

Modeling and analysis of mechanisms underlying high-resolution functional MRI of cortical columns

Dissertation

Zur Erlangung des Grades eines
Doktors der Naturwissenschaften

der Mathematisch-Naturwissenschaftlichen Fakultät
und
der Medizinischen Fakultät
der Eberhard-Karls-Universität Tübingen

vorgelegt
von

Denis Chaimow
aus Bernau bei Berlin, Deutschland

im Dezember 2016

Tag der mündlichen Prüfung: 17. Mai 2017

Dekan der Math.-Nat. Fakultät: Prof. Dr. W. Rosenstiel

Dekan der Medizinischen Fakultät: Prof. Dr. I. B. Autenrieth

1. Berichterstatter: Dr. Amir Shmuel

2. Berichterstatter: Prof. Dr. Klaus Scheffler

Prüfungskommission: Prof. Dr. Andreas Bartels

Dr. Amir Shmuel

Prof. Dr. Klaus Scheffler

Prof. Dr. Martin Giese

Erklärung / Declaration:

Ich erkläre, dass ich die zur Promotion eingereichte Arbeit mit dem Titel:

„Modeling and analysis of mechanisms underlying high-resolution functional MRI of cortical columns“

selbständig verfasst, nur die angegebenen Quellen und Hilfsmittel benutzt und wörtlich oder inhaltlich übernommene Stellen als solche gekennzeichnet habe. Ich versichere an Eides statt, dass diese Angaben wahr sind und dass ich nichts verschwiegen habe. Mir ist bekannt, dass die falsche Abgabe einer Versicherung an Eides statt mit Freiheitsstrafe bis zu drei Jahren oder mit Geldstrafe bestraft wird.

I hereby declare that I have produced the work entitled “.....”, submitted for the award of a doctorate, on my own (without external help), have used only the sources and aids indicated and have marked passages included from other works, whether verbatim or in content, as such. I swear upon oath that these statements are true and that I have not concealed anything. I am aware that making a false declaration under oath is punishable by a term of imprisonment of up to three years or by a fine.

Berlin, den 20. Dezember 2016

Denis Chaimow

A C K N O W L E D G E M E N T S

First and foremost I thank my advisor Amir Shmuel. He supported and mentored me through many years despite that we found ourselves working at different institutions and on different continents. Without his guidance and feedback this thesis would not have been achievable.

Furthermore I thank Kamil Uğurbil for his advice and support. I worked in his research groups at the Center for Magnetic Resonance Research (CMRR) in Minneapolis at the University of Minnesota and at the Max Planck Institute (MPI) for Biological Cybernetics in Tübingen. Both institutions supported me financially. The people at CMRR and the MPI provided stimulating discussions, whether science related or otherwise. In particular I thank my collaborators Essa Yacoub and Ute Goerke.

I also thank the Graduate Training Center of Neuroscience in particular Horst Herbert, Tina Lampe and Katja Thielges.

Last but not least I thank my family and friends. My parents supported me morally and financially through my studies and beyond. It would not have been possible without them. I thank my wife Claire for her patience, love and assistance, and my daughter Nora for being a source of joy in the final phase of this work.

C O N T E N T S

Summary	9
Synopsis	
Modeling and analysis of mechanisms underlying high-resolution functional MRI of cortical columns	13
Study 1	
Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns	37
Study 2	
Spatial specificity of the functional MRI blood oxygenation response relative to neuronal activity	55
Study 3	
A more accurate account of the effect of k-space sampling and signal decay on the effective spatial resolution in functional MRI	105
Study 4	
Optimization of functional MRI for detection, decoding and imaging the response patterns of cortical columns	141

S U M M A R Y

High spatial resolution functional MRI (fMRI) and advanced multivariate analysis techniques are promising tools for studying the cortical basis of human cognitive processes at the level of columns and layers. However the true spatial specificity of high-resolution fMRI has not been quantified, and the basis for decoding from fine scale structures using large voxels and relatively low magnetic field strength is unknown. It is also not yet known what method and voxel size is optimal for decoding and what voxel size is optimal for high-resolution imaging. In this thesis we present four studies that answer part of these questions using a model-based approach of imaging cortical columns.

We started our investigation of model-based analysis of high-resolution fMRI of cortical columns by addressing the specific problem of how it is possible to decode information thought to be mediated by cortical columns using large voxels at low field strength. 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 a 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.

Having developed a model of imaging ODCs we then used this model to estimate the spatial specificity of BOLD fMRI, specifically at high field (7 T). Previous attempts at characterizing the spatial specificity of the blood oxygenation level dependent functional MRI (BOLD fMRI) response by estimating its point-spread function (PSF) have conventionally relied on spatial representations of visual stimuli in area V1. Consequently, their estimates were confounded by the width and scatter of receptive fields of V1 neurons. Here, we circumvent these limits by instead using the inherent cortical spatial organization of ocular dominance columns (ODCs) to determine the PSF for both Gradient Echo (GE) and Spin Echo (SE) BOLD imaging at 7 Tesla. By applying Markov Chain Monte Carlo sampling on a probabilistic generative model of imaging ODCs, we quantified the PSFs that best predict the spatial structure and magnitude of differential ODCs' responses. Prior distributions for the ODC model parameters were determined by analyzing published data of cytochrome oxidase patterns from post-mortem histology of human V1 and of neurophysiological ocular dominance indices. The most probable PSF full-widths at half-maximum were 0.82 mm (SE) and 1.02 mm (GE). Our results provide a quantitative basis for the spatial specificity of BOLD fMRI at ultra-high fields, which can be used for planning and interpretation of high-resolution differential fMRI of fine-scale cortical organizations.

Our BOLD fMRI PSF findings show that the PSF is considerably smaller than what was reported previously. This in turn raised the question of the role of the imaging PSF, which now has become relevant. Next we show that the commonly used magnitude point-spread function fails to accurately represent the true effects of k-space sampling and signal decay, and propose an alternative model that accounts more accurately for these effects. The effects of k-space sampling and signal decay on the effective spatial resolution of MRI and functional MRI (fMRI) are commonly assessed by means of the magnitude point-spread function (PSF), defined as the absolute values (magnitudes) of the complex MR imaging PSF. It is commonly assumed that this magnitude PSF signifies blurring, which can be quantified by its full-width at half-maximum (FWHM). Here we show that the magnitude PSF fails to accurately represent the true effects of k-space sampling and signal decay. Firstly, a substantial part of the width of the magnitude PSF is due to MRI sampling per se. This part is independent of any signal decay and its effect depends on the spatial frequency composition of the imaged object. Therefore, it cannot always be expected to introduce blurring. Secondly, MRI reconstruction is typically followed by taking the absolute values (magnitude image) of the reconstructed complex image. This introduces a non-linear stage into the process of image formation. The *complex* imaging PSF does not fully describe this process, since it does not reflect the stage of taking the magnitude image. Its corresponding *magnitude* PSF fails to correctly describe this process, since convolving the original pattern with the magnitude PSF is different from the true process of taking the absolute following a convolution with the complex imaging PSF. Lastly, signal decay can have not only a blurring, but also a high-pass filtering effect. This

cannot be reflected by the strictly positive width of the magnitude PSF. As an alternative, we propose to model the imaging process by decomposing it into a signal decay-independent MR sampling part and an approximation of the signal decay effect. We approximate the latter as a convolution with a Gaussian PSF or, if the effect is that of high-pass filtering, as reversing the effect of a convolution with a Gaussian PSF. We show that for typical high-resolution fMRI at 7 Tesla, signal decay in Spin-Echo has a moderate blurring effect (FWHM = 0.89 voxels, corresponds to 0.44 mm for 0.5 mm wide voxels). In contrast, Gradient-Echo acts as a moderate high-pass filter that can be interpreted as reversing a Gaussian blurring with FWHM = 0.59 voxels (0.30 mm for 0.5 mm wide voxels). Our improved approximations and findings hold not only for Gradient-Echo and Spin-Echo fMRI but also for GRASE and VASO fMRI. Our findings support the correct planning, interpretation, and modeling of high-resolution fMRI.

In our first study we used our model to analyze imaging of cortical columns under a very specific scenario. We studied a best case scenario for decoding the stimulated eye from ODCs imaged at 3T using large voxels. In order to do so, we formalized available knowledge about fMRI of cortical columns. In particular, the ability of fMRI to resolve cortical columnar organization depends on several interdependent factors, e.g. the spatial scale of the columnar pattern, the point-spread of the BOLD response, voxel size and the signal-to-noise ratio. In our fourth study we aim to analyze how these factors contribute and combine in imaging of arbitrary cortical columnar patterns at varying field strengths and voxel sizes. In addition, we compared different pattern imaging approaches. We show how detection, decoding and reconstruction of a fine scale organization depend on the parameters of the model, and we predict optimal voxel sizes for each approach under various scenario.

The capacity of fMRI to resolve cortical columnar organizations depends on several factors, e.g. the spatial scale of the columnar pattern, the point-spread of the fMRI response, the voxel size, and the SNR considering thermal and physiological noise. How these factors combine, and what is the voxel size that optimizes fMRI of cortical columns remain unknown. Here we combine current knowledge into a quantitative model of fMRI of patterns of cortical columns. We compare different approaches for imaging patterns of cortical columns, including univariate and multivariate based detection, multi-voxel pattern analysis (MVPA) based decoding, and reconstruction of the pattern of cortical columns. We present the dependence of their performance on the parameters of the imaged pattern and the data acquisition, and predict voxel sizes that optimize fMRI under various scenarios. To this end, we modeled differential imaging of realistic patterns of cortical columns with different spatial scales and degrees of irregularity. We quantified the capacity to detect and decode stimulus-specific responses by analyzing the distribution of voxel-wise differential responses relative to noise. We quantified the accuracy with which the spatial pattern of cortical columns can be reconstructed as the correlation between the underlying columnar pattern and the imaged pattern. For regular patterns, optimal voxel widths for detection, decoding and reconstruction were close to half the main cycle length of the organization. Optimal voxel widths for irregu-

lar patterns were less dependent on the main cycle length, and differed between univariate detection, multivariate detection and decoding, and reconstruction. We compared the effects of different factors of Gradient Echo fMRI at 3 Tesla (T), Gradient Echo fMRI at 7T and Spin-Echo fMRI at 7T, and found that for all measures (detection, decoding, and reconstruction), the width of the fMRI point-spread has the most significant effect. In contrast, different response amplitudes and noise characteristics played a comparatively minor role. We recommend specific voxel widths for optimal univariate detection, for multivariate detection and decoding, and for reconstruction under these three data-acquisition scenarios. Our study supports the planning, optimization, and interpretation of fMRI of cortical columns and the decoding of information conveyed by these columns.

SYNOPSIS

Modeling and analysis of mechanisms
underlying high-resolution functional
MRI of cortical columns

Introduction

The cerebral cortex and human cognition

How does our brain allow us to perceive the world? How does it enable us to remember, to plan actions and to communicate? It is widely believed that the cerebral cortex plays a major role in our cognitive abilities. The cerebral cortex is a thin, 3 mm thick sheet, extending approximately 50 cm by 50 cm. It consists of individual areas, that can be distinguished based on cytoarchitecture, patterns of neuronal connections and involvement in cognitive tasks.

The neo-cortex comprises six layers that differ in type and density of neurons. Neurons across these six layers form so called microcircuits or microcolumns, a basic reoccurring computational structure that might form the basis of cortical processing (Bastos et al., 2012; Douglas and Martin, 2004). Connections between areas originate and terminate in distinct layers. These connectivity patterns provide a basis for distinguishing two types of connections. They, respectively, give rise to feed-forward and feed-back flow of information relative to the input that the brain receives from the sensory organs. These different types of connections impose a hierarchy of cortical areas, suggesting a distributed and hierarchical processing of information (Felleman and van Essen, 1991).

Neurons in the same microcolumn tend to share the same stimulus preferences in their responses. For example, neurons in a microcolumn in primary visual cortex (V1) will all respond preferentially to bars of the same orientation if presented at a certain position in the visual field through their preferred eye of origin. Traversing a distance of approximately 0.5 mm to 1 mm along the cortical sheet in a single direction, one will find neurons with a different preference with respect to some stimulus dimension. The results of this organization are columnar maps in which neurons form semi-periodic clusters, between 0.5 mm and 1 mm wide, in which they share the preference for some stimulus attribute. As neurons are generally sensitive to multiple stimulus attributes, several such maps may be present within an area. Prominent examples are maps of orientation columns and ocular dominance columns (ODCs) in V1.

To summarize, neurons in a cortical column respond similarly to the same stimulus while neurons in the same cortical layer receive input from similar origins within the local microcircuit and from other areas along the cortical hierarchy. This organization makes the spatial scale of columns and layers suitable to provide information about how the cortex processes information locally and between areas.

Studying the neuronal basis of cognition in humans

Most of what is known about the structure and physiology of the cerebral cortex has been the result of invasive physiological and anatomical studies in animals. Most notably, macaque monkeys have served as a model organism for cortical visual processing in humans. However, for a number of reasons it is desirable to study such processes in humans directly. For example, animals need to be trained, higher cognitive tasks like language can only be studied in

humans and it is not guaranteed that results in animals will apply completely to humans.

However, for ethical reasons, research methods for human subjects are limited compared to animals. Available methods to study the human brain include the analysis of brain injuries, purely behavioral experiments, electroencephalography (EEG), magnetoencephalography (MEG), transcranial magnetic stimulation (TMS) and positron emission tomography (PET).

Brain injuries provide very crude spatial and no temporal resolution. Behavioral experiments do not allow us to observe cortical processes directly. EEG and MEG provide excellent temporal resolution but it is very difficult to infer the spatial origin of observed responses. TMS allows us to manipulate responses in the brain, but similar to EEG and MEG these responses cannot be localized precisely. Finally, positron emission tomography measures metabolic activity of the brain using radioactive tracers. In addition to its invasive nature, its temporal resolution is low and its spatial resolution is not sufficiently high to resolve layers and columns. None of these techniques allow us to study the human brain non-invasively at the resolution of columns and layers.

Functional magnetic resonance imaging

In 1990, Ogawa et al. (1990) showed that magnetic resonance imaging (MRI) can be sensitive to the oxygenation level of blood. This led to the development of blood-oxygenation-level dependent (BOLD) fMRI (Bandettini et al., 1992; Kwong et al., 1992; Ogawa et al., 1992). BOLD fMRI soon became arguably the most important tool of cognitive neuroscience.

BOLD fMRI is based on the fact that following increases in neuronal activity local arterial cerebral blood flow (CBF) increases are relatively larger than the increases in oxygen consumption. This results in higher levels of oxygenated hemoglobin and consequently reduced levels of deoxygenated hemoglobin (deoxyhemoglobin). Deoxyhemoglobin is paramagnetic, and thus it interferes with the homogeneity of the magnetic field in and around blood vessels (intra- and extravascular effects), causing the spins of diffusing hydrogen atoms to dephase. The end result of reduced deoxyhemoglobin levels due to increased neuronal activity are signal increases in MRI sequences that are sensitive to dephasing-dependent changes in transversal relaxation rates.

The most commonly used BOLD fMRI pulse sequence and contrast is gradient echo (GE). GE BOLD is sensitive to the intra- and extravascular effects of activation-induced changes in the deoxyhemoglobin content of blood. At standard magnetic field strengths (1.5 T, 3 T) the signal is dominated by contributions from larger blood vessels (Jochimsen et al., 2004; Uludağ et al., 2009). At higher magnetic field strengths, the strong intravascular component of these large blood vessels decreases, while the extravascular signal changes around capillaries and smaller vessels increase (Uludağ et al., 2009; Yacoub et al., 2001). Additional weighting towards the microvasculature can be achieved by using spin echo (SE) BOLD imaging, which suppresses extravascular signal contributions from larger blood vessels (Uludağ et al., 2009; Yacoub et al., 2003).

High spatial resolution fMRI

The majority of fMRI studies were done using GE BOLD fMRI at 1.5 T and 3 T at relatively low spatial resolution and were concerned with identifying which cortical areas were involved in which cognitive tasks. However, a number of studies used fMRI at higher magnetic field strengths with higher spatial resolution to successfully image columnar structures in the human cortex

(Cheng et al., 2001; Goodyear and Menon, 2001; Yacoub et al., 2008; 2007; Zimmermann et al., 2011). More recently, also the possibility to image layer dependent activity was demonstrated (Chaimow et al., 2012; 2011; De Martino et al., 2015; Fracasso et al., 2016; Muckli et al., 2015).

However, high-resolution studies are methodologically challenging. In addition to requiring high magnetic fields they place high demands on the gradient system in order to achieve a high nominal resolution. Furthermore they can suffer from relatively low SNR and motion artifacts. Finally the attainable spatial resolution is ultimately limited by the spatial specificity of the BOLD response.

The first study to quantify the spatial specificity of the BOLD response (Engel et al., 1997) used an elegant phase-encoding paradigm that induced traveling waves of retinotopic neural activity in the primary visual cortex (V1). Assuming a shift-invariant linear response, Engel et al. (1997) estimated the point-spread function (PSF), which represents the spatial response that would be elicited by a small point stimulus. They found the full-width at half-maximum (FWHM) of the GE BOLD PSF to be 3.5 mm at 1.5 T. Similar values (3.9 mm for GE BOLD and 3.4 mm for SE BOLD) have been reported at 3 T (Parkes et al., 2005) using a paradigm similar to that used in Engel et al. (1997). To estimate the GE BOLD PSF at 7 T, we previously measured the spatiotemporal spread of the fMRI response in grey matter regions around the V1 representation of edges of visual stimuli (Shmuel et al., 2007). To reduce contributions from macroscopic veins, we excluded voxels that showed vessel-like response features. The mean measured and estimated FWHMs were 2.34 ± 0.20 mm and < 2 mm, respectively. The spatial specificity of SE BOLD fMRI at ultra-high magnetic fields has not yet been quantified.

All previous studies of the spatial specificity of the BOLD response had relied on spatial representations of visual stimuli in V1. Because of the width and scatter of V1 receptive fields, such estimates are confounded. They reflect the combined spread of neural and BOLD response. Therefore, they only measure the capacity of the BOLD response to resolve retinotopic representations but not its ability to resolve more fine-grained neural activity.

Multivariate pattern analysis - decoding

A different approach towards probing fine structures relies on the application of multivariate pattern analysis (MVPA) methods, rather than increasing the spatial resolution of fMRI. The traditional approach to fMRI analysis can be summarized as massive univariate statistical analysis using a general linear model (GLM). The time-course of every voxel is separately analyzed for activation caused by experimental conditions.

MVPA considers the activity of multiple voxels as multivariate data points. Furthermore MVPA most often uses a decoding approach in which a machine learning algorithm is trained to predict the experimental condition from the measured activity. Successful prediction (decoding) of the experimental condition from an “unseen” subset of the data, that was not available to the algorithm for training, is then taken as evidence for the presence of condition specific information in the measured activity.

A number of studies (e.g. Haxby et al., 2001; Haynes and Rees, 2005; Kamitani and Tong, 2005) have demonstrated that multivariate machine learning algorithms can decode visual stimuli from fMRI data. Using GE BOLD fMRI data obtained at 3 T, these algorithms decoded information thought to be mediated by cortical columns. This result seemed 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 ODCs and orientation columns in humans). This result was even more surprising considering the relatively wide PSF of GE BOLD fMRI signals at 3T (Engel et al., 1997; Parkes et al., 2005).

Important questions

High spatial resolution BOLD fMRI and advanced multivariate analysis techniques promise to provide a window into activity of the human cerebral cortex at the scale of columns and layers. The non-invasiveness and applicability of these methods to the living human brain make them well suited for studying the cortical basis of human cognitive processes.

However the true spatial specificity of high-resolution BOLD fMRI has not been quantified. Furthermore the basis for decoding from fine scale structure using large voxels and relatively low magnetic field strength is not known. It is also not known what method and voxel size is optimal for decoding and what voxel size is optimal for high-resolution imaging. Finally it is not clear how to interpret imaging results from unknown cortical-columns-like patterns of activity.

To optimally plan high-resolution BOLD fMRI studies and to correctly interpret their results, it is necessary to know the inherent limits of spatial specificity of BOLD fMRI relative to the sites where changes in neuronal activity occur. It is also important to quantitatively understand how factors such as spatial specificity, voxel size, noise and the structure of interest contribute and interact in the process of imaging or decoding.

In this thesis we present four studies that answer part of these question using a model-based approach of imaging cortical columns.

Study 1: Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns

We started our investigation of model-based analysis of fMRI from cortical columns by addressing the specific problem of how it is possible to decode information thought to be mediated by cortical columns using large voxels acquired at low field strength.

Several mechanisms had 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.

Out of these five mechanisms proposed to underlie decoding of stimulus orientation or the stimulated eye we sought to assess the plausibility of two of the suggested mechanisms: (I) aliasing and (II) local irregularities, by creating a model of imaging ODC patterns. The aim was to evaluate quantitatively what levels of functional contrast and classification performance 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, the process of voxel sampling, and noise.

The model that we developed served as a basis for this and the remaining three studies of this thesis (Study 1, Fig. 1). The first component of the model, i.e. the modeling of realistic patterns of ODCs, followed Rojer and Schwartz (1990). It consisted of band-pass filtering of spatial white noise followed by applying a sigmoidal point-wise non-linearity, which controlled the smoothness of transitions between left and right eye preference regions. The spatial BOLD response was modeled as a convolution of the ODC pattern with a Gaussian PSF. MRI k-space sampling was modeled by restricting the spatial frequency space to its central part in accordance with the modeled field of view, sampling matrix and voxel size (Haacke et al., 1999).

The modeled voxel differential responses followed 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 differential contrast values that exist in the specific distribution. We defined the term *contrast range* as this standard deviation of the distribution of differential responses. Contrast range, the number of voxels and the level of noise allowed us to calculate expected classification performance.

First we studied how contrast range and classification performance depended on the voxel width and BOLD point spread width. The BOLD PSF and the voxel width acted as low-pass filters, reducing information conveyed by high frequencies, and were found to be important factors in determining the functional contrast and classification performance. Nonetheless we found that for large point spreads, the voxel width had almost no effect on functional contrast. In contrast, for large voxel widths, increasing BOLD PSF still decreased the functional contrast. 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 low frequencies.

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. Not only are these high-frequency components 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. Increasing these irregularities improves classification performance.

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 low in amplitude though, insufficient for accounting for classification performance reported at 3 Tesla. We expect that lower frequencies, 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 possibly complex spatio-temporal filtering of neuronal activity play significant roles in decoding.

Study 2: Spatial specificity of the functional MRI blood oxygenation response relative to neuronal activity

Having developed a model of imaging ODCs, we then used this model to provide a better answer to the question of spatial specificity of BOLD fMRI, specifically at high field (7 T).

All previous attempts at characterizing the spatial specificity of the BOLD fMRI response (Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007) relied on an implicit assumption that neuronal responses to small visual stimuli are point-like. To estimate the spatial specificity of the BOLD response, these studies have conventionally relied on spatial representations of visual stimuli in area V1. Unlike the implicit assumption of point-like responses, the receptive fields of neurons in V1 have non-zero spatial extents (Hubel and Wiesel, 1968). In addition, electrode measurements in macaque V1, oriented orthogonally relative to the surface of cortex, have demonstrated substantial scatter in the center of receptive fields (Hubel and Wiesel, 1974). Therefore, the pattern of neural activity parallel to the cortical surface is a blurred representation of

the visual stimulus. This implies that receptive field size and scatter pose a lower limit on any BOLD fMRI PSF width that is estimated using spatial representations of visual stimuli in V1. Consequently, the previously computed estimates of the spatial specificity of the fMRI response were confounded by the width and scatter of receptive fields of V1 neurons. Such estimates are limited in that they solely measure the capacity of the BOLD response to resolve retinotopic representations; they do not measure its ability to resolve more fine-grained neural activity. Yet only this latter resolvability matters for functional imaging at the spatial scale of cortical columns.

Our model allowed us to circumvent the limits posed by the retinotopic representation of visual stimuli by instead using the inherent cortical spatial organization in the form of ODCs. We extended our model to a probabilistic model in order to quantify the widths of the PSF that best predict the spatial structure and magnitude of ODC responses acquired at 7 T using GE and SE BOLD imaging.

The key idea is to treat the Gaussian white noise, that was spatially filtered in order to model realistic columnar patterns, as a high dimensional parameter with a multivariate normal prior probability distribution. By applying Markov Chain Monte Carlo sampling on our probabilistic generative model of imaging ODCs, we quantified the PSFs that best predict the spatial structure and magnitude of differential ODCs' responses. We determined prior distributions for the ODC model parameters by analyzing published data of cytochrome oxidase patterns from post-mortem histology of human V1 (Adams et al., 2007) and of neurophysiological ocular dominance indices (Berens et al., 2008; Hubel and Wiesel, 1968). We then reanalyzed fMRI data (Yacoub et al., 2007) from human ODCs. We estimated responses to left and right eye stimulations. Their voxel-wise difference resulted in a differential ODC map while the distribution of their average response and of their estimation standard error made it possible to constrain the overall BOLD amplitude and the measurement noise of the model.

Using the data (differential responses), the probabilistic model and the parameter priors, we applied Markov chain Monte Carlo (MCMC) sampling in order to sample from the posterior distribution of FWHM parameters. The most probable PSF FWHMs averaged over subjects were 0.82 mm (SE) and 1.02 mm (GE).

Across all modeled ODC patterns, the GE PSF was almost always wider than the SE PSF. We calculated the resulting posterior distribution of differences between GE and SE point-spread widths. Averaged over subjects, the most probable difference was 0.25 mm.

In addition, we used our 7 T GE result and how it related to previous results confounded by receptive fields to make a rough estimate of what the non-confounded PSF widths at lower fields would be. The width of the BOLD PSF has been estimated to be 3.5 mm for 1.5 T GE BOLD (Engel et al., 1997), 3.9 mm for 3 T GE BOLD and 3.4 mm for 3 T SE BOLD (Parkes et al., 2005). Assuming that on average the receptive field effect can be modeled as another convolution with a Gaussian, we estimated a receptive field effect with an

FWHM of 1.72 mm and non-confounded PSF estimates of 3 mm (1.5 T GE BOLD), 2.8 mm (3 T GE BOLD) and 3.5 mm (3 T SE BOLD).

Our results facilitate planning and interpretation of high-resolution fMRI studies of fine scale cortical organizations.

Study 3: A more accurate account of the effect of k-space sampling and signal decay on the effective spatial resolution in functional MRI

Our BOLD fMRI PSF findings (Study 2) showed that the PSF is considerably smaller than what was reported previously. This in turn raised the question of the role of the *imaging* PSF, which has now become relevant. The imaging PSF describes the effect of the MR imaging process, specifically discrete k-space sampling and T_2/T_2^* signal decay during sampling, on the effective spatial resolution. The imaging PSF contributes to the overall BOLD fMRI PSF.

Given the relatively narrow widths of the overall BOLD fMRI PSF that we found (0.82 mm for 7 T SE and 1.02 mm for 7T GE), results from previous imaging PSF studies (e.g. FWHM of ~ 1 mm or larger for partial Fourier 7T SE using 0.5 mm voxels; Kemper et al., 2015) would suggest that the contribution of the imaging PSF to the overall BOLD fMRI PSF could be quite substantial. Paradoxically, its width might be even larger than that of the overall BOLD fMRI PSF. This necessitates a careful analysis of how the imaging PSF had been modeled.

The imaging PSF is commonly studied in form of the magnitude point-spread function, defined as the absolute values (magnitudes) of the complex MR imaging PSF. It is commonly assumed that this magnitude PSF signifies blurring, which can be quantified by its full-width at half-maximum (FWHM). However, in Study 3 we show that the magnitude PSF fails to accurately represent the true effects of k-space sampling and signal decay.

Firstly, a substantial part of the width of the magnitude PSF is due to MRI sampling per se. This part is independent of any signal decay and its effect depends on the spatial frequency composition of the imaged object. Therefore, it cannot always be expected to introduce blurring. Secondly, MRI reconstruction is typically followed by taking the absolute values (magnitude image) of the reconstructed complex image. This introduces a non-linear stage into the process of image formation. The complex imaging PSF does not fully describe this process, since it does not reflect the stage of taking the magnitude image. Its corresponding magnitude PSF fails to correctly describe this process, since convolving the original pattern with the magnitude PSF is different from the true process of taking the absolute following a convolution with the complex imaging PSF. Lastly, signal decay can have not only a blurring, but also a high-pass filtering effect. This cannot be reflected by the width of the magnitude PSF.

We simulated T_2/T_2^* decay of GE and SE EPI at 7T and MR sampling of simulated realistic columnar patterns. We found that for MR imaging with no signal decay, the imaged pattern was very similar to the original pattern. How-

ever, the FWHM of the magnitude PSF corresponding to MRI with no signal decay was 0.6 mm. GE imaging had the effect of a moderate high-pass filter. Also this result could not be expected by simply considering the FWHM of the magnitude PSF, which was 0.67 mm. The SE image was slightly blurred. However, convolutions of the original pattern with Gaussian PSFs of the same widths as those computed from the SE magnitude PSFs and convolutions with the SE magnitude PSF itself resulted in stronger blurring.

In order to characterize the complete non-linear MRI process using a PSF, we propose an alternative method. First we approximated the MRI process linearly. For small deviations from a spatially constant baseline pattern, the resulting linear approximation is identical to a convolution with the real component of the complex imaging PSF.

In a second step, we fitted Gaussians to the MTF of the real component of the imaging PSF (and to its inverse). This allowed us to approximate the MRI process as a convolution with a Gaussian PSF or, if the effect is that of high-pass filtering, as reversing the effect of a convolution with a Gaussian PSF, followed by MR sampling with no decay. We found that signal decay in SE has a moderate blurring effect (FWHM = 0.89 voxels, corresponds to 0.44 mm for 0.5 mm wide voxels). In contrast, GE acts as a moderate high-pass filter that can be interpreted as reversing a Gaussian blurring with FWHM = 0.59 voxels (0.30 mm for 0.5 mm wide voxels).

The magnitude PSF describes the absolute level of influence that neighboring positions in the original pattern have on each other's value in the image. However, it fails to characterize the nature of this influence (e.g. blurring or high-pass filtering). While our approach is an approximation, it results in an intuitive characterization of the effect of signal decay and makes it possible to compare it to previously reported FWHMs of PSFs associated with the entire BOLD fMRI process (BOLD PSF).

Study 4: Optimization of functional MRI for detection, decoding and imaging the response patterns of cortical columns

In Study 1 we used our model to analyze imaging of cortical columns under a very specific scenario. We studied a best case scenario for decoding the stimulated eye from ODCs imaged at 3T using large voxels. In order to do so, we formalized the currently available knowledge about fMRI of cortical columns. In particular, the ability of fMRI to resolve cortical columnar organization depends on several interdependent factors, e.g. the spatial scale of the columnar pattern, the point-spread of the BOLD response, voxel size and the signal-to-noise ratio.

In Study 4 we aimed to analyze how these factors contribute and combine in imaging of arbitrary cortical columnar patterns at varying field strengths and voxel sizes. We compared different pattern imaging approaches and show how detection, decoding and reconstruction of a fine scale organization depend on the parameters of the model and predict optimal voxel sizes for each approach under various scenarios.

Extending our model employed in study 1, we modeled generic isotropic columnar patterns of varying spatial scale and degree of irregularity. We demonstrate how contrast-to-noise ratio (CNR) can be determined from the spatial frequency spectrum of the columnar pattern, the BOLD point-spread function, the voxel width, and noise.

CNR and the number of voxels were used to estimate expected decoding accuracy for a linear classifier and the probability to detect a condition-specific response in single voxels (univariate) and to detect a pattern response (multivariate). The probability to achieve statistically significant decoding and the probability to detect a pattern response were virtually identical. The accuracy of reconstructing fine scale organizations was quantified as the correlation between the underlying columnar patterns and the imaged pattern.

We simulated univariate and multivariate detection probability and pattern correlation as a function of voxel size and show how optimal voxel size arises as a trade off between functional contrast and noise.

We further estimated how optimal voxel width changes as a function of the spatial organization of the pattern and as a function of the BOLD point-spread function. For regular patterns, optimal voxel widths for univariate detection, multivariate detection and decoding, and reconstruction were close to half the main cycle length of the organization. Optimal voxel widths for irregular patterns were less dependent on the main cycle length, and differed between univariate detection, multivariate detection and decoding, and reconstruction.

We compared the effects of different factors of GE fMRI at 3 T, GE fMRI at 7T and SE fMRI at 7T, and found that for all measures (detection, decoding, and reconstruction), the width of the fMRI point-spread has the most significant effect. In contrast, different response amplitudes and noise characteristics played a comparatively minor role.

Finally, we recommend specific voxel widths for optimal univariate detection, for multivariate detection and decoding, and for reconstruction under these three data-acquisition scenarios.

Implications for columnar imaging

Resolving cortical columns using 7T BOLD fMRI

In Study 2 we quantified the BOLD PSF in human subjects relative to neuronal activity, avoiding the confounding effects of scatter and size of visual receptive fields which were not eliminated in previous estimations (Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007). The most probable PSF FWHMs were 0.82 mm (SE) and 1.02 mm (GE). Our PSF widths provide a quantitative basis for resolving cortical columns and allow to model the spatial spread of the BOLD response.

Table 2 of Study 4 shows that columnar patterns with cycle lengths of 1.4 mm or larger are expected to be well imaged at 7 T using GE or SE BOLD. Assuming differential analysis, their imaged patterns are expected to have correlations of 0.7071 or higher relative to the original pattern. This minimal cycle length of 1.4 mm corresponds in differential analysis to column widths of 0.7

mm, which is narrower than human ODCs (~ 0.9 mm; Adams et al., 2007) and orientation columns (~ 0.8 mm; Yacoub et al., 2008).

The imaging PSF, which accounts for k-space sampling effects (discrete finite sampling of k-space and signal decay during sampling) is part of the overall BOLD PSF. We have shown in study 3 that the effect of signal decay is relatively small compared to the overall BOLD PSF. Its effect is proportional to the voxel width and therefore its contribution to the overall BOLD PSF decreases with decreasing voxel width. The effect of k-space sampling per se is not expected to cause blurring if the voxel width is sufficiently small for sampling the larger part of the original spatial frequencies.

GE vs. SE BOLD fMRI

In Study 2 we found the SE BOLD PSF (FWHM = 0.82 mm) to be narrower than the GE BOLD PSF (FWHM = 1.02 mm). This is expected because the re-focusing pulse in SE imaging suppresses the extravascular signal around larger blood vessels while leaving the signal around the microvasculature intact (Uludağ et al., 2009; Yacoub et al., 2003). However, we found that the difference between the PSF widths of GE and SE was relatively small. We believe that this is due to the fact that the influence of larger blood vessels was reduced by using a differential imaging paradigm, even when using 7T GE BOLD fMRI.

Consequently, both GE and SE BOLD imaging techniques seem capable of resolving cortical columns when applying differential imaging analysis. However, GE maps are more susceptible to confounds introduced in voxels containing blood vessels which may not be fully suppressed in differential imaging. Therefore, obtaining results of high spatial specificity using GE depends on the region of interest and on methods to mask out blood vessels or reduce their effect.

SE is less susceptible to large-vessel confounds, that may not be suppressed by differential imaging. The response amplitude of SE is lower than that of GE. However, for imaging of highly granular structures such as ODC's at such high resolutions, the *differential* contrast is similar for GE and SE fMRI. Overall, we believe that SE or the related GRASE pulse sequence (Feinberg et al., 2008; Oshio and Feinberg, 1991) are the methods of choice for mapping fine structures, especially when relying on single-condition analysis. However, which data acquisition method is optimal depends on the goal of the study and the spatial scale of the neuronal architecture under investigation.

Probing columns at lower fields by means of multivariate pattern analysis

Small PSF widths are necessary to reconstruct true fine-scale organizations. Yet, such small PSF widths can only be expected at high field. At standard magnetic field strengths, we inferred the width of the BOLD PSF with no receptive field confounds to be 2.8 mm (1.5 T GE BOLD), 3.3 mm (3 T GE BOLD) and 2.7 mm (3 T SE BOLD). These PSF widths are considerably larger than our estimates for 7T GE BOLD and 7T SE BOLD. The reason for this is that the BOLD signal (both GE and SE BOLD) at lower field strengths is domi-

nated by intravascular signals from draining veins (Jochimsen et al., 2004; Uludağ et al., 2009). At higher field strengths, the contributions from intravascular signals are reduced due to a shortening of the venous blood T2. In parallel, the relative contributions of extravascular signals around small vessels increase (Duong et al., 2003; Uludağ et al., 2009; Yacoub et al., 2003; 2001).

In Study 4 we show that such wide PSFs as obtained for magnetic field strengths of 3T, make it impossible to resolve fine-scale organizations at the scale of cortical columns. But, as our study shows, there is a difference between resolving columnar patterns, thus reconstructing their true organization, and detection of stimulus specific information encoded in their responses. Making this distinction explicit allowed us to estimate optimal imaging parameters for each of the approaches, and to demonstrate that optimal parameters are not necessarily identical between approaches.

In particular, reconstruction generally requires smaller voxel size than detection. This is especially true for irregular patterns, where detection can benefit from low-frequency components using relatively large voxels with high SNR, whereas accurate reconstruction requires high spatial-frequency content to be captured, which can only be achieved with smaller voxels. As a result and in line with our Study 1, relatively large voxels at lower field strength may pick up information present in lower spatial frequencies due to the inherent irregularities of columnar patterns. However, Studies 1 and 4 also show that these low-spatial-frequency contributions are expected to be relatively small. The fact that a number of studies decoded information thought to be mediated by cortical columns with more precision than our modeling predicted suggests that additional factors not considered by our model may be at play. Such factors include: (1) underestimated contributions of irregularities, (2) large-scale bias, (3) biased draining regions, and possibly (4) spatiotemporal complex kernels.

(1) The role of irregularities

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 (Blasdel et al., 1995; Rojer and Schwartz, 1990) that the spatial frequency spectra of real ODC columns correspond to a more heavy-tailed filter function than the Gaussian filter. In other words, ODC organizations are expected to include higher contributions of relatively lower 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 lower spatial frequencies, and would imply improved classification performance.

Experimental evidence for significant, local rather than global, contributions of lower 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 (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 (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 millimeter scale (Swisher et al., 2010).

(2) Large scale bias

In addition to moderately low spatial frequencies reflecting irregularities of the columnar pattern, very low spatial frequencies, reflecting large-scale components of the organization were proposed to play a role (Op de Beeck, 2010). These include the oblique and radial effects (Furmanski and Engel, 2000; Sasaki et al., 2006) 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). 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. Freeman et al. (2011; 2013) confirmed the existence of a coarse scale topographic organization of orientation preference in human V1 and found that this organization suffices to explain successful decoding using larger voxels at low field. Beckett et al. (2012) found a large scale retinotopy dependent organization for direction of motion and showed that it is sufficient to allow decoding.

While there seems to be ample evidence that large scale organizations contribute to decoding, it is not clear whether they provide a sufficient explanation. Furthermore their exact nature and spatial scales are still being debated (Alink et al., 2013; Carlson, 2014; Carlson and Wardle, 2015; Maloney, 2015; Mannion et al., 2009; Pratte et al., 2016; Wang et al., 2014).

(3) Biased draining regions

Alternatively, draining regions that cover cortical maps and columns non-homogeneously may cause selective responses of their corresponding blood vessels (Gardner, 2010; Gardner et al., 2006; Kamitani and Tong, 2006; 2005; Kriegeskorte and Bandettini, 2007; Shmuel, 2010). In this scenario 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. (Gardner et al., 2006), Shmuel et al. (Shmuel et al., 2010) and Thompson et al. (2011). However, while random biases in draining are likely to play a role, Adams et al. (Adams et al., 2014) and Blinder et al. (2013) found no systematic specialization of the vascular supply for cortical columns.

(4) Variations in PSF or spatiotemporal complex kernels

Lastly, Kriegeskorte et al. (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. 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. (Gardner et al., 2006) and Shmuel et al. (Shmuel et al., 2010) constitute a specific scenario of the more general concept of interpreting fMRI sampling as spatio-temporal filtering of neuronal activity.

What does decoding at low field strength tell us?

Irrespective of the exact mechanisms, all proposed mechanisms reflect neuronal selectivity. Even though the exact spatial information is lost, the signals are expected to originate at the neuronal level. MVPA allows us to detect the presence of information about a stimulus or an experimental condition within an area. This information, when not detectable with conventional univariate analysis, indicates the possibility that fine scale structures such as cortical columns are involved in its processing. Studies at low magnetic field strengths can therefore help to find what kind of columnar level information processing an area is involved in. But it is unlikely that such studies would be sensitive to all aspects of columnar activity and ultimately they are not capable of resolving the activity in individual columns. Such a resolvability, however, may be necessary in order to understand local information processing.

Limitations of our current approach and questions for future research

There is a number of issues that were beyond the scope of this work but that we believe need to be addressed in order to better understand and advance the use of high spatial resolution fMRI to study cortical processes.

Imaging of layer dependent activity

Our investigation focused mainly on columnar patterns. The other fine-scale cortical structure of interest is the organization of the cortical sheet into layers. As we have discussed above, layer dependent activity may play a role in differentiating the origin of information flow that causes local responses as well as differences in how the local microcircuit processes information.

The PSF that we studied in our Study 2 describes the spread along the cortical manifold averaged over all cortical depths. It can be termed the *average tangential* PSF. There are two additional aspects that need to be studied in order to understand the spatial specificity with respect to cortical layers. The first aspect is the depth dependence of the tangential PSF. It is not known how the *layer specific tangential* PSF varies quantitatively as a function of depth and how this variation may depend on imaging methods and parameters. The fact that it does vary has been shown by Polimeni et al. (2010) who demonstrated

that the spatial specificity with respect to a retinotopic stimulus was lowest towards the pial surface

The second aspect is the spatial specificity *across* layers, orthogonal to the spread that the tangential PSF describes. It should be noted that there are some differences in cerebrovascular organization with respect to radial and angular direction (Duvernoy et al., 1981). The largest blood vessels are the pial surface veins that extend in various orientations along the tangential plane. Somewhat smaller are cortical- penetrating veins that are organized radially, traversing the different cortical layers. The smallest vessels, the capillaries, form a fine mesh that locally appears to be isotropic. However, their density varies with cortical layers (Weber et al., 2008). For these reasons, we cannot directly apply our PSF to the imaging of cortical layers. In addition, the distinctiveness of layers and the bounded nature of the cortical depth continuum appear to make a PSF convolution model ill-suited for fMRI of cortical layers. Nonetheless, the laminar organization of the vasculature (Adams et al., 2014) may turn out to benefit the resolvability of layer specific activity using fMRI, provided that the signal contributions of penetrating veins and pial draining veins can be minimized.

Recent results (Chaimow et al., 2012; 2011; De Martino et al., 2015; Fracasso et al., 2016; Muckli et al., 2015; Olman et al., 2012) suggest that it is possible to differentially resolve layer-specific signals on the scale of 1 mm or finer. For example, in a study which is not part of this thesis (Chaimow et al., 2012; 2011) we measured layer responses in areas V1 and MT to visual stimuli of coherently and incoherently moving random dot patterns as well as to flickering checkerboards. We found different layer response profiles for different types of stimuli. In particular coherently moving dots caused smaller responses in superficial layers of V1 than incoherently moving dots. This suppression of responses may reflect feedback from MT, where the detection of coherent motion leads to suppression of downstream areas compatible with the concept of predictive coding.

We believe layer studies to be a promising direction of further research, especially within the context of theories about feed-forward and feed-back flow of information such as predictive coding.

Arbitrary orientations of voxels and curved cortex

Our studies focused on two dimensional imaging. We assumed that fMRI data were acquired from a slice that overlapped tangentially with a flat region of cortical gray matter whose thickness was similar to the thickness of the visual cortex (2.5 mm). This approach has been successfully implemented in several studies of columnar imaging from V1 (Cheng et al., 2001; Goodyear and Menon, 2001; Shmuel et al., 2010; Yacoub et al., 2008; 2007). Still, it poses limitations on the area of interest and on the cortical anatomy of subjects.

A more general approach would consider curved cortex as well as arbitrary positions and orientations of voxels relative to the cortex. This more general situation would result in potential partial volume effects, affect the alignment of the voxel integration functions relative to the cortical manifold and could introduce differences in column-to-column distances along the pial surface of

the cortex compared to distances along the gray matter/white matter boundary.

As we discussed in Study 4, our optimization results are likely to generalize approximately for the scenario of arbitrary voxel orientations and sampling of curved cortex. It is however possible, that under these circumstances, detection, decoding performance, and reconstruction quality are somewhat reduced, and that optimal voxels will be somewhat smaller in order to reduce effects related to partial volume and curved cortex. However, the exact consequences of considering such a general scenario need to be examined through additional modeling and empirical studies preferably taking laminar aspects into account.

Variations in PSF, blood vessel and response variation

Our estimated PSF is an average “first order” model. It has been proposed (Kriegeskorte et al., 2010) and demonstrated (Polimeni et al., 2010) that the BOLD response depends on the cortical site, suggesting that it is more complex than a convolution with a single prototypical Gaussian. The spatial extent and the magnitude of the response may vary due to local variations in vascular geometry. As a consequence, a convolutional model with a single Gaussian function could only be an approximating simplification. Modeling the spatial BOLD response as a Gaussian PSF is a useful approximation for comparing fMRI contrasts, quantitative modeling, interpretation, and planning of high-resolution fMRI studies.

But it is of interest to quantify not only the first approximation but also its higher order variations. Is there a more precise model that describes the spatial specificity of fMRI methods (in tangential and radial direction with respect to cortical organization). As an example, Aquino et al. (2012) modeled the BOLD response as a travelling wave evolving in time and found that deconvolution of neural dynamics using such a model resulted in physiologically more plausible spatiotemporal patterns than when using a model separable in space and time (Aquino et al., 2014).

If the PSF varies as a function of location, to which extent does it vary? What does this variation depend on? Does it vary in amplitude, width and shape? The relatively wide distribution of average responses across voxels in Fig.4 of our Study 2 (distribution of l/r averg. resp.) indicates that spatial variation in response amplitude may be a factor that needs to be considered.

Finally what consequences does the variability of spatial specificity have for our expected ability to detect stimulus specific columnar response patterns, to decode from these responses and to reconstruct their fine scale organization? PSF variations may not matter much for reconstruction of a pattern as they are unlikely to influence the average correspondence between the true and the imaged pattern. Such variations could however have an effect on decoding algorithms that employ feature selection and would benefit disproportionately from voxels with narrow or biased PSFs. However, the true effect of the variability needs to be studied empirically by future studies which will apply high-resolution fMRI to large volumes of the brain.

Empirical studies

In addition to estimating the spatial specificity of BOLD fMRI at 7T, we used other current knowledge about columnar patterns, MR imaging, signal-to-noise ratios in fMRI and analysis methods, and developed quantitative models in order to make predictions about imaging of columns. Ultimately, empirical studies need to test how well our predictions work and to inform us about possible gaps in our knowledge.

Conclusions

We believe that high-resolution fMRI is an important tool for studying the cortical basis of cognitive processes in humans. In this thesis we have used model-based analysis of columnar imaging in order to study the spatial specificity of BOLD fMRI and to analyze factors involved in decoding from and imaging of fine-scale columnar patterns. We found that 7T BOLD imaging is capable of imaging typical columnar patterns and we quantified its spatial specificity and its potential to reconstruct and detect patterns. This quantitative knowledge may help to build models, plan studies and interpret results. Imaging approaches at lower field strengths relying on MVPA may exploit low-spatial-frequency information present due to irregularities in columnar patterns. Several additional mechanisms may explain unexpectedly high classification performances in the literature and are still under debate. Additional studies will be needed in order to further advance the potential of high-resolution fMRI to contribute to the understanding of the cortical basis of human cognitive processing.

References

- Adams, D.L., Piserchia, V., Economides, J.R., Horton, J.C., 2014. Vascular Supply of the Cerebral Cortex is Specialized for Cell Layers but Not Columns. *Cereb Cortex* 25, 3673–3681. doi:10.1093/cercor/bhu221
- Adams, D.L., Sincich, L.C., Horton, J.C., 2007. Complete pattern of ocular dominance columns in human primary visual cortex. *J Neurosci* 27, 10391–10403. doi:10.1523/JNEUROSCI.2923-07.2007
- Alink, A., Krugliak, A., Walther, A., Kriegeskorte, N., 2013. fMRI orientation decoding in V1 does not require global maps or globally coherent orientation stimuli. *Front. Psychology* 4, 1–14. doi:10.3389/fpsyg.2013.00493/abstract
- Aquino, K.M., Robinson, P.A., Schira, M.M., Breakspear, M., 2014. Deconvolution of neural dynamics from fMRI data using a spatiotemporal hemodynamic response function. *Neuroimage* 94, 203–215. doi:10.1016/j.neuroimage.2014.03.001
- Aquino, K.M., Schira, M.M., Robinson, P.A., Drysdale, P.M., Breakspear, M., 2012. Hemodynamic traveling waves in human visual cortex. *PLoS Comput Biol* 8, e1002435. doi:10.1371/journal.pcbi.1002435
- Bandettini, P.A., Wong, E.C., Hinks, R.S., Tikofsky, R.S., Hyde, J.S., 1992. Time course EPI of human brain function during task activation. *Magn Reson Med* 25, 390–397.
- Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., Friston, K.J., 2012. Canonical microcircuits for predictive coding. *Neuron* 76, 695–711.

- doi:10.1016/j.neuron.2012.10.038
- Beckett, A., Peirce, J.W., Sanchez-Panchuelo, R.-M., Francis, S., Schluppeck, D., 2012. Contribution of large scale biases in decoding of direction-of-motion from high-resolution fMRI data in human early visual cortex. *Neuroimage* 63, 1623–1632. doi:10.1016/j.neuroimage.2012.07.066
- Berens, P., Keliris, G.A., Ecker, A.S., Logothetis, N.K., Tolias, A.S., 2008. Comparing the feature selectivity of the gamma-band of the local field potential and the underlying spiking activity in primate visual cortex. *Front Syst Neurosci* 2, 2. doi:10.3389/neuro.06.002.2008
- Blasdel, G., Obermayer, K., Kiorpes, L., 1995. Organization of ocular dominance and orientation columns in the striate cortex of neonatal macaque monkeys. *Vis Neurosci* 12, 589–603.
- Blinder, P., Tsai, P.S., Kaufhold, J.P., Knutsen, P.M., Suhl, H., Kleinfeld, D., 2013. The cortical angiome: an interconnected vascular network with noncolumnar patterns of blood flow. *Nat Neurosci* 16, 889–897. doi:10.1038/nn.3426
- Carlson, T.A., 2014. Orientation decoding in human visual cortex: new insights from an unbiased perspective. *J Neurosci* 34, 8373–8383. doi:10.1523/JNEUROSCI.0548-14.2014
- Carlson, T.A., Wardle, S.G., 2015. Sensible decoding. *Neuroimage* 110, 217–218. doi:10.1016/j.neuroimage.2015.02.009
- Chaimow, D., Yacoub, E., Feinberg, D.A., Goerke, U., Uğurbil, K., Shmuel, A., 2012. Layer specific fMRI correlates of motion processing in human cortical areas V1 and MT, in: Presented at the 18th Annual Meeting of the Organization for Human Brain Mapping, Beijing.
- Chaimow, D., Yacoub, E., Goerke, U., Uğurbil, K., Shmuel, A., 2011. Response amplitude in supra-granular layers of V1 differs for coherent and incoherent visual motion, in: Presented at the 17th Annual Meeting of the Organization of Human Brain Mapping, Quebec City.
- Cheng, K., Waggoner, R.A., Tanaka, K., 2001. Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron* 32, 359–374.
- De Martino, F., Moerel, M., Uğurbil, K., Goebel, R., Yacoub, E., Formisano, E., 2015. Frequency preference and attention effects across cortical depths in the human primary auditory cortex. *Proceedings of the National Academy of Sciences* 112, 16036–16041. doi:10.1073/pnas.1507552112
- Douglas, R.J., Martin, K.A.C., 2004. Neuronal circuits of the neocortex. *Annu Rev Neurosci* 27, 419–451. doi:10.1146/annurev.neuro.27.070203.144152
- Duong, T.Q., Yacoub, E., Adriany, G., Hu, X., Uğurbil, K., Kim, S.-G.G., 2003. Microvascular BOLD contribution at 4 and 7 T in the human brain: Gradient-echo and spin-echo fMRI with suppression of blood effects. *Magn Reson Med* 49, 1019–1027. doi:10.1002/mrm.10472
- Duvernoy, H.M., Delon, S., Vannson, J.L., 1981. Cortical blood vessels of the human brain. *Brain Res Bull* 7, 519–579.
- Engel, S.A., Glover, G.H., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7, 181–192. doi:10.1093/cercor/7.2.181
- Feinberg, D.A., Harel, N., Ramanna, S., Uğurbil, K., Yacoub, E., 2008. Sub-millimeter Single-shot 3D GRASE with Inner Volume Selection for T2 weighted fMRI applications at 7 Tesla. *Proc. Intl. Soc. Mag. Reson. Med.* 16, 1–1.
- Felleman, D.J., van Essen, D.C., 1991. Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cereb Cortex* 1, 1–47. doi:10.1093/cercor/1.1.1
- Fracasso, A., Petridou, N., Dumoulin, S.O., 2016. Systematic variation of population

- receptive field properties across cortical depth in human visual cortex. *Neuroimage* 139, 427–438. doi:10.1016/j.neuroimage.2016.06.048
- Freeman, J., Brouwer, G.J., Heeger, D.J., Merriam, E.P., 2011. Orientation decoding depends on maps, not columns. *J Neurosci* 31, 4792–4804. doi:10.1523/JNEUROSCI.5160-10.2011
- Freeman, J., Heeger, D.J., Merriam, E.P., 2013. Coarse-scale biases for spirals and orientation in human visual cortex. *J Neurosci* 33, 19695–19703. doi:10.1523/JNEUROSCI.0889-13.2013
- Furmanski, C.S., Engel, S.A., 2000. An oblique effect in human primary visual cortex. *Nat Neurosci* 3, 535–536. doi:10.1038/75702
- Gardner, J.L., 2010. Is cortical vasculature functionally organized? *Neuroimage* 49, 1953–1956. doi:10.1016/j.neuroimage.2009.07.004
- Gardner, J.L., Sun, P., Tanaka, K., Heeger, D.J., Cheng, K., 2006. Classification analysis with high spatial resolution fMRI reveals large draining veins with orientation specific responses, in: pp. 614–640.
- Goodyear, B.G., Menon, R.S., 2001. Brief visual stimulation allows mapping of ocular dominance in visual cortex using fMRI. *Hum Brain Mapp* 14, 210–217.
- Haacke, M.E., Brown, R.W., Thompson, M.R., Venkatesan, R., 1999. Magnetic resonance imaging: physical principles and sequence design.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293, 2425–2430. doi:10.1126/science.1063736
- Haynes, J., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8, 686–691.
- Hubel, D., Wiesel, T.N., 1968. Receptive fields and functional architecture of monkey striate cortex. *J Physiol (Lond)* 195, 215–243.
- Hubel, D.H., Wiesel, T.N., 1974. Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor. *J Comp Neurol* 158, 295–305. doi:10.1002/cne.901580305
- Jochimsen, T.H., Norris, D.G., Mildner, T., Möller, H.E., 2004. Quantifying the intra- and extravascular contributions to spin-echo fMRI at 3 T. *Magn Reson Med* 52, 724–732. doi:10.1002/mrm.20221
- Kamitani, Y., Tong, F., 2006. Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16, 1096–1102. doi:10.1016/j.cub.2006.04.003
- Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8, 679–685. doi:10.1038/nn1444
- Kemper, V.G., De Martino, F., Vu, A.T., Poser, B.A., Feinberg, D.A., Goebel, R., Yacoub, E., 2015. Sub-millimeter T2 weighted fMRI at 7 T: comparison of 3D-GRASE and 2D SE-EPI. *Front Neurosci* 9, 163. doi:10.3389/fnins.2015.00163
- Kriegeskorte, N., Bandettini, P.A., 2007. Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage* 38, 649–662.
- 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. doi:10.1016/j.neuroimage.2009.09.059
- Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences* 89, 5675–5679.
- Maloney, R.T., 2015. The basis of orientation decoding in human primary visual cortex: fine- or coarse-scale biases? *J Neurophysiol* 113, 1–3. doi:10.1152/jn.00196.2014

- Mannion, D., McDonald, J., Clifford, C.W.G., 2009. Discrimination of the local orientation structure of spiral Glass patterns early in human visual *Neuroimage*.
- Menon, R.S., Goodyear, B.G., 1999. Submillimeter functional localization in human striate cortex using BOLD contrast at 4 Tesla: implications for the vascular point-spread function. *Magn Reson Med* 41, 230–235.
- Muckli, L., De Martino, F., Vizioli, L., Petro, L.S., Smith, F.W., 2015. Contextual Feedback to Superficial Layers of V1. *Curr Biol*.
- Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences* 87, 9868–9872. doi:10.1073/pnas.87.24.9868
- Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S.-G.G., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences* 89, 5951–5955. doi:10.1073/pnas.89.13.5951
- Olman, C.A., Harel, N., Feinberg, D.A., He, S., Zhang, P., Ugurbil, K., Yacoub, E., 2012. Layer-specific fMRI reflects different neuronal computations at different depths in human V1. *PLoS ONE* 7, e32536–9. doi:10.1371/journal.pone.0032536
- Op de Beeck, H.P., 2010. Probing the mysterious underpinnings of multi-voxel fMRI analyses. *Neuroimage* 50, 567–571. doi:10.1016/j.neuroimage.2009.12.072
- Oshio, K., Feinberg, D.A., 1991. GRASE (Gradient- and spin-echo) imaging: a novel fast MRI technique. *Magn Reson Med* 20, 344–349.
- Parkes, L.M., Schwarzbach, J.V., Bouts, A.A., Deckers, R.H.R., Pullens, P., Kerskens, C.M., Norris, D.G., 2005. Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. *Magn Reson Med* 54, 1465–1472. doi:10.1002/mrm.20712
- Polimeni, J.R., Fischl, B., Greve, D.N., Wald, L.L., 2010. Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1. *Neuroimage* 52, 1334–1346. doi:10.1016/j.neuroimage.2010.05.005
- Pratte, M.S., Sy, J.L., Swisher, J.D., Tong, F., 2016. Radial bias is not necessary for orientation decoding. *Neuroimage* 127, 23–33. doi:10.1016/j.neuroimage.2015.11.066
- Roger, A., Schwartz, E., 1990. Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biol Cybern* 62, 381–391.
- Sasaki, Y., Rajimehr, R., Kim, B.W., Ekstrom, L.B., Vanduffel, W., Tootell, R.B.H., 2006. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron* 51, 661–670. doi:10.1016/j.neuron.2006.07.021
- Shmuel, A., 2010. Locally Measured Neuronal Correlates of Functional MRI Signals 63–82. doi:10.1007/978-3-540-87919-0_4
- 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. doi:10.1016/j.neuroimage.2009.08.040
- Shmuel, A., Yacoub, E., Chaimow, D., Logothetis, N.K., Ugurbil, K., 2007. Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 Tesla. *Neuroimage* 35, 539–552. doi:10.1016/j.neuroimage.2006.12.030
- Swisher, J.D., Gatenby, J.C., Gore, J.C., Wolfe, B.A., Moon, C.-H., Kim, S.-G.G., Tong, F., 2010. Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *J Neurosci* 30, 325–330. doi:10.1523/JNEUROSCI.4811-09.2010
- Thompson, R., Correia, M., Cusack, R., 2011. Vascular contributions to pattern analysis: Comparing gradient and spin echo fMRI at 3T. *Neuroimage* 56, 643–650. doi:10.1016/j.neuroimage.2010.03.061

- Tychsen, L., Burkhalter, A., 1997. Nasotemporal asymmetries in V1: ocular dominance columns of infant, adult, and strabismic macaque monkeys. *J Comp Neurol* 388, 32–46.
- Uludağ, K., Müller-Bierl, B., Uğurbil, K., 2009. An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. *Neuroimage* 48, 150–165. doi:10.1016/j.neuroimage.2009.05.051
- Wang, H.X., Merriam, E.P., Freeman, J., Heeger, D.J., 2014. Motion Direction Biases and Decoding in Human Visual Cortex. *J Neurosci* 34, 12601–12615. doi:10.1523/JNEUROSCI.1034-14.2014
- Weber, B., Keller, A., Reichold, J., Logothetis, N.K., 2008. The Microvascular System of the Striate and Extrastriate Visual Cortex of the Macaque. *Cereb Cortex* 18, 2318–2330. doi:10.1093/cercor/bhm259
- Yacoub, E., Duong, T.Q., van de Moortele, P.F., Lindquist, M., Adriany, G., Kim, S.-G.G., Uğurbil, K., Hu, X.P., 2003. Spin-echo fMRI in humans using high spatial resolutions and high magnetic fields. *Magn Reson Med* 49, 655–664. doi:10.1002/mrm.10433
- Yacoub, E., Harel, N., Uğurbil, K., 2008. High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences* 105, 10607–10612. doi:10.1073/pnas.0804110105
- Yacoub, E., Shmuel, A., Logothetis, N.K., Uğurbil, 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., Shmuel, A., Pfeuffer, J., van de Moortele, P.F., Adriany, G., Andersen, P., Vaughan, J.T., Merkle, H., Uğurbil, K., Hu, X., 2001. Imaging brain function in humans at 7 Tesla. *Magn Reson Med* 45, 588–594. doi:10.1002/mrm.1080
- Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P.-F., Feinberg, D., Adriany, G., Chaimow, D., Shmuel, A., Uğurbil, 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. doi:10.1371/journal.pone.0028716

STUDY 1

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

Denis Chaimow, Essa Yacoub, Kâmil Uğurbil, Amir Shmuel


The manuscript has been published in *NeuroImage* 56 (2011), pp. 627—642.

Author contributions

Denis Chaimow developed and implemented the model, and wrote the manuscript. Essa Yacoub contributed to discussions. Kâmil Uğurbil contributed to discussions. Amir Shmuel advised on the developed of the model, and edited the manuscript.

Acknowledgements

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.



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

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

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

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

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

ARTICLE INFO

Article history:

Received 30 December 2009

Revised 14 September 2010

Accepted 16 September 2010

Available online 22 September 2010

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.

© 2010 Elsevier Inc. All rights reserved.

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

* Corresponding author. MNI, 3801 University St., Room 753A, Montreal, QC, Canada H3A 2B4.

E-mail address: amir.shmuel@mcgill.ca (A. Shmuel).

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.

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	
ϵ	Pattern branchiness, variations parallel to ODC bands	
α	Sharpness parameter of the sigmoidal non-linearity in ODC	$f(x) = \frac{1}{1 + e^{-\alpha x}}$
σ_{BOLD}	Bold point spread width	$\text{FWHM} = 2.35 \cdot \sigma_{\text{BOLD}}$
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)	$\sigma = \frac{1}{\text{tSNR}}$
\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	$\frac{\sqrt{nc}}{\sigma}$
n	Number of voxels	
t	Number of averaged volumes	

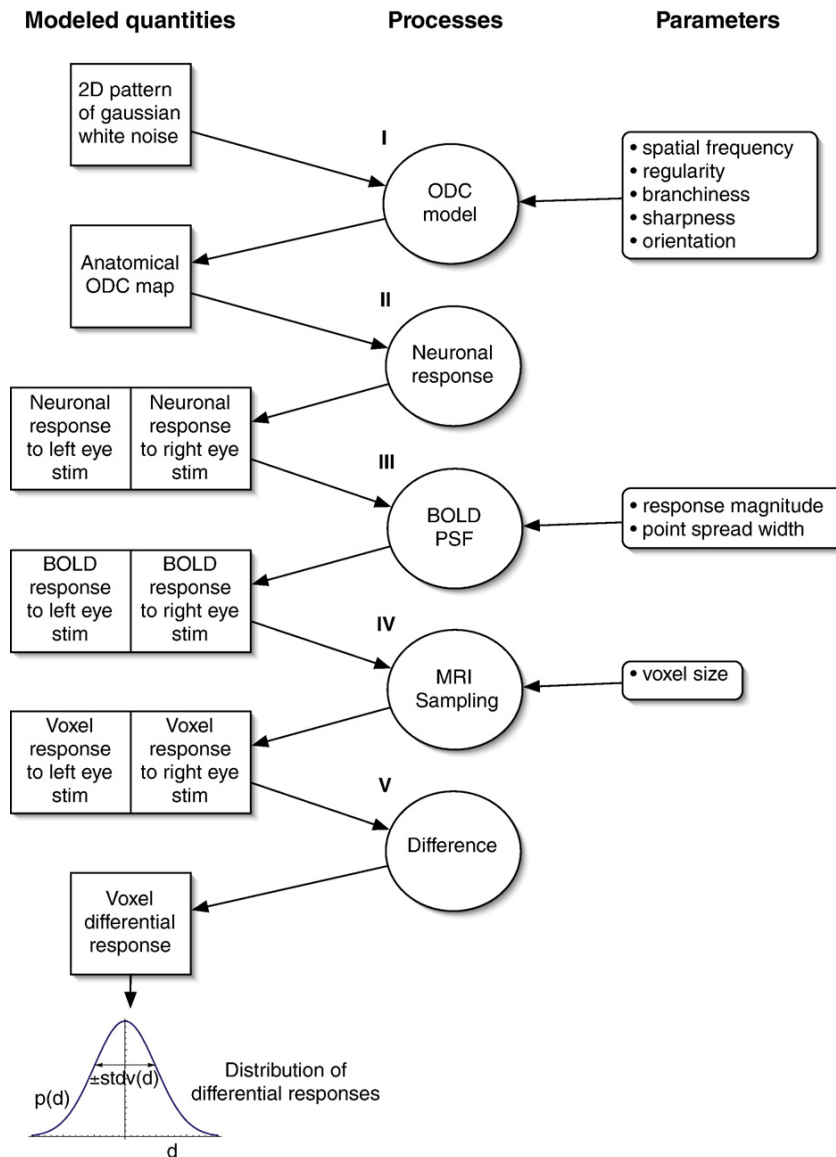


Fig. 1. Model overview. In stage 1, ODC maps are modeled by spatial filtering of white noise. In stage 2 we simulate the neuronal response to right- or left eye stimulation. In stage 3 the neuronal response is convolved with a BOLD point-spread function. In stage 4 the BOLD response is transformed into 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").

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).

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

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×48 mm and 192 mm×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

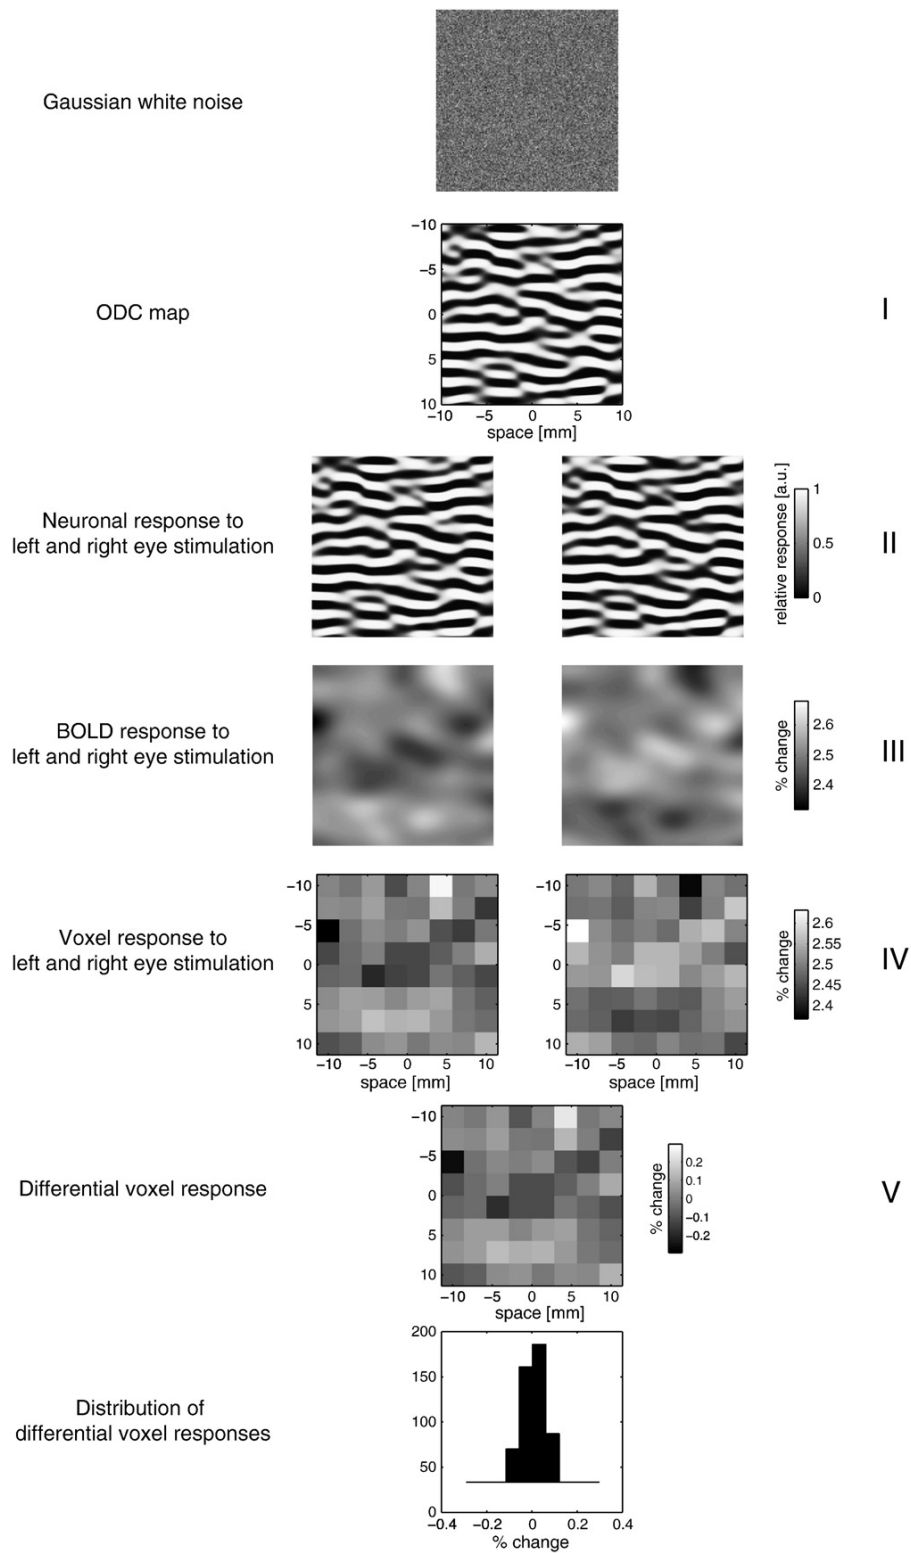


Fig. 2. Numerical realization of the model. The figure presents a numerical simulation of the model. Each row shows the results of one single stage of the model: the Gaussian white noise input, the ODC map, the neuronal response, the BOLD response, the voxel response, and the differential voxel response. 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.

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

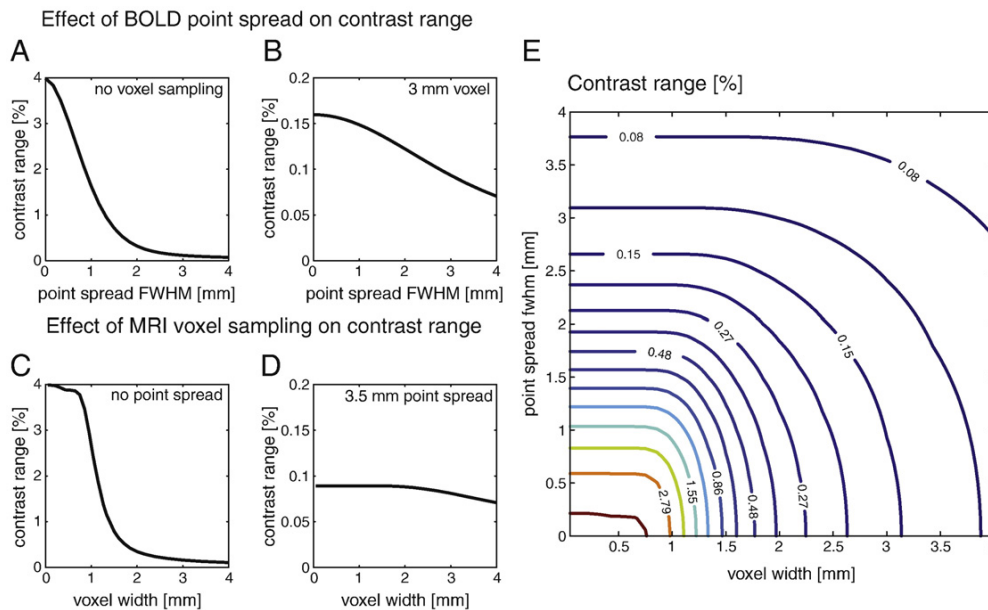


Fig. 3. Dependence of contrast range on voxel width and BOLD point spread. Contrast range is defined as the standard deviation of the distribution of differential responses (percent change relative to baseline). Contrast range in a set of imaged voxels is presented as a function of FWHM of BOLD point 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.

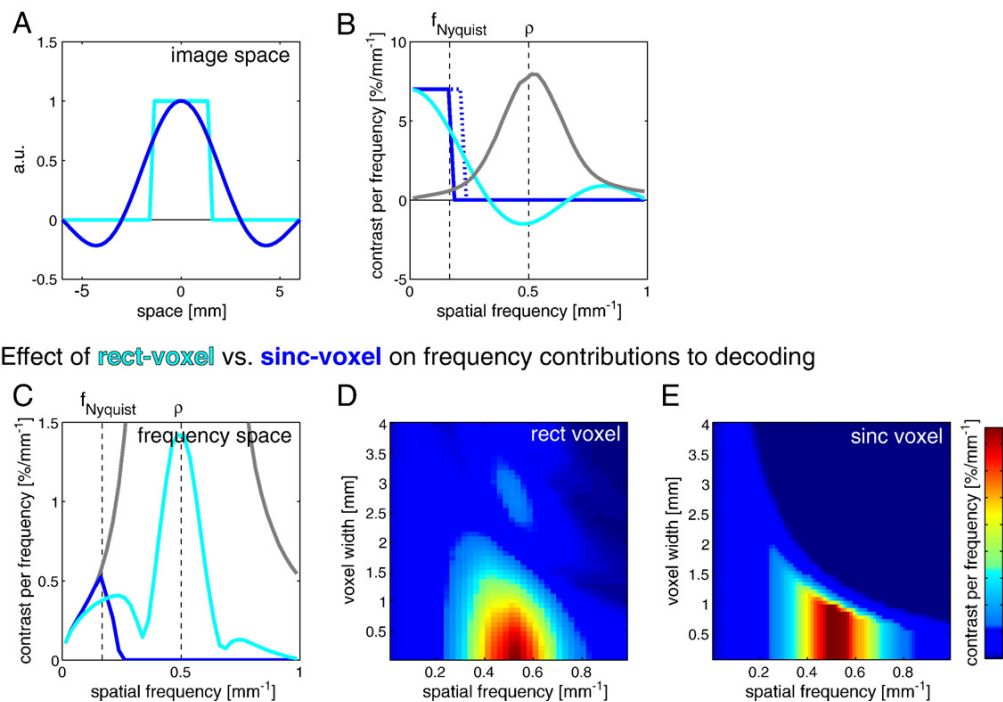
Voxel as **rect-function** vs. voxel as **sinc-function**

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 wide 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.

(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 ($0.167 \cdot \sqrt{2}$). 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);

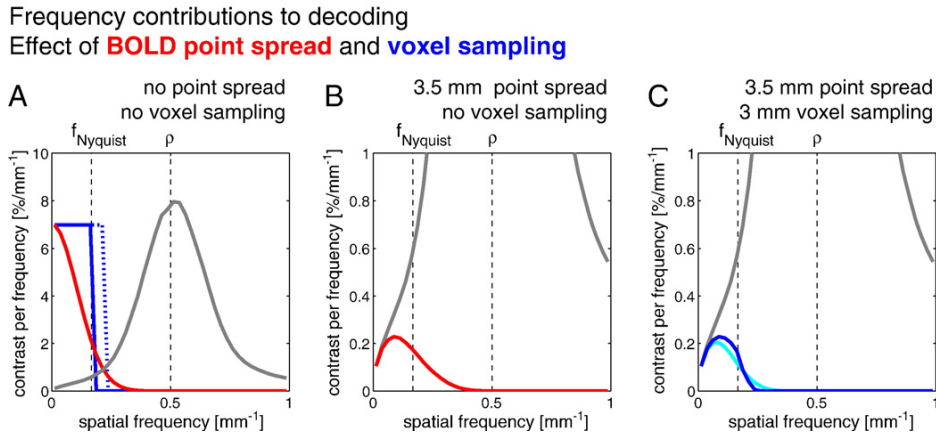


Fig. 5. Contributions of pattern spatial frequency components to the contrast range. The figure presents contributions of spatial frequency components in the ODC pattern to the range of contrasts in the set of imaged voxels, in a format similar to that used in Fig. 4. In A–C, different subsets of the 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.

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

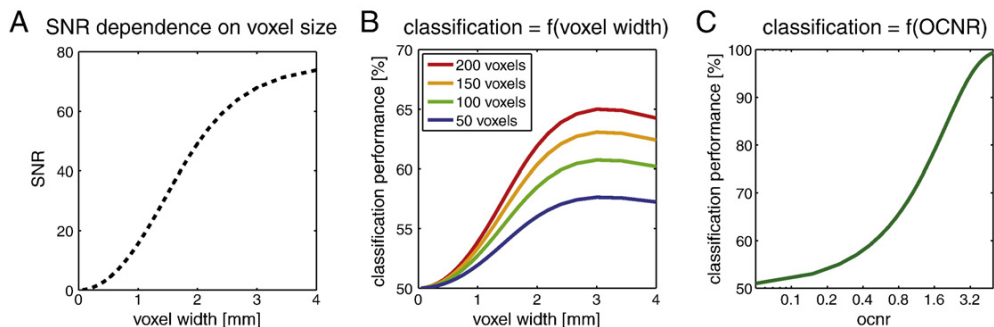


Fig. 6. Classification performance. (A) Time-course SNR (tSNR) as a function of voxel width at 3 Tesla. Noise levels were computed following Triantafyllou et al. (2005), using TR = 2 s (see section C in the Appendix). B presents classification performance as a function of in-plane voxel width. The slice 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.

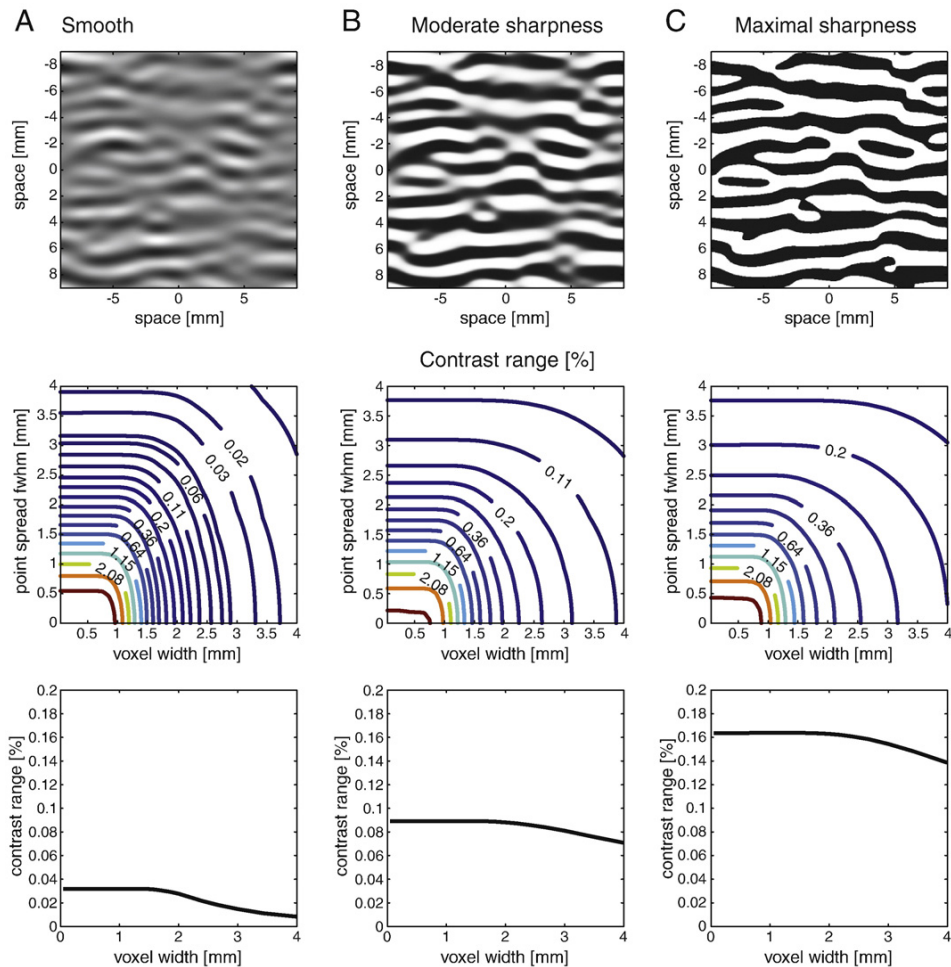


Fig. 7. The effect of varying ODC pattern sharpness on the sampled contrast range. The sharpness of transitions between ODCs can be modeled by a sigmoidal non-linearity with a degree controlled by parameter α . The figure shows the effect of this non-linearity on the model results by using different 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.

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.

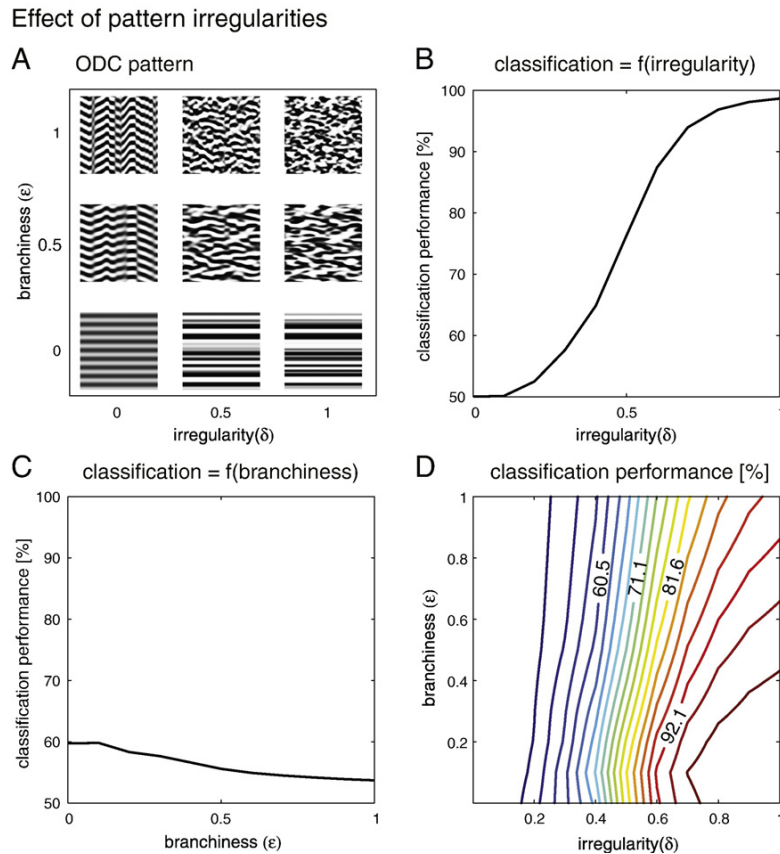


Fig. 8. The effect of pattern irregularities on classification rate. The irregularity of the ODC pattern is varied to study its effect on classification performance. Classification performance is predicted for 100 voxels and a TR of 2 s. (A) demonstrates the effects of the irregularity parameter (δ) 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.

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.

Conflict of interest statement

The authors declare that they have no conflicts of interest.

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}_i . Each vector \vec{x}_i is sampled from one of two normal distributions $\mathcal{N}(\vec{\mu}_A, \Sigma_A)$ and $\mathcal{N}(\vec{\mu}_B, \Sigma_B)$ corresponding to the two stimulation 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 $\vec{w} = \vec{d} / \|\vec{d}\|$ pointing in the direction of the line connecting the means of the two distributions. This results in one-dimensional variables y_t , given by:

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

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 ($y_t < 0$).

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

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

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

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

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

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 $OCNR = \frac{\sqrt{ntc}}{\sigma}$.

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:

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

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):

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

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 $y_B(\vec{r})$.

$$y_A(\vec{r}) = \frac{1}{2} + \frac{x(\vec{r})}{2}$$

$$y_B(\vec{r}) = \frac{1}{2} - \frac{x(\vec{r})}{2}.$$

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 $x(\vec{r}^*)$ is then given by:

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

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

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

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 $\tilde{x}_{\text{stim}}(\vec{r}^*) = b + b \cdot x(\vec{r}^*)$ during stimulation and $\tilde{x}_{\text{bl}}(\vec{r}^*) = b$ during baseline.

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

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

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):

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

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):

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

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

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

Inserting this result into equation 3 we get:

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

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

References

- Adams, D.L., Sincich, L.C., Horton, J.C., 2007. Complete pattern of ocular dominance columns in human primary visual cortex. *J. Neurosci.* 27, 10391–10403.
- Blasdel, G., Obermayer, K., Kiorpes, L., 1995. Organization of ocular dominance and orientation columns in the striate cortex of neonatal macaque monkeys. *Vis. Neurosci.* 12, 589–603.
- Boynton, G.M., 2005. Imaging orientation selectivity: decoding conscious perception in V1. *Nat. Neurosci.* 8, 541–542.
- Boynton, G.M., Engel, S., Glover, G., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* 16, 4207–4221.
- Duda, R.O., Hart, P.E., Strok, D.G., 2006. *Pattern classification*, 2nd ed. Wiley-Interscience.
- Engel, S., Glover, G., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb. Cortex* 7, 181–192.
- Furmanski, C.S., Engel, S.A., 2000. An oblique effect in human primary visual cortex. *Nat. Neurosci.* 3, 535–536.
- Gardner, J.L., Sun, P., Tanaka, K., Heeger, D.J., Cheng, K., 2006. Classification analysis with high spatial resolution fMRI reveals large draining veins with orientation specific responses. Society for Neuroscience Meeting, Atlanta, GA, USA.
- Greenspan, H., 2009. Super-Resolution in Medical Imaging. *Comput. J.* 52, 43–63.
- Haacke, M.E., Brown, R.W., Thompson, M.R., 1999. *Magnetic resonance imaging: physical principles and sequence design*. Wiley-Liss.
- Haxby, J., Gobbini, M., Furey, M., Ishai, A., Schouten, J., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293, 2425–2430.
- Haynes, J., Rees, G., 2005a. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci.* 8, 686–691.
- Haynes, J., Rees, G., 2005b. Predicting the stream of consciousness from activity in human visual cortex. *Curr. Biol.* 15, 1301–1307.
- Haynes, J., Rees, G., 2006. Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* 7, 523–534.
- Horton, J.C., Dagi, L., McCrane, E., de Monasterio, F., 1990. Arrangement of ocular dominance columns in human visual cortex. *Arch. Ophthalmol.* 108, 1025–1031.
- 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.
- Kriegeskorte, N., Bandettini, P.A., 2007. Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage* 38, 649–662.
- 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.
- Krüger, G., Kastrup, A., Glover, G.H., 2001. Neuroimaging at 1.5 T and 3.0 T: comparison of oxygenation-sensitive magnetic resonance imaging. *Magn. Reson. Med.* 45, 595–604.
- Mayer, G.S., Vrscaj, E.R., 2007. Measuring information gain for frequency-encoded super-resolution MRI. *Magn. Reson. Imaging* 25, 1058–1069.
- Obermayer, K., Blasdel, G.G., 1993. Geometry of orientation and ocular dominance columns in monkey striate cortex. *J. Neurosci.* 13, 4114–4129.
- Op de Beeck, H.P., 2010. Against hyperacuity in brain reading: spatial smoothing does not hurt multivariate fMRI analyses? *Neuroimage* 49, 1943–1948.
- Parkes, L.M., Schwarzbach, J.V., Bouts, A.A., Deckers, R.H.R., Pullens, P., Kerskens, C.M., Norris, D.G., 2005. Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. *Magn. Reson. Med.* 54, 1465–1472.
- Pereira, F., Mitchell, T., Botvinick, M., 2009. Machine learning classifiers and fMRI: a tutorial overview. *Neuroimage* 45, S199–S209.
- Roger, A., Schwartz, E., 1990. Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biol. Cybern.* 62, 381–391.
- Sasaki, Y., Rajimehr, R., Kim, B.W., Ekstrom, L.B., Vanduffel, W., Tootell, R.B.H., 2006. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron* 51, 661–670.
- 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.
- Shmuel, A., Yacoub, E., Chaimow, D., Logothetis, N.K., Ugurbil, K., 2007. Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 Tesla. *Neuroimage* 35, 539–552.
- Swisher, J.D., Gatenby, J.C., Gore, J.C., Wolfe, B.A., Moon, C.-H., Kim, S.-G., Tong, F., 2010. Multiscale pattern analysis of orientation-selective activity in the primary visual cortex. *J. Neurosci.* 30, 325–330.
- Triantafyllou, C., Hoge, R., Krueger, G., Wiggins, C., Potthast, A., Wiggins, G., Wald, L., 2005. Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage* 26, 243–250.
- Tychsen, L., Burkhalter, A., 1997. Nasotemporal asymmetries in V1: ocular dominance columns of infant, adult, and strabismic macaque monkeys. *J. Comp. Neurol.* 388, 32–46.
- Yacoub, E., Shmuel, A., Logothetis, N.K., 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.

STUDY 2

Spatial specificity of the functional MRI blood oxygenation response relative to neuronal activity

Denis Chaimow, Essa Yacoub, Kâmil Uğurbil, Amir Shmuel

The manuscript is under revision for publication in NeuroImage.

Author contributions

Denis Chaimow developed and implemented the model, analyzed the data, and wrote the manuscript. Essa Yacoub acquired the data, contributed to discussions and commented on the manuscript. Kâmil Uğurbil contributed to discussions and commented on the manuscript. Amir Shmuel acquired the data, advised on the development of the model, and edited the manuscript.

Acknowledgements

This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada (AS, NSERC Discovery grants RGPIN 375457-09 and RGPIN-2015-05103).

Abstract

Previous attempts at characterizing the spatial specificity of the blood oxygenation level dependent functional MRI (BOLD fMRI) response by estimating its point-spread function (PSF) have conventionally relied on spatial representations of visual stimuli in area V1. Consequently, their estimates were confounded by the width and scatter of receptive fields of V1 neurons. Here, we circumvent these limits by instead using the inherent cortical spatial organization of ocular dominance columns (ODCs) to determine the PSF for both Gradient Echo (GE) and Spin Echo (SE) BOLD imaging at 7 Tesla. By applying Markov Chain Monte Carlo sampling on a probabilistic generative model of imaging ODCs, we quantified the PSFs that best predict the spatial structure and magnitude of differential ODCs' responses. Prior distributions for the ODC model parameters were determined by analyzing published data of cytochrome oxidase patterns from post-mortem histology of human V1 and of neurophysiological ocular dominance indices. The most probable PSF full-widths at half-maximum were 0.82 mm (SE) and 1.02 mm (GE). Our results provide a quantitative basis for the spatial specificity of BOLD fMRI at ultra-high fields, which can be used for planning and interpretation of high-resolution differential fMRI of fine-scale cortical organizations.

Introduction

Functional magnetic resonance imaging (fMRI) of the human brain is increasingly being used to investigate fine-scale structures such as cortical columns (Cheng et al., 2001; De Martino et al., 2015; Goodyear and Menon, 2001; Menon et al., 1997; Shmuel et al., 2010; Yacoub et al., 2008; 2007; Zimmermann et al., 2011). To optimally plan high-resolution fMRI studies and to correctly interpret their results it is necessary to know the inherent limits of the fMRI spatial specificity relative to the sites where changes in neuronal activity occur.

The most commonly used fMRI approach relies on gradient echo (GE) blood oxygenation level dependent (BOLD) contrast (Bandettini et al., 1992; Kwong et al., 1992; Ogawa et al., 1990; 1992). GE BOLD is sensitive to the intra- and extravascular effects of activation-induced changes in the deoxy-hemoglobin content of blood. At standard magnetic field strengths (1.5 T, 3 T) the signal is dominated by contributions from larger blood vessels. At higher magnetic field strengths the strong intravascular component of these large blood vessels decreases, while the extravascular signal changes around capillaries and smaller vessels increase (Uludağ et al., 2009; Yacoub et al., 2001). Additional weighting towards the microvasculature can be achieved by using spin echo (SE) BOLD imaging, which suppresses extravascular signal contributions from larger blood vessels (Uludağ et al., 2009; Yacoub et al., 2003).

The first study to quantify the spatial specificity of the BOLD response (Engel et al., 1997) used an elegant phase-encoding paradigm that induced traveling waves of retinotopic neural activity in the primary visual cortex (V1). Assuming a shift-invariant linear response, Engel et al. (1997) estimated the point-spread function (PSF), which represents the spatial response that would

be elicited by a small point stimulus. They found the full-width at half-maximum (FWHM) of the GE BOLD PSF to be 3.5 mm at 1.5 T. Similar values (3.9 mm for GE BOLD and 3.4 mm for SE BOLD) have been reported at 3 T (Parkes et al., 2005) using a paradigm similar to that used in Engel et al. (1997). To estimate the GE BOLD PSF at 7 T, we previously measured the spatiotemporal spread of the fMRI response in grey matter regions around the V1 representation of edges of visual stimuli (Shmuel et al., 2007). To reduce contributions from macroscopic veins, we excluded voxels that showed vessel-like response features. The mean measured and estimated FWHMs were 2.34 ± 0.20 mm and < 2 mm, respectively. The spatial specificity of SE BOLD fMRI at ultra-high magnetic fields has not yet been quantified.

All previous attempts at characterizing the spatial specificity of the BOLD fMRI response (Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007) relied on an implicit assumption that neuronal responses to small visual stimuli are point-like. However, to estimate the spatial specificity of the BOLD response, these studies have conventionally relied on spatial representations of visual stimuli in area V1. Unlike the implicit assumption of point-like responses, the receptive fields of neurons in V1 have non-zero spatial extents (Hubel and Wiesel, 1968). In addition, electrode measurements in macaque V1, oriented orthogonally relative to the surface of cortex have demonstrated substantial scatter in the center of receptive fields (Hubel and Wiesel, 1974). Therefore, the pattern of neural activity parallel to the cortical surface is a blurred representation of the visual stimulus. This implies that receptive field size and scatter pose a lower limit on any BOLD fMRI PSF width that is estimated using spatial representations of visual stimuli in V1. Consequently, the previously computed estimates of the spatial specificity of the fMRI response were confounded by the width and scatter of receptive fields of V1 neurons. Such estimates are limited in that they solely measure the capacity of the BOLD response to resolve retinotopic representations; they do not measure its ability to resolve more fine-grained neural activity. Yet only this latter resolvability matters for functional imaging at the spatial scale of cortical columns.

Here, we estimate and compare the PSF widths of GE and SE BOLD imaging at 7 T using a novel approach. We circumvent the limits posed by the retinotopic representation of visual stimuli by instead using the inherent cortical spatial organization of ocular dominance columns (ODCs). To this end, we fit a model of ODCs imaging (Chaimow et al., 2011) to ODCs responses acquired at 7 T (Yacoub et al., 2007). We quantify the width of the PSF that best predicts the spatial structure and magnitude of differential ODC responses. Since we do not have access to the underlying anatomical ODC patterns and neurophysiological responses, we use a probabilistic modeling approach. We constrain the model ODC parameters by estimating features of real ODC patterns taken from post-mortem cytochrome oxidase (CO) maps of human ODCs (Adams et al., 2007) and neurophysiological response distributions in primates (Berens et al., 2008; Hubel and Wiesel, 1968). We then fit our model by Markov Chain Monte Carlo (MCMC) sampling. Our results provide a quantitative basis for the spatial specificity of differential BOLD fMRI at ultra-high fields.

Methods

Overview

We developed a probabilistic generative model of imaging ODCs in order to estimate widths of GE and SE BOLD PSFs that would best explain our previously obtained fMRI data of differential ODC maps (Yacoub et al., 2007).

The measured fMRI maps consisted of voxel estimates of the difference in BOLD responses to left and right eye stimulation. We modeled these responses as the sum of predictions from an ODC imaging model and measurement noise. The predictions from the ODC imaging model were completely determined by a set of parameters. Accounting for the effect of the measurement noise allowed us to first express the probability of observing the measured fMRI maps as a function of model parameters. In a second stage, we derived the posterior probability of the model parameters given the observed data and the prior probability of parameters.

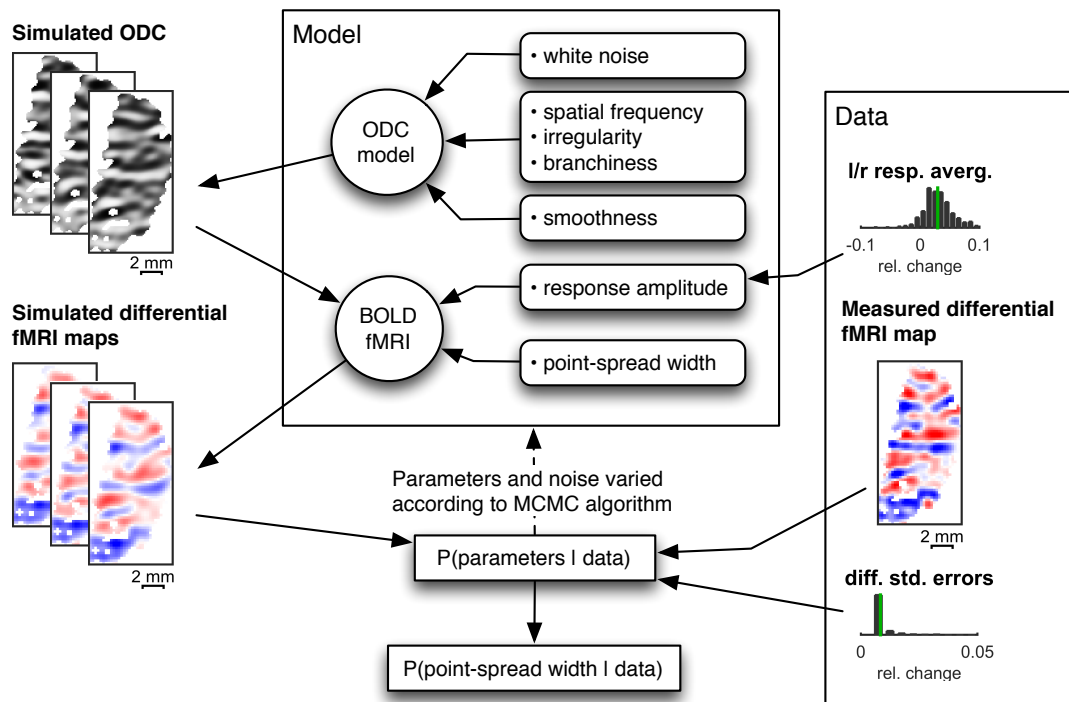


Fig. 1 Overview of Markov Chain Monte Carlo fitting. The model was fitted to the fMRI data using Markov Chain Monte Carlo (MCMC) sampling. For an arbitrary given set of parameters, the model generated a differential fMRI map (left). This map was compared to the measured fMRI map (right) and the likelihood of parameters given the data was calculated. The MCMC algorithm uses this likelihood together with parameter priors to further traverse the parameter space. After sufficiently many iterations the resulting parameter samples are distributed according to their joint posterior probability distribution.

Model of imaging ODCs

We implemented a model of imaging ODCs (Chaimow et al., 2011; see Fig. 1 for an overview; see Appendix A in the Suppl. Material for the detailed equations). The first component of the model, i.e. the modeling of realistic ODCs, followed (Rojer and Schwartz, 1990). It consisted of band-pass filtering of spatial white

noise using an anisotropic filter. The filtering was followed by applying a sigmoidal point-wise non-linearity, which controlled the smoothness of transitions between left and right eye preference regions.

The spatial BOLD response was modeled as a convolution of the ODCs pattern with a Gaussian PSF. We modeled separate BOLD responses for the GE and SE maps. However, as these maps were obtained from the exact same region in area V1 in each subject, the model constituted one ODC map underlying both the GE and SE responses. MRI k-space sampling was modeled by restricting the spatial frequency space to its central part in accordance with the modeled field of view, sampling matrix and voxel size (Haacke et al., 1999).

All model parameters are listed in Table 1. Three parameters controlled general ODC properties: the main (peak) spatial frequency ρ , the degree of irregularity δ and the branchiness ϵ . The spatial white noise served as a high-dimensional parameter determining the specific manifestation of the ODC pattern. A parameter ω controlled the smoothness of transitions between regions showing left and right eye preference. The PSFs of the BOLD responses were parameterized by amplitudes β_{GE} and β_{SE} and by their FWHM $fwhm_{GE}$ and $fwhm_{SE}$ (the parameters of interest).

Relative to our previously published model, we made 2 slight modifications. First, to simplify derivation of the gradient we subsequently used for implementing the MCMC sampling, we modified the formulation of the band-pass filtering kernel (see Appendix A). Second, due to consideration of step size determined by the MCMC algorithm, we defined and used a smoothness parameter ω instead of using its inverse, the sharpness parameter α which we used previously.

Model implementation

The model was implemented in MATLAB (The MathWorks Inc., Natick, MA, USA). All model computations were carried out on a Cartesian grid of $0.125 \times 0.125 \text{ mm}^2$ resolution. Spatial filtering used for the ODC and BOLD PSF modeling was carried out in the frequency domain using discrete Fourier transforms. The discrete Fourier transform assumes signals to be periodic, thereby forcing opposite edges of the grid to be continuous. In order to minimize this effect on modeling, the simulated area was extended relative to the data by doubling the length of each dimension.

Prior estimation

Table 1 presents an overview of employed priors for all parameters. The spatial Gaussian white noise had an independent multivariate normal distribution with a standard deviation of 1 as its prior. Priors for the ODC model parameters were estimated from anatomical and neurophysiological data as described in the following sections.

Table 1 Model Parameters. Free parameters are probabilistic parameters. We used MCMC to sample from their joint distribution. Fixed parameters were estimated directly from the data and were held constant during MCMC sampling.

Free parameters	Description	Prior	MCMC starting value
$n_{i,j}$	Spatial white noise to be filtered, determines specific ODC pattern	Multivariate standard normal distribution	Up-sampled differential fMRI map, normalized to have a standard deviation of 1
ρ	Main spatial frequency of ODC	Limited normal distribution, parameters estimated from anatomical data	Set to mean of estimates from anatomical data
δ	Irregularity	Normal distribution, parameters estimated from anatomical data	Set to mean of estimates from anatomical data
ϵ	Branchiness	Normal distribution, parameters estimated from anatomical data	Set to mean of estimates from anatomical data
ω	Smoothness of transitions	Uniform between 0.3 and 2 based on neurophysiology data	Set to mean of limits
θ	Orientation of ODC	Flat	$\pi/2$ (corresponds to ODC bands parallel to medio-lateral direction)
$fwhm_{GE}$	GE point-spread-function width (full width at half maximum)	Flat	2 mm
$fwhm_{SE}$	SE point-spread-function width (full width at half maximum)	Flat	2 mm
Fixed parameters	Description	Estimation	
β_{GE}	GE response amplitude	Twice the median across voxels of the average between GE responses to left and right eye stimulation	
β_{SE}	SE response amplitude	Twice the median across voxels of the average between SE responses to left and right eye stimulation	
$\hat{\sigma}_{GE}^2$	GE measurement noise variance	Mean across voxels of estimation variance of differential GE response	
$\hat{\sigma}_{SE}^2$	SE measurement noise variance	Mean across voxels of estimation variance of differential SE response	

Estimation of priors from cytochrome oxidase data

Four single hemisphere images of complete patterns of ODCs in the human brain taken from Adams et al. (2007) were reanalyzed. These images were originally obtained by postmortem staining for CO activity in human subjects who had lost one eye. The goal of this analysis was to find model parameters that gave rise to modeled ODC maps whose spatial power spectra most closely resemble those of real human ODCs (Fig. 2).

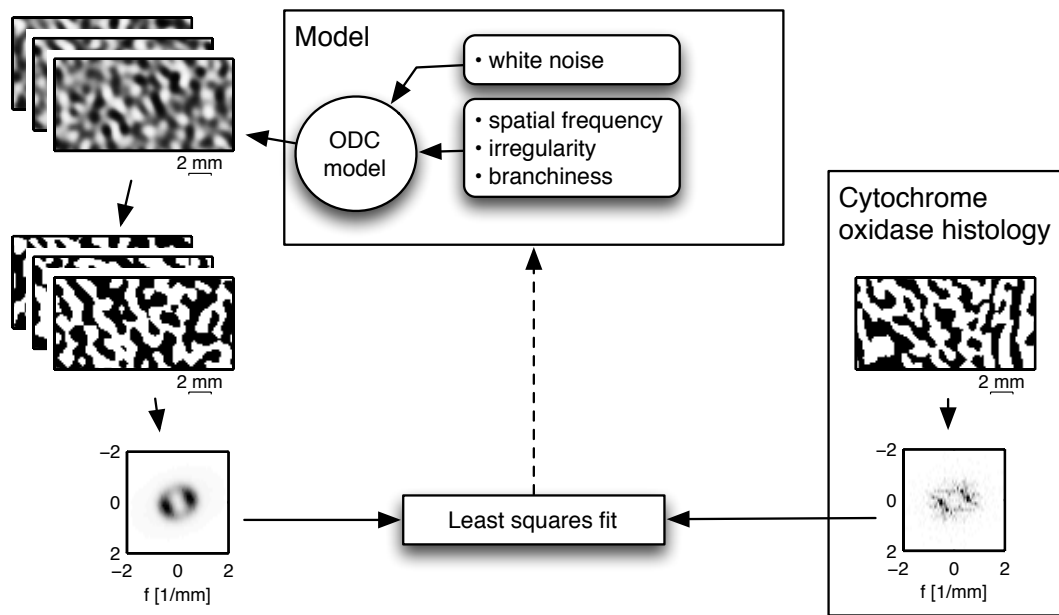


Fig. 2 Overview of cytochrome oxidase fitting. In order to constrain the parameters of the ODC model (main spatial frequency, irregularity and branchiness), cytochrome oxidase (CO) maps of human ODCs (Adams et al., 2007) were analyzed. Model parameters were optimized so that the spatial power spectra of binary ODC maps generated by the model (left) resembled those of the CO maps (right).

Two rectangular regions from each image, which corresponded in cortical location and extent to our fMRI ODC maps, were selected. To this end, V1 boundaries and the representation of the fovea were delineated. The eccentricity for each point in V1 was computed from the cortical distance to the fovea d_{fovea} as $0.75 + \exp\left(\frac{d_{fovea}}{17.3}\right)$ (Horton and Hoyt, 1991). For every point in the map, the two locations on the upper and lower V1 boundary (representing the vertical meridians) with eccentricity equal to that of the considered point were identified. The angular distances (along points with the same eccentricity) from the point under consideration to each of those two points on the boundary were calculated. The horizontal meridian was defined as the set of all V1 points for which those two distances were equal (green line in Fig. 3A).

The two regions to be selected, corresponding to the upper and lower banks of the calcarine sulcus, were then defined using the following criteria. First, the spatial extent was set to 15.7 mm x 8 mm, so that the area was equal to the mean area of our fMRI regions of interest (ROI) and the aspect ratio was equal to the mean aspect ratio of our fMRI ROIs. Second, ROIs had to be 5 mm away from the horizontal meridian and centered within an eccentricity range of 3° to 10° , corresponding to the expected location of the flat regions of the calcarine sulcus (Cheng et al., 2001).

The pattern of the CO map was binarized, in order to obtain the pattern of absolute ocular dominance (i.e. left or right eye preference). Then, for each map, we fitted the parameters of the ODC part of the model such that the spatial power spectra of the simulated binarized ODC maps were similar to those of the CO maps of ODCs (Fig. 3B, measured ODC; Fig. 3C, simulated ODC). For a model that consists of the filtering of spatial white noise only (i.e. with-

out a sigmoidal point-wise non-linearity or binarization), the frequency spectrum of the output is expected to resemble the filter shape. Here, the spatial frequency spectra of the binarized CO maps were used to obtain first estimates of the ODC filter parameters. To that end, spectra were resampled to polar coordinates and their radial and angular averages were computed. Model equations for the radial and angular filter components (see Appendix A) were fitted to these averages using the MATLAB Curve Fitting Toolbox, enabling the extraction of parameter estimates (ρ , δ , θ , ϵ). Random binary ODC patterns (400 in each step) were simulated using these estimates as initial values. Their power spectra were averaged, and the sum of squared differences between the data spectrum and the average simulated spectrum was computed. An optimization algorithm in MATLAB (fminsearch; Lagarias et al., 1998) was used to find parameters that minimized this sum of squared differences.

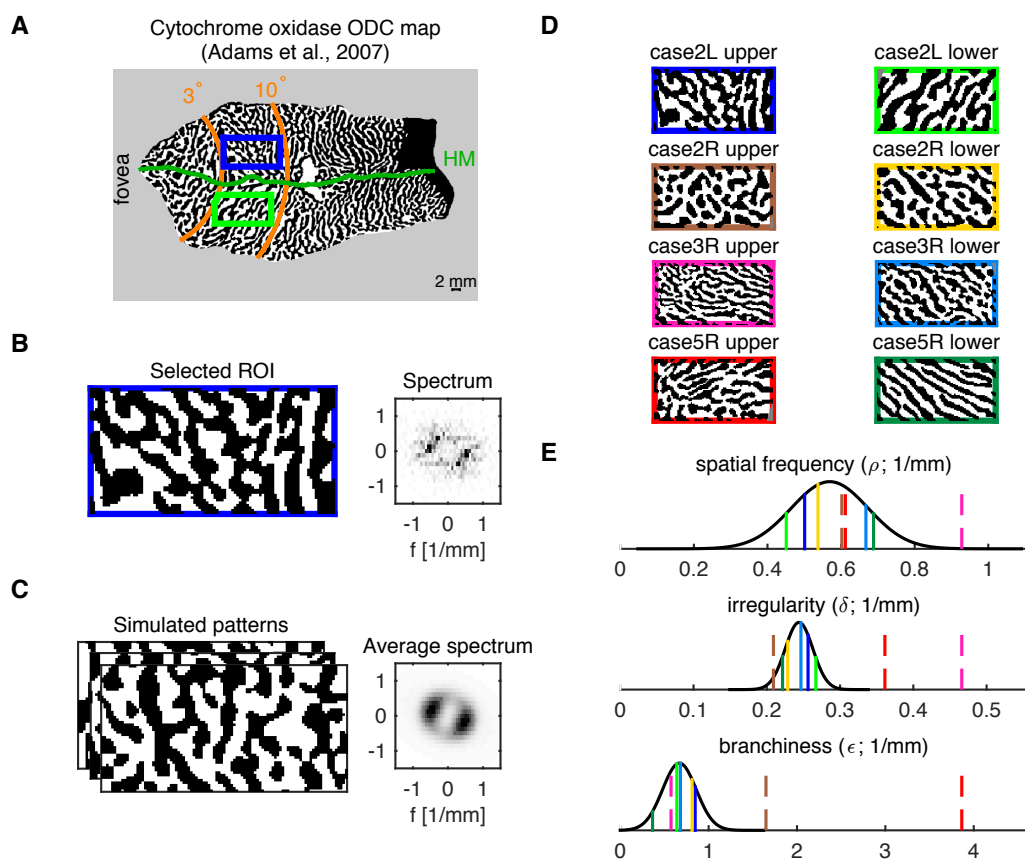


Fig. 3 Results of cytochrome oxidase ODC map analysis. A Cytochrome oxidase ODC maps from human V1 were imported from Adams et al. (2007). Regions of interest (ROIs) were selected to be comparable in size and location to our fMRI data. B The ODC pattern from the selected upper bank region (left) is shown next to its spatial power spectrum (right). C Model parameters were optimized to produce simulated patterns (left) whose average spectrum (right) was comparable to the spectrum of the data in B. D The patterns from the upper and lower bank regions of all four cases are shown with color-coded surrounds. E The estimated values of model parameters from all patterns are shown as vertical lines using the same colors as in D. Dashed lines indicate values from patterns that were classified as outliers in the distribution of spatial frequency, irregularity, or branchiness. The mean and standard deviation of the remaining values were used to define Gaussian distributions (black) to be used as priors for fitting the model to the fMRI data.

Three maps and their fitted parameters showed outlier features. For example, ‘case5R upper’ (Fig 3D, red surround) showed very thin bands in one region (upper left) immediately adjacent to a region of thicker bands (bottom right), resulting in outlier estimates of irregularity and branchiness (red lines in Fig. 3E, middle and bottom panels). Such abrupt changes may have resulted from the processing of anatomical specimens or from the presence of curved boundaries between locally flat regions.

In order to avoid atypical parameters estimates, parameter values whose absolute deviation from the median exceeded 3.7 times the median absolute deviation (corresponding to 2.5 standard deviations of a normal distribution; Leys et al., 2013) were marked as outliers. Maps in which at least one parameter was marked as an outlier were excluded from further analysis. This procedure resulted in the exclusion of three maps (cases 2R, 3R and 5R upper) from further analysis.

For each parameter, model priors were defined as normal distributions with means and standard deviations equal to the sample means and standard deviations of parameter values from the remaining maps. In order to further discourage extreme parameter values, we set the prior for ρ outside two standard deviations from the mean to zero.

Estimation of a prior for the smoothness parameter ω

A prior for the smoothness parameter ω was constructed on the basis of ocular dominance indices (ODIs) as reported in the neurophysiological literature. We assumed that ODI distributions in humans are similar to those in the macaque.

ODIs are defined as $ODI = \frac{y_{left} - y_{right}}{y_{left} + y_{right}}$, where y_{left} and y_{right} denote the response values to stimuli presented to the left and right eye respectively (e.g. Berens et al., 2008). ODIs were calculated from differential ODC maps generated by our model and fit to ODI distributions from the literature (Berens et al., 2008; Hubel and Wiesel, 1968). ω was allowed to vary while all other parameters were fixed as the mean of their anatomical data estimates. The value of ω that resulted in the smallest Kullback-Leibler divergence between the modeled and the target distribution was selected.

In order to fit modeled ODI distributions to the seven-class classification of Hubel and Wiesel (1968), classes 1 and 7 were collapsed into an exclusively responding class (left or right eye). Classes 2, 3, 5 and 6 were collapsed into an intermediately responding class, with class 4 responding indifferently. Modeled ODI distributions were also transformed into these three classes. The range of absolute ODIs that were assigned to the exclusively responding class and to the indifferently responding class were defined by a class width parameter. The values of ω and of the width parameter that together resulted in the smallest Kullback-Leibler divergence between the modeled and the target distribution were selected.

fMRI Data acquisition

7 T BOLD fMRI data from Yacoub et al. (2007) were reanalyzed. The data were obtained from three subjects in six sessions each, using GE (three sessions) and SE (three sessions) imaging. The target ROI of one subject was un-

sually and densely covered by large blood vessels. We therefore excluded the data from this subject and used the two other datasets. A single slice was imaged; it was selected such that it was parallel to and maximally overlapping with a flat gray matter region of the calcarine sulcus. The in-plane resolution was $0.5 \times 0.5 \text{ mm}^2$ and the slice thickness was 3 mm. Each run included a baseline epoch, in which a blank gray image was presented, and alternating epochs of left or right eye stimulation. Detailed descriptions of the methods used for data acquisition can be found in Yacoub et al. (2007).

fMRI Data processing

Data reconstruction

The measured k-space data were preprocessed using dynamic off-resonance in k-space (DORK) to remove respiration-induced fluctuations in resonance frequency (Pfeuffer et al., 2002). Subsequently, a Fourier transformation was applied in order to transform the data to the image space. Three datasets (subject 2, SE) were acquired using partial Fourier and were reconstructed using a homodyne reconstruction algorithm (Noll et al., 1991).

Motion correction

Residual head motion was corrected using AFNI's 3dvolreg (Cox and Jesmanowicz, 1999). This algorithm requires multiple slices; therefore, identical copies of the slice were concatenated from above and below. The additional slices were later discarded from the output of the algorithm. The reference volume in each run was set to the volume with the highest average correlation to all other volumes. Each run was motion-corrected using the two-passes option and Fourier interpolation. All volumes for which any voxel was displaced more than 1 mm relative to the reference volume were marked as motion outliers and were later excluded from the general linear model (GLM) analysis. All resulting transformation matrices were saved.

Between-run motion correction was carried out by first averaging all within-run corrected volumes from each run. Next, the series created from concatenating these single-run averages was corrected in the same manner as described above. Again, all transformation matrices were saved.

Finally, the within-run corrected data was transformed by applying the saved between-run transformations. The resulting combination of two interpolations (within- and between-runs) was created for intermediate use only. For our quantitative analysis of the PSF, only one interpolation was applied. This one interpolation accounted for all alignments and registrations of the data (see below).

Outlier volume detection

For every volume, the measured fMRI signal in each voxel was compared to the entire time-course of that voxel by computing the z-score of the measured fMRI signal relative to the voxel's time series across all other volumes. Volumes that were already marked as motion outliers were not included in this calculation. The volume under consideration was marked as an outlier volume if the

average of z-scores across all voxels was larger than 2. Outlier volumes were later excluded from the GLM analysis.

Between-days and between-modality registration

ROIs for each session were imported from our previous analysis of the data (Yacoub et al., 2007). For each modality (GE or SE), between-days registration was performed using mean intensity images averaged over all runs of a given day. First, the single-day ROIs were aligned according to their centers of mass. Next, each day's mean intensity image was cropped such that all images contained the same amount of space around their ROIs, the ROIs were in the same position, and all cropped images were of equal size. Weight masks were calculated for each day by assigning a weight of one to all voxels inside the ROI and zero to all voxels further away from the ROI than 25 mm. Voxels outside of the ROI but closer than 25 mm were assigned an intermediate weight that varied smoothly between one and zero according to the function $0.5 \left(\cos \left(\pi \frac{d_{ROI}}{25 \text{ mm}} \right) + 1 \right)$, where d_{ROI} is the shortest distance to the ROI. Out of the three days, the one whose mean intensity image had the highest average correlation to those of all other days (weighted by the mask) was selected as the reference day. Each day was registered to the reference day using FSL's flirt 2D registration without large-scale search (Jenkinson et al., 2002; Jenkinson and Smith, 2001), using the weight masks and normalized correlation as the cost function. All transformations were saved.

Initial between-modality (GE and SE) registration was carried out using the registered GE and SE images averaged over all days. A procedure similar to that used for within-modality registration was employed, except that the correlation ratio served as the cost function. For the data of subject 1, AFNI's 3dvolreg (using the same options as in motion correction) produced a better registration than FSL's flirt based on visual inspection, and was therefore used.

All registration results were visually inspected. Residual misalignments found in one day of subject 1 and one day of subject 2 were manually corrected.

Data resampling

In order to avoid smoothing of the data due to multiple interpolations, all transformation matrices (within-run motion correction, between-run motion correction, between-days within-modality registration and between-modalities registration) were combined. All unprocessed data was transformed using one single Fourier interpolation per volume using AFNI's 3drotate (Cox and Jesmanowicz, 1999).

GLM analysis

For each run, a GLM was fit to each single voxel time-course. The model consisted of a constant predictor and the two stimulation paradigms (left and right eye stimulation) convolved with a standard hemodynamic response function. Volumes that were previously determined to be outliers due to extensive head motion or imaging artifacts were excluded from the fit. Relative responses were calculated by dividing the estimated stimulus response magnitudes by the

estimated constant baseline. Differential responses were calculated as the difference between left and right eye responses. Unspecific responses were calculated as the average between left and right eye responses. In addition, standard errors of all estimates were calculated. For visualization purposes only, estimated response maps as well as modeled response maps were interpolated to $0.25 \times 0.25 \text{ mm}^2$ resolution by zero-padding in spatial frequency space and were band-pass filtered for eliminating cycles shorter than 1.25 mm and longer than 12 mm.

Multi-run and multi-day averaging

Standard errors were comparable between runs and days. Accordingly, single-day responses were calculated by averaging the GLM estimates of the single-run responses. Likewise, responses for each subject and imaging modality (GE or SE) were estimated by averaging single-day responses.

Standard errors for these averaged responses were estimated as standard errors of the mean from the distribution of single-day responses. Averaged response maps from all three days were used for further processing for subject 1. For subject 2, between-days correlation of a pair of SE sessions was significantly lower than those obtained from all other pairs in our data. We therefore averaged only the two most reproducible SE sessions (highest correlation of differential responses) and, separately, the two most reproducible GE sessions in order to achieve equal processing between SE and GE.

Optimization of between-modality registration using differential maps

Between-modality registration was further optimized. The GE and SE differential maps were shifted relative to each other vertically and horizontally by multiples of a quarter voxel up to three voxels in each direction and the set of shifts that resulted in the highest correlation between differential GE and SE maps was saved. To avoid multiple interpolations, this shift was combined with all previously found transformations (i.e. motion correction and registration) into a single transformation and interpolation. All unprocessed data were transformed over again as described above, followed by GLM analysis, multi-run and multi-day averaging.

Quantities used for MCMC fitting

Image artifacts, noise and blood vessels may result in some voxels with extreme differential responses that would have a disproportionate effect on fitting the model. For this reason, all voxels with a differential response showing absolute deviation from the median exceeding 3.7 times the median absolute deviation (corresponding to 2.5 standard deviations of a normal distribution; Leys et al., 2013) were excluded (percentage of excluded voxels in subject 1: 5.0% GE and 4.1% SE, in subject 2: 2.8% GE and 2.3% SE). Note that although this procedure may have removed some voxels with large vessel contributions, it was not meant to systematically remove all voxels with such contributions (see discussion). We then calculated the median unspecific response and the root mean square (RMS) of the differential response standard errors from the remaining voxels. We set the maximum response amplitudes β_{GE} and

β_{SE} to twice the median of the left/right averaged GE and SE responses, respectively, as defined by the model.

MCMC fitting

The posterior probability of GE and SE PSF widths given the data and priors over parameters was estimated using MCMC sampling (Fig. 1). MCMC sampling was implemented using a Hamiltonian Monte Carlo algorithm (Duane et al., 1987; see Neal, 2011 for a more recent review).

The algorithm requires input in the form of a function that computes the negative log posterior probability (the potential energy of the model; see Appendix B in the Suppl. Material) and its gradient (see Appendix C in the Suppl. Material). The log posterior probability depends on model parameter values, their prior probabilities, the data and the uncertainty of the data. The data in this sense were the maps of measured differential GE and SE responses within the ROI that were not excluded as outliers. The uncertainty of the data was characterized by the RMS of differential response standard errors, calculated separately for GE and SE. The exact form of the log posterior probability and derivations of the formulae for efficient computation of its gradient are described in the Appendix.

Two parameters determine the dynamics of parameter space exploration. The first parameter, the number of leapfrog steps per iteration, was set to a value of 20. The second parameter, the step size, was initially set to 0.005 and was adjusted adaptively so that the acceptance probability stayed close to the theoretical optimum of 0.651 (Neal, 2011). In addition, the step size was varied randomly within a range of $\pm 20\%$ to avoid periodicity in the trajectories (Neal, 2011).

Initial values used for all model parameters can be found in Table 1.

The MCMC algorithm was run for 512,000 iterations, of which every 256th sample was retained. The set of all retained samples is an approximation to the joint posterior probability distribution of all parameters given the data and the model, while taking prior distributions into account.

See the attached video, that demonstrates the initial stages of the fitting.

Analysis of MCMC sampling results

The samples of PSF widths were binned into 0.04 mm wide intervals. We then identified the bin that contained the highest number of samples, which is the maximum a posteriori probability estimate obtained from the marginal distribution for the PSF width. Highest marginal posterior density credible intervals at the 95% level were computed by selecting the narrowest intervals containing 95% of the PSF width samples.

MCMC sampling diagnostics

The quality of the MCMC sampling process was assessed by visual inspection of parameter sample traces, autocorrelation estimates of the samples traces and the Geweke diagnostic, which is a z-test for difference between sample means in the first 10% and last 50% of samples (Geweke, 1991).

Estimation of T_2/T_2^* blurring or sharpening effect

Imaging modulation transfer functions (MTF) were estimated from the last volume of each run in which the phase-encode gradients were switched off (Kemper et al., 2015). This resulted in read-out lines that were expected to vary in amplitude only, according to the phase-encode direction imaging MTF (reflecting T_2/T_2^* decay as experienced by all normally acquired volumes). First, the peak position along the read-out direction was found from the average of the absolute magnitudes computed over all read-out lines. Next, the imaging MTF was estimated by combining the absolute magnitudes of all read-out lines at the peak position into a vector.

We estimated the blurring or sharpening due to T_2/T_2^* decay as a separate effect from the effect of finite and discrete MR sampling (for more details see Chaimow and Shmuel, 2016, in preparation). To this end, the complex imaging PSF was computed by applying a discrete Fourier transform to the estimated imaging MTF.

Convolution of the original pattern with the real component of the complex imaging PSF is an approximation to the full MRI acquisition (Chaimow and Shmuel, 2016, in preparation), including the last stage of taking the absolute of the complex values obtained at the end of the reconstruction.

We computed the inverse discrete Fourier transform of the real component of the complex imaging PSF, resulting in its MTF. Two Gaussian functions with zero means were separately fitted to the MTF of the real component of the complex PSF and to its inverse. Then, we compared the goodness of fit (R^2) obtained by the two fitted Gaussians. If the better fit was obtained by fitting a Gaussian to the MTF of the real component of the complex PSF, the effect of T_2/T_2^* could be described as Gaussian blurring. If the better fit was obtained by fitting a Gaussian to the inverse of the MTF of the real component, the effect of T_2/T_2^* could be described as a sharpening that could reverse a specific Gaussian blurring.

We therefore computed the FWHM of the Gaussian with the better goodness of fit (obtained by fitting to either the MTF of the real component or to its inverse).

Note that in the case of sharpening, the computed FWHM characterizes the Gaussian ‘used’ for blurring which would be reversed by the sharpening effect of the T_2/T_2^* decay. FWHM estimates for each modality were first averaged over all runs of each session (day) and then over all sessions of each subject.

Inclusion of T_2/T_2^* blurring in the model

A version of our model that included the effect of T_2/T_2^* decay was fitted to our data. Depending on whether the T_2/T_2^* decay effect resulted in blurring or sharpening, the BOLD MTF (Appendix A, BOLD response) was changed to: $MTF_{k,l}(fwhm, \beta) = \beta \cdot e^{-2\pi^2 \sigma_{BOLD}^2 r(k,l)} \cdot e^{-2\pi^2 \sigma_{IMG}^2 l^2}$ (for blurring) or $MTF_{k,l}(fwhm, \beta) = \beta \cdot 1/e^{-2\pi^2 \sigma_{BOLD}^2 r(k,l)} \cdot e^{-2\pi^2 \sigma_{IMG}^2 l^2}$ (for sharpening), where $\sigma_{IMG} = fwhm_{IMG}/2\sqrt{2\log 2}$ and $fwhm_{IMG}$ is the estimated FWHM of the Gaussian blurring kernel that models the effect of T_2/T_2^* decay. We assumed

the second dimension (associated with index l) to be the phase-encode dimension.

Results

Our goal was to fit a probabilistic generative model to maps of ODCs obtained with GE- and SE-based BOLD fMRI (Fig. 1). We aimed to infer posterior probability distributions of model parameters, specifically the width of the GE and SE BOLD fMRI PSF.

Parameter priors obtained from real human ODC

The ODC imaging model (Chaimow et al., 2011) consisted of simulating realistic ODCs by the filtering of spatial white noise (Rojer and Schwartz, 1990) followed by a spatial BOLD response and MRI k-space sampling.

Before we fitted our model to fMRI data, we determined priors for the ODC parameters by incorporating statistical information obtained from real human ODC patterns (Fig. 2). To this end, we analyzed CO maps of ODCs from human V1 taken from Adams et al. (2007).

It should be noted that CO labeling intensities are expected to provide a fairly accurate estimate of the preferred eye. However, there are multiple, potentially non-linear transformations between neuronal activity, staining intensity and the final processed image. These make it unlikely that the CO intensities quantitatively reflect the relative ocular dominance. Therefore, we only used binarized versions of these maps, thresholded to represent the absolute preference to either left or right eye stimulation). We eventually determined priors on ocular dominance from neurophysiological recordings (Berens et al., 2008; Hubel and Wiesel, 1968).

We first restricted the maps of the entire V1 to small regions (Fig. 3A) whose size and location were similar to those of our fMRI data (with origins in flat regions of the calcarine sulcus). Then, for each map, we fitted the parameters of the ODC part of the model such that the spatial power spectra of the simulated binarized ODC maps were similar to those of the CO maps of ODCs (Fig. 3B, measured ODC; Fig. 3C, simulated ODC). Simulated maps generated using these parameters looked qualitatively similar to the true CO maps (Fig. 3C). The set of all imported maps is shown in Figure 3D and the estimated parameters from all maps are shown color-coded in Figure 3E.

For each parameter we defined Gaussian priors that fit the distribution of all remaining parameter estimates (black curves, Fig. 3E). In particular, the prior for the main pattern frequency ρ had a mean of 0.57 cycles/mm with a standard deviation of 0.1 cycles/mm, which corresponds to an average column width of 0.87 mm.

Smoothness of ODC maps

In order to construct a prior for the smoothness parameter ω we analyzed distributions of ocular dominance indices (ODIs) as reported in the neurophysiological literature. ODIs quantify the relative contributions of each eye to meas-

ured responses, and their distribution is tightly linked to the smoothness parameter ω . Small values of ω result in sharp transitions between columns associated with ODIs close to +1 or -1. Large values of ω result in smooth transitions, with few locations reaching absolute monocular responses and most ODIs being close to 0.

We analyzed ODI distributions taken from Hubel and Wiesel (1968) and Berens et al. (2008) by fitting ODI distributions computed from our model as a function of smoothness ω . We found a value of $\omega = 1.5$ to best explain ODI distributions corresponding to the data in Berens et al. (2008), whereas the data in Hubel and Wiesel (1968) were best fitted with $\omega = 0.36$. Both datasets came from macaque monkeys. Berens et al. (Berens et al., 2008) used multi-unit activity, a measure whose ODIs are expected to be blurred relative to single neuron responses and are therefore expected to match a higher ω . Data from Hubel and Wiesel (Hubel and Wiesel, 1968) presented single-unit responses but were less quantitative. We therefore chose a uniform prior distribution for ω , limited by 0.3 from below and 2 from above, effectively reflecting the range of uncertainty associated with ω .

GE and SE BOLD maps of ODC

Having constructed a generative model with realistic priors, the next step was to process the fMRI data and to extract all quantities needed to fit the model. We reanalyzed fMRI data from two subjects (Yacoub et al., 2007) using a general linear model (GLM) to estimate responses to left and right eye stimulation (Fig. 4). The single-eye response maps were dominated by global unspecific responses and superimposed band-shaped modulations (Fig. 4A).

We separated these two components by first calculating the voxel-wise difference between left and right eye responses, yielding the differential ODC maps (Fig. 4C). Here, the band-shaped organization is clearly visible. The range of differential contrasts as defined by their standard deviation was 1.8% (GE) and 1.5% (SE) for subject 1, and 1.0% (GE) and 1.0% (SE) for subject 2.

In addition, we calculated voxel-wise averages of left and right eye responses (Fig. 4B). According to our model, which assumes antagonistic patterns of neuronal responses, this average response is expected to be independent of the local ocular preference. Furthermore, it is expected to be equal to a spatially homogeneous response with half the amplitude of the highest possible ocular dominance (with no response to the non-preferred eye).

We calculated the median of this left/right average response over all voxels. It was 3.0% (GE) and 1.9% (SE) for subject 1, and 3.7% (GE) and 2.0% (SE) for subject 2. In accordance with the model, we then set the amplitudes of the model PSFs to twice these values.

Finally, we estimated the measurement noise level of the differential maps as the root mean square (RMS) of all standard errors estimated by the GLM (Fig. 4D). It was 0.9% (GE) and 0.9% (SE) for subject 1, and 0.6% (GE) and 0.8% (SE) for subject 2.

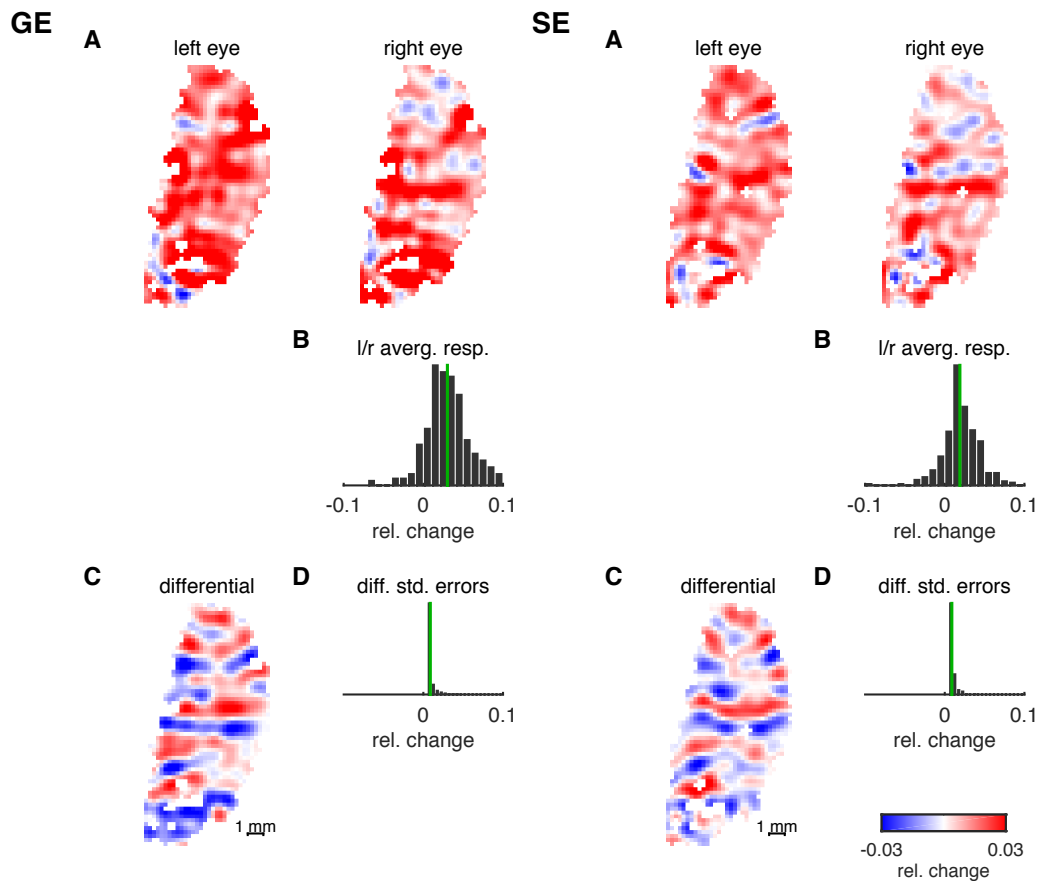


Fig. 4 fMRI ODC data. Results from the GLM analysis of fMRI data from subject 1 for GE (left) and SE (right). A Responses to left and right eye stimulation relative to baseline. B The response maps to the left and right eyes from A were averaged. B shows the distribution of the average response. Its median (in green) was used to set the overall amplitude of the BOLD response model. C The difference between left and right eye responses yields the differential ODC map. D The distribution of standard errors of all differential responses. From this distribution we estimated the noise level used by the model. The color look-up-table applies to all response maps.

Estimation of GE and SE point-spread widths

We went on to estimate the probability distributions of GE and SE PSF widths given our data. Theoretically, this requires integrating the posterior probability distribution of model parameters over all other parameters (including the high-dimensional spatial noise parameter). However, exact integration over this high dimensional space is not feasible. We therefore used MCMC to sample from the posterior probability distribution. Every sample contains all the parameters necessary to simulate one anatomical ODC map and the resulting GE and SE fMRI differential maps. The algorithm draws parameter samples with a probability proportional to how likely these parameters are to have generated the measured data given all the priors.

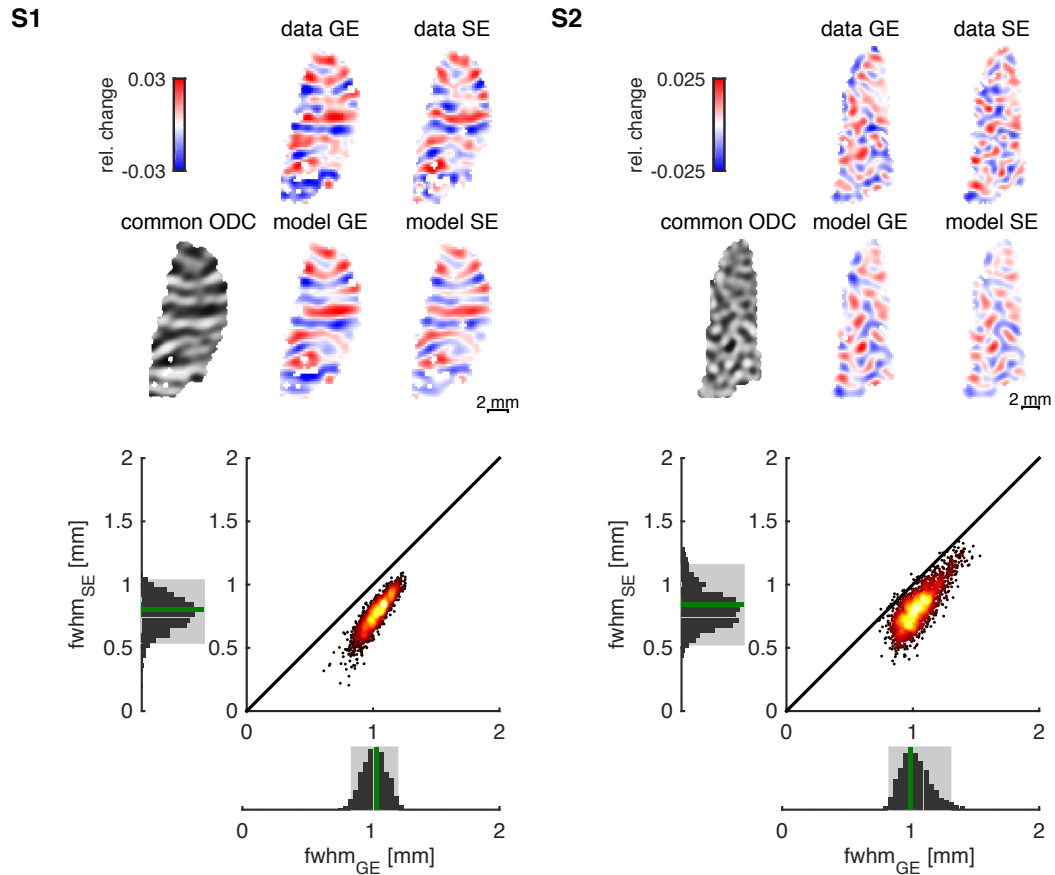


Fig. 5 Results of point-spread width estimation. The probability distribution of PSF given the data was estimated using Markov Chain Monte Carlo Sampling. The model simulated GE and SE maps each with their own BOLD parameters and with a common underlying ODC map. Results are shown for both subjects. The first row shows the measured differential ocular dominance map from the GE (left) and SE (right) experiments. The second row shows the modeled underlying ODC maps (left) and the modeled differential fMRI maps from the maximum a posteriori (MAP) sample. The bottom part of the figure shows the joint and marginal distributions of GE and SE point-spread full-widths at half-maximum (FWHM). The gray rectangles show the 95% credible intervals (highest posterior density interval). The scatter plots show that the vast majority of individual GE PSF samples were wider than their SE PSF counterparts. The MAP estimates (green bars) obtained from the marginal distributions of the FWHMs of the PSFs were 1.04 and 1.0 mm (GE), and 0.8 and 0.84 mm (SE) for subjects 1 and 2 respectively.

Figure 5 (second row, common ODC) shows one of many possible ODC patterns generated by our model. It was generated using the parameter sample with the highest posterior probability. Differential BOLD fMRI maps modeled as arising from this shared ODC pattern (Fig. 5 second row, model GE and model SE) resemble the data closely (Fig. 5 first row, data GE and data SE). The distribution of PSF widths from all samples (Fig. 5 bottom) is an estimate of the true probability distribution of PSF widths for that data (see the attached video that demonstrates the initial iterations of the fitting procedure).

Table 2 Summary of estimated PSF full-widths at half-maximum (FWHM). Maximum a posteriori estimates (MAP) with 95% credible intervals (CI), means with standard deviations and medians were calculated from MCMC samples of the marginal posterior distributions of GE and SE PSFs FWHM and from the distribution of their differences. The results labeled with 'Accounting for T_2/T_2^* decay' were obtained from fitting a separate model, which accounted for the effect of these decays, to the data.

	Subject 1			Subject 2			Mean over Subjects		
GE PSF (FWHM)	MAP [95% CI]	1.04 mm [0.85 mm...1.21 mm]		1.00 mm [0.84 mm...1.31 mm]		1.02 mm			
	Mean \pm SD	1.03 mm \pm 0.09 mm		1.04 mm \pm 0.12 mm		1.04 mm			
	Median	1.03 mm		1.03 mm		1.03 mm			
Accounting for T_2^* decay	MAP [95% CI]	1.04 mm [0.88 mm...1.23 mm]		1.04 mm [0.86 mm...1.34 mm]		1.04 mm			
	Mean \pm SD	1.06 mm \pm 0.09 mm		1.08 mm \pm 0.12 mm		1.07 mm			
	Median	1.06 mm		1.06 mm		1.06 mm			
SE PSF (FWHM)	MAP [95% CI]	0.80 mm [0.54 mm...1.03 mm]		0.84 mm [0.53 mm...1.15 mm]		0.82 mm			
	Mean \pm SD	0.77 mm \pm 0.13 mm		0.82 mm \pm 0.16 mm		0.80 mm			
	Median	0.78 mm		0.81 mm		0.80 mm			
Accounting for T_2/T_2^* decay	MAP [95% CI]	0.76 mm [0.49 mm...1.01 mm]		0.68 mm [0.41 mm...1.14 mm]		0.72 mm			
	Mean \pm SD	0.76 mm \pm 0.14 mm		0.74 mm \pm 0.18 mm		0.75 mm			
	Median	0.76 mm		0.73 mm		0.75 mm			
Difference GE - SE	MAP [95% CI]	0.24 mm [0.14 mm...0.38mm]		0.24 mm [0.05 mm...0.42 mm]		0.24 mm			
	Mean \pm SD	0.26 mm \pm 0.06 mm		0.22 \pm 0.10 mm		0.24 mm			
	Median	0.25 mm		0.22 mm		0.24 mm			
Accounting for T_2/T_2^* decay	MAP [95% CI]	0.28 mm [0.17 mm...0.43 mm]		0.32 mm [0.11 mm...0.55 mm]		0.30 mm			
	Mean \pm SD	0.30 mm \pm 0.07 mm		0.34 mm \pm 0.11 mm		0.32 mm			
	Median	0.30 mm		0.34 mm		0.32 mm			

Figure 5 (bottom) and Table 2 present the results of PSF widths. For subject 1, the most probable (maximum a posteriori estimate obtained from the marginal distribution of FWHMs) GE PSF width was 1.04 mm (FWHM), with 95% of the values falling between 0.85 mm and 1.21 mm. The most probable SE PSF width was 0.80 mm, with 95% of the values falling between 0.54 mm and 1.03 mm. For subject 2, the most probable GE PSF width was 1.00 mm, with 95% of the values falling between 0.84 mm and 1.31 mm. The most probable SE PSF width was 0.84 mm, with 95% of the values falling between 0.53 mm and 1.15 mm.

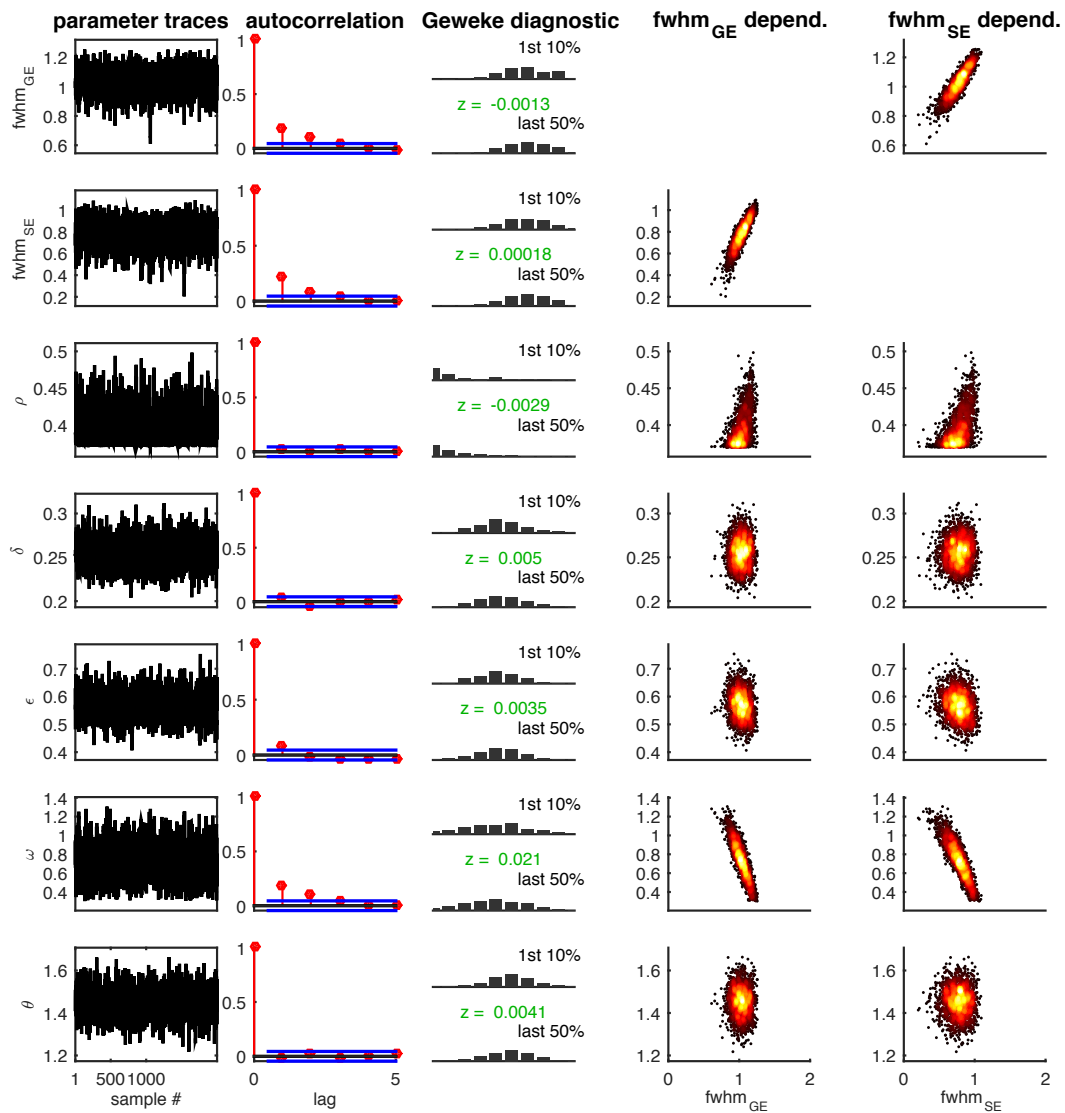
Furthermore, the samples of GE and SE PSF widths were correlated. This means that ODC model parameters that resulted in a relatively higher GE PSF width also resulted in a relatively higher SE PSF width. Across all modeled underlying anatomical ODC patterns, the GE PSF was almost always wider than the SE PSF. We calculated the resulting posterior distribution of differences between GE and SE PSF widths. The bottom part of Table 2 summarizes the estimated differences for the two subjects. For subject 1, the most probable difference was 0.24 mm, with 95% of the values falling between 0.14 mm and 0.38 mm. The most probable difference obtained for subject 2 was 0.24 mm, with 95% of the values falling between 0.05 mm and 0.42 mm.

Evaluation of model fit

The validity of our results depends on how well the MCMC samples approximate the target distribution. The MCMC sampling distribution approaches the target distribution when the number of iterations goes to infinity (e.g. see Neal, 1993). For sufficiently large number of iterations, MCMC effectively samples from the target distribution.

While there cannot be proof that the target distribution has been reached, there are a number of indications that are considered reliable. The first is that the traces of samples of all parameters have settled into a stationary distribution, with no slow drifts over iterations. This can be seen in the single parameter trace plots (Fig. 6, first column) and their autocorrelation plots (Fig. 6, second column). In addition, the Geweke diagnostic (Geweke, 1991) shows that for all single parameters the mean of the first 10% of the samples was not significantly different from the last 50% of the samples ($|z| < 1.96$). The Geweke diagnostics for the high-dimensional noise follow a standard normal distribution (Fig. 6, bottom, distribution of z-scores), as would be expected by chance under the hypothesis that the means are not different. Figure 6 also shows the dependences between PSF widths and ODC parameters (last two columns). As can be seen, higher levels of smoothness parameter (ω) values and to a lesser extent lower levels of the main spatial frequency parameter (ρ) values made a narrower PSF more likely.

A Parameters



B Noise values

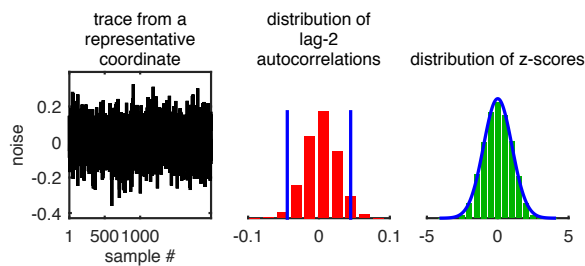


Fig. 6 Convergence diagnostics of Markov Chain Monte Carlo sampling. Markov Chain Monte Carlo needs to run for a sufficient number of iterations in order to yield samples from the modeled probability distribution. Indications for convergence are: (1) stationarity of the parameter sampling distributions, and (2) sample autocorrelations decrease rapidly with increasing lag, relative to the total number of samples. This figure examines convergence for subject 1. The upper part of the figure (A) shows diagnostics for the standard model parameters. The bottom part (B) shows diagnostics for the white noise values that act as parameters to determine the ODC pattern. The first column (A and B) shows traces of the sampled parameters. For the noise values (B), one exemplary trace is shown from the center of the map. The second column (A and B) shows sample autocorrelations as a function of lag. The horizontal blue lines (A) indicate the 95%-confidence bounds around 0 for a white noise process. Consecutive samples (lag=1) show low autocorrelation. However, samples of lag 2 (and higher) show autocorrelation esti-

mates that are comparable to those obtained from uncorrelated white noise. In B, the autocorrelation from the noise samples are summarized by the histogram of lag-2 autocorrelations from all coordinates together. Here, 95%-confidence bounds for a white noise process are indicated by vertical blue lines. The third column (A and B) presents the Geweke convergence (stationarity) diagnostic, which is a z-test (z-scores shown in green) for testing whether the means of the first 10% and last 50% of samples are different. In A, 2 histograms per each parameter show how similar their respective distributions are. In B, the z-scores from the noise samples are shown as a histogram together with a blue plot of the standard normal probability density representing the null-hypothesis of $z=0$. The last two columns (A) show the sample covariation between each parameter (vertical axis) and the GE and SE point-spread function FWHM (horizontal axis).

Discussion

Possible confounds: our estimates are upper bounds of BOLD fMRI spatial specificity

The PSF widths that we estimated (1.02 mm for GE BOLD, 0.82 mm for SE BOLD) reflect the realistically achievable spatial specificity of BOLD signals at ultra-high field strength (7T). However, they are only upper bounds for the true BOLD PSF widths. Subjects' head motion, data interpolation and intra-acquisition T_2/T_2^* decay can all introduce additional blurring (but see section below on the effect of T_2/T_2^* decay), causing the estimated PSF to be wider than the true PSF.

In order to minimize head motion, data was acquired from trained subjects using a bite bar. Before each scan, the position of the region of interest (ROI) was checked and the slices repositioned if necessary. We corrected the data for residual head motion and discarded any problematic volumes. We aligned data from multiple days and checked the alignment carefully. In order to further optimize between-modality registrations, we also took the differential fMRI response patterns into account, making use of the fact that they emerged from the same underlying neuronal ODC pattern.

Motion correction and between-day registration required spatial interpolation of the data. We minimized any blurring effects by applying all spatial transformations combined using one single Fourier interpolation (Cox and Jesmanowicz, 1999).

All high spatial resolution BOLD fMRI experiments will be influenced by these effects to a similar degree as ours, making our reported PSF widths good estimates for the practically relevant compound effect.

Possible confounds: the contribution of the imaging PSF to the total BOLD fMRI PSF

In addition to the effects of the hemodynamic and metabolic responses on the spatial specificity of fMRI, the MRI acquisition process influences the effective resolution of the acquired images. Specifically, the sampling of k-space by means of temporal gradient encoding defines the spatial resolution. However, the effective spatial resolution along the phase encoding direction in EPI acquisitions can be subject to blurring or sharpening, because of T_2/T_2^* decay while the k-space is being sampled. This can potentially contribute to the overall measured spread of the BOLD fMRI signal.

In order to minimize this effect, our data were acquired using a reduced field-of-view (in SE) and multiple segments. These measures limited the total read-out duration per segment (25.6 ms for GE and 24 ms for SE) to approximately the T_2^* of the tissue (Uludağ et al., 2009) and are expected to result in only minor blurring or sharpening (Haacke et al., 1999)).

We estimated the blurring or sharpening and their contributions to the total BOLD PSF. In general, MRI data acquisition using EPI has two distinct effects on the effective spatial resolution. The first is the effect of the finite and discrete MR sampling with no decay. However, in the current study MR sampling was part of the model and therefore has already been accounted for. The second effect is the already mentioned T_2/T_2^* decay blurring or sharpening. This effect is limited to the phase encode direction (vertical direction in all presented maps). In order to characterize it, we estimated the imaging modulation transfer functions (MTF) along the phase encoding direction from a reference volume obtained in each run, in which the phase-encode gradients were switched off (Kemper et al., 2015).

Using a model of Gaussian convolution and MRI sampling (Chaimow and Shmuel, 2016, in preparation) we obtained Gaussian functions that can model the separate effect of T_2/T_2^* decay. We then used MCMC to fit a separate version of our model to our data, where we applied the decay effect to the simulated ODC patterns by modulating the values acquired in the simulated k-space. The results presented in Table 2 show that the effect of the signal decay on the total BOLD PSF was small. For GE, while accounting for the signal decay, we obtained PSFs wider than the effective PSF obtained directly from the BOLD fMRI response. This indicates that signal decay in the GE fMRI used for obtaining our data has a sharpening effect. In contrast, the signal decay in the SE fMRI used for obtaining our data has a blurring effect. These results demonstrate that the physiological BOLD response measured with GE fMRI (that theoretically does not include signal decay) is less spatially specific than the same physiological BOLD response measured with SE fMRI (with no signal decay). This difference in spatial specificity of the GE and SE BOLD responses (with no signal decay) is even slightly larger than the corresponding effective difference obtained from the overall measured fMRI responses with signal decay.

What do our estimated point-spread function widths describe?

The BOLD PSF describes the spatial specificity of the BOLD fMRI signal by characterizing the spatial response that would be elicited by a small point stimulus. Specifically, our BOLD PSF width measures the spread of the BOLD fMRI response (I) elicited by a small spot of neuronal activity, (II) along the cortical manifold, (III) using a differential response analysis, (IV) assuming that in a differential analysis paradigm the average spread can be described by a Gaussian function, and (V) considering a relatively long time scale.

(I) BOLD PSF relative to the local neuronal activity

To the best of our knowledge, our PSF estimates are the first to quantify the BOLD spread in human subjects relative to local neuronal activity. We previ-

ously estimated the FWHM of the 7 T GE BOLD PSF to be smaller than 2 mm by measuring the spread of the BOLD fMRI response around the V1 representation of edges of visual stimuli (Shmuel et al., 2007). We expect that our previous estimates as well as others' (Engel et al., 1997; Parkes et al., 2005) included contributions from non-zero extent of receptive fields and the scatter of receptive field position of neurons in V1.

Hubel and Wiesel (1974) reported that in the macaque "... a 2 mm \times 2 mm block of cortex contains the machinery needed to analyze a region of visual field roughly equal to the local field size plus scatter".

These observations suggest that visual stimuli will result in neuronal activity that is blurred on the surface of human V1. All PSF widths that have been estimated using spatial representations of visual stimuli included this neuronal spread by nature of their experimental design. In the current study, we instead used a spatial structure of neuronal responses that is inherent to the cortex—ODC patterns. This allowed us to estimate a PSF that does not contain contributions from the spatial spread of responses to visual stimuli.

There are a number of measures of neuronal activity that a BOLD PSF could potentially relate to, notably single-unit activity (SUA), multi-unit activity (MUA) and local field potentials (LFP). Under specific circumstances, these measures can show very different activity. Under most conditions, however, they are highly correlated. This is likely to be true when mapping a cortical columnar organization. The main difference is that the spatial extent (that influences the smoothness of the spatial response pattern) of these signals increases from SUA to MUA to LFP. We estimated a smoothness prior using ODI distributions of SUA and MUA. Consequently, our PSF is based on these signals. The BOLD PSF from LFP would be narrower than our estimate because of the wider cortical spread of LFP compared to MUA activity (Xing et al., 2009).

(II) Spatial BOLD response along the cortical manifold

It has been demonstrated (Polimeni et al., 2010) that the PSF consists of different radial and tangential components relative to the cortical surface. The radial component describes the spread across cortical layers while the tangential component describes the spread parallel to the cortical surface. Here we investigated the tangential PSF, averaged over all layers. This is the component that is most relevant for imaging the representation of cortical columns parallel to the cortical surface. Accordingly, the location and orientations of voxels, the ROI, and the voxel size were all optimized to sample gray matter tangentially and to obtain an average from all layers.

It should be noted that there are some differences in cerebrovascular organization with respect to radial and angular direction (Duvernoy et al., 1981). The largest blood vessels are the pial surface veins that extend in various orientations along the tangential plane. Somewhat smaller are cortical-penetrating veins that are organized radially, traversing the different cortical layers. The smallest vessels, the capillaries, form a fine mesh that locally appears to be isotropic. However, their density varies with cortical layers (Weber et al., 2008). For these reasons, we cannot directly apply our PSF to the imaging of cortical layers. In addition, the distinctiveness and finite extent of layers

appear to make a PSF convolution model ill-suited for fMRI of cortical layers. However, some recent results (De Martino et al., 2015; Fracasso et al., 2016; Muckli et al., 2015; Olman et al., 2012) suggest it is possible to differentially resolve layer-specific signals on the scale of 1 mm or less.

(III-IV) On modeling the average differential BOLD response as a Gaussian PSF We assumed the average (over space) PSF to be a Gaussian function. However, the shape of the spread in specific cortical locations may be more complex and location-dependent (Kriegeskorte et al., 2010; Polimeni et al., 2010). Also, its width as well as the magnitude of the response may vary due to local variations in vascular geometry. In fact, the relatively wide distribution of average responses in our data (Fig. 4, distribution of l/r averg. resp.) supports this latter intuition. Therefore, a convolutional model with a single Gaussian function can only be an approximating simplification. Nevertheless, we believe that such a simplifying approach provides a useful approximation for planning and interpretation of high-resolution fMRI studies and for quantitative modeling.

As part of our pre-processing before fitting the model to the data, we removed 1.5 - 5.0 % of voxels that had extreme differential values (see methods section for precise criterion for exclusion). Part of these voxels were located in areas that were previously shown to contain blood vessels (Shmuel et al., 2010). However, for our current analysis we did not explicitly and systematically remove voxels that were affected by larger blood vessels. Our reasoning was that a consistent removal of all voxels suspected to be influenced by larger blood vessels would have reduced contiguous areas of ODCs, which would have made the model fitting more difficult.

We expect the influence of geometric variations in local vasculature to be higher for veins and venules than for capillaries because of their respective diameters and densities. Consequently, the GE BOLD signal, which is more sensitive to larger pial surface veins will be more affected by these local variations. As a result, GE BOLD imaging does not only suffer from a slightly wider PSF than SE BOLD, but it is also subject to local distortions when larger blood vessels are present.

However, although draining veins may show responses with a preference to a subset of features encoded in a columnar organization (Shmuel et al., 2010), differential analysis reduces contributions from macroscopic vessels because of their tendency to drain blood from a region larger than that of a small number of columns. Taken together, a Gaussian PSF model by itself is likely not a good model for single-condition imaging when influenced by large blood vessels (e.g. in GE BOLD imaging). In contrast, we expect that a Gaussian PSF is a good model in a differential analysis paradigm, which reduces contributions from macroscopic vessels. The BOLD PSFs we report here reflect the spatial specificity that can be achieved in a *differential* paradigm. They do not reflect the spatial specificity expected from *single-condition* imaging that involves contributions from macroscopic vessels, such as single-condition GE fMRI and to a lesser extent, single-condition SE fMRI.

(V) Spatial specificity as a function of stimulus duration

It has been shown that the early phase of the positive BOLD response (up until ~ 4 s after stimulus onset) is spatially more specific than the later phase (Goodyear and Menon, 2001; Shmuel et al., 2007). On the other hand, stimulation paradigms that use very brief stimulation durations suffer from a highly reduced contrast-to-noise ratio, because the response does not develop to its highest potential amplitude.

We found previously that after 4 s the spatial BOLD response remained stable and that the entire spatiotemporal response could be well approximated by the first spatial principle component (Shmuel et al., 2007). Aquino et al. (2012) modeled the BOLD response as a travelling wave evolving in time and found that deconvolution of neural dynamics using such a model resulted in physiologically more plausible spatiotemporal patterns than when using a model separable in space and time (Aquino et al., 2014). The spatial profile alone, however, was very similar for both models.

Taken together, long stimulation paradigms are an efficient way of high-resolution imaging and their spatial PSF can be well described by a single time-independent component. The stimulation periods for our data were 48 s long, thereby making our PSF most applicable to long stimulation paradigms.

Spatial specificity of the BOLD response

Constraints on the spatial specificity of BOLD

The positive BOLD signal depends on decreases in deoxyhemoglobin content in the capillaries which then propagate downstream to draining venules and veins. These decreases are caused by elevated cerebral blood flow (CBF) and only smaller fractional increases in the oxygen consumption rate, following increases in neuronal activity. CBF is regulated at a sub-millimeter scale: (Duong et al., 2001). Similarly, Vazquez et al. (2014) reported a spread of cerebral blood volume (CBV) of $103 - 175 \mu\text{m}$ (FWHM) in mice using optical imaging. Although this measure is not directly comparable to the CBF spread in a different species (human subjects), it demonstrates that hemodynamic signals can show very high spatial specificity. The CBF response is the ultimate lower limit for the spatial specificity of any BOLD-based technique.

The deoxyhemoglobin content changes in the draining venules and veins are ultimately diluted downstream, because the draining veins pool blood not only from active but also from non-active regions. For an activated area of 100 mm^2 , Turner et al. (2002) estimated the maximal extent of undiluted oxygenation changes along a vein to be 4.2 mm. For these reasons, we can expect the PSF width of any BOLD-based imaging technique to fall in this range; that is, less than 1 mm (Duong et al., 2001) to approximately 4.2 mm (Turner, 2002). The values will be determined by how much weighting towards the microvasculature can be achieved and on the actual presence of larger draining veins in the region of interest.

PSF dependence on field strength

At standard magnetic fields, the width of the BOLD PSF has been estimated to be 3.5 mm for 1.5 T GE BOLD (Engel et al., 1997), 3.9 mm for 3 T GE BOLD

and 3.4 mm for 3 T SE BOLD (Parkes et al., 2005). These estimates of PSF widths were confounded by the above described receptive field and scatter effects. We can make a rough estimate of what the non-confounded PSF widths at lower fields would be. We assume that on average the receptive field effect can be modeled as another convolution with a Gaussian. It follows that the square of the confounded PSF width is equal to the sum of squares of the receptive field effect width and the non-confounded PSF width. For the receptive field effect we get an FWHM of 2.12 mm when using 2.35 mm as the 7 T GE BOLD confounded PSF width (Shmuel et al., 2007) and 1.02 mm as the corresponding non-confounded PSF width (results from our current study). This in turn results in non-confounded estimates of 2.8 mm (1.5 T GE BOLD), 3.3 mm (3 T GE BOLD) and 2.7 mm (3 T SE BOLD).

These PSF widths are considerably larger than the estimates from the current study (1.02 mm for 7T GE BOLD, 0.82 mm for 7T SE BOLD). The reason for this is that the BOLD signal (both GE and SE BOLD) at lower field strengths is dominated by intravascular signals from draining veins (Jochimsen et al., 2004; Uludağ et al., 2009). At higher field strengths, the contributions from intravascular signals are reduced due to a shortening of the venous blood T_2 . In parallel, the relative contributions of extravascular signals around small vessels increase (Duong et al., 2003; Uludağ et al., 2009; Yacoub et al., 2003; 2001).

All PSF widths from field strengths of up to 3 T appear to fall close to the wider end of possible PSF widths. In contrast, PSF widths using SE and GE at 7 T appear close to their theoretical minimum.

T_2^* and T_2 based imaging methods: GE, SE and GRASE

We found the SE BOLD PSF to be narrower than the GE BOLD PSF. This is expected because the refocusing pulse in SE imaging suppresses the extravascular signal around larger blood vessels while leaving the signal around the microvasculature intact. As a result, compared to GE BOLD fMRI, SE BOLD signals obtained at 7T have relatively larger contributions from the spatially more specific microvasculature, whereas at lower field strength the signal of either SE or GE BOLD fMRI is dominated by intravascular contributions of large blood vessels.

However, the suppression of extravascular signal around larger blood vessels by SE at high fields is not perfect. Only the k-space data that is sampled at the exact echo time will result in absolute suppression (pure T_2 weighting as compared to T_2^* weighting). The extent to which sampled k-space data is affected by T_2^* weighting increases with increasing total read-out time. Consequently longer total read-out times in SE result in decreased spatial specificity

(Goense and Logothetis, 2006) and are expected to have a wider point-spread function (though still narrower than GE).

Other T_2 based functional imaging methods such as GRASE (Oshio and Feinberg, 1991) and 3D-GRASE (Feinberg et al., 2008) are expected to have similar spatial specificity as SE. Whether their PSFs are slightly wider or narrower will mainly depend on the T_2^* weighting component associated with such methods (i.e. echo train lengths of gradient recalled echoes employed in between successive 180° pulses), in addition to their T_2 component. In fact, Kem-

per et al. (2015) have reported that 3D-GRASE had a smaller bias towards pial surface veins owing to the smaller T_2^* contribution when a reduced field of view is employed in zoomed 3D-GRASE compared to the longer in-plane echo-train of 2D-SE EPI.

Although we found a wider PSF for GE BOLD than for SE BOLD fMRI, the difference was relatively small (1.02 mm for GE BOLD, 0.82 mm for SE BOLD). We believe that this is due to the fact that the influence of larger blood vessels can be reduced by using a differential imaging paradigm, even when using 7T GE BOLD fMRI. Consequently, both GE and SE BOLD imaging techniques seem capable of resolving cortical columns when applying differential imaging analysis.

However, GE maps are more susceptible to confounds introduced in voxels containing blood vessels which may not be fully suppressed in differential imaging. Therefore, obtaining results of high spatial specificity using GE depends on the region of interest and on methods to mask out blood vessels.

SE is less susceptible to large-vessel confounds, that may not be suppressed by differential imaging. The response amplitude of SE is lower than that of GE. However, for imaging of highly granular structures such as ODC's at such high resolutions, the *differential* contrast is similar for GE and SE fMRI. Overall, we believe that SE is the method of choice for mapping finer structures, especially when relying on single-condition analysis. However, which data acquisition method is optimal depends on the goal of the study and the spatial scale of the neuronal architecture under investigation.

The application of probabilistic models of cortical columns and MR imaging

We have extended our quantitative model for imaging ODCs to a probabilistic generative model and used it to infer the PSF widths by means of MCMC sampling.

A critical component to the successful application of MCMC to our model is the Hamiltonian Monte Carlo (HMC) algorithm (Duane et al., 1987), which makes use of the gradient of the model posterior probability. Importantly, we were able to derive an efficient way to compute this gradient (Appendices C and D). HMC has the advantage of very efficiently exploring the parameter space. However, for high-dimensional problems such as ours, every step may take a long time because the gradient components for all variables need to be computed. Because of the specific form of the computations in our model (convolutions and a point-wise non-linearity), it was possible to compute the gradient efficiently as a combination of convolutions and point-wise non-linearities as well. In principle, such efficient computation should be possible for a wide range of similar models, making HMC a powerful method for fitting such models.

We believe that the novel approach we introduce to the field of imaging cortical columns, of fitting a model of imaging columns to corresponding measured data, will be useful beyond our current study. For example, when imaging an unknown columnar structure, questions about its organization

(e.g. isotropy, spatial frequency, irregularity) can be addressed via inference on model parameters.

Conclusion

We have quantified the BOLD PSF in human subjects relative to neuronal activity, avoiding the confounding effects of scatter and size of visual receptive fields which were not eliminated in previous estimations (Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007). As a result, our BOLD PSF estimates characterize the spatial specificity when employing imaging of fine scale cortical organizations such as cortical columns. Previous studies have shown that BOLD fMRI at 4 T and 7 T is capable of resolving cortical columns on the sub-millimeter scale when differential analysis is employed (Cheng et al., 2001; Menon and Goodyear, 1999; Yacoub et al., 2008; 2007; Zimmermann et al., 2011). Our results provide a quantitative basis for this resolvability and facilitate planning and interpretation of high-resolution fMRI studies of fine scale cortical organizations.

References

- Adams, D.L., Sincich, L.C., Horton, J.C., 2007. Complete pattern of ocular dominance columns in human primary visual cortex. *J Neurosci* 27, 10391–10403. doi:10.1523/JNEUROSCI.2923-07.2007
- Aquino, K.M., Robinson, P.A., Schira, M.M., Breakspear, M., 2014. Deconvolution of neural dynamics from fMRI data using a spatiotemporal hemodynamic response function. *Neuroimage* 94, 203–215. doi:10.1016/j.neuroimage.2014.03.001
- Aquino, K.M., Schira, M.M., Robinson, P.A., Drysdale, P.M., Breakspear, M., 2012. Hemodynamic traveling waves in human visual cortex. *PLoS Comput Biol* 8, e1002435. doi:10.1371/journal.pcbi.1002435
- Bandettini, P.A., Wong, E.C., Hinks, R.S., Tikofsky, R.S., Hyde, J.S., 1992. Time course EPI of human brain function during task activation. *Magn Reson Med* 25, 390–397.
- Berens, P., Keliris, G.A., Ecker, A.S., Logothetis, N.K., Tolias, A.S., 2008. Comparing the feature selectivity of the gamma-band of the local field potential and the underlying spiking activity in primate visual cortex. *Front Syst Neurosci* 2, 2. doi:10.3389/neuro.06.002.2008
- Chaimow, D., Shmuel, A., 2016. The effect of MR sampling and signal decay on the effective spatial resolution in BOLD fMRI of cortical columns (in preparation).
- Chaimow, D., Yacoub, E., Ugurbil, K., Shmuel, A., 2011. Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns. *Neuroimage* 56, 627–642. doi:10.1016/j.neuroimage.2010.09.037
- Cheng, K., Waggoner, R., Tanaka, K., 2001. Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron* 32, 359–374.
- Cox, R.W., Jesmanowicz, A., 1999. Real-time 3D image registration for functional MRI. *Magn Reson Med* 42, 1014–1018.
- De Martino, F., Moerel, M., Ugurbil, K., Goebel, R., Yacoub, E., Formisano, E., 2015. Frequency preference and attention effects across cortical depths in the human primary auditory cortex. *Proceedings of the National Academy of Sciences* 112, 16036–16041. doi:10.1073/pnas.1507552112
- Duane, S., Kennedy, A.D., Pendleton, B.J., Roweth, D., 1987. Hybrid Monte Carlo.

Physics letters B.

- Duong, T.Q., Kim, D.S., Ugurbil, K., Kim, S.G., 2001. Localized cerebral blood flow response at submillimeter columnar resolution. *Proceedings of the National Academy of Sciences* 98, 10904–10909. doi:10.1073/pnas.191101098
- Duong, T.Q., Yacoub, E., Adriany, G., Hu, X., Ugurbil, K., Kim, S.-G., 2003. Microvascular BOLD contribution at 4 and 7 T in the human brain: Gradient-echo and spin-echo fMRI with suppression of blood effects. *Magn Reson Med* 49, 1019–1027.
- Duvernoy, H.M., Delon, S., Vannson, J.L., 1981. Cortical blood vessels of the human brain. *Brain Res Bull* 7, 519–579.
- Engel, S., Glover, G., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7, 181–192.
- Feinberg, D.A., Harel, N., Ramanna, S., Ugurbil, K., Yacoub, E., 2008. Sub-millimeter Single-shot 3D GRASE with Inner Volume Selection for T2 weighted fMRI applications at 7 Tesla. *Proc. Intl. Soc. Mag. Reson. Med.* 16, 1–1.
- Fracasso, A., Petridou, N., Dumoulin, S.O., 2016. Systematic variation of population receptive field properties across cortical depth in human visual cortex. *Neuroimage* 139, 427–438. doi: 10.1016/j.neuroimage.2016.06.048.
- Geweke, J., 1991. Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments.
- Goense, J.B.M., Logothetis, N.K., 2006. Laminar specificity in monkey V1 using high-resolution SE-fMRI. *Magn Reson Imaging* 24, 381–392.
- Goodyear, B.G., Menon, R.S., 2001. Brief visual stimulation allows mapping of ocular dominance in visual cortex using fMRI. *Hum Brain Mapp* 14, 210–217.
- Haacke, M.E., Brown, R.W., Thompson, M.R., 1999. *Magnetic resonance imaging: physical principles and sequence design* 914.
- Horton, J.C., Hoyt, W.F., 1991. The representation of the visual field in human striate cortex. A revision of the classic Holmes map. *Arch Ophthalmol* 109, 816–824.
- Hubel, D., Wiesel, T., 1968. Receptive fields and functional architecture of monkey striate cortex. *J Physiol (Lond)* 195, 215–243.
- Hubel, D.H., Wiesel, T.N., 1974. Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor. *J Comp Neurol* 158, 295–305. doi:10.1002/cne.901580305
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17, 825–841.
- Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis* 5, 143–156.
- Jochimsen, T.H., Norris, D.G., Mildner, T., Möller, H.E., 2004. Quantifying the intra- and extravascular contributions to spin-echo fMRI at 3 T. *Magn Reson Med* 52, 724–732. doi:10.1002/mrm.20221
- Kemper, V.G., De Martino, F., Vu, A.T., Poser, B.A., Feinberg, D.A., Goebel, R., Yacoub, E., 2015. Sub-millimeter T2 weighted fMRI at 7 T: comparison of 3D-GRASE and 2D SE-EPI. *Front Neurosci* 9, 163. doi:10.3389/fnins.2015.00163
- 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. doi:10.1016/j.neuroimage.2009.09.059
- Kwong, K.K., Belliveau, J.W., Chesler, D.A., Goldberg, I.E., Weisskoff, R.M., Poncelet, B.P., Kennedy, D.N., Hoppel, B.E., Cohen, M.S., Turner, R., 1992. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences* 89, 5675–5679.

- Lagarias, J.C., Reeds, J.A., Wright, M.H., Wright, P.E., 1998. Convergence Properties of the Nelder--Mead Simplex Method in Low Dimensions. *SIAM J. Optim.* 9, 112–147. doi:10.1137/S1052623496303470
- Leys, C., Ley, C., Klein, O., Bernard, P., Licata, L., 2013. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology* 49, 764–766. doi:10.1016/j.jesp.2013.03.013
- Menon, R.S., Goodyear, B.G., 1999. Submillimeter functional localization in human striate cortex using BOLD contrast at 4 Tesla: implications for the vascular point-spread function. *Magn Reson Med* 41, 230–235.
- Menon, R.S., Ogawa, S., Strupp, J., Ugurbil, K., 1997. Ocular dominance in human V1 demonstrated by functional magnetic resonance imaging. *J Neurophysiol* 77, 2780–2787.
- Muckli, L., De Martino, F., Vizioli, L., Petro, L.S., Smith, F.W., 2015. Contextual Feedback to Superficial Layers of V1. *Curr Biol*.
- Neal, R., 2011. MCMC using Hamiltonian dynamics, in: *Handbook of Markov Chain Monte Carlo*, pp. 113–162.
- Neal, R.M., 1993. Probabilistic inference using Markov chain Monte Carlo methods.
- Noll, D.C., Nishimura, D.G., Macovski, A., 1991. Homodyne detection in magnetic resonance imaging. *Ieee Transactions on Medical Imaging* 10, 154–163. doi:10.1109/42.79473
- Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences* 87, 9868–9872.
- Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S.G., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences* 89, 5951–5955.
- Olman, C.A., Harel, N., Feinberg, D.A., He, S., Zhang, P., Ugurbil, K., Yacoub, E., 2012. Layer-specific fMRI reflects different neuronal computations at different depths in human V1. *PLoS ONE* 7, e32536–9. doi:10.1371/journal.pone.0032536
- Oshio, K., Feinberg, D.A., 1991. GRASE (Gradient- and spin-echo) imaging: a novel fast MRI technique. *Magn Reson Med* 20, 344–349.
- Parkes, L.M., Schwarzbach, J.V., Bouts, A.A., Deckers, R.H.R., Pullens, P., Kerskens, C.M., Norris, D.G., 2005. Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. *Magn Reson Med* 54, 1465–1472. doi:10.1002/mrm.20712
- Pfeuffer, J., van de Moortele, P.F., Ugurbil, K., Hu, X., Glover, G., 2002. Correction of physiologically induced global off resonance effects in dynamic echo-planar and spiral functional imaging. *Magn Reson Med* 47, 344–353.
- Polimeni, J.R., Fischl, B., Greve, D.N., Wald, L.L., 2010. Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1. *Neuroimage* 52, 1334–1346. doi:10.1016/j.neuroimage.2010.05.005
- Roger, A., Schwartz, E., 1990. Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biol Cybern* 62, 381–391.
- 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. doi:10.1016/j.neuroimage.2009.08.040
- Shmuel, A., Yacoub, E., Chaimow, D., Logothetis, N.K., Ugurbil, K., 2007. Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 Tesla.

- Neuroimage 35, 539–552. doi:10.1016/j.neuroimage.2006.12.030
- Turner, R., 2002. How much cortex can a vein drain? Downstream dilution of activation-related cerebral blood oxygenation changes. *Neuroimage* 16, 1062–1067. doi:10.1006/nimg.2002.1082
- Uludağ, K., Müller-Bierl, B., Ugurbil, K., 2009. An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. *Neuroimage* 48, 150–165. doi:10.1016/j.neuroimage.2009.05.051
- Vazquez, A.L., Fukuda, M., Crowley, J.C., Kim, S.-G., 2014. Neural and hemodynamic responses elicited by forelimb- and photo-stimulation in channelrhodopsin-2 mice: insights into the hemodynamic point spread function. *Cereb Cortex* 24, 2908–2919. doi:10.1093/cercor/bht147
- Weber, B., Keller, A., Reichold, J., Logothetis, N.K., 2008. The Microvascular System of the Striate and Extrastriate Visual Cortex of the Macaque. *Cereb Cortex* 18, 2318–2330. doi:10.1093/cercor/bhm259
- Xing, D., Yeh, C.-I., Shapley, R.M., 2009. Spatial spread of the local field potential and its laminar variation in visual cortex. *J Neurosci* 29, 11540–11549. doi:10.1523/JNEUROSCI.2573-09.2009
- Yacoub, E., Duong, T.Q., van de Moortele, P.F., Lindquist, M., Adriany, G., Kim, S.-G., Ugurbil, K., Hu, X., 2003. Spin-echo fMRI in humans using high spatial resolutions and high magnetic fields. *Magn Reson Med* 49, 655–664.
- Yacoub, E., Harel, N., Ugurbil, K., 2008. High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences* 105, 10607–10612. doi:10.1073/pnas.0804110105
- Yacoub, E., Shmuel, A., Logothetis, N.K., 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., Shmuel, A., Pfeuffer, J., van de Moortele, P.F., Adriany, G., Andersen, P., Vaughan, J.T., Merkle, H., Ugurbil, K., Hu, X., 2001. Imaging brain function in humans at 7 Tesla. *Magn Reson Med* 45, 588–594.
- 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. doi:10.1371/journal.pone.0028716.

Appendix A. Model of imaging ocular dominance columns

Preliminaries

Simulations before MRI sampling were carried out on a grid of size $N_1^{sim} \times N_2^{sim}$. Simulations of MRI data were of size $N_1^{MRI} \times N_2^{MRI}$. We use i, j and k, l as indices of 2-dimensional image and spatial frequency space, respectively. Furthermore, $r(k, l)$ is the absolute spatial frequency and $\phi(k, l)$ the orientation that the point with indices k, l represents. The two-dimensional discrete Fourier transform and its inverse (D.3 and D.4) are denoted as $dft2$ and $idft2$.

Overview over model computations

The ocular dominance columns (ODCs) imaging model can be described as a function

$$f(q) = (f_{GE}(q), f_{SE}(q)),$$

that takes the the list of model parameter values $q = ((n_{i,j}), \rho, \delta, \epsilon, \theta, \omega, fwhm_{GE}, fwhm_{SE})$ and the fixed parameters β_{GE} and β_{SE} (see Table 1) as input and generates differential fMRI maps of ODC $f_{GE}(q) = (mri_{i,j}^{GE})$ and $f_{SE}(q) = (mri_{i,j}^{SE})$ in a number of steps:

$$(\widetilde{odc}_{i,j}) = idft2 \left[\left(dft2 \left[(n_{i,j}) \right]_{k,l} \cdot F_{k,l}^{ODC}(\rho, \delta, \epsilon, \theta) \right) \right], \quad (A.1)$$

$$odc_{i,j} = s(\widetilde{odc}_{i,j}, \omega), \quad (A.2)$$

$$(bold_{i,j}^{GE}) = idft2 \left[\left(dft2 \left[(odc_{i,j}) \right]_{k,l} \cdot MTF_{k,l}(fwhm_{GE}, \beta_{GE}) \right) \right], \quad (A.3)$$

$$(bold_{i,j}^{SE}) = idft2 \left[\left(dft2 \left[(odc_{i,j}) \right]_{k,l} \cdot MTF_{k,l}(fwhm_{SE}, \beta_{SE}) \right) \right], \quad (A.4)$$

$$(mri_{i,j}^{GE}) = idft2 \left[\left(dft2 \left[(bold_{i,j}^{GE}) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \right] \frac{N_1^{MRI} N_2^{MRI}}{N_1^{sim} N_2^{sim}}, \quad (A.5)$$

$$(mri_{i,j}^{SE}) = idft2 \left[\left(dft2 \left[(bold_{i,j}^{SE}) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \right] \frac{N_1^{MRI} N_2^{MRI}}{N_1^{sim} N_2^{sim}}. \quad (A.6)$$

(A.1): The ODC pattern $(odc_{i,j})$ was modeled by filtering two-dimensional Gaussian white noise $(n_{i,j})$ using a non-isotropic filter $(F_{k,l}^{ODC})$ and followed by (A.2) a point-wise sigmoidal non-linearity s that controlled the smoothness of transitions between left and right eye preference columns. (A.3 and A.4): The BOLD response was modeled as a convolution with a Gaussian point-spread function. It was implemented as multiplication in spatial frequency space with its Fourier transform the modulation transfer function $(MTF_{k,l})$. (A.5 and A.6): MRI sampling was simulated by restricting the spatial frequency space representation to its central part (indices given by index functions I_1^{MRI} and I_2^{MRI}) in accordance with the voxel size. The last factor corrects for the reduction in scale caused by applying $idft2$ to the reduced grid size.

We can combine operations A.3 and A.5, as well as A.4 and A.6:

$$(mri_{i,j}^{GE}) = idft2 \left[\left(dft2 \left[(odc_{i,j}^{GE}) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \cdot F_{k,l}^{BOLD} (fwhm_{GE}, \beta_{GE}) \right) \right],$$

$$(mri_{i,j}^{SE}) = idft2 \left[\left(dft2 \left[(odc_{i,j}^{SE}) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \cdot F_{k,l}^{BOLD} (fwhm_{SE}, \beta_{SE}) \right) \right],$$

where $(F_{k,l}^{BOLD})$ is of size $N_1^{MRI} \times N_2^{MRI}$ with:

$$F_{k,l}^{BOLD} (fwhm, \beta) = MTF_{I_1^{MRI}(k), I_2^{MRI}(l)} (fwhm, \beta) \frac{N_1^{MRI} N_2^{MRI}}{N_1^{sim} N_2^{sim}}$$

ODC filter

An unnormalized ODC filter $\tilde{F}_{k,l}^{ODC}$ was defined in spatial frequency space as the product of radial and angular components:

$$\tilde{F}_{k,l}^{ODC} = \tilde{F}_{k,l}^{ODC_{rad}} \cdot \tilde{F}_{k,l}^{ODC_{ang}}.$$

The radial component $\tilde{F}_{k,l}^{ODC_{rad}}$ is the sum of two Gaussian functions centered on $+\rho$ and $-\rho$, where ρ is the main spatial frequency of the pattern.

$$\begin{aligned}\tilde{F}_{k,l}^{ODC_{rad}} &= \tilde{F}_{k,l}^{ODC_{rad}^+} + \tilde{F}_{k,l}^{ODC_{rad}^-}, \\ \tilde{F}_{k,l}^{ODC_{rad}^+} &= e^{-\frac{(r(k,l)-\rho)^2}{2\delta^2}}, \\ \tilde{F}_{k,l}^{ODC_{rad}^-} &= e^{-\frac{(r(k,l)-(-\rho))^2}{2\delta^2}}.\end{aligned}$$

The angular component is the sum of two von Mises distribution functions centered on $+\theta$ and $-\theta$, where θ is the orientation of the pattern:

$$\begin{aligned}\tilde{F}_{k,l}^{ODC_{ang}} &= \tilde{F}_{k,l}^{ODC_{ang}^1} + \tilde{F}_{k,l}^{ODC_{ang}^2}, \\ \tilde{F}_{k,l}^{ODC_{ang}^1} &= e^{\frac{\cos(\phi(k,l)-\theta)}{\epsilon^2}}, \\ \tilde{F}_{k,l}^{ODC_{ang}^2} &= e^{\frac{\cos(\phi(k,l)-(\theta+\pi))}{\epsilon^2}} = e^{\frac{-\cos(\phi(k,l)-\theta)}{\epsilon^2}}.\end{aligned}$$

In order for the filter output to have the same variance as the input (independent of filter parameter values) we normalized the filter:

$$F_{k,l}^{ODC} = \frac{\tilde{F}_{k,l}^{ODC}}{C^{ODC}},$$

where

$$C^{ODC} = \sqrt{\frac{\sum_{k,l=1}^{N_1^{sim} \times N_2^{sim}} (\tilde{F}_{k,l}^{ODC})^2}{N_1^{sim} N_2^{sim}}}.$$

Sigmoidal non-linearity

The point-wise sigmoidal non-linearity $s(x, \omega)$ was defined as:

$$s(x, \omega) = 2 \left(\tilde{s} \left(\frac{x}{\omega} \right) - 0.5 \right),$$

with the standard sigmoidal function defined as:

$$\tilde{s}(x) = \frac{1}{1 + e^{-x}}.$$

BOLD response

The BOLD response modulation transfer function was defined as

$$MTF_{k,l}(fwhm, \beta) = \beta \cdot e^{-2\pi^2 \sigma_{BOLD}^2 \cdot r(k,l)^2},$$

with:

$$\sigma_{BOLD} = \frac{fwhm}{2\sqrt{2 \log 2}} \approx \frac{fwhm}{2.35}.$$

It is the Fourier transform of a Gaussian point-spread function with a full-width at half-maximum of $fwhm$. It is scaled such that a spatially extended neuronal response of 1 results in a BOLD response of amplitude β .

MRI sampling

MRI sampling was simulated by restricting the spatial frequency representation according to the following index functions:

$$I_1^{MRI}(k) = \begin{cases} k & \text{if } k \leq \frac{N_1^{MRI}}{2} \\ k - N_1^{MRI} + N_1^{sim} & \text{if } k > \frac{N_1^{MRI}}{2} \end{cases},$$

$$I_2^{MRI}(l) = \begin{cases} l & \text{if } l \leq \frac{N_2^{MRI}}{2} \\ l - N_2^{MRI} + N_2^{sim} & \text{if } l > \frac{N_2^{MRI}}{2} \end{cases}.$$

Appendix B. Posterior probability and potential energy

Posterior probability

Let $\mathcal{D} = \{(data_{i,j}^{GE}), (data_{i,j}^{SE})\}$ be the data of differential fMRI maps imaged using GE and SE BOLD fMRI, respectively. The likelihood - the probability to observe the data \mathcal{D} given a specific set of model parameter values q is:

$$P[\mathcal{D}|q] = P \left[(data_{i,j}^{GE}) = f_{GE}(q) + (v_{i,j}^{GE}), (data_{i,j}^{SE}) = f_{SE}(q) + (v_{i,j}^{SE}) | q \right],$$

where $(v_{i,j}^{GE})$ and $(v_{i,j}^{SE})$ are patterns of measurement noise.

We assume the measurement noise to be independent between voxels and imaging modalities and to be distributed normally with (estimated) variances $\hat{\sigma}_{GE}^2$ and $\hat{\sigma}_{SE}^2$. Furthermore we define $(d_{i,j}^{GE}) = (data_{i,j}^{GE}) - (mri_{i,j}^{GE})$ and $(d_{i,j}^{SE}) = (data_{i,j}^{SE}) - (mri_{i,j}^{SE})$ to be the patterns of deviations of the data from the model prediction. The likelihood can then be expressed as:

$$P[\mathcal{D}|q] = \frac{1}{(2\pi\hat{\sigma}_{GE}\hat{\sigma}_{SE})^{N_1^{MRI}N_2^{MRI}}} \prod_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} e^{-\frac{(d_{i,j}^{GE})^2}{2\hat{\sigma}_{GE}^2} - \frac{(d_{i,j}^{SE})^2}{2\hat{\sigma}_{SE}^2}}.$$

From this result we can calculate the posterior probability of parameters q given the data \mathcal{D} :

$$\begin{aligned} P[q|\mathcal{D}] &= \frac{P[\mathcal{D}|q]P[q]}{P[\mathcal{D}]} \\ &= A \cdot P[q] \cdot \prod_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} e^{-\frac{(d_{i,j}^{GE})^2}{2\hat{\sigma}_{GE}^2} - \frac{(d_{i,j}^{SE})^2}{2\hat{\sigma}_{SE}^2}}, \end{aligned} \quad (\text{B.1})$$

where $A = \frac{1}{P[\mathcal{D}] \cdot (2\pi\hat{\sigma}_{GE}\hat{\sigma}_{SE})^{N_1^{MRI}N_2^{MRI}}}$ is a constant factor that is independent of q and $P[q]$ is the prior probability over parameters.

Prior probability

The prior probability over all parameters q was defined as:

$$P[q] = \begin{cases} 0 & \text{if } \omega < 0 \\ & \text{or } fwhm_{GE} < 0 \\ & \text{or } fwhm_{SE} < 0 \\ P[(n_{i,j})] \cdot P[\rho] \cdot P[\delta] \cdot P[\epsilon] \cdot P[\omega] & \text{otherwise} \end{cases},$$

with individual parameter priors were defined as:

$$P[(n_{i,j})] \propto e^{-\frac{\sum_{i,j=1}^{N_1^{sim} \times N_2^{sim}} n_{i,j}^2}{2}},$$

$$P[\rho] \propto \begin{cases} 0 & \text{if } |\rho - \mu_\rho| > 2\sigma_\rho \\ e^{-\frac{(\rho - \mu_\rho)^2}{2\sigma_\rho^2}} & \text{otherwise} \end{cases},$$

$$P[\delta] \propto \begin{cases} 0 & \text{if } \delta < 0 \\ e^{-\frac{(\delta - \mu_\delta)^2}{2\sigma_\delta^2}} & \text{otherwise} \end{cases},$$

$$P[\epsilon] \propto \begin{cases} 0 & \text{if } \epsilon < 0 \\ e^{-\frac{(\epsilon - \mu_\epsilon)^2}{2\sigma_\epsilon^2}} & \text{otherwise} \end{cases},$$

$$P[\omega] \propto \begin{cases} 0 & \text{if } \omega < \omega_{min} \\ 0 & \text{if } \omega > \omega_{max} \\ 1 & \text{otherwise} \end{cases},$$

where $\mu_\rho, \mu_\delta, \mu_\epsilon$ and $\sigma_\rho, \sigma_\delta, \sigma_\epsilon$ are the means and standard deviations of the ODC priors that were estimated from cytochrome oxidase data and ω_{min} and ω_{max} are lower and upper limits for ω that were set based on results from the neurophysiological literature.

Potential energy

The potential energy of a state of parameter values q is its negative log-posterior probability plus an arbitrary constant C :

$$E(q) = -\log P[q|\mathcal{D}] + C.$$

We apply B.1 and express the energy as the sum of three parts:

$$= E_{GE}(q) + E_{SE}(q) + E(q)_{prior},$$

with:

$$E_{GE}(q) = \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} \frac{(d_{i,j}^{GE})^2}{2\hat{\sigma}_{GE}^2},$$

$$E_{SE}(q) = \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} \frac{(d_{i,j}^{SE})^2}{2\hat{\sigma}_{SE}^2}$$

and

$$E_{prior}(q) = \begin{cases} \infty & \begin{array}{l} \text{if } |\rho - \mu_\rho| > 2\sigma_\rho \\ \text{or } \delta < 0 \\ \text{or } \epsilon < 0 \\ \text{or } \omega > \omega_{min} \\ \text{or } \omega > \omega_{max} \\ \text{or } fwhm_{GE} < 0 \\ \text{or } fwhm_{SE} < 0 \end{array} \\ \frac{\sum_{i,j=1}^{N_1^{sim} \times N_2^{sim}} n_{i,j}^2}{2} + \frac{(\rho - \mu_\rho)^2}{2\sigma_\rho^2} + \frac{(\delta - \mu_\delta)^2}{2\sigma_\delta^2} + \frac{(\epsilon - \mu_\epsilon)^2}{2\sigma_\epsilon^2} & \text{otherwise} \end{cases}.$$

Appendix C. Potential energy gradient

We start by deriving derivatives for functions used by the model.

Derivatives of the ODC filter

The derivatives of the unnormalized radial and angular filter component parts with respect to their parameters are:

$$\frac{\partial \tilde{F}_{k,l}^{ODC^+_{rad}}}{\partial \rho} = \tilde{F}_{k,l}^{ODC^+_{rad}} \cdot \frac{r(k,l) - \rho}{\delta^2},$$

$$\frac{\partial \tilde{F}_{k,l}^{ODC^-_{rad}}}{\partial \rho} = \tilde{F}_{k,l}^{ODC^-_{rad}} \cdot -\frac{r(k,l) - (-\rho)}{\delta^2},$$

$$\frac{\partial \tilde{F}_{k,l}^{ODC^+_{rad}}}{\partial \delta} = \tilde{F}_{k,l}^{ODC^+_{rad}} \cdot \frac{(r(k,l) - \rho)^2}{\delta^3},$$

$$\frac{\partial \tilde{F}_{k,l}^{ODC^-_{rad}}}{\partial \delta} = \tilde{F}_{k,l}^{ODC^-_{rad}} \cdot \frac{(r(k,l) - (-\rho))^2}{\delta^3},$$

$$\frac{\partial \tilde{F}_{k,l}^{ODC^1_{ang}}}{\partial \theta} = \tilde{F}_{k,l}^{ODC^1_{ang}} \cdot \frac{\sin(\phi(k,l) - \theta)}{\epsilon^2},$$

$$\frac{\partial \tilde{F}_{k,l}^{ODC^2_{ang}}}{\partial \theta} = \tilde{F}_{k,l}^{ODC^2_{ang}} \cdot -\frac{\sin(\phi(k,l) - \theta)}{\epsilon^2},$$

$$\frac{\partial \tilde{F}_{k,l}^{ODC^1_{ang}}}{\partial \epsilon} = \tilde{F}_{k,l}^{ODC^1_{ang}} \cdot -2 \frac{\cos(\phi(k,l) - \theta)}{\epsilon^3},$$

$$\frac{\partial \tilde{F}_{k,l}^{ODC^2_{ang}}}{\partial \epsilon} = \tilde{F}_{k,l}^{ODC^2_{ang}} \cdot 2 \frac{\cos(\phi(k,l) - \theta)}{\epsilon^3}.$$

We combine these derivatives to form derivatives of the full radial and angular components:

$$\begin{aligned}\frac{\partial \tilde{F}_{k,l}^{ODC_{rad}}}{\partial \rho} &= \frac{\partial \tilde{F}_{k,l}^{ODC_{rad}^+}}{\partial \rho} + \frac{\partial \tilde{F}_{k,l}^{ODC_{rad}^-}}{\partial \rho}, \\ \frac{\partial \tilde{F}_{k,l}^{ODC_{rad}}}{\partial \delta} &= \frac{\partial \tilde{F}_{k,l}^{ODC_{rad}^+}}{\partial \delta} + \frac{\partial \tilde{F}_{k,l}^{ODC_{rad}^-}}{\partial \delta}, \\ \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}}}{\partial \theta} &= \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}^1}}{\partial \theta} + \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}^2}}{\partial \theta}, \\ \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}}}{\partial \epsilon} &= \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}^1}}{\partial \epsilon} + \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}^2}}{\partial \epsilon},\end{aligned}$$

and in turn to form the derivatives of the full unnormalized filter:

$$\begin{aligned}\frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \rho} &= \frac{\partial \tilde{F}_{k,l}^{ODC_{rad}}}{\partial \rho} \cdot \tilde{F}_{k,l}^{ODC_{ang}}, \\ \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \delta} &= \frac{\partial \tilde{F}_{k,l}^{ODC_{rad}}}{\partial \delta} \cdot \tilde{F}_{k,l}^{ODC_{ang}}, \\ \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \theta} &= \tilde{F}_{k,l}^{ODC_{rad}} \cdot \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}}}{\partial \theta}, \\ \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \epsilon} &= \tilde{F}_{k,l}^{ODC_{rad}} \cdot \frac{\partial \tilde{F}_{k,l}^{ODC_{ang}}}{\partial \epsilon}.\end{aligned}$$

The derivatives of the normalization constant are:

$$\begin{aligned}\frac{\partial C^{ODC}}{\partial \rho} &= \frac{1}{2\sqrt{\frac{\sum_{k,l=1}^{N_1^{sim} \times N_2^{sim}} (\tilde{F}_{k,l}^{ODC})^2}{N_1^{sim} N_2^{sim}}}} \cdot \frac{1}{N_1^{sim} N_2^{sim}} \cdot \sum_{k,l=1}^{N_1^{sim} \times N_2^{sim}} 2 \cdot \tilde{F}_{k,l}^{ODC} \cdot \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \rho} \\ &= \frac{\sum_{k,l=1}^{N_1^{sim} \times N_2^{sim}} \tilde{F}_{k,l}^{ODC} \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \rho}}{C^{ODC} N_1^{sim} N_2^{sim}}, \\ \frac{\partial C^{ODC}}{\partial \delta} &= \frac{\sum_{k,l=1}^{N_1^{sim} \times N_2^{sim}} \tilde{F}_{k,l}^{ODC} \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \delta}}{C^{ODC} N_1^{sim} N_2^{sim}}, \\ \frac{\partial C^{ODC}}{\partial \theta} &= \frac{\sum_{k,l=1}^{N_1^{sim} \times N_2^{sim}} \tilde{F}_{k,l}^{ODC} \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \theta}}{C^{ODC} N_1^{sim} N_2^{sim}},\end{aligned}$$

$$\frac{\partial C^{ODC}}{\partial \epsilon} = \frac{\sum_{k,l=1}^{N_1^{sim} \times N_2^{sim}} \tilde{F}_{k,l}^{ODC} \frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \epsilon}}{C^{ODC} N_1^{sim} N_2^{sim}},$$

resulting in the following derivatives of the complete normalized filter:

$$\frac{\partial F_{k,l}^{ODC}}{\partial \rho} = \frac{\frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \rho} \cdot C^{ODC} - \tilde{F}_{k,l}^{ODC} \cdot \frac{\partial C^{ODC}}{\partial \rho}}{(C^{ODC})^2},$$

$$\frac{\partial F_{k,l}^{ODC}}{\partial \delta} = \frac{\frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \delta} \cdot C^{ODC} - \tilde{F}_{k,l}^{ODC} \cdot \frac{\partial C^{ODC}}{\partial \delta}}{(C^{ODC})^2},$$

$$\frac{\partial F_{k,l}^{ODC}}{\partial \theta} = \frac{\frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \theta} \cdot C^{ODC} - \tilde{F}_{k,l}^{ODC} \cdot \frac{\partial C^{ODC}}{\partial \theta}}{(C^{ODC})^2},$$

$$\frac{\partial F_{k,l}^{ODC}}{\partial \epsilon} = \frac{\frac{\partial \tilde{F}_{k,l}^{ODC}}{\partial \epsilon} \cdot C^{ODC} - \tilde{F}_{k,l}^{ODC} \cdot \frac{\partial C^{ODC}}{\partial \epsilon}}{(C^{ODC})^2}.$$

Derivatives of the sigmoidal non-linearity

The derivative of the standard sigmoidal function $\tilde{s}(x)$ is:

$$\tilde{s}'(x) = \tilde{s}(x)(1 - \tilde{s}(x)).$$

Using this result we get the following derivatives for our sigmoidal non-linearity $s(x)$:

$$\begin{aligned} \frac{\partial s(x)}{\partial x} &= \frac{\partial s(x)}{\partial \tilde{s}(x/\omega)} \cdot \frac{\partial \tilde{s}(x/\omega)}{\partial (x/\omega)} \cdot \frac{\partial (x/\omega)}{x} \\ &= \frac{2}{\omega} \tilde{s}(x/\omega)(1 - \tilde{s}(x/\omega)), \end{aligned}$$

$$\begin{aligned} \frac{\partial s(x)}{\partial \omega} &= \frac{\partial s(x)}{\partial \tilde{s}(x/\omega)} \cdot \frac{\partial \tilde{s}(x/\omega)}{\partial (x/\omega)} \cdot \frac{\partial (x/\omega)}{\partial \omega} \\ &= \frac{-2x}{\omega^2} \tilde{s}(x/\omega)(1 - \tilde{s}(x/\omega)). \end{aligned}$$

Derivative of the BOLD modulation transfer function

The derivative of the BOLD modulation transfer function is:

$$\begin{aligned} \frac{\partial MTF_{k,l}(fwhm, \beta)}{\partial fwhm} &= \frac{\partial MTF_{k,l}(fwhm, \beta)}{\partial \sigma_{BOLD}} \frac{\partial \sigma_{BOLD}}{\partial fwhm} \\ &= \frac{-4MTF_{k,l} r(k, l)^2 \pi^2 fwhm}{8 \log 2} \end{aligned}$$

The derivative of the combined BOLD-MRI filter is:

$$\frac{\partial F_{k,l}^{BOLD} (fwhm, \beta)}{\partial fwhm} = \frac{\partial MTF_{I_1^{MRI}(k), I_2^{MRI}(l)}(fwhm, \beta)}{\partial fwhm} \frac{N_1^{MRI} N_2^{MRI}}{N_1^{sim} N_2^{sim}}.$$

Potential energy gradient components

The energy gradient is composed of the following derivatives:

$$\begin{aligned}
\frac{\partial E(q)}{\partial n_{i,j}} &= \frac{\partial E_{GE}(q)}{\partial n_{i,j}} + \frac{\partial E_{SE}(q)}{\partial n_{i,j}} + \frac{\partial E_{prior}(q)}{\partial n_{i,j}}, \\
\frac{\partial E(q)}{\partial \rho} &= \frac{\partial E_{GE}(q)}{\partial \rho} + \frac{\partial E_{SE}(q)}{\partial \rho} + \frac{\partial E_{prior}(q)}{\partial \rho}, \\
\frac{\partial E(q)}{\partial \delta} &= \frac{\partial E_{GE}(q)}{\partial \delta} + \frac{\partial E_{SE}(q)}{\partial \delta} + \frac{\partial E_{prior}(q)}{\partial \delta}, \\
\frac{\partial E(q)}{\partial \epsilon} &= \frac{\partial E_{GE}(q)}{\partial \epsilon} + \frac{\partial E_{SE}(q)}{\partial \epsilon} + \frac{\partial E_{prior}(q)}{\partial \epsilon}, \\
\frac{\partial E(q)}{\partial \theta} &= \frac{\partial E_{GE}(q)}{\partial \theta} + \frac{\partial E_{SE}(q)}{\partial \theta}, \\
\frac{\partial E(q)}{\partial \omega} &= \frac{\partial E_{GE}(q)}{\partial \omega} + \frac{\partial E_{SE}(q)}{\partial \omega}, \\
\frac{\partial E(q)}{\partial fwhm_{GE}} &= \frac{\partial E_{GE}(q)}{\partial fwhm_{GE}}, \\
\frac{\partial E(q)}{\partial fwhm_{SE}} &= \frac{\partial E_{SE}(q)}{\partial fwhm_{SE}}.
\end{aligned}$$

The gradient is not defined for $|\rho - \mu_\rho| > 2\sigma_\rho$ or $\delta < 0$ or $\epsilon < 0$ or $\omega < \omega_{min}$ or $\omega > \omega_{max}$ or $fwhm_{GE} < 0$ or $fwhm_{SE} < 0$ (regions where $E_{prior}(q) = \infty$).

Energy gradient with respect to noise variables

The derivatives of the *GE* energy component are:

$$\begin{aligned}
&\frac{\partial E_{GE}(q)}{\partial n_{i,j}} \\
&= \frac{1}{\tilde{\sigma}_{GE}^2} \sum_{i',j'=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i',j'}^{GE} \frac{\partial mri_{i',j'}}{\partial n_{i,j}} \\
&= \frac{1}{\tilde{\sigma}_{GE}^2} \sum_{i',j'=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i',j'}^{GE} \sum_{i'',j''=1}^{N_1^{sim} \times N_2^{sim}} \frac{\partial mri_{i',j'}}{\partial odc_{i'',j''}} \frac{\partial odc_{i'',j''}}{\partial \widetilde{odc}_{i'',j''}} \frac{\partial \widetilde{odc}_{i'',j''}}{\partial n_{i,j}} \\
&= \frac{1}{\tilde{\sigma}_{GE}^2} \sum_{i',j'=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i',j'}^{GE} \sum_{i'',j''=1}^{N_1^{sim} \times N_2^{sim}} \frac{\partial mri_{i',j'}}{\partial odc_{i'',j''}} s'(\widetilde{odc}_{i'',j''}) \frac{\partial \widetilde{odc}_{i'',j''}}{\partial n_{i,j}} \\
&= \frac{1}{\tilde{\sigma}_{GE}^2} \sum_{i'',j''=1}^{N_1^{sim} \times N_2^{sim}} s'(\widetilde{odc}_{i'',j''}) \frac{\partial \widetilde{odc}_{i'',j''}}{\partial n_{i,j}} \sum_{i',j'=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i',j'}^{GE} \frac{\partial mri_{i',j'}}{\partial odc_{i'',j''}}
\end{aligned}$$

For both sums we apply D.8 (see D.2 for the definition of the zero-padding operation $\text{zp}_{I_1^{MRI}, I_2^{MRI}}^{N_1^{sim}, N_2^{sim}}$):

$$= \frac{1}{\tilde{\sigma}_{GE}^2} \text{dft2} \left[\left(F_{k,l}^{ODC} \cdot \text{idft2} \left[\left(s'(\widetilde{odc}_{i'',j''}) \cdot \text{zp}_{I_1^{MRI}, I_2^{MRI}}^{N_1^{sim}, N_2^{sim}} \left[\text{dft2} \left[\left(F_{k',l'}^{BOLD} \cdot \text{idft2} \left[\left((d_{i',j'}^{GE})_{k',l'} \right) \right]_{i'',j''} \right) \right]_{k,l} \right) \right] \right) \right]_{k,l} \right)$$

Similarly for the *SE* energy component we get:

$$\frac{\partial E_{SE}(q)}{\partial n_{i,j}}$$

$$= \frac{1}{\sigma_{SE}^2} \text{dft2} \left[\left(F_{k,l}^{ODC} \cdot \text{idft2} \left[\left(s'(\widetilde{odc}_{i'',j''}) \cdot \text{zp}_{I_1^{MRI}, I_2^{MRI}}^{N_1^{sim}, N_2^{sim}} \left[\text{dft2} \left[\left(F_{k',l'}^{BOLD} \cdot \text{idft2} \left[\left(d_{i',j'}^{SE} \right) \right]_{k',l'} \right) \right]_{i'',j''} \right) \right]_{k,l} \right) \right] \right)$$

The contribution of the prior energy component is:

$$\frac{\partial E_{\text{prior}}(q)}{\partial n_{i,j}} = n_{i,j}.$$

Energy gradient with respect to ODC filter parameters

The derivatives of the *GE* energy component are:

$$\begin{aligned} & \frac{\partial E_{GE}(q)}{\partial \rho} \\ &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \frac{\partial mri_{i,j}}{\partial \rho} \\ &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \sum_{i',j'=1}^{N_1^{sim} \times N_2^{sim}} \frac{\partial mri_{i,j}}{\partial odc_{i',j'}} \frac{\partial odc_{i',j'}}{\partial \rho} \\ &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \sum_{i',j'=1}^{N_1^{sim} \times N_2^{sim}} \frac{\partial mri_{i,j}}{\partial odc_{i',j'}} \frac{\partial odc_{i',j'}}{\partial \widetilde{odc}_{i',j'}} \frac{\partial \widetilde{odc}_{i',j'}}{\partial \rho} \\ &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \sum_{i',j'=1}^{N_1^{sim} \times N_2^{sim}} \frac{\partial mri_{i,j}}{\partial odc_{i',j'}} s'(\widetilde{odc}_{i',j'}) \frac{\partial \widetilde{odc}_{i',j'}}{\partial \rho} \end{aligned}$$

The last two factors together can be regarded as a pattern indexed by i', j' . We apply D.7:

$$= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{odc}_{i',j'}) \cdot \frac{\partial \widetilde{odc}_{i',j'}}{\partial \rho} \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \right]_{i,j}$$

$\frac{\partial \widetilde{odc}_{i',j'}}{\partial \rho}$ is the derivative of the ODC filter output $\widetilde{odc}_{i',j'}$ with respect to a filter parameter. We apply D.9:

$$\begin{aligned} &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \\ &\cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{odc}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[\left(n_{i'',j''} \right) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \rho} \right) \right]_{i',j'} \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \right] \end{aligned}$$

The derivatives of $E_{GE}(q)$ with respect to the remaining ODC filter parameters δ , ϵ and θ differ only in the derivative of $F_{k',l'}^{ODC}$:

$$\frac{\partial E_{GE}(q)}{\partial \delta} = \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE}$$

$$\cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{\text{odc}}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[(n_{i'',j''}) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \delta} \right) \right]_{i',j'} \right) \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}}$$

$$\frac{\partial E_{GE}(q)}{\partial \epsilon} = \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE}$$

$$\cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{\text{odc}}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[(n_{i'',j''}) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \epsilon} \right) \right]_{i',j'} \right) \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}}$$

$$\frac{\partial E_{GE}(q)}{\partial \theta} = \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE}$$

$$\cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{\text{odc}}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[(n_{i'',j''}) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \theta} \right) \right]_{i',j'} \right) \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}}$$

Similarly for the *SE* energy components we get:

$$\frac{\partial E_{SE}(q)}{\partial \rho} = \frac{1}{\sigma_{SE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{SE}$$

$$\cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{\text{odc}}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[(n_{i'',j''}) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \rho} \right) \right]_{i',j'} \right) \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}}$$

$$\frac{\partial E_{SE}(q)}{\partial \delta} = \frac{1}{\sigma_{SE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{SE}$$

$$\cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{\text{odc}}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[(n_{i'',j''}) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \delta} \right) \right]_{i',j'} \right) \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}}$$

$$\frac{\partial E_{SE}(q)}{\partial \epsilon} = \frac{1}{\sigma_{SE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{SE}$$

$$\cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{\text{odc}}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[(n_{i'',j''}) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \epsilon} \right) \right]_{i',j'} \right) \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}}$$

$$\frac{\partial E_{SE}(q)}{\partial \theta} = \frac{1}{\sigma_{SE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{SE} \cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(s'(\widetilde{odc}_{i',j'}) \cdot \text{idft2} \left[\left(\text{dft2} \left[(n_{i',j''}) \right]_{k',l'} \cdot \frac{\partial F_{k',l'}^{ODC}}{\partial \theta} \right) \right]_{i',j'} \right) \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right)$$

The contributions of the prior energy components are:

$$\frac{\partial E_{prior}(q)}{\partial \rho} = \frac{\rho - \mu_\rho}{\sigma_\rho^2},$$

$$\frac{\partial E_{prior}(q)}{\partial \delta} = \frac{\delta - \mu_\delta}{\sigma_\delta^2},$$

$$\frac{\partial E_{prior}(q)}{\partial \epsilon} = \frac{\epsilon - \mu_\epsilon}{\sigma_\epsilon^2}.$$

Energy gradient with respect to the smoothness parameter

The derivatives of the GE component are:

$$\begin{aligned} \frac{\partial E_{GE}(q)}{\partial \omega} &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \frac{\partial mri_{m,n}^{GE}}{\partial \omega} \\ &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \sum_{i',j'=1}^{N_1^{sim} \times N_2^{sim}} \frac{\partial mri_{i,j}^{GE}}{\partial odc_{i',j'}} \frac{\partial odc_{i',j'}}{\partial \omega} \\ &= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \sum_{i',j'=1}^{N_1^{sim} \times N_2^{sim}} \frac{\partial mri_{i,j}^{GE}}{\partial odc_{i',j'}} \frac{\partial s(\widetilde{odc}_{i',j'})}{\partial \omega} \end{aligned}$$

for the second sum we apply D.7:

$$= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(\frac{\partial s(\widetilde{ODC}_{i',j'})}{\partial \omega} \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \right]_{i,j}.$$

Similarly for the SE components we get:

$$\frac{\partial E_{SE}(q)}{\partial \omega} = \frac{1}{\sigma_{SE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{SE} \cdot \text{idft2} \left[\left(F_{k,l}^{BOLD} \cdot \text{dft2} \left[\left(\frac{\partial s(\widetilde{ODC}_{i',j'})}{\partial \omega} \right) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \right) \right]_{i,j}.$$

Energy gradient with respect to BOLD PSF width

The derivatives of the GE component are:

$$\frac{\partial E_{GE}(q)}{\partial fwhm_{GE}} = \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \frac{\partial mri_{i,j}^{GE}}{\partial fwhm_{GE}}$$

We apply D.9:

$$= \frac{1}{\sigma_{GE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{GE} \cdot \text{idft2} \left[\left(\text{dft2} \left[(odc_{i',j'}) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \cdot \frac{\partial F_{k,l}^{BOLD}}{\partial fwhm_{GE}} \right) \right]_{i,j}.$$

Similarly for the SE components we get:

$$\frac{\partial E_{SE}(q)}{\partial fwhm_{SE}} = \frac{1}{\sigma_{SE}^2} \sum_{i,j=1}^{N_1^{MRI} \times N_2^{MRI}} d_{i,j}^{SE} \cdot \text{idft2} \left[\left(\text{dft2} \left[(odc_{i',j'}) \right]_{\substack{k=I_1^{MRI}(1), \dots, I_1^{MRI}(N_1^{MRI}), \\ l=I_2^{MRI}(1), \dots, I_2^{MRI}(N_2^{MRI})}} \cdot \frac{\partial F_{k,l}^{BOLD}}{\partial fwhm_{SE}} \right) \right]_{i,j}.$$

Appendix D. Efficient computation of some expressions consisting of derivatives of filtering operations

Let us assume we have a pattern $(x_{i,j})$ of size $N_1^{in} \times N_2^{in}$ (the input), whose discrete Fourier transform is $(X_{k,l})$. We filter (or equivalently: convolve) that pattern using a linear filter $(F_{k,l})$ defined in frequency space and of size $N_1^{out} \times N_2^{out}$. The filtering is carried out as:

$$Y_{k,l} = F_{k,l} \cdot X_{I_1(k), I_2(l)}, \quad (\text{D.1})$$

where $(Y_{k,l})$ is the discrete Fourier transform of the filtered, $N_1^{out} \times N_2^{out}$ sized pattern $(y_{m,n})$ (the output). I_1 and I_2 are index functions that allow to assign a specific subset of elements of $(X_{k,l})$ to $(Y_{k,l})$ (e.g. when restricting the spatial frequency space representation in order to simulate MRI sampling). Note that for $N_1^{out} = N_1^{in}$, $N_2^{out} = N_2^{in}$, $I_1(k) = k$ and $I_2(l) = l$ this formalism describes a normal filtering operation in which the size and structure of the spatial frequency space is not altered.

Our goal here is to derive alternate formulations for some derivative expressions that result in a more efficient computation of gradients for models that contain such filtering operations.

Definitions

We define a zero-padding operation $\text{zp}_{I_1, I_2}^{N_1^{in}, N_2^{in}}$ that allows us to up-sample patterns with size equal to the output pattern to match the size of the input pattern:

$$(\bar{Y}_{k,l}) = \text{zp}_{I_1, I_2}^{N_1^{in}, N_2^{in}} [(Y_{k,l})], \quad (\text{D.2})$$

where $(\bar{Y}_{k,l})$ is of size $N_1^{in} \times N_2^{in}$ such that:

$$\bar{Y}_{k,l} = \begin{cases} Y_{k',l'} & \text{if there exist } (k', l') \text{ with } I_1(k') = k \text{ and } I_2(l') = l \\ 0 & \text{otherwise} \end{cases}.$$

We use the following definitions of the two-dimensional discrete Fourier transform $(X_{k,l}) = \text{dft2} [(x_{i,j})]$ and its inverse $(x_{i,j}) = \text{idft2} [(X_{k,l})]$, where $(x_{i,j})$ and $(X_{k,l})$ are of size $N_1 \times N_2$:

$$X_{k,l} = \sum_{i,j=1}^{N_1 \times N_2} x_{i,j} e^{-i2\pi \frac{(i-1)(k-1)}{N_1^{in}}} e^{-i2\pi \frac{(j-1)(l-1)}{N_2^{in}}}, \quad (\text{D.3})$$

$$x_{i,j} = \frac{1}{N_1 N_2} \sum_{k,l=1}^{N_1^{in} \times N_2^{in}} X_{k,l} e^{i2\pi \frac{(k-1)(i-1)}{N_1^{in}}} e^{i2\pi \frac{(l-1)(j-1)}{N_2^{in}}}, \quad (D.4)$$

with partial derivatives given as:

$$\frac{\partial X_{k,l}}{\partial x_{i,j}} = e^{-i2\pi \frac{(i-1)(k-1)}{N_1^{in}}} e^{-i2\pi \frac{(j-1)(l-1)}{N_2^{in}}}, \quad (D.5)$$

$$\frac{\partial x_{i,j}}{\partial X_{k,l}} = \frac{1}{N_1^{in} N_2^{in}} e^{i2\pi \frac{(k-1)(i-1)}{N_1^{in}}} e^{i2\pi \frac{(l-1)(j-1)}{N_2^{in}}}. \quad (D.6)$$

Weighted sums of derivatives with respect to input

The first expression of interest is the weighted sum (weights given as $(g_{i,j})$) of the derivatives of an arbitrary output element $y_{i,j}$ with respect to all input elements $x_{i',j'}$. We derive a formula that allows to compute this expressions for all output elements $y_{i,j}$ simultaneously using one dft2 and one idft2 operation.

$$\begin{aligned} & \sum_{i',j'=1}^{N_1^{in} \times N_2^{in}} g_{i',j'} \frac{\partial y_{i,j}}{\partial x_{i',j'}} \\ &= \sum_{i',j'=1}^{N_1^{in} \times N_2^{in}} g_{i',j'} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} \frac{\partial y_{i,j}}{\partial Y_{k,l}} \sum_{k',l'=1}^{N_1^{in} \times N_2^{in}} \frac{\partial Y_{k,l}}{\partial X_{k',l'}} \frac{\partial X_{k',l'}}{\partial x_{i',j'}}. \end{aligned}$$

Because of D.1, the second to last term is equal to $F_{k,l}$, if $(k',l') = (I_1(k), I_2(l))$ and 0 otherwise:

$$= \sum_{i',j'=1}^{N_1^{in} \times N_2^{in}} g_{i',j'} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} \frac{\partial y_{i,j}}{\partial Y_{k,l}} F_{k,l} \frac{\partial X_{I_1(k), I_2(l)}}{\partial x_{i',j'}}$$

Furthermore making use of D.5 and D.6 we get:

$$\begin{aligned} &= \sum_{i',j'=1}^{N_1^{in} \times N_2^{in}} g_{i',j'} \frac{1}{N_1^{out} N_2^{out}} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} e^{i2\pi \frac{(k-1)(i-1)}{N_1^{out}}} e^{i2\pi \frac{(l-1)(j-1)}{N_2^{out}}} e^{-i2\pi \frac{(i'-1)(I_1(k)-1)}{N_1^{in}}} e^{-i2\pi \frac{(j'-1)(I_2(l)-1)}{N_2^{in}}} \\ &= \frac{1}{N_1^{out} N_2^{out}} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} e^{i2\pi \frac{(k-1)(i-1)}{N_1^{out}}} e^{i2\pi \frac{(l-1)(j-1)}{N_2^{out}}} \sum_{i',j'=1}^{N_1^{in} \times N_2^{in}} g_{i',j'} e^{-i2\pi \frac{(i'-1)(I_1(k)-1)}{N_1^{in}}} e^{-i2\pi \frac{(j'-1)(I_2(l)-1)}{N_2^{in}}} \\ &= \frac{1}{N_1^{out} N_2^{out}} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} F_{k,l} e^{i2\pi \frac{(k-1)(i-1)}{N_1^{out}}} e^{i2\pi \frac{(l-1)(j-1)}{N_2^{out}}} \text{dft2} \left[(g_{i',j'}) \right]_{I_1(k), I_2(l)} \\ &= \text{idft2} \left[\left(F_{k,l} \cdot \text{dft2} \left[(g_{i',j'}) \right]_{I_1(k), I_2(l)} \right) \right]_{i,j}. \quad (D.7) \end{aligned}$$

A related expression is the weighted sum of the derivatives of all output elements $y_{i',j'}$ with respect to an arbitrary but specific input element $x_{i,j}$. Again, we derive a formula that allows to compute this expressions for all input elements $x_{i,j}$ simultaneously using one dft2 and one idft2 operation.

$$\sum_{i',j'=1}^{N_1^{out} \times N_2^{out}} g_{i',j'} \frac{\partial y_{i',j'}}{\partial x_{i,j}}$$

$$= \sum_{i',j'=1}^{N_1^{out} \times N_2^{out}} \mathcal{G}_{i',j'} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} \frac{\partial y_{i',j'}}{\partial Y_{k,l}} \sum_{k',l'=1}^{N_1^{in} \times N_2^{in}} \frac{\partial Y_{k,l}}{\partial X_{k',l'}} \frac{\partial X_{k',l'}}{\partial x_{i,j}}$$

Because of D.1, the second to last term is equal to $F_{k,l}$, if $(k', l') = (I_1(k), I_2(l))$ and 0 otherwise:

$$= \sum_{i',j'=1}^{N_1^{out} \times N_2^{out}} \mathcal{G}_{i',j'} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} \frac{\partial y_{i',j'}}{\partial Y_{k,l}} F_{k,l} \frac{\partial X_{I_1(k), I_2(l)}}{\partial x_{i,j}}.$$

Furthermore making use of D.5 and D.6 we get:

$$\begin{aligned} &= \sum_{i',j'=1}^{N_1^{out} \times N_2^{out}} \mathcal{G}_{i',j'} \frac{1}{N_1^{out} N_2^{out}} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} F_{k,l} e^{i2\pi \frac{(k-1)(i'-1)}{N_1^{out}}} e^{i2\pi \frac{(l-1)(j'-1)}{N_2^{out}}} e^{-i2\pi \frac{(i-1)(I_1(k)-1)}{N_1^{in}}} e^{-i2\pi \frac{(j-1)(I_2(l)-1)}{N_2^{in}}} \\ &= \frac{1}{N_1^{out} N_2^{out}} \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} F_{k,l} e^{-i2\pi \frac{(i-1)(I_1(k)-1)}{N_1^{in}}} e^{-i2\pi \frac{(j-1)(I_2(l)-1)}{N_2^{in}}} \sum_{i',j'=1}^{N_1^{out} \times N_2^{out}} \mathcal{G}_{i',j'} e^{i2\pi \frac{(k-1)(i'-1)}{N_1^{out}}} e^{i2\pi \frac{(l-1)(j'-1)}{N_2^{out}}} \\ &= \sum_{k,l=1}^{N_1^{out} \times N_2^{out}} F_{k,l} \text{idft2} \left[(\mathcal{G}_{i',j'}) \right]_{k,l} e^{-i2\pi \frac{(i-1)(I_1(k)-1)}{N_1^{in}}} e^{-i2\pi \frac{(j-1)(I_2(l)-1)}{N_2^{in}}} \end{aligned}$$

Instead of summing over the output indices $k, l = 1, \dots, N_1^{out} \times N_2^{out}$, we can zero-pad $F_{k,l} \text{idft2} \left[(\mathcal{G}_{i',j'}) \right]_{k,l}$ and sum over the input indices $k', l' = 1, \dots, N_1^{in} \times N_2^{in}$. The additional elements are 0 due to the zero-padding and will not contribute to the sum.

$$\begin{aligned} &= \sum_{k',l'=1}^{N_1^{in} \times N_2^{in}} \text{zp}_{I_1, I_2}^{N_1^{in}, N_2^{in}} \left[\left(F_{k,l} \cdot \text{idft2} \left[(\mathcal{G}_{i',j'}) \right]_{k,l} \right) \right]_{k',l'} e^{-i2\pi \frac{(i-1)(k'-1)}{N_1^{in}}} e^{-i2\pi \frac{(j-1)(l'-1)}{N_2^{in}}} \\ &= \text{dft2} \left[\text{zp}_{I_1, I_2}^{N_1^{in}, N_2^{in}} \left[\left(F_{k,l} \cdot \text{idft2} \left[(\mathcal{G}_{i',j'}) \right]_{k,l} \right) \right] \right]_{i,j}. \end{aligned} \quad (\text{D.8})$$

Derivative with respect to a filter parameter

Let us assume that the filter representation $(F_{k,l})$ depends on some parameter q . We derive a formula that allows to calculate the partial derivatives of all output elements $y_{m,n}$ with respect to q simultaneously using one dft2 and one idft2 operation.

$$\frac{\partial y_{i,j}}{\partial q} = \sum_{k',l'}^{N_1^{out} \times N_2^{out}} \frac{\partial y_{i,j}}{\partial Y_{k',l'}} \sum_{k'',l''}^{N_1^{out} \times N_2^{out}} \frac{\partial Y_{k',l'}}{\partial F_{k'',l''}} \frac{\partial F_{k'',l''}}{\partial q}$$

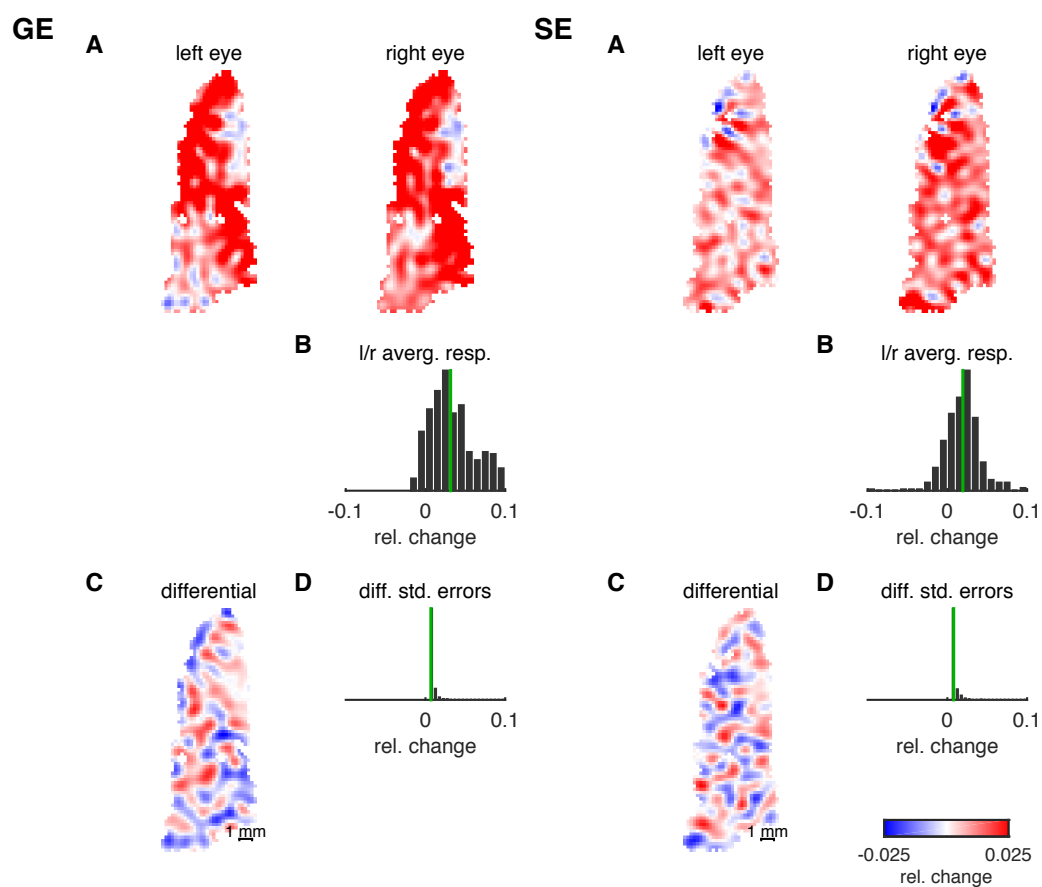
Because of D.1, the second to last term is equal to $X_{I_1(k'), I_2(l')}$ if $(k'', l'') = (I_1(k'), I_2(l'))$ and 0 otherwise:

$$= \sum_{k',l'}^{N_1^{out} \times N_2^{out}} \frac{\partial y_{i,j}}{\partial Y_{k',l'}} X_{I_1(k'), I_2(l')} \frac{\partial F_{k',l'}}{\partial q}$$

We apply D.6:

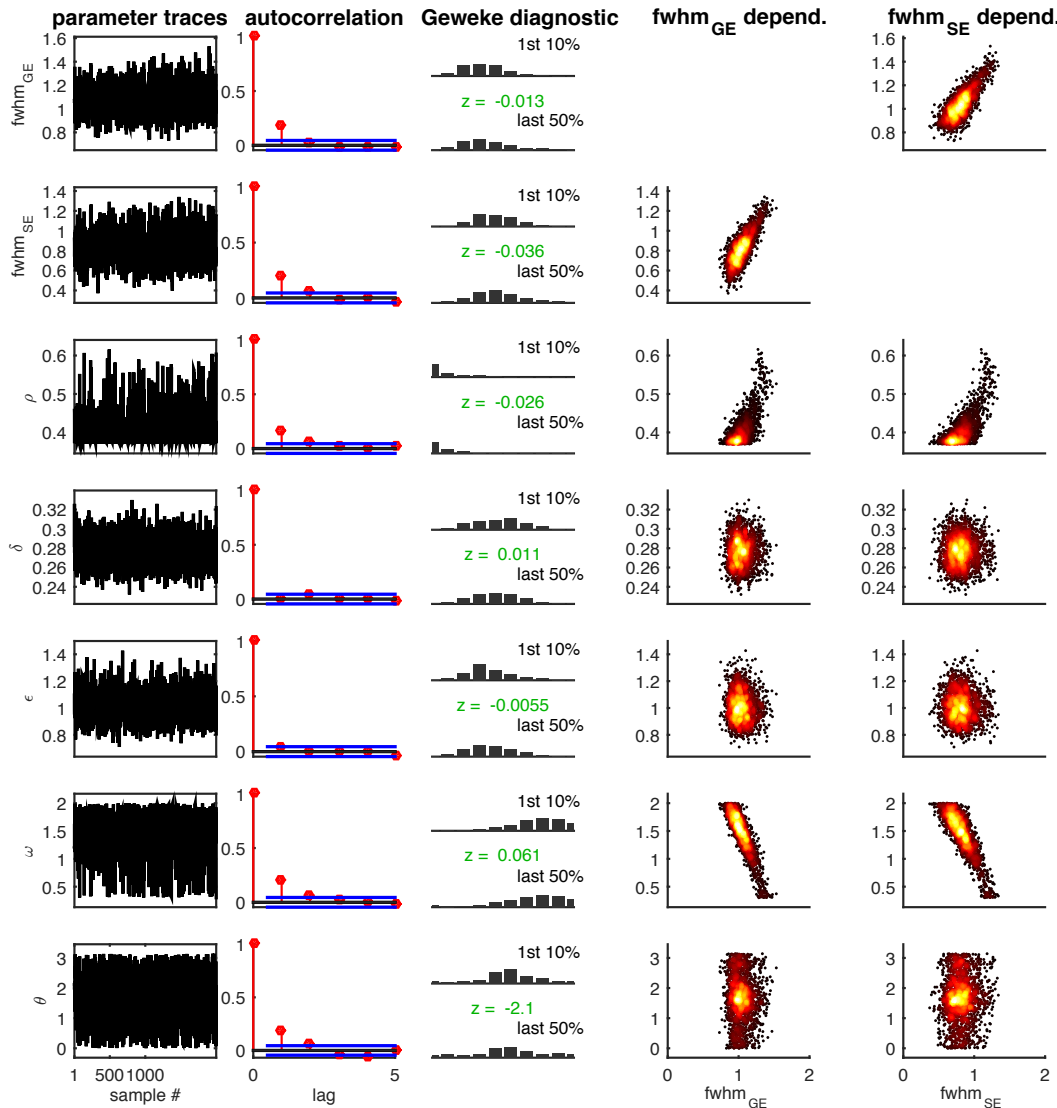
$$\begin{aligned} &= \sum_{k',l'}^{N_1^{out} \times N_2^{out}} \frac{1}{N_1^{out} N_2^{out}} e^{i2\pi \frac{(k'-1)(i-1)}{N_1^{out}}} e^{i2\pi \frac{(l'-1)(j-1)}{N_2^{out}}} X_{I_1(k'), I_2(l')} \frac{\partial F_{k',l'}}{\partial q} \\ &= \text{idft2} \left[\left(\text{dft2} \left[(x_{i',j'}) \right]_{I_1(k'), I_2(l')} \cdot \frac{\partial F_{k',l'}}{\partial q} \right) \right]_{i,j}. \end{aligned} \quad (\text{D.9})$$

Supplementary Material

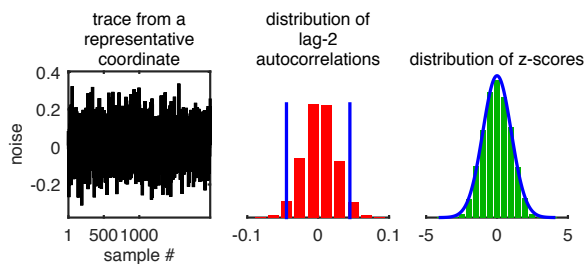


Sup. Fig. 1 fMRI ODC data. Results from the GLM analysis of fMRI data from subject 2 for GE (left) and SE (right). **A** Responses to left and right eye stimulation relative to baseline. **B** The response maps to the left and right eyes from **A** were averaged. **B** shows the distribution of the average response. Its median (in green) was used to set the overall amplitude of the BOLD response model. **C** The difference between left and right eye responses yields the differential ODC map. **D** The distribution of standard errors of all differential responses. From this distribution we estimated the noise level used by the model. The color look-up-table applies to all response maps.

A Parameters



B Noise values



Sup. Fig. 2 Convergence diagnostics of Markov Chain Monte Carlo sampling. Markov Chain Monte Carlo needs to run for a sufficient number of iterations in order to yield samples from the modeled probability distribution. Indications for convergence are: (1) stationarity of the parameter sampling distributions, and (2) sample autocorrelations decrease rapidly with increasing lag, relative to the total number of samples. This figure examines convergence for subject 2. The upper part of the figure (A) shows diagnostics for the standard model parameters. The bottom part (B) shows diagnostics for the white noise values that act as parameters to determine the ODC pattern. The first column (A and B) shows traces of the sampled parameters. For the noise values (B), one exemplary trace is shown from the center of the map. The second column (A and B) shows sample autocorrelations as a function of lag. The horizontal blue lines (A) indicate the 95%-confidence bounds around 0 for a white noise process. Consecutive samples

(lag=1) show low autocorrelation. However, samples of lag 2 (and higher) show autocorrelation estimates that are comparable to those obtained from uncorrelated white noise. In B, the autocorrelation from the noise samples are summarized by the histogram of lag-2 autocorrelations from all coordinates together. Here, 95%-confidence bounds for a white noise process are indicated by vertical blue lines. The third column (A and B) presents the Geweke convergence (stationarity) diagnostic, which is a z-test (z-scores shown in green) for testing whether the means of the first 10% and last 50% of samples are different. In A, 2 histograms per each parameter show how similar their respective distributions are. In B, the z-scores from the noise samples are shown as a histogram together with a blue plot of the standard normal probability density representing the null-hypothesis of $z=0$. The last two columns (A) show the sample covariation between each parameter (vertical axis) and the GE and SE point-spread function FWHM (horizontal axis).

STUDY 3

A more accurate account of the effect of k-space sampling and signal decay on the effective spatial resolution in functional MRI

Denis Chaimow and Amir Shmuel

The manuscript is about to be submitted for publication.

Author contributions

Denis Chaimow developed and implemented the model, and wrote the manuscript. Amir Shmuel advised on the development of the model and edited the manuscript.

Acknowledgements

We thank Avery Berman for his comments on an earlier version of the manuscript, and Nicky Tam for English editing. This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada (AS, NSERC Discovery grants RGPIN 375457-09 and RGPIN-2015-05103).

Abstract

The effects of k-space sampling and signal decay on the effective spatial resolution of MRI and functional MRI (fMRI) are commonly assessed by means of the magnitude point-spread function (PSF), defined as the absolute values (magnitudes) of the complex MR imaging PSF. It is commonly assumed that this magnitude PSF signifies blurring, which can be quantified by its full-width at half-maximum (FWHM). Here we show that the magnitude PSF fails to accurately represent the true effects of k-space sampling and signal decay.

Firstly, a substantial part of the width of the magnitude PSF is due to MRI sampling per se. This part is independent of any signal decay and its effect depends on the spatial frequency composition of the imaged object. Therefore, it cannot always be expected to introduce blurring. Secondly, MRI reconstruction is typically followed by taking the absolute values (magnitude image) of the reconstructed complex image. This introduces a non-linear stage into the process of image formation. The *complex* imaging PSF does not fully describe this process, since it does not reflect the stage of taking the magnitude image. Its corresponding *magnitude* PSF fails to correctly describe this process, since convolving the original pattern with the magnitude PSF is different from the true process of taking the absolute following a convolution with the complex imaging PSF. Lastly, signal decay can have not only a blurring, but also a high-pass filtering effect. This cannot be reflected by the strictly positive width of the magnitude PSF.

As an alternative, we propose to model the imaging process by decomposing it into a signal decay-independent MR sampling part and an approximation of the signal decay effect. We approximate the latter as a convolution with a Gaussian PSF or, if the effect is that of high-pass filtering, as reversing the effect of a convolution with a Gaussian PSF. We show that for typical high-resolution fMRI at 7 Tesla, signal decay in Spin-Echo has a moderate blurring effect (FWHM = 0.89 voxels, corresponds to 0.44 mm for 0.5 mm wide voxels). In contrast, Gradient-Echo acts as a moderate high-pass filter that can be interpreted as reversing a Gaussian blurring with FWHM = 0.59 voxels (0.30 mm for 0.5 mm wide voxels). Our improved approximations and findings hold not only for Gradient-Echo and Spin-Echo fMRI but also for GRASE and VASO fMRI. Our findings support the correct planning, interpretation, and modeling of high-resolution fMRI.

Introduction

The spatial specificity of functional MRI (fMRI) based on the Blood Oxygenation Level Dependent (BOLD) signal depends on the spatial properties of the hemodynamic response. Specifically, it depends on the relative contributions of the micro-vascular and macro-vascular components of the hemodynamic response to the fMRI signal. In addition to the effects of the hemodynamic response on the spatial specificity of fMRI, the MRI acquisition process influences the effective resolution of the acquired images. Specifically, the sampling of k-space by means of temporal gradient encoding defines the spatial resolution.

However, the effective spatial resolution can be compromised in the presence of T_2 and/or T_2^* decay, which potentially contribute to the overall measured spread of the BOLD fMRI signal.

The BOLD point-spread function (PSF) is a measure used to approximate the spatial spread of the BOLD response to a localized increase in neuronal activity. A convolution of the pattern of neuronal activity with a single BOLD PSF kernel is not a precise model of the spatial specificity of the BOLD response, because of the variability in the vascular components as a function of space (Polimeni et al., 2010). However, it provides a useful measure, based on the average BOLD PSF across space, for comparing the spatial specificity between different fMRI contrasts and techniques.

The full-width at half-maximum (FWHM) of the gradient echo (GE) BOLD PSF at 1.5 T was found to be 3.5 mm (Engel et al., 1997). Similar values of 3.9 mm for GE BOLD and 3.4 mm for Spin-Echo (SE) BOLD have been reported at 3 T (Parkes et al., 2005). We previously estimated the FWHM of the GE BOLD PSF to be below 2 mm at 7 T (Shmuel et al., 2007). Narrower BOLD PSFs at higher field strength are thought to result from reduced intravascular contributions from larger blood vessels and increases in extravascular signal changes around capillaries and smaller vessels (Yacoub et al., 2001). Additional relative weighting towards the microvasculature, and thus further increases in spatial specificity, can be achieved by using SE BOLD imaging, which suppresses extravascular signal contributions from larger blood vessels (Uludağ et al., 2009; Yacoub et al., 2003).

The use of high field strengths and developments in pulse sequences that lead to decreases in the BOLD fMRI PSF allow investigation into the function of ever finer structures such as cortical columns (Cheng et al., 2001; Goodyear and Menon, 2001; Menon et al., 1997; Nasr et al., 2016; Shmuel et al., 2010; Yacoub et al., 2008; 2007; Zimmermann et al., 2011;). Consequently, in order to optimize such experiments and to understand their inherent limitations, it becomes important to assess the contribution of the MRI sampling process and of T_2/T_2^* decay to the overall BOLD fMRI PSF.

The MR imaging process can be described by means of a complex valued MR imaging PSF (Haacke et al., 1999). The magnitude PSF (formed by the absolute values (magnitudes) of the complex MR imaging PSF) and the corresponding FWHM of the magnitude PSF, have been used to assess the effective spatial resolution and to quantify the blurring that the MR sampling process introduces (Constable and Gore, 1992; Farzaneh et al., 1990; Haacke, 1987; Kemper et al., 2015; Oshio and Singh, 1989; Qin, 2012).

Here, we simulate T_2/T_2^* decay and MR imaging of realistic columnar patterns to show that the FWHM of the magnitude PSF is neither a meaningful nor an accurate measure for quantifying the effect of MR sampling on the effective spatial resolution, especially in the context of functional MRI. As an alternative, we propose to decompose the modeling of the imaging process into two components: one component accounts for MR sampling, independent of the signal decay; a second component, formulated as a convolution with a Gaussian kernel, approximates the blurring effect due to the T_2/T_2^* decay.

Methods

Discrete representations of simulated spaces

All simulations were implemented in MATLAB (The MathWorks Inc., Natick, MA, USA). Spatial dimensions were considered relative to an arbitrary voxel width. A field-of-view (FOV) of 32 voxels was simulated, represented by 256 equally spaced points (resolution 8 times finer than the voxel width). The spatial frequency space (k-space) was simulated on a corresponding grid of 256 equally spaced points representing a spatial frequency range between -128 and +127 cycles per 32 voxels. Spatial frequencies sampled by MRI (see below) are represented by the central part of this simulated k-space, covering the spatial frequency range between -16 and +15 cycles per 32 voxels.

Modeling of signal decay

GE signal decay $f_{GE}(t)$ and SE signal decay $f_{SE}(t)$ were modeled according to (Haacke et al., 1999):

$$f_{GE}(t) = e^{-t/T_2^*}$$

$$f_{SE}(t) = \begin{cases} e^{-t/T_2^*} & 0 < t < T_E/2 \\ e^{-t/T_2} e^{-(T_E-t)/T_2'} & T_E/2 < t < T_E \\ e^{-t/T_2} e^{-(t-T_E)/T_2'} & t > T_E \end{cases}$$

where T_E represents the echo time. Relaxation time constants for gray matter at 7T were used (see Uludağ et al., 2009 for a review of relaxation times). T_2 was set to 50 ms. T_2^* of gray matter at 7T was 27.8 ms. In order to account for additional macroscopic inhomogeneities, a volumetric T_2^* value of 17 ms was used (Kemper et al., 2015).

Calculation of Modulation Transfer Functions

If not stated otherwise, a total read-out time of 27.8 ms (equal to T_2^*) for the full k-space acquisition using 32 lines (phase-encode steps) was assumed. For partial Fourier acquisition, the total read-out time was shortened, so that it was proportional to the reduction in k-space coverage, resulting in a total read-out time of 20.85 ms for the acquisition of 24 lines (3/4 partial Fourier).

For GE, the echo time T_E was set to the true T_2^* of 27.8 ms. The SE echo time was set to 55 ms (Yacoub et al., 2003). The modulation transfer functions (MTFs) of MR imaging are sums of Dirac delta functions (rect-function-windowed Dirac comb function), where each Dirac delta function is modulated by a signal decay factor (Appendix A). Discrete representations of MTFs were computed by first calculating the sampling time of each k-space line relative to excitation and then setting the line's MTF value to the signal decay value for this time, or to zero if it fell outside the range of sampled lines.

Calculation of point-spread functions

PSFs were calculated by taking an inverse discrete Fourier transform of the discrete representation of the corresponding MTF.

Simulation of responses of cortical columns

MR imaging simulations were applied to simulated realistic ocular dominance column (ODC) patterns and to simulated general isotropic columnar patterns.

ODC response patterns were simulated by anisotropic filtering of Gaussian white noise (Rojer and Schwartz, 1990). Detailed modeling equations can be found in (Chaimow et al., 2011). We modified the mathematical form of the band-pass filtering kernel relative to our previously published model, expressing it as a product of radial and angular components. The non-normalized filter as a function of absolute spatial frequency r and orientation ϕ is given as

$$\tilde{F}_{ODC}(r, \phi) = \left(e^{-\frac{(r-\rho)^2}{2\delta^2}} + e^{-\frac{(r-(-\rho))^2}{2\delta^2}} \right) \cdot \left(e^{\frac{\cos \phi - \theta}{\epsilon^2}} + e^{\frac{\cos \phi - (\theta + \pi)}{\epsilon^2}} \right).$$

Unless stated otherwise, the main spatial frequency parameter ρ was set to 0.5 cycles/mm, corresponding to an average column width of 1 mm (Yacoub et al., 2007). Parameters δ and ϵ controlled the degree of irregularities orthogonal and parallel to the main axis of elongation of the ODC bands, respectively. We specified those parameters by defining *relative irregularity* parameters as $\delta_{rel} = \frac{\delta}{\rho}$ and $\epsilon_{rel} = \frac{\epsilon}{\rho}$, that determined the level of irregularities independent of the chosen main spatial frequency. Unless stated otherwise, relative irregularity parameters δ_{rel} and ϵ_{rel} were set to 0.5 and 1 respectively. The orientation parameter θ was set to $\pi/2$, so that the main axis of elongation of the ODC bands was orthogonal to the phase-encode direction.

For 1D modeling (along the phase-encode direction), the 2D model was reduced by only considering the radial component of the filter (first factor), setting the angular component (second factor) to 1. This 1D model can be regarded as a general one dimensional columnar model, valid not only for anisotropic organizations such as ODC, but also for isotropic columnar patterns.

The sharpness parameter α was set to 1.4, resulting in a mid-level of sharpness. The default maximum response amplitude was set to 5%. Single condition columnar response patterns were added to a background signal intensity of 1. For the definition of the sharpness parameter and the response amplitude, see (Chaimow et al., 2011).

Simulation of MR imaging

MR imaging was modeled by multiplying the k-space representation of a pattern with the MTF and applying an inverse discrete Fourier transformation. This was followed by taking the magnitude of the resulting complex values. The size of the MTF was identical to that of the pattern k-space representation. However, the MTF was zero for k-space components not sampled by MRI imaging. Therefore the result of the discrete Fourier transform was a high-resolution representation of the voxel-size dependent MR sampled signal. It is equivalent to an interpolation using zero-filling in k-space. This high-resolution representation was then down-sampled to correspond to the actual voxel size, resulting in single voxel signals in accordance with the MR imaging equation (Appendix A).

Different MTFs introduce differences in the overall amplitude scaling of resulting images. To allow the comparison of the image patterns obtained by

considering different MTFs, we normalized the magnitude images in the last stage of the simulation. To this end, we divided each of the magnitude images by a constant equal to the result of simulating the entire MRI sampling process with the specific MTF applied to a constant pattern of value 1. Note that 1 is also the value of the homogeneous background onto which we superimposed the simulated ODC pattern with the maximal amplitude of 5%.

Simulation of Partial Fourier

Partial Fourier imaging and reconstruction was simulated by setting the first or last $\frac{1}{4}$ of the MTF components within the full acquisition sampling range to either zero (zero-filling reconstruction) or to their conjugate symmetric counterparts (conjugate symmetry reconstruction). The most negative k-value of the full acquisition sampling range has no positive k-space counterpart due to the slight asymmetric sampling of an even number of k-space lines. Therefore, in early omission partial Fourier and conjugate symmetry reconstruction, the most negative k-value was set to zero.

Implementation of approximating MRI models

Convolutions between simulated patterns and various kernel functions were implemented as multiplication of their respective discrete Fourier transforms or MTF representations, followed by inverse discrete Fourier transform back into image space. For Gaussian blurring, Gaussian kernels were computed as $e^{-x^2/2\sigma^2}$ where $\sigma = FWHM/2.355$.

Convolutions followed by MR sampling were modeled by first setting the MTF outside the range of the sampled lines to zero. Then, this modified MTF was multiplied with the k-space representation of the pattern, and an inverse discrete Fourier transformation was applied. Lastly, the magnitudes of the resulting complex values were computed. This high-resolution representation was then down-sampled according to the voxel size, resulting in single voxel signals according to the MR imaging equation (Appendix A).

Contrast range

The contrast range of each of 1000 simulated individual response patterns, superimposed on a background signal intensity level of 1, was computed as the standard deviation (SD) of all responses while taking into account the theoretical mean ($1 + \beta/2$, where β is the maximum response level) used for simulating the original pattern. All individual contrast range estimates were averaged resulting in the average contrast range.

Frequency spectra

Spatial frequency spectra were computed by taking the absolute value of the discrete Fourier transform separately for each of 1,000 individual response patterns. Average spatial frequency spectra were computed by averaging all individual spectra.

Evaluation of a linear approximation of the MR imaging process

We approximated the MR imaging process linearly as a convolution with the real component of the complex imaging PSF (see Appendix B). This linear approximation of the MR imaging process was evaluated by simulating 1000 different cortical columnar response patterns for each combination of a main spatial frequency (8 values from 1 cycle per Field of View (FOV) to 1 cycle per 2 voxels), a relative irregularity parameter (10 values from 0.1 to 1), and a range of maximum response amplitudes (1% - 10% in steps of 1% and 10% - 100% in steps of 10%). The complete MR imaging of each of these patterns was simulated. In addition, for each of these patterns we computed the linear approximation of the MR imaging process. The results of the linear approximations of the MR imaging process were compared to the complete MR imaging simulations by computing the root-mean-squared-errors (RMSE) relative to the standard deviation of the simulated patterns of the complete MR imaging process.

In addition, for a response amplitude of 5%, the same patterns were also convolved with the magnitude PSF. The results of convolving with the magnitude PSF were compared to the complete MR imaging simulations by computing the root-mean-squared-errors (RMSE) relative to the standard deviation of the simulated patterns of the complete MR imaging process.

Definition and fitting of a Gaussian point-spread function model for signal decay

The MTF corresponding to the real component of the complex imaging PSF (MTF_k^{real}) was calculated by transforming the real component of the complex imaging PSF back into the spatial frequency domain. The inverse MTF of this real component ($MTF_k^{real^{-1}}$) was calculated as $MTF_k^{real^{-1}} = 1/MTF_k^{real}$ for all k-space lines within the sampling range (and zero outside the sampling range).

Gaussian functions of the form $a \cdot e^{-2\pi^2\sigma^2k^2}$ were fitted to (MTF_k^{real}) and ($MTF_k^{real^{-1}}$) within the range of $-(k-1)$ to $(k-1)$. The amplitude a was constrained to be equal to the center component MTF_0^{real} or $MTF_0^{real^{-1}}$. This resulted in Gaussian PSFs of the form $PSF(x) = e^{-x^2/2\sigma^2}$, whose FWHM = $\sigma \cdot 2.355$. To fit the Gaussian functions, we used the MATLAB function 'fit' from the Curve Fitting Toolbox (The MathWorks Inc., Natick, MA, USA). The Gaussian fit (either to the real component of the complex imaging PSF or to its inverse) with the higher R^2 was considered for further characterizing the effects of signal decay.

Evaluation of the Gaussian point-spread function model of signal decay

The Gaussian PSF model of signal decay was evaluated by simulating 1000 different cortical columnar response patterns with a main spatial frequency of 1 cycle per 4 voxels and a relative irregularity of 0.5. The complete MR imaging of these patterns was simulated for different total read-out durations and partial Fourier acquisition schemes (including full k-space acquisition).

In addition, for each of these patterns and acquisition parameters (different total read-out durations and partial Fourier acquisition schemes), we comput-

ed the convolution of the pattern with a Gaussian PSF model for the signal decay (see definition and fitting described in the previous section). This was followed by MR sampling with no decay. The results of these approximations of the MR imaging process were compared to the complete MR imaging simulations by computing the root-mean-squared-errors (RMSE) relative to the standard deviation of the simulated patterns of the complete MR imaging process.

Estimation of a pattern specific Gaussian point-spread function model of signal decay

One thousand different cortical columnar response patterns with a main spatial frequency of 1 cycle per 4 voxels and a relative irregularity of 0.5 were simulated. MR imaging of these patterns was simulated for different total-read out durations and partial Fourier acquisition schemes (including full k-space acquisition). For each pattern, we computed convolutions with PSFs corresponding to Gaussian MTFs and inverse of Gaussian MTFs, while considering the FWHM as a free parameter. This was followed by MR sampling with no decay. Using MATLAB's `fminsearch` (The MathWorks Inc., Natick, MA, USA), the FWHM (and the choice of Gaussian MTF or inverse of Gaussian MTF) was optimized such that the mean-squared-error between the approximation and the full MR imaging simulation was minimized.

Results

The MR imaging point-spread function

We first summarize how the MRI acquisition process of a pattern can be described using PSFs. Appendix A provides detailed equations. The theory follows Haacke et al. (1999).

MRI with no signal decay

Echo-planar imaging (EPI) samples the two-dimensional k-space representation of the pattern by sequentially sampling individual lines along the first dimension (read-out direction), each separated by a step in the phase-encode direction. This results in a grid of sampled k-space points from which the original image is reconstructed using an inverse discrete Fourier transform. Since the dimensions in the Fourier transform are separable, we can focus on one dimension: the phase-encode dimension.

Assuming no signal decay takes place, we can formulate the MRI sampling process as an inverse Fourier transform of the product between the k-space representation of the pattern and a rect-function-windowed Dirac comb function (Figure 1, no decay, MTF). This function describes the effect of a linear system as multiplication in spatial frequency space and is commonly termed a modulation transfer function (MTF). Typically an even number of points, $N = 2n$ is sampled, resulting in a slightly asymmetric coverage of k-space over the region $[-n\Delta k, (n - 1)\Delta k]$, where Δk is the step size in the k-space.

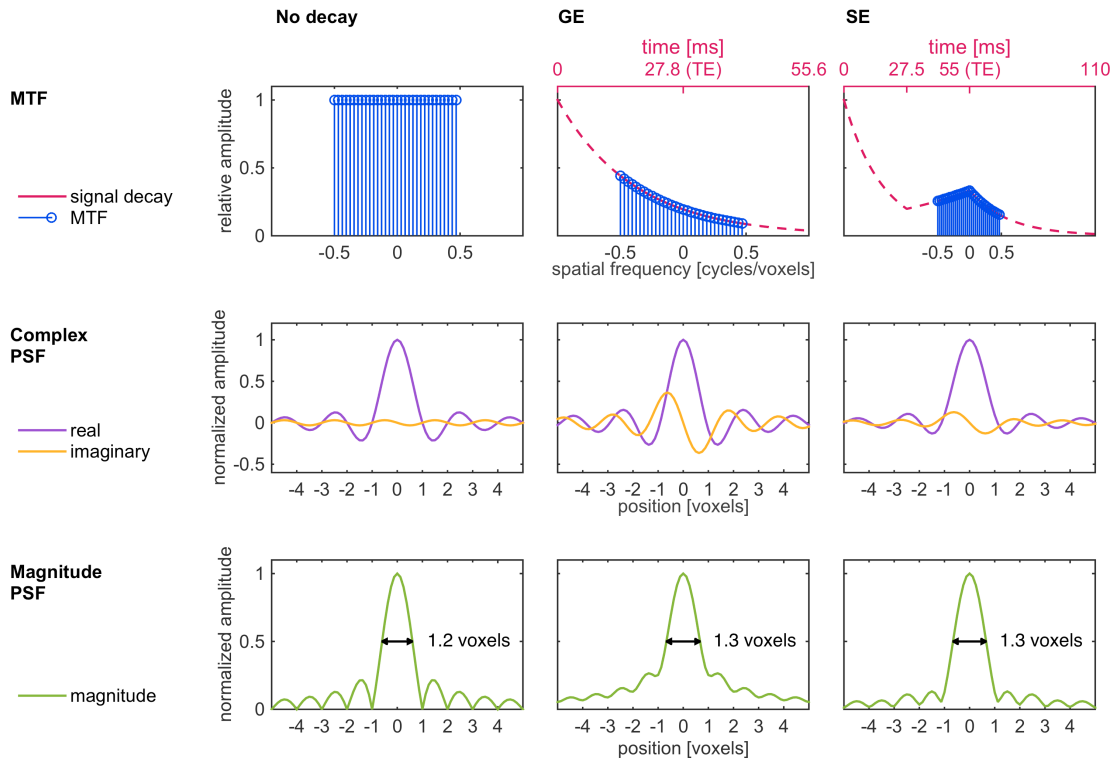


Fig. 1 MR complex imaging PSF and its absolute values. This figure demonstrates how MR imaging can be described using a PSF. The three columns illustrate the scenarios of no signal decay, signal decay in GE imaging, and signal decay in SE imaging, respectively. The first row shows the modulation transfer function (MTF) of the imaging process. In the case of no decay, the MTF is a Dirac comb function that corresponds to sampled k-space data points that are sampled in time from the lowest to the highest k-value. In GE imaging, this comb function is modulated by the T_2^* signal decay (red line). In SE a refocusing pulse results in a decay curve with a peak at the echo time. The second row shows the corresponding complex PSFs defined as the Fourier transform of the MTF. These functions show the complex signal one would obtain along the phase encoding direction from an infinitesimally small point-like structure. One can model the complex signal obtained by MRI along the phase-encoding direction as a convolution of the original pattern with the complex PSF. In contrast, in the general case one cannot model the magnitude of the signal obtained along the phase-encoding direction as the result of a convolution process. The third row shows the absolute (magnitude) values of the complex imaging PSF. Its FWHM (black arrows) has been previously used as a common measure to describe the spatial specificity of the MR imaging process. For relaxation time constants measured at 7 Tesla and a total readout duration of 27.8 ms, the FWHM of the magnitude PSF is 1.20 voxels for no decay, 1.34 voxels for GE, and 1.32 voxels for SE.

Multiplying the k-space data with an MTF is equivalent to convolving the image space data with the MTF's Fourier transform, which is the imaging point-spread function (PSF) (Figure 1, no decay, complex PSF). On its own, the PSF describes the image one would obtain from an infinitesimally small point-like structure.

The non-zero imaginary component of the no-decay PSF is a result of the above mentioned asymmetry in the MTF. MTF asymmetries are also caused by signal decay as described below. MTF asymmetries and other phase-influencing artifacts result in reconstructed images that are generally complex valued with non-zero phase. Commonly, the absolute values (magnitudes) of the complex image are considered for further analysis. Likewise, the spatial resolution of MRI and fMRI is often characterized by measuring the full-width at half-maximum of the magnitude PSF, obtained by taking the absolute values (mag-

nitudes) of the complex PSF (Figure 1, no decay, magnitude PSF). Here, without any signal decay effects the FWHM is 1.20 voxels.

However, neither the complex PSF nor the magnitude PSF can describe the spatial resolution of the MRI process correctly under all circumstances. The magnitude PSF in itself is not the PSF of the imaging process, because convolution with the absolute values of the complex PSF is not equivalent to taking the absolute values after convolution with the complex PSF (which is the common practice in MRI reconstruction). In contrast, the complex PSF does describe the imaging process (excluding the operation of taking the absolute values of the complex image). However, given its complex nature, how to use the complex PSF for quantifying the effective spatial resolution of the absolute values image is not obvious.

Signal decay in Gradient-Echo and Spin-Echo functional MRI

So far, we have not considered the change in signal strength with time following excitation. In GE imaging, the signal decays with a time constant of T_2^* (Figure 1, GE, MTF), which subsumes tissue dependent spin-spin relaxation (time constant T_2) and additional dephasing due to magnetic field inhomogeneities (time constant T_2'). In SE imaging, the signal similarly decays with time constant T_2^* . At half the echo time, however, a refocusing pulse causes reversal of the accumulated T_2' decay, while T_2 decay continues. After the echo time is reached, the signal returns to a decay with time constant T_2^* (Figure 1, SE, MTF).

During the decay, the k-space is sampled for the total acquisition time from the smallest (most negative) k-space value to the highest k-space value, such that the center of k-space ($k=0$) is sampled at the echo time. Note that this common sampling order, also called linear ordering, is not the only one possible. For example, in centric ordering, the sampling trajectory starts at the center of k-space and alternates between increasingly positive and decreasingly negative k-space coordinates. Here, we only consider linear ordering. For each k-space step in the phase-encode direction, an entire line along the read-out direction is acquired while the signal decays only minimally. Consequently, the effect of signal decay on the read-out direction can be neglected. However, in the phase-encode direction, the signal decay modulates the sampled data, causing different weighting of different spatial frequency components. The MTFs of GE and SE along the phase-encoding direction reflect this weighting (Figure 1, GE and SE, MTF). Therefore, the complex PSFs of GE and SE MRI differ from the complex PSF of the imaging process with no signal decay (Figure 1, GE and SE, complex PSF). The time constants used in our simulations reflect imaging at 7 Tesla; unless specifically mentioned otherwise, we considered a total readout duration of 27.8 ms (equal to T_2^* of gray matter). The FWHM of the magnitude PSFs were 1.34 voxels for GE and 1.32 voxels for SE, both larger than the FWHM of the magnitude PSF with no decay (1.20 voxels).

The Effect of the MRI process on the effective spatial resolution for imaging cortical columns

In the previous section we have discussed why the magnitude PSF does not correctly describe the MR imaging process. Does it follow then, that the FWHM of the magnitude PSF fails to accurately describe the effective spatial resolution of fMRI?

To address this question specifically in the context of imaging cortical columns, we simulated fMRI sampling of BOLD responses of patterns of ocular dominance columns (Chaimow et al., 2011; Rojer and Schwartz, 1990).

Figure 2 (simulated pattern, 2D pattern) shows a 2D modeled pattern and an excerpt from a 1D pattern (simulated pattern, 1D pattern excerpt). The 1D pattern follows the horizontal direction of the 2D pattern, here considered as the phase-encode direction.

These 2D and 1D modeled patterns represent a BOLD pattern, consisting of a spatially constant baseline signal of 1 and a superimposed ODC pattern-dependent BOLD response. The BOLD response can vary between 0% and 5% relative to the baseline value. We did not model the spatial spread of the BOLD response (meaning we assumed no spread) in order to not distract from the effects of the imaging process.

The main spatial frequency of the simulated ODC BOLD pattern was 0.5 cycles/mm (Yacoub et al., 2007), reflected as the maximum in the spatial frequency spectrum (Figure 2, simulated pattern, average spatial frequency spectrum, vertical blue lines). The irregularity of the pattern is reflected in a distribution of additional spatial frequency contributions around the two maxima.

In order to quantify the functional contrast of true or imaged responses, we defined the contrast range as the standard deviation around the average response (the average response is defined as one half of the maximum response relative to the background intensity, namely 1.025). The contrast range averaged over 1,000 simulated one-dimensional patterns, was 1.30% (relative to the baseline of 1).

MRI with no signal decay

First we analyzed the effect of MR imaging with no signal decay. We simulated MRI sampling of the simulated original pattern (Figure 2, Simulated pattern) using a voxel width of 0.5 mm. Figure 2 (No decay, 2D pattern) shows the corresponding imaged two-dimensional pattern. Figure 2 (No decay, 1D pattern excerpt) shows the values of the imaged and interpolated one-dimensional pattern (red). The MR imaged pattern is shown as a high-resolution representation, which is equivalent to an interpolation using zero-filling in k-space. This was done in order to facilitate visual comparison to the original pattern (the alternative presentation format would consist of a discrete function due to voxelization). The imaged pattern was very similar to the original pattern (blue). The average contrast range computed over 1,000 imaged patterns decreased only slightly from 1.30% (average SD of 1,000 original patterns) to 1.29% (average SD of 1,000 no-decay MRI patterns).

In addition, we compared the average frequency spectrum of the imaged patterns to the average spectrum computed over the original patterns (Figure

2, no decay, average spatial frequency spectrum). Within the sampled k-space range (spatial frequencies below 1 cycle/mm), the two spectra were identical. Outside of this range, the average spectrum of the imaged patterns was zero.

Note, however, that the FWHM of the magnitude PSF corresponding to MRI with no signal decay was 1.20 voxels (= 0.6 mm in our specific simulation). This result could be wrongly interpreted to imply that MR image formation, ignoring decay, is comparable to blurring with a kernel (e.g. a Gaussian) of the same width.

Figure 2 (no decay, 1D pattern excerpt, orange curve) shows that in contrast to actual MRI sampling, such blurring would have resulted in a reduced amplitude (contrast range = 0.95%) and a shift in the spatial frequency spectrum to 0.4375 cycles/mm (Figure 2, no decay, average spatial frequency spectrum, orange lines). Using the magnitude PSF as a convolution kernel in itself (although, in fact, it is not a convolution kernel of the MRI process) resulted in an even larger reduction in contrast (0.45%, Figure 2, no decay, 1D pattern excerpt, green) and a shift of the spatial frequency distribution towards lower frequencies (maximum at 0.3750 cycles/mm, Figure 2, no decay, average spatial frequency spectrum, green).

MR imaging in the presence of signal decay: Gradient-Echo imaging

Next, we analyzed the effect of signal decay. Figure 2 (GE, 2D pattern) shows the GE imaged two-dimensional pattern. There is no noticeable blur relative to the no-decay image. In fact, the interpolated one-dimensional imaged pattern (Figure 2, GE, 1D pattern excerpt, red) showed a higher amplitude compared to the original pattern (blue). The average contrast range increased from 1.30 % (original) to 1.41% (GE).

This increase in contrast did not result in a noticeable difference in the peak spatial frequency, which remained at 0.5 cycles/mm (Figure 2, GE, average spatial frequency spectrum; the resolution we employed in our simulated k-space was 0.0625 cycles/mm). However, relative to the average spectrum of the original patterns, the amplitude at the peak spatial frequency increased, while the amplitudes at spatial frequencies close to 0 cycles/mm remained constant. This shows that GE imaging had the effect of a moderate high-pass filter. Similar to our conclusion for the imaging with no signal decay, this result could not be expected by simply considering the positive FWHM of the magnitude PSF, which was 1.34 voxels (= 0.67 mm in our specific simulation).

A convolution with a Gaussian of the same width resulted in contrast reduction (Figure 2, GE, 1D pattern excerpt, orange). The average contrast range dropped from 1.30% to 0.89%. Also, the peak in the spatial frequency spectrum shifted to a lower frequency of 0.375 cycles/mm. Convolution of the original pattern with the magnitude PSF resulted in a larger reduction in contrast (average contrast range 0.33%) and a shift of the peak spatial frequency towards lower frequencies (0.375 cycles/mm).

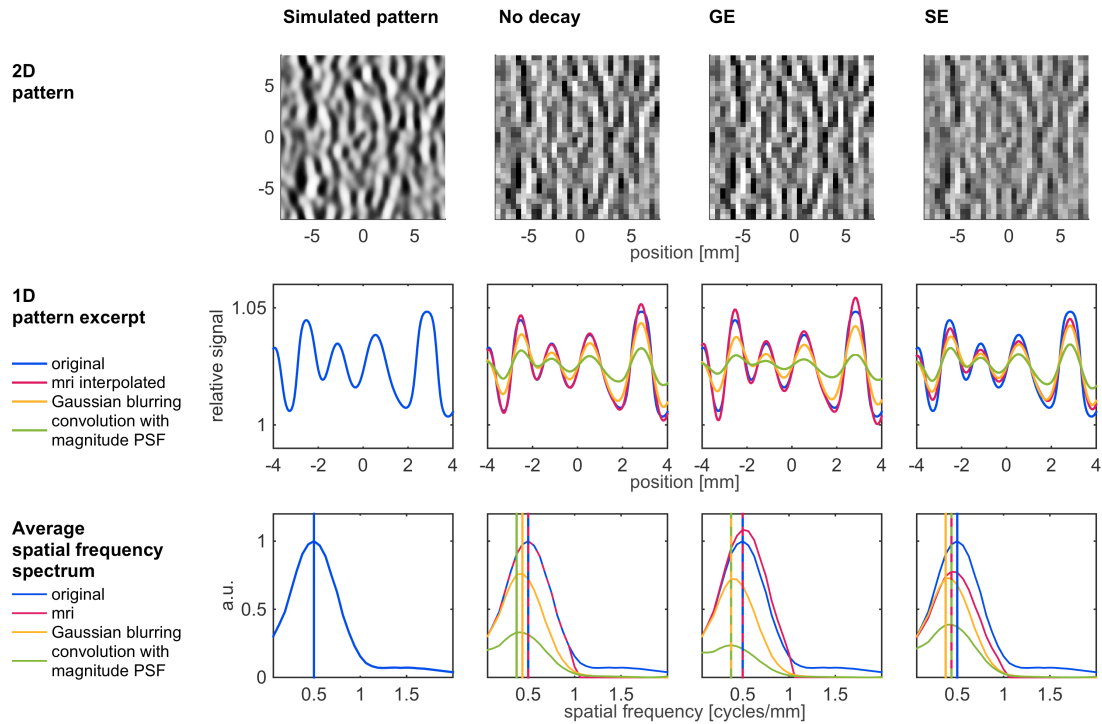


Fig. 2 The effect of imaging PSF on MRI of a columnar pattern. We simulated and analyzed columnar ocular dominance patterns (first column) and their MR imaging with no signal decay (second column), GE imaging (third column) and SE imaging (fourth column). The first row shows the simulated original 2D pattern and the resulting MR images using a voxel size of 0.5 mm. The second row shows an excerpt from a 1D simulated pattern (blue) and its no-decay, GE, and SE imaged counterparts (red). The 1D imaged patterns (which could be presented as non-continuous functions due to voxelization) were interpolated using zero-filling in k-space in order to facilitate comparison. The no-decay imaged pattern was very similar to the original pattern. The GE and SE imaged patterns showed slight increases and decreases, respectively, in contrast. In addition, we simulated convolutions (in orange) of the original pattern with Gaussian PSFs of the same widths as those computed from the magnitude PSFs (black arrows in Fig. 1) and convolutions with the magnitude PSF itself (in green). They all show lower contrast than that in the corresponding MR simulations. The third row shows spatial frequency spectra averaged from 1,000 simulated 1D patterns. The spatial frequency showing the maximal amplitude in each spectrum is marked with a vertical line with corresponding color. In cases where segments of the spectra obtained from the original pattern and the pattern obtained by MRI sampling were identical, we present alternating dashed red (for MRI sampling) and blue (for the original pattern) curves. Similarly, we present alternating dashed blue and red vertical lines in cases for which the frequencies showing the maximal amplitude were identical across the original pattern and MRI sampling. Imaging with no decay did not change the frequency spectrum within the sampled range (up to 1 cycle/mm). For GE, no change in the frequency showing the maximal amplitude was detected at the resolution we applied. SE imaging resulted in a slightly lower spatial frequency showing the maximal amplitude. In contrast, all spectra obtained by the Gaussian convolution (orange) and by the magnitude PSF convolution (green) showed lower spatial frequencies associated with the maximal amplitudes.

MR imaging in the presence of signal decay: Spin-Echo imaging

Figure 2 (SE, 2D pattern) shows the SE imaged two-dimensional pattern. The SE image is slightly blurred relative to the image obtained with no signal decay. Similarly, the interpolated one-dimensional pattern excerpt (Figure 2, SE, 1D pattern excerpt, red) shows a lower amplitude compared to the original pattern (blue). The average contrast range decreased from 1.30% (original) to 1.01% (SE) and the peak spatial frequency shifted to a lower frequency of 0.4375 cycles/mm (Figure 2, SE, average spatial frequency spectrum).

The reduction in contrast and peak spatial frequency suggests that SE imaging had a blurring effect on the original pattern, which could be consistent with the spatial extent of its magnitude PSF (FWHM of 0.66 mm). However, convolution with a Gaussian of the same width resulted in even larger contrast reductions (Figure 2, SE, 1D pattern excerpt, orange). The average contrast range decreased further to 0.90% and the peak in the spatial frequency spectrum shifted to 0.375 cycles/mm. True also for SE, convolution with the magnitude PSF reduced the contrast (0.52%) more than the MRI simulation and the convolution with a Gaussian kernel (of the same width as the width of the magnitude PSF) did. The peak spatial frequency remained at 0.4375 relative to the SE MR simulation.

An alternative approach to quantifying the effect of MR imaging on the effective spatial resolution

Convolution with the real component of the complex imaging PSF linearly approximates the complete MRI process

We have shown that neither the linear process of convolution with the complex PSF (without taking the absolute) nor the linear process of convolution with the magnitude PSF can faithfully describe the entire non-linear MR imaging and reconstruction process. To characterize the complete MRI process using a PSF, we propose an alternative, optimal, linear approximation.

The best linear approximation of a function around a point x_0 is the derivative of the function at x_0 . For a function of a single variable $f(x)$, the derivative $f'(x)$ represents a tangent line which can be interpreted as a linear approximation to $f(x)$ by mapping small deviations $x_0 + \Delta x$ onto $f(x_0) + \Delta x \cdot f'(x_0) \approx f(x_0 + \Delta x)$.

In the case of the MR imaging and reconstruction process, the function under consideration is not a function of a single variable but a functional, which maps a pattern onto a set of imaged voxel values. The derivative of this functional is a linear transformation that itself depends on a baseline pattern (the *point* at which the derivative is evaluated) in the same manner that a tangent depends on the point (x_0) at which it is defined. Similar to a tangent line, this linear transformation maps a pattern with small deviations from the baseline pattern onto a set of voxel values approximating the true imaged pattern.

In Appendix B, we computed this derivative for a spatially constant baseline pattern. This derivative is a linear transformation that approximates the complete MR process for other patterns, with small response deviations relative to this spatially constant baseline pattern. In appendix B we show that this linear approximation is identical to a convolution with the real component of the complex imaging PSF (presented in the middle row of Figure 1, in purple).

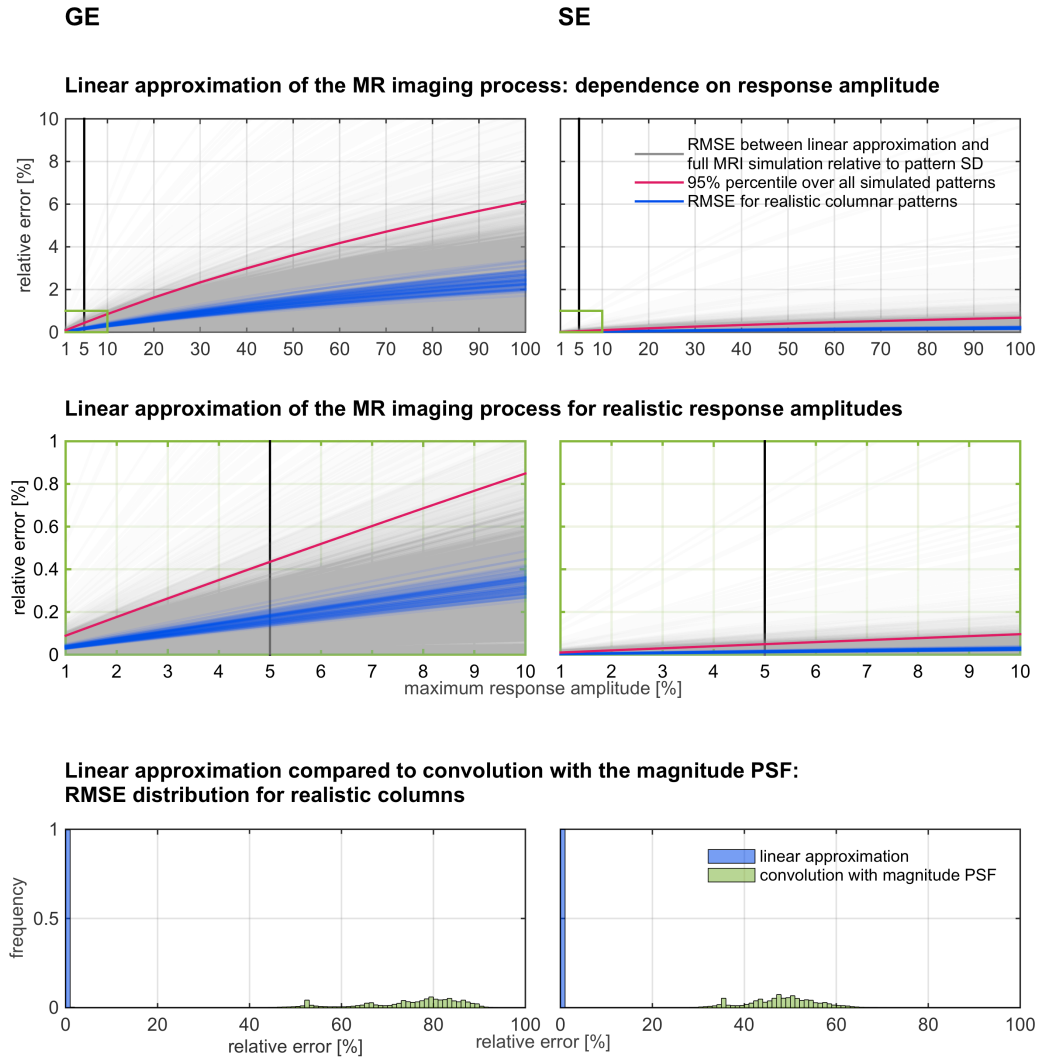


Fig. 3 Evaluation of the linear approximation of the complete MRI process. The upper row shows the dependence of the linear approximation for GE imaging (left) and SE imaging (right) as a function of maximum response amplitude. Cortical columnar response patterns for a wide range of spatial frequency parameters, irregularity parameters, and maximum response amplitudes were simulated (1,000 for each parameter combination; results from 100 patterns were used for visualization). MR imaging of these patterns was simulated. In addition, a linear approximation of the MR imaging process consisting of a convolution of the pattern with the real component of the complex imaging PSF was computed. The results of the linear approximations were compared to the complete MR imaging simulations by means of the root-mean-squared-errors (RMSE) relative to the standard deviation of the images obtained by the complete MRI process. The gray curves show the distribution of relative RMSEs from all simulated patterns; the red curve presents their 95 percentile. The blue curves show results from realistic columnar parameters (intermediate irregularity and spatial frequency). All curves show increased errors with increasing response amplitude. For the default, realistic response amplitude level of 5% (indicated by a vertical black line), most relative errors were well below 1%. The middle row of panels presents a magnified view of the errors for such realistic response amplitudes. The bottom part of the figure compares the distribution of relative RMSEs obtained by the linear approximation with those obtained by a convolution with the magnitude PSF. The default, realistic response amplitude level of 5% was used for this comparison. The convolution with the magnitude PSF resulted in substantially higher errors.

Next, we evaluated how well this convolution approximates the true MR imaging and reconstruction process. In particular, because the linear approximation is expected to be valid for small deviations around the spatially constant baseline pattern, we quantified the dependence of the quality of the approximation on the response amplitude. To this end, we simulated a wide

range of columnar patterns with different response amplitudes. We then compared the results obtained by simulating the full MRI process for each of these patterns to those obtained by a convolution of the pattern with the real component of the complex imaging PSF (Fig. 3). Note that as expected, for large amplitude deviations from the spatially constant baseline pattern, the linear approximation can result in relatively large errors (e.g. 95th percentile of the RMSE were 6.13% and 0.68% for 100% response amplitude imaged with GE and SE fMRI, respectively; Figure 3, upper row). However, the root-mean-squared error for realistic response amplitudes was small. For example, for the maximal response amplitude of 5% imaged with GE fMRI, the 95th percentile was 0.43% of the standard deviation of the pattern obtained by the full MRI simulation (Figure 3, middle row). For SE imaging, the relative RMSE was even lower; for the maximal response amplitude of 5% imaged with SE fMRI it was 0.05% (Figure 3, middle row). In contrast, convolution of the same response patterns (with response amplitude of 5%) with the magnitude PSF resulted in median relative root-mean-squared errors of 78% (GE) and 48% (SE) (Figure 3, bottom row).

Quantifying the effect of signal decay by fitting a two-component model consisting of convolution with a Gaussian PSF followed by MR sampling with no decay

In itself, the real component of the complex imaging PSF, in particular its width, is not suited to characterize the effective spatial resolution of the MR imaging process. The reason is that it represents not only signal decay, but also the MR sampling process. As we have shown, the latter is pattern-dependent and irrelevant if the voxels are sufficiently small to sample the spatial frequency spectrum, such that the imaged pattern is similar to the original pattern. Furthermore, it is not possible to easily discriminate the blurring characteristics of SE from the high-pass filtering characteristics of GE on the basis of the real component of the complex imaging PSFs.

However, instead of considering the real components of the complex imaging PSFs, we can consider their spatial frequency representations. Figure 4 (second row) shows the MTF of the real component of the complex imaging PSF for GE and SE (blue). The MTF of the real component is equal to the average of the positive and negative components of the original imaging MTF (presented in Figure 4, first row), assigned in a mirror-symmetric manner around the center ($k=0$) of the k -space.

The MTF of the real component of the complex imaging PSF can be regarded as the product of two factors representing two different processes. The first factor is the rect-function-windowed Dirac comb function that describes the MRI sampling with no signal decay. The second factor is a modulation of the sampled signal due to signal decay. In order to model this second factor we can choose a function that fits the modulation within the sampled k -space range. The Fourier transform of this function will be a convolution kernel in the image space, which describes the effect of signal decay. It does not describe the effect of MRI sampling with no signal decay, which is qualitatively different and depends on the imaged pattern.

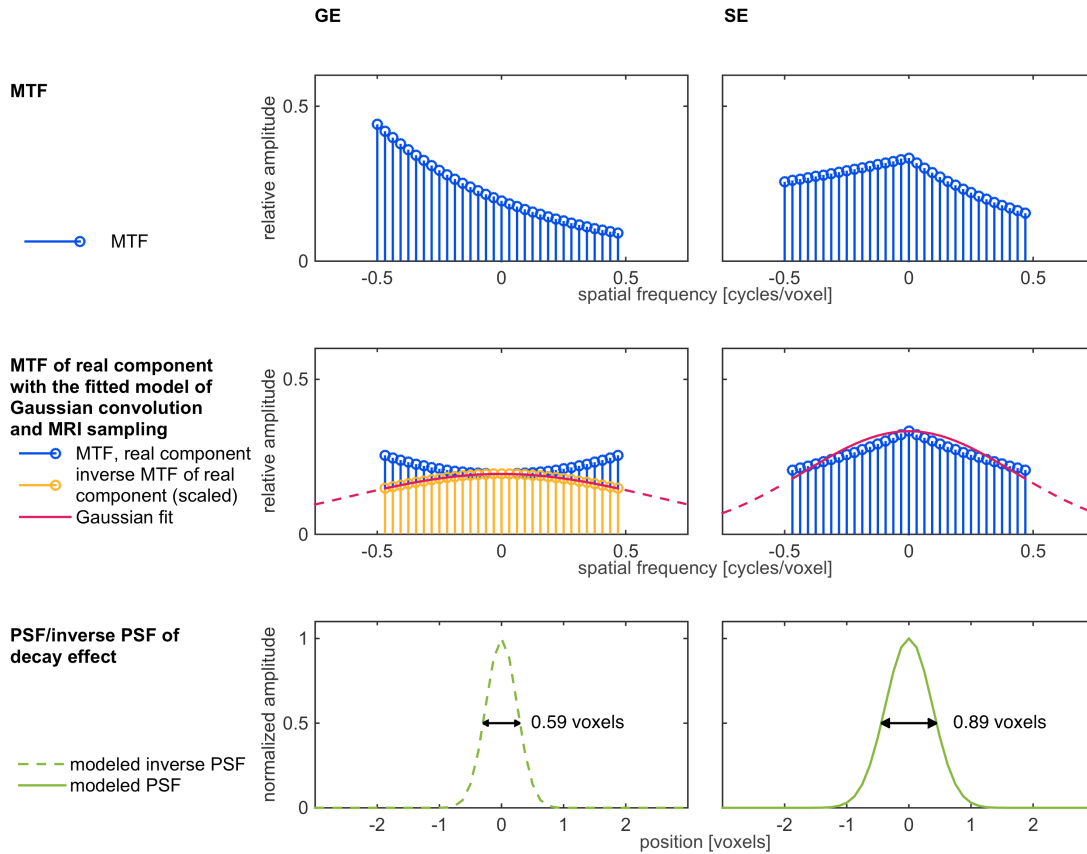


Fig. 4 Fitting of a two-component model consisting of convolution with a Gaussian PSF that accounts for signal decay followed by MR sampling with no decay. This figure shows the fitting of a Gaussian convolution and MRI sampling model for GE imaging (first column) and SE imaging (second column). Gaussian functions (second row, red) were fitted to the MTF of the real component of the SE image PSF (second row, right side column, in blue) and to the scaled inverse MTF of the real component of the GE image (second row, left side column, in orange). Outside the sampled k-space range, the continuation of the Gaussian fit is shown as a dashed red line. The Fourier transforms of these functions are Gaussian PSFs (bottom row, in green). For SE, this PSF describes the blurring due to the signal decay. For GE, it describes a hypothetical blurring that would be reversed by the high-pass filter properties of the T_2^* decay effect. Their respective FWHM (black arrows) are 0.59 voxels (GE) and 0.89 voxels (SE).

We can then apply MRI sampling with no signal decay. The image we get from the convolution that accounts for the signal decay and the MRI sampling is either identical to the complete MRI process (depending on the function fitted to the k-space representation in its sampled range) or approximates it.

For SE imaging, we chose a Gaussian function and fitted it to the signal decay-dependent modulation of the MTF of the real component of the complex imaging PSF (= the second factor; Figure 4 second row, SE, red; See discussion for justification of modeling the second factor as a Gaussian function). An inverse Fourier transformation of this fitted Gaussian results in a Gaussian PSF in the image space, allowing for the interpretation of the signal decay effect as Gaussian blurring. The FWHM of this Gaussian PSF was 0.89 voxels (for a total read duration of 27.8 ms at 7T).

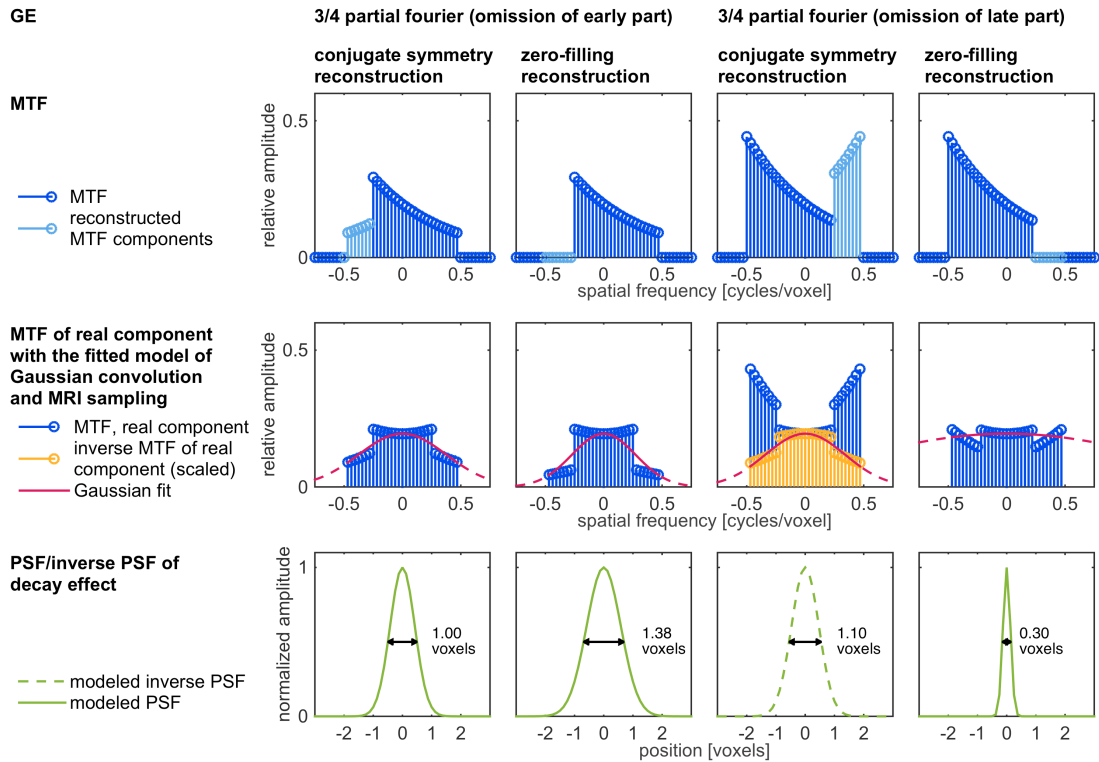


Fig. 5 Fitting of the two-component model for GE partial Fourier acquisition. This figure shows the fitting of a two-component model for partial Fourier acquisition using GE imaging. Omission of the first $\frac{1}{4}$ (columns 1 and 2) and last $\frac{1}{4}$ (columns 3 and 4) of phase-encode steps were simulated. Furthermore, a reconstruction that exploits conjugate symmetry (columns 1 and 3) was compared to a zero-filling reconstruction (columns 2 and 4). The first row shows the imaging MTF resulting from measurement components (dark blue) and reconstruction components (light blue). Gaussian functions (second row, in red) were fitted to the MTF of the real component of the imaging PSF (early-omission partial Fourier, columns 1 and 2, and late-omission partial Fourier using zero-filling reconstruction, column 4, in blue) and to the inverse MTF (here scaled for clarity of presentation) of the real component of the imaging PSF (late-omission partial Fourier using conjugate symmetry reconstruction, column 3, in orange). Outside the sampled k-space range, the continuation of the Gaussian fit is presented as a dashed red line. The Fourier transforms of these functions are Gaussian PSFs (bottom row, in green). For early-omission partial Fourier (columns 1 and 2) and late-omission partial Fourier using zero-filling reconstruction (column 4), these PSFs describe the blurring due to the signal decay. For late-omission partial Fourier using conjugate symmetry reconstruction (column 3), the PSF describes a hypothetical blurring that would be reversed by the high-pass filter properties of the T_2^* decay effect.

For GE imaging, a Gaussian function is not a good fit, since the MTF of the real component of the complex imaging PSF shows increasing amplitudes with increasing spatial frequency (Figure 4, second row, GE, blue), consistent with its high-pass filtering properties we have shown above (Figure 2). However, instead of this MTF, we can consider its inverse (1 divided by the MTF; Figure 4, second row, GE, orange). The inverse MTF describes the process that would be reversed by a convolution with the real component of the complex PSF. We fitted a Gaussian to the signal decay-dependent modulation of this inverse MTF (which resulted in a higher R^2 , than the Gaussian fit to the non-inverted MTF) and calculated its corresponding Gaussian PSF in the image space. This allows for interpreting GE imaging as reversing (deconvolving) a Gaussian blur. The FWHM of this Gaussian was 0.59 voxels (for a total read duration of 27.8 ms at 7T).

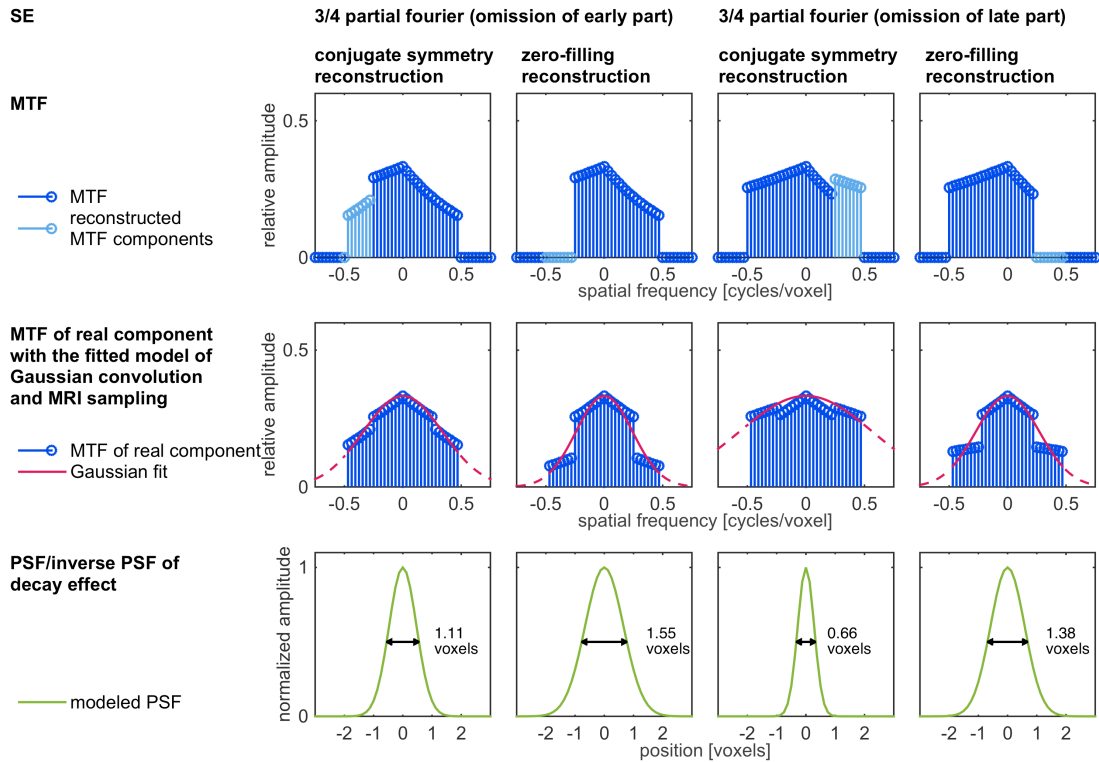


Fig. 6 Fitting of the two-component model for SE partial Fourier acquisition. This figure shows the fitting of a two-component model for partial Fourier acquisition using SE imaging. Omission of the first $\frac{1}{4}$ (columns 1 and 2) and last $\frac{1}{4}$ (columns 3 and 4) of phase encode steps were considered. Furthermore, a reconstruction that exploited conjugate symmetry (columns 1 and 3) was compared to a zero-filling reconstruction (columns 2 and 4). The first row shows the imaging MTF resulting from measurement components (dark blue) and reconstruction components (light blue). Gaussian functions (second row, in red) were fitted to the MTF of the real component of the imaging PSF for all columns 1-4. Outside the sampled k-space range, the continuation of the Gaussian fit is presented as a dashed red line. The Fourier transforms of these functions are Gaussian PSFs (bottom row, in green) that have a blurring effect. These PSFs describe the blurring due to the signal decay.

Partial Fourier acquisition

In addition to the standard EPI acquisition described so far, one can shorten the total read-out duration by only acquiring parts of the conjugate symmetric k-space. This is known as partial Fourier acquisition.

In order to study how signal decay affects partial Fourier acquisition, we simulated MTFs resulting from partial Fourier imaging in which either the first $\frac{1}{4}$ or last $\frac{1}{4}$ of the phase-encode steps were omitted (the total read out duration was shortened accordingly from 27.8 ms to 20.85 ms). We then applied the same modeling methodology described above for the full k-space acquisitions (Figures 5 and 6 for GE and SE, respectively). In addition, we compared a reconstruction that exploited the conjugate symmetry (Figures 5 and 6, columns 1 and 3) to a reconstruction with simple zero-filling (Figures 5 and 6, columns 2 and 4).




Fig. 7 FWHMs of Gaussian PSFs that model the effect of signal-decay as a function of total read-out duration. This figure shows the results of fitting our two component model that accounts for MR sampling and signal decay for different GE (left) and SE (right) imaging scenarios and for different total read-out durations. The FWHM of the fitted Gaussian PSFs are presented as a function of total-read out duration. For Partial Fourier acquisition, the total-read out duration axis is labeled at the top. Partial Fourier acquisition total-read out durations were shortened according to the fraction of omitted k-space (1/4) relative to the corresponding full k-space total read-out duration (bottom axis). The vertical black line represents a total-read out duration of 27.8 ms (20.85 ms for partial Fourier). Negative FWHM values indicate that the Gaussian PSFs resulted from the inverse of the MTF of the real component of the complex imaging PSF. Such negative values represent a hypothetical blurring that is reversed by the high-pass filter properties of the T_2^* decay effect.

For GE, partial Fourier with early omission resulted in blurring (Figure 5 columns 1 and 2) as opposed to the high-pass filtering observed in full k-space acquisition (Figure 4). The blurring was more substantial in zero-filling reconstruction (FWHM = 1.38 voxels) than in conjugate symmetry reconstruction (FWHM = 1.00 voxels).

Partial Fourier with late omission using conjugate symmetry reconstruction resulted in high-pass filtering (Figure 5, column 3). This high-pass filtering effect (reverse kernel FWHM = 1.10 voxels) increased relative to that of full k-space acquisition (reverse kernel FWHM = 0.59 voxels). Partial Fourier with late omission using zero filling reconstruction (Figure 5, column 4) resulted in moderate low-pass filtering (FWHM = 0.30 voxels).

For SE, partial Fourier with early omission (Figure 6, columns 1 and 2) resulted in increased blurring relative to that shown by the full k-space acquisition (FWHM = 0.89 voxels). The blurring was more substantial in zero-filling reconstruction (FWHM = 1.55 voxels) than in conjugate symmetry reconstruction (FWHM = 1.10 voxels). Partial Fourier with late omission also resulted in blurring (Figure 6, columns 3 and 4). Here, conjugate symmetry and zero-filling reconstructions resulted in decreased (FWHM = 0.66 voxels) and in-

creased (FWHM = 1.38 voxels) blurring, respectively, relative to the blurring obtained from the full k-space acquisition (FWHM = 0.89 voxels).

All results so far were based on an assumed total readout duration of 27.8 ms. In order to extend our results and to study the dependence of signal decay blurring on the total read out duration, we repeated the simulation of signal decay-dependent MTFs and the fitting of our model of Gaussian convolution and MRI sampling for a range of total readout durations (Figure 7).

For almost all imaging scenarios, including GE and SE imaging, the type of decay-dependent effect (blurring or high-pass filtering) was independent of total readout duration. However, the effect's strength increased with increasing total readout duration. The only exception was the zero-filling reconstruction of GE partial Fourier imaging that omits the late acquisitions. Here, short total readout durations resulted in blurring while longer total readout durations resulted in high-pass filtering.

Evaluation of modeling the complete MR process as a convolution with a Gaussian PSF that accounts for signal decay followed by MR sampling with no decay

Lastly, we evaluated how well our simplifying model approximated a complete MRI acquisition model (Figure 8). We compared complete simulations of fMRI of columnar patterns including signal decay to simulations that used our Gaussian PSF model of signal decay followed by MR sampling with no decay. We then quantified their deviations by calculating the root-mean-squared errors relative to the standard deviation of the patterns obtained by the simulation of the complete MRI process. The relative errors differed for different imaging methods and total readout durations. For full k-space acquisition (Figure 8, upper row), the median error (red curve) obtained for a total readout duration of 27.8 ms was 3.91% and 8.39% using GE and SE fMRI, respectively.

The relative errors obtained for the majority of partial Fourier acquisition schemes were substantially higher. For conjugate symmetry reconstruction (Figure 8, middle row), the median error computed for GE fMRI and a total readout duration of 20.85 ms (that corresponds to 27.8 ms for full k-space acquisition) was 17.55% and 11.6% for omissions of the early and late part of the k-space, respectively. The corresponding errors for SE fMRI were 5.79% and 9.51%.

For zero filling reconstruction (Figure 8, bottom row), the median error computed for GE fMRI and a total readout duration of 20.85 ms was 31.15% and 14.19% for omissions of the early and late part of the k-space, respectively. The corresponding errors for SE fMRI were 18.14% and 16.37%.

Our Gaussian PSF model of signal decay approximates the signal decay effect in MR imaging as a pattern-independent linear process. As we have just demonstrated, the differences between the results of this approximation and the true imaging process are low for full k-space acquisitions, but are higher for partial Fourier acquisitions. In order to obtain an even better approximation, we can define an alternative approximation. Given a specific pattern, we can determine a Gaussian PSF (or its inverse process) that accounts for signal decay, such that convolution with this specific Gaussian followed by MR sam-

pling with no decay results in the best Gaussian-based approximation of the complete MR imaging process.

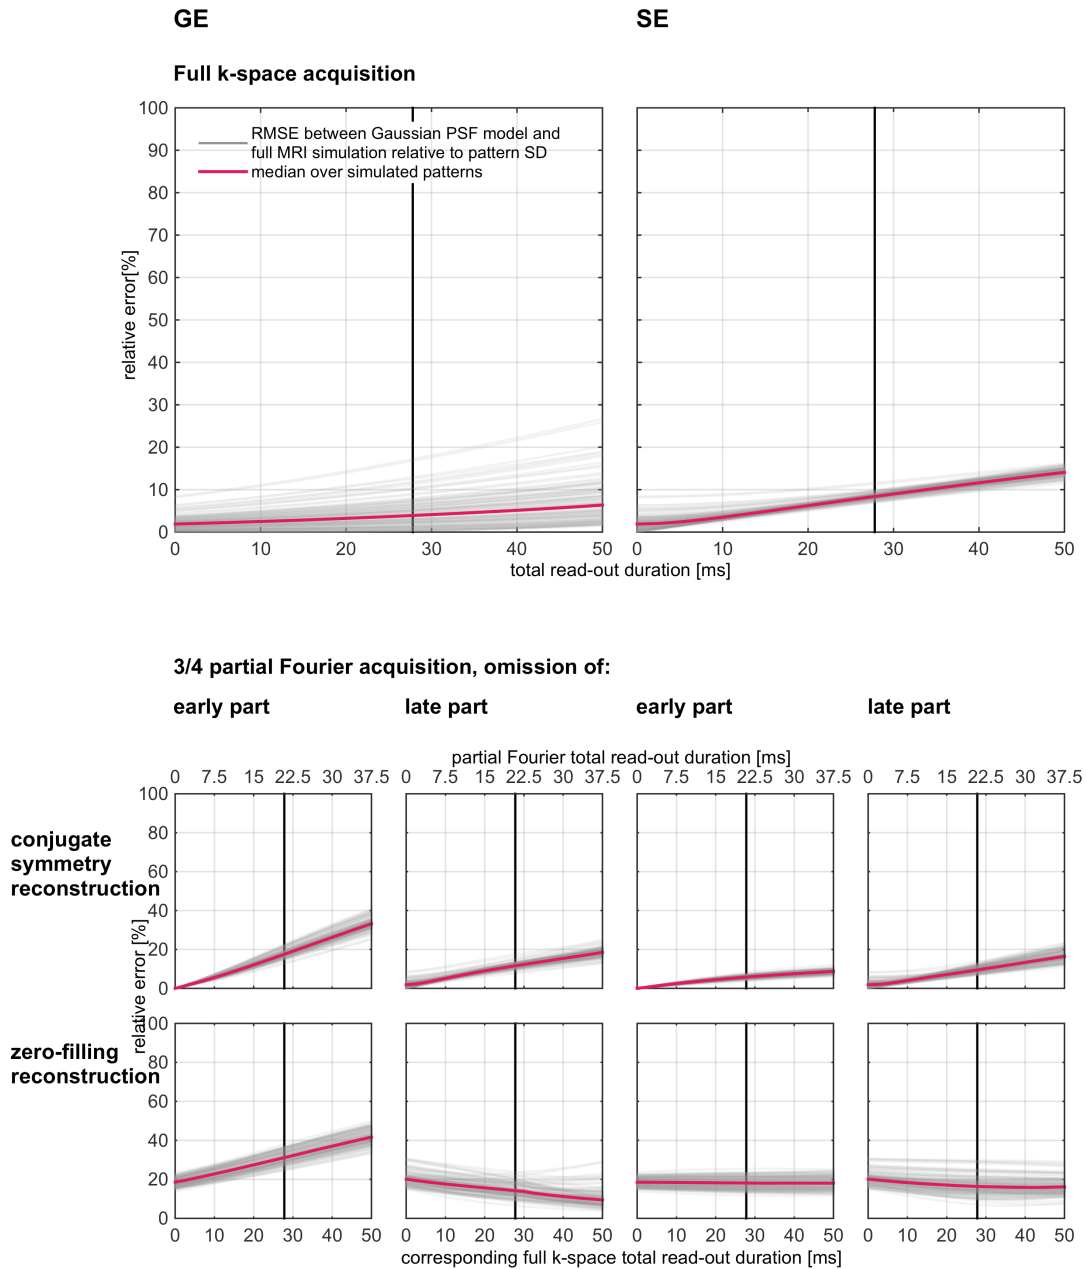


Fig. 8 Evaluation of modeling the complete MR process as a convolution with a Gaussian PSF that accounts for signal decay followed by MR sampling with no decay. 1,000 cortical columnar response patterns (with intermediate irregularity and spatial frequency) were simulated (results from 100 patterns were used for visualization). For GE (left) and SE (right) full k-space acquisition (top) and partial Fourier acquisition (bottom), convolution with the fitted Gaussian PSF (or its inverse model, see text) followed by MR sampling with no decay was compared to the complete MR imaging process. To this end, the root-mean-squared-errors (RMSE) relative to the standard deviation of the complete MR images were calculated. The gray curves show the distribution of relative RMSEs from all simulated patterns as a function of total read-out duration. The vertical black line represents a total-read out duration of 27.8 ms (20.85 ms for partial Fourier). The medians of RMSEs are shown in red. For Partial Fourier acquisition, the total-read out duration axis is labeled at the top. Partial Fourier acquisition total-read out durations were shortened according to the fraction of omitted k-space (1/4) relative to the corresponding full k-space total read-out duration (bottom axis). RMSEs were generally low for full k-space acquisition but became higher for the majority of partial Fourier acquisition schemes.

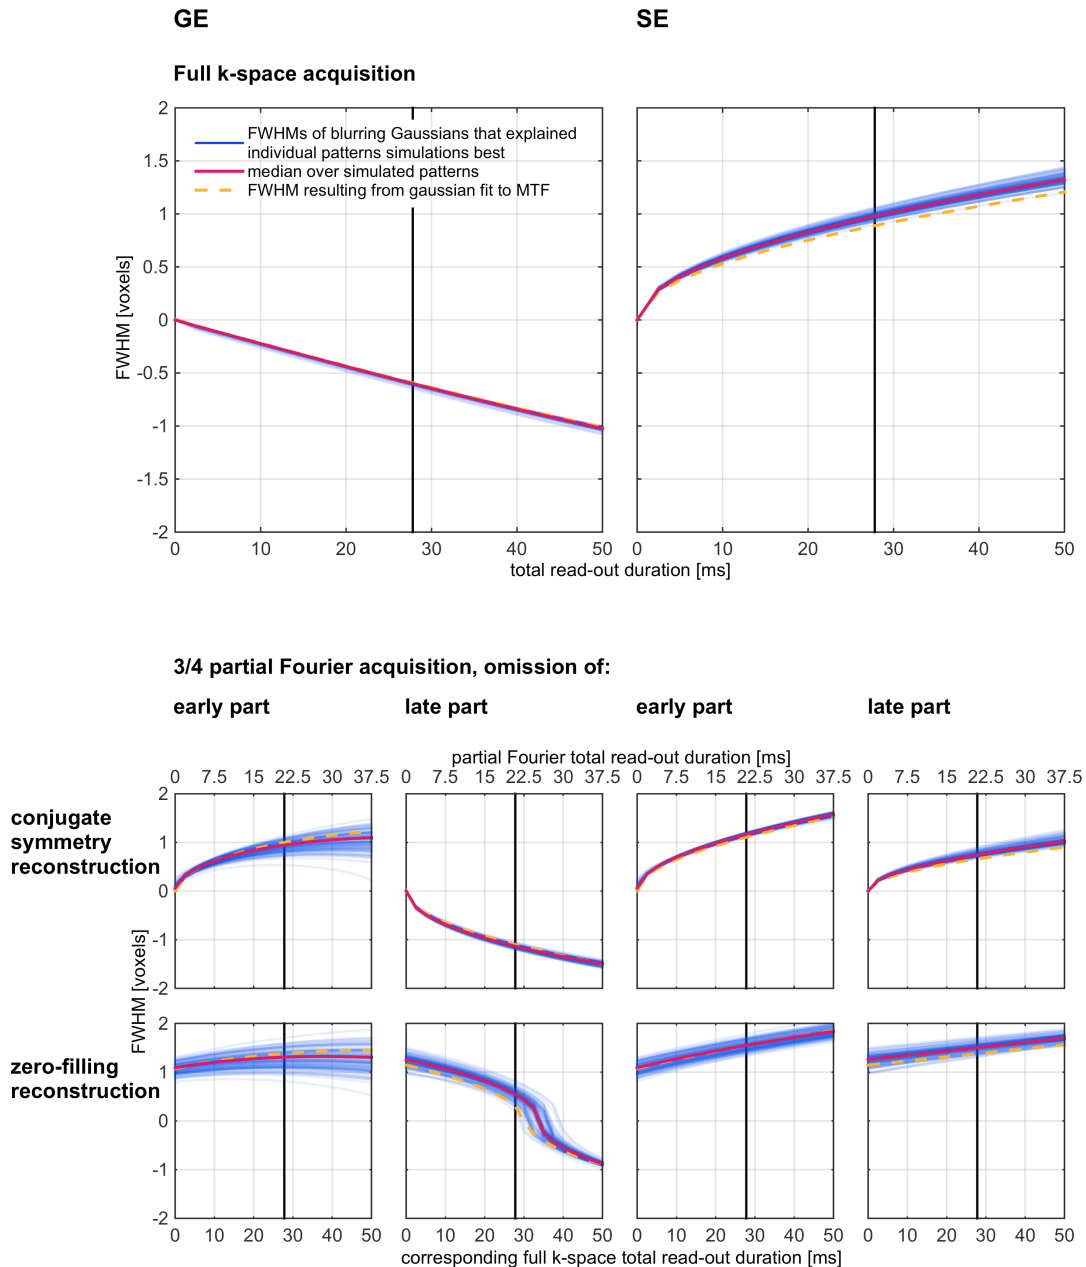


Fig. 9 Estimation of pattern specific Gaussian PSFs that model the effect of signal-decay. 1,000 cortical columnar response patterns (with intermediate irregularity and spatial frequency) were simulated (results from 100 patterns were used for visualization). We present results from GE (left) and SE (right) full k-space acquisition (top) and partial Fourier acquisition (bottom). For each pattern, complete MR imaging was simulated and compared to the result of convolution of the pattern with Gaussians (or their inverse models, see text) followed by MR sampling with no decay. For each pattern, the pattern specific Gaussian that resulted in the smallest mean squared error relative to the complete MR imaging simulation were determined. The blue curves show the distribution of FWHMs of these Gaussians as a function of total read-out duration. The vertical black line represents a total-read out duration of 27.8 ms (20.85 ms for partial Fourier). The medians of FWHMs are shown in red. For Partial Fourier acquisition, the total read-out duration axis is labeled at the top. Partial Fourier acquisition total read-out durations were shortened according to the fraction of omitted k-space (1/4) relative to the corresponding full k-space total read-out duration (bottom axis). The FWHMs of our pattern independent Gaussian PSF model are presented in dashed orange for comparison.

Figure 9 presents the resulting Gaussian PSF FWHMs (blue, median in red) as a function of total read-out duration for full k-space acquisitions (upper row) and for different partial Fourier acquisition schemes (middle and bottom rows). In addition, we present the previously estimated pattern-independent Gaussian PSF FWHMs (dashed orange curves). Note that the patterns used for this evaluation were the results of different columnar patterns, but they all shared the same statistical properties (main pattern frequency = 1 cycle per 4 voxels, relative irregularity = 0.5).

The pattern-dependent PSF widths (blue curves) for the full k-space acquisitions did not vary much across patterns (they were relatively independent of the specific pattern). They corresponded reasonably well to those obtained from our pattern-independent model (dashed orange curve). For the majority of partial Fourier acquisition schemes, the variability of pattern-specific PSF widths was somewhat higher. In addition, the differences between the median over simulated patterns and the pattern independent model were on average higher than those obtained for the full k-space acquisitions. This suggests that using a pattern-independent, single PSF width does not fully characterize the effective spatial resolution under partial Fourier acquisitions. We note, however, that for both full k-space and partial Fourier acquisitions, the estimations obtained from our proposed pattern-independent model (in dashed orange curves) matched those obtained from the pattern-dependent simulations reasonably well.

Discussion

The Imaging PSF and effective spatial resolution

The PSF of an imaging system is defined as the image obtained from an infinitesimally small point-like object. If the imaging system is linear and shift invariant, its response to an arbitrary object can be described as a convolution with the imaging PSF as the convolution kernel. This latter property is what makes the PSF useful in describing the spatial characteristics of an imaging system.

The effect of a system with an imaging PSF with its maximum at zero and whose strictly positive values do not increase with increasing distance from zero can be intuitively understood as spatial spreading or blurring. Such systems produce smoothed image versions of any object or pattern. The smoothing can be quantified by measures of the imaging PSF width (such as the FWHM). However, other PSF shapes may characterize certain systems, resulting in more complex effects that require careful evaluation and may not be intuitive.

Note that the width of the PSF is not the only possible measure of effective spatial resolution. Another measure that has been proposed is the area under the PSF divided by the value of the PSF at the origin (Haacke et al., 1999). In the case of MR imaging without signal decay, it is identical to the actual voxel width.

The magnitude PSF is not a good measure for the effective spatial resolution of MRI and fMRI

We have shown that the FWHM of the magnitude PSF is not a good measure for the effect of T_2 and T_2^* signal decay on spatial resolution. There are three reasons as to why this is the case.

The larger part of the FWHM of the magnitude PSF is due to MR sampling and not due to signal decay

The magnitude PSF of MR imaging, even with no signal decay, has a FWHM of 1.2 voxels compared to 1.4 voxels and 1.3 voxels (examples given for total readout of 30 ms at 7T) for GE and SE decay, respectively. However this MR sampling effect with no signal decay cannot be regarded as a simple spread of signal. In fact, it acts as a hard low-pass filter that discards all spatial frequency components higher than the voxel size-dependent highest sampled spatial frequency, and leaves all other spatial frequency components unchanged. For a pattern dominated by the latter spatial frequency components, MR sampling with no signal decay has virtually no blurring effect. For a pattern with spatial frequencies limited to those sampled by the MRI process, MR sampling with no signal decay has no blurring effect at all. Thus, the FWHM of the magnitude PSF is a pattern-independent measure, whereas the effective resolution of MRI sampling with no signal decay does depend on the imaged pattern.

The MR imaging process is non-linear

The last step of the reconstruction in an MR imaging process typically involves taking the absolute values (magnitude image) of the complex image values. The complex image values are the results of a process that can be described as a convolution of the original pattern with the complex imaging PSF. Taking the absolute values of the complex image values makes the MRI process non-linear. In general, the result of convolving a pattern with a complex kernel, then taking the magnitude image, is not equal to convolving the same pattern with the magnitude values of the complex kernel. This general statement applies in the specific case of MRI: the image obtained by the MRI process is different from that obtained by convolving the original pattern with the magnitude PSF.

To illustrate the effect of the non-linearity of the MRI process, we will describe two scenarios involving the imaging of a point-like structure. While the FWHM of the magnitude PSF can describe the effective spatial resolution of imaging an infinitesimally small point-like structure with no background, it cannot describe the MRI sampling of any arbitrary structure. For example, MRI of a pattern composed of a similar infinitesimally small point-like structure superimposed on a spatially constant baseline with an amplitude higher than the amplitude of the magnitude PSF would include negative side lobes relative to the baseline. Thus, the magnitude PSF fails to correctly describe the MRI process.

Note that even convolving the original pattern with the complex imaging PSF does not fully describe the MRI process, since it does not reflect the operation of taking the magnitude image. Therefore, even the complex imaging PSF

on its own does not completely reflect the typical MRI sampling and reconstruction process.

Signal decay may cause blurring or high-pass filtering with identical FWHM of the magnitude PSF

Signal decay does not always blur the pattern; it can also cause high-pass filtering. This is the case for the GE simulations we have conducted. Whether signal decay results in blurring or high-pass filtering depends on the shape of the decay curve and on the ordering of k-space acquisitions. Specific decay curves and ordering of k-space acquisitions may result in imaging PSFs that are different, but share the same FWHM. Indeed, the magnitude PSFs that are associated with imaging processes that result in blurring or high-pass filtering can be different but may have identical widths. This clearly limits the interpretation based on the magnitude PSF's width measure.

What does the magnitude PSF describe?

In the 3 previous sub-sections, we have shown that the magnitude PSF is not a good measure for characterizing the effective spatial resolution of the MRI process. The magnitude PSF carries less information relative to the complex imaging PSF. Simply relying on the FWHM of the magnitude PSF further reduces the available information. The magnitude PSF cannot differentiate between reduced effective spatial resolution due to the MRI sampling per se or due to signal decay. The effect of MRI sampling with no signal decay depends on the original pattern, whereas the FWHM of the magnitude PSF does not. While the magnitude PSF does describe the magnitude MR image of a small point-like structure with no background signal, it is not a convolution kernel of the MR imaging process. The FWHM of the magnitude PSF does not differentiate between blurring and high-pass filtering effects.

What, then, does the magnitude PSF describe, and can it be used for any characterization of the MRI process? The magnitude PSF has some general value in that it describes the absolute level of influence that neighboring positions in the original pattern have on each other's value in the image. The problem is that it fails to characterize the nature of this influence (e.g. blurring or high-pass filtering), which depends on the signs of the components, the phase and the overall shape of the underlying complex PSF.

Applicability of our simplified model

Approximation of fMRI as a convolution with the real component of the complex imaging PSF

We have shown that the MR imaging process can be approximated by a convolution with the real component of the complex imaging PSF. This approximation works well if the original pattern constitutes a low amplitude spatially varying pattern superimposed on a constant, spatially homogeneous background of a higher amplitude. For response amplitudes of 5%, we found the typical RMSE relative to the standard deviation of the pattern imaged by the complete MRI process to be well below 1%. This scenario holds for fMRI, where gray matter has a relatively uniform baseline intensity and the focus is

on superimposed signal changes of approximately 1%-5%. In contrast, this approximation may not be as appropriate for structural MRI where the low signal background, as well as objects of varying size and intensity, are of interest.

Separation of MRI sampling and signal decay

The separation of MR sampling and signal decay makes it possible to consider their respective effects separately. MR sampling with no decay does not automatically result in a blurred image. If the voxel width is sufficiently small for sampling the larger part of the original spatial frequencies, the MRI sampling will have no blurring effect.

For example, the spatial frequency content of the BOLD response of cortical columns is limited due to the smoothness of the neuronal columnar organizations and the spreads of the neurophysiological and hemodynamic responses. Therefore, the blurring effect of fMRI sampling with adequately small voxels can be neglected. The condition is that the voxels are sufficiently small, such that they can capture the main (peak) frequency of a columnar organization and the frequencies showing elevated power around it.

When untangled from MR sampling, signal decay can be described as a blurring process which we model as Gaussian blurring. This is a simplifying model, with precision that varies with the actual imaging MTF. However, the deviations of this simplifying model relative to the complete fMRI process are small for typical signal changes (fMRI response) and noise levels in fMRI.

Importantly, it results in a simple and intuitive characterization of the blurring associated with signal decay that makes it possible to compare it to previously reported FWHMs of PSFs associated with the entire BOLD fMRI process (BOLD PSF) (Chaimow et al., 2016; Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007). It further makes it possible to decompose the PSF of fMRI into two components: one caused by k-space sampling and signal decay, and the other caused by a physiological fMRI contrast-dependent spread.

Our proposed model applies to both isotropic and anisotropic cortical columns

We have used the example of ocular dominance columns (ODC) in order to show the limitations of using the magnitude PSF for characterizing the effective spatial resolution of fMRI (Figure 2). However, our characterization of the MR imaging process as Gaussian blurring followed by MR sampling with no signal decay does not depend on this specific example. In particular because we only considered the phase-encode dimension, we were able to use one dimensional columnar patterns. These one dimensional patterns can be regarded as general models of columnar patterns, characterized by a main pattern frequency and a level of irregularity.

Applicability of our proposed model to other fMRI and MRI methods

In addition to GE and SE BOLD fMRI, other methods such as 3D GRASE (Feinberg et al., 2008), VASO (Lu et al., 2003) and ASL (Detre et al., 1992) have been used for high resolution imaging (Duong et al., 2001; Huber et al., 2015; Zimmermann et al., 2011). Do our findings generalize and account for MRI sampling and signal decay in these methods? Our analysis of the imaging PSF depends on two conditions. The first is the shape of signal decay during linear

sampling of k-space. This condition holds when using different echo times or preparatory pulses that only affect the absolute magnitude of the magnetization to which the modeled T_2^* or T_2/T_2^* decay is applied. Our derivations are valid for such scenarios. The second is the validity of a linear approximation of the response pattern, which holds if the response is composed of a small amplitude response pattern relative to a larger amplitude spatially homogenous baseline.

In 3D GRASE, multiple refocusing pulses and subsequent EPI acquisitions (partitions) follow a single excitation pulse. Each partition represents a step in 3D k-space in the direction orthogonal to what is commonly considered the slice planes. The signal decay within each partition is proportional to that of a single SE EPI acquisition (Fig. 1, first row, SE). Across partitions, the amplitude changes according to T_2 decay.

Consequently, the entire 3D k-space modulation due to signal decay can be separated into a product of decay between and within partitions. Furthermore, the separability of dimensions in the Fourier transform implies that these respective components determine the spatial filtering due to signal decay across slices (between partition decay) and within slices (within partition decay).

As a result, our SE findings are valid for the effective in-plane resolution of 3D GRASE acquisitions. In order to approximate the effective in-plane spatial resolution of 3D GRASE, one needs to consider the total readout for a single partition and refer to the SE results in Figure 7. However, the effective spatial resolution across slices is determined by a decay similar to GE acquisition, but with a time constant of T_2 instead of T_2^* , and with a total read-out duration that covers the acquisition of all partitions.

VASO, a method that indirectly measures changes in cerebral blood volume, applies an inversion recovery pulse prior to a normal GE or SE EPI sequence. The effect of the inversion recovery pulse is that all signals from blood are nulled when the excitation pulse occurs. This, in addition to the often used very short echo time, results in a change of scaling but not in a change of the shape of T_2^* decay (for GE) or T_2/T_2^* decay (for SE). In addition, VASO signal changes are small, on the order of -1% (Lu and van Zijl, 2012). Together, these features allow us to apply our results directly to VASO imaging. Therefore, Figure 7 provides the effective spatial resolution of VASO acquired by means of GE or SE imaging.

For cerebral blood flow imaging, e.g. using ASL, the situation is different. If EPI acquisition is used, the signal decay follows our analysis, therefore complying with the first condition as described above. However, cerebral blood flow changes are on the order of 20%-60% which does not follow the second condition, making our linear approximation much less accurate.

How does the effective spatial resolution influence functional imaging?

Signal amplitude reductions

In the current study, we focus on signal decay-dependent modulation of amplitudes of spatial frequency components relative to each other. However, signal decay also causes overall amplitude decreases that bring about a reduced signal to thermal noise ratio (SNR). Such amplitude decreases depend on total

read-out duration. In addition, read-out duration per read-out line determines receiver bandwidth, with higher bandwidth (shorter read-out duration) resulting in increased noise. Both factors need to be considered in order to find an optimal total read-out duration that maximizes SNR (Qin, 2012), all within the constraints (and potential effects) of matrix size, field of view, echo time, and gradient strength.

SNR associated with a spatial frequency

SNR and effective spatial resolution can be considered together by evaluating the SNR at a spatial frequency. For detecting or decoding stimulus-specific responses, SNR needs to be high for at least part of the spatial frequencies that contribute to stimulus-specific responses, independent of whether the overall image is blurred. However, if one aims to obtain a precise reconstruction of the response pattern, then both high SNR and an image with an undistorted frequency spectrum are necessary.

In this context, it is of interest to discuss the difference between partial Fourier reconstructions employing conjugate symmetry and simple zero-filling. It may appear that the conjugate symmetry reconstruction is always superior to zero-filling due to its reduced blurring effect. Indeed, this is the case if the aim is to image a pattern precisely. However, the situation is different if we consider the SNR at each spatial frequency. Compared to conjugate symmetry reconstruction, zero-filling reduces the amplitude of high spatial frequencies, as the contribution of their omitted components is set to zero. However, the noise at these spatial frequencies is reduced proportionally. Therefore, the spatial frequency-specific SNR is equal between conjugate symmetry and zero filling reconstructions of partial Fourier acquisition. The spatial frequency specific SNR obtained by full k-space acquisition is higher than those obtained by both conjugate symmetry and zero-filling reconstructions of data obtained by partial Fourier acquisition. This is because full k-space acquisition benefits from averaging of independent noise across the negative and positive k-space parts. The result is an expected increase of spatial frequency-specific SNR by a factor of $\sqrt{2}$ compared to partial Fourier acquisitions, independent of the employed reconstruction method.

Significance of signal decay blurring relative to the overall BOLD point spread function

The absolute width of the PSF due to signal decay is proportional to voxel width. At ultra-high magnetic field, the in-plane voxel width used for high-resolution fMRI can be as small as 0.5 mm. In Table 1, we compare the FWHM of the resulting PSF due to signal decay to the overall BOLD fMRI PSF as estimated in (Chaimow et al., 2016). Furthermore, we estimated the BOLD PSF while accounting for the effect of signal decay. We considered that consecutive Gaussian convolutions result in a Gaussian convolution with a total width equal to the square root of the sum of squares of individual convolution widths. The results show that the contribution of fMRI signal decay to the overall BOLD fMRI PSF is relatively small. This is especially true, considering that signal decay blurring acts only on the phase-encode direction.

Table 1 Comparison of decay dependent imaging PSF and BOLD PSF. This table compares the FWHM of the PSF due to signal decay (fMRI signal decay PSF, as estimated in the current study) to that of the overall BOLD fMRI PSF as estimated in (Chaimow et al., 2016). A voxel width of 0.5 mm is assumed. The bottom row shows the FWHM of the BOLD PSF when accounting for the contribution of fMRI signal decay. These numbers show that the contribution of fMRI signal decay is relatively small. Note that for GE, accounting for the PSF due to signal decay widens the PSF function, due to the high-pass filtering effect of the signal decay in GE BOLD fMRI.

	GE	SE
Overall BOLD fMRI PSF (Chaimow et al., 2016)	1.02 mm	0.82 mm
fMRI signal decay PSF	-0.30 mm (high pass filter)	0.44 mm
BOLD PSF, accounting for fMRI signal decay	0.98 mm	0.69 mm

Conclusion

We have demonstrated that the FWHM of the absolute values of the complex imaging PSF (magnitude PSF) is a poor and potentially misleading measure for the effect of signal decay on the effective spatial resolution. Instead, we propose to first linearly approximate the typically non-linear process of MR sampling and reconstruction and then to separately consider the effects of two components of the imaging process. The first component is the MR sampling with no signal decay, which acts as a hard low-pass filter. It discards all spatial frequencies higher than the voxel size-dependent maximal sampled spatial frequency and leaves all other frequencies untouched. The second component depends on the signal decay. We have shown that the effect of this second component can be approximated by either Gaussian blurring or high-pass filtering that reverses the effect of a Gaussian blurring. For typical SE parameters at 7 Tesla, we found that the Gaussian blurring attributed to signal decay has a PSF FWHM of 0.89 voxels (0.44 mm for 0.5 mm wide voxels). In contrast, GE at 7 Tesla has a high-pass filter effect, reversing a Gaussian blurring with a PSF FWHM of 0.59 voxels (0.30 mm for 0.5 mm wide voxels). We conclude that signal decay in SE fMRI with full k-space acquisition at 7 Tesla has a more moderate blurring effect compared to the effect implied by the commonly used FWHM of the magnitude PSF. We further conclude that signal decay in GE fMRI at 7 Tesla (and also at other field strengths, not shown) has a high-pass filtering effect, opposite to what can be expected from describing the effect of signal decay on GE fMRI by the FWHM of the corresponding magnitude PSF.

References

Chaimow, D., Yacoub, E., Uğurbil, K., Shmuel, A., 2016. Spatial specificity of the functional MRI blood oxygenation response relative to metabolic activity, in: Presented at the 22nd Annual Meeting of the Organization for Human Brain Mapping,

Geneva.

- Chaimow, D., Yacoub, E., Uğurbil, K., Shmuel, A., 2011. Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns. *Neuroimage* 56, 627–642. doi:10.1016/j.neuroimage.2010.09.037
- Cheng, K., Waggoner, R.A., Tanaka, K., 2001. Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron* 32, 359–374.
- Constable, R.T., Gore, J.C., 1992. The loss of small objects in variable TE imaging: implications for FSE, RARE, and EPI. *Magn Reson Med* 28, 9–24.
- Detre, J.A., Leigh, J.S., Williams, D.S., Koretsky, A.P., 1992. Perfusion imaging. *Magn Reson Med* 23, 37–45.
- Duong, T.Q., Kim, D.S., Uğurbil, K., Kim, S.-G.G., 2001. Localized cerebral blood flow response at submillimeter columnar resolution. *Proceedings of the National Academy of Sciences* 98, 10904–10909. doi:10.1073/pnas.191101098
- Engel, S.A., Glover, G.H., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7, 181–192. doi:10.1093/cercor/7.2.181
- Farzaneh, F., Riederer, S.J., Pelc, N.J., 1990. Analysis of T2 limitations and off-resonance effects on spatial resolution and artifacts in echo-planar imaging. *Magn Reson Med* 14, 123–139.
- Feinberg, D.A., Harel, N., Ramanna, S., Uğurbil, K., Yacoub, E., 2008. Sub-millimeter Single-shot 3D GRASE with Inner Volume Selection for T2 weighted fMRI applications at 7 Tesla. *Proc. Intl. Soc. Mag. Reson. Med.* 16, 1–1.
- Goodyear, B.G., Menon, R.S., 2001. Brief visual stimulation allows mapping of ocular dominance in visual cortex using fMRI. *Hum Brain Mapp* 14, 210–217.
- Haacke, M.E., 1987. The effects of finite sampling in spin-echo or field-echo magnetic resonance imaging. *Magn Reson Med* 4, 407–421.
- Haacke, M.E., Brown, R.W., Thompson, M.R., 1999. *Magnetic resonance imaging: physical principles and sequence design.*
- Huber, L., Goense, J., Kennerley, A.J., Trampel, R., Guidi, M., Reimer, E., Ivanov, D., Neef, N., Gauthier, C.J., Turner, R., Möller, H.E., 2015. Cortical lamina-dependent blood volume changes in human brain at 7 T. *Neuroimage* 107, 23–33. doi:10.1016/j.neuroimage.2014.11.046
- Kemper, V.G., De Martino, F., Vu, A.T., Poser, B.A., Feinberg, D.A., Goebel, R., Yacoub, E., 2015. Sub-millimeter T2 weighted fMRI at 7 T: comparison of 3D-GRASE and 2D SE-EPI. *Front Neurosci* 9, 163. doi:10.3389/fnins.2015.00163
- Lu, H., Golay, X., Pekar, J.J., van Zijl, P.C.M., 2003. Functional magnetic resonance imaging based on changes in vascular space occupancy. *Magn Reson Med* 50, 263–274. doi:10.1002/mrm.10519
- Lu, H., van Zijl, P.C.M., 2012. A review of the development of Vascular-Space-Occupancy (VASO) fMRI. *Neuroimage* 62, 736–742. doi:10.1016/j.neuroimage.2012.01.013
- Menon, R.S., Ogawa, S., Strupp, J., Uğurbil, K., 1997. Ocular dominance in human V1 demonstrated by functional magnetic resonance imaging. *J Neurophysiol* 77, 2780–2787.
- Nasr, S., Polimeni, J.R., Tootell, R.B., 2016. Interdigitated Color- and Disparity-Selective Columns within Human Visual Cortical Areas V2 and V3. *J Neurosci.* 36:1841-57.
- Oshio, K., Singh, M., 1989. A computer simulation of T2 decay effects in echo planar imaging. *Magn Reson Med* 11, 389–397.
- Parkes, L.M., Schwarzbach, J.V., Bouts, A.A., Deckers, R.H.R., Pullens, P., Kerskens, C.M., Norris, D.G., 2005. Quantifying the spatial resolution of the gradient echo

- and spin echo BOLD response at 3 Tesla. *Magn Reson Med* 54, 1465–1472. doi:10.1002/mrm.20712
- Polimeni, J.R., Fischl, B., Greve, D.N., Wald, L.L., 2010. Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1. *Neuroimage* 52, 1334–1346. doi:10.1016/j.neuroimage.2010.05.005
- Qin, Q., 2012. Point spread functions of the T2 decay in k-space trajectories with long echo train. *Magn Reson Imaging* 30, 1134–1142. doi:10.1016/j.mri.2012.04.017
- Rojer, A., Schwartz, E., 1990. Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biol Cybern* 62, 381–391.
- Shmuel, A., Chaimow, D., Raddatz, G., Uğurbil, 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. doi:10.1016/j.neuroimage.2009.08.040
- Shmuel, A., Yacoub, E., Chaimow, D., Logothetis, N.K., Uğurbil, K., 2007. Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 Tesla. *Neuroimage* 35, 539–552. doi:10.1016/j.neuroimage.2006.12.030
- Uludağ, K., Müller-Bierl, B., Uğurbil, K., 2009. An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. *Neuroimage* 48, 150–165. doi:10.1016/j.neuroimage.2009.05.051
- Yacoub, E., Duong, T.Q., van de Moortele, P.F., Lindquist, M., Adriany, G., Kim, S.-G.G., Uğurbil, K., Hu, X.P., 2003. Spin-echo fMRI in humans using high spatial resolutions and high magnetic fields. *Magn Reson Med* 49, 655–664. doi:10.1002/mrm.10433
- Yacoub, E., Harel, N., Uğurbil, K., 2008. High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences* 105, 10607–10612. doi:10.1073/pnas.0804110105
- Yacoub, E., Shmuel, A., Logothetis, N.K., Uğurbil, 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., Shmuel, A., Pfeuffer, J., van de Moortele, P.F., Adriany, G., Andersen, P., Vaughan, J.T., Merkle, H., Uğurbil, K., Hu, X., 2001. Imaging brain function in humans at 7 Tesla. *Magn Reson Med* 45, 588–594. doi:10.1002/mrm.1080
- Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P.-F., Feinberg, D., Adriany, G., Chaimow, D., Shmuel, A., Uğurbil, 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. doi:10.1371/journal.pone.0028716

Appendix A. Modeling and simulating MR imaging

This section provides detailed equations of the MR imaging model as used in our simulations. The theory follows Haacke et al. (1999).

We consider the phase-encode dimension only and analyze it separately from the read-out dimension. This is justified because of the separability of the Fourier transform. Let $y(x)$ be a spatial pattern and $s(k) = \mathcal{F} [y(x)]$ its k-space representation obtained by Fourier transform. Furthermore, let L be the field-of-view and $N = 2n$ the matrix size with voxel width $\Delta x = L/N$.

MR sampling with no signal decay

MRI samples k-space in steps of $\Delta k = 1/L$ from $-n\Delta k$ to $(n-1)\Delta k$. The reconstructed imaged pattern (y_q^{MRI}), where $q \in [-n, \dots, n-1]$, is obtained by taking the absolute values of an inverse discrete Fourier transform, such that

$$y_q^{MRI} = \left| \Delta x \Delta k \sum_{p=-n}^{n-1} s(p\Delta k) e^{i\pi p q / n} \right|. \quad (\text{A.1})$$

Defining $MTF(k) = \Delta x \Delta k \sum_{p=-n}^{n-1} \delta(k - p\Delta k)$, A.1 can be rewritten as:

$$y_q^{MRI} = \left| \int_{-\infty}^{\infty} MTF(k) s(k) e^{i2\pi k q \Delta x} dk \right| = \left| \mathcal{F}^{-1} [MTF(k) s(k)] (q \Delta x) \right|. \quad (\text{A.2})$$

This shows that $MTF(k)$ is the modulation-transfer function of the linear part of the MRI process (up to taking absolute values). As such it has an associated point-spread function $psf(x) = \mathcal{F}^{-1} [MTF(k)]$, allowing us to express A.2 as a convolution:

$$y_q^{MRI} = \left| (y * psf) (q \Delta x) \right|. \quad (\text{A.3})$$

MR sampling in the presence of signal decay

Let $t(k)$ be the time that an individual k-space point k is being acquired and let $f(k) = f(t(k))$ be the relative signal decay amplitude at that time. Such an acquisition in the presence of signal decay results in an effective k-space representation $f(k)s(k)$, changing A.1 to

$$y_q^{MRI} = \left| \Delta x \Delta k \sum_{p=-n}^{n-1} f(p\Delta k) s(p\Delta k) e^{i\pi p q / n} \right|. \quad (\text{A.4})$$

In this situation the MR imaging equations A.2 and A.3 still apply, if we absorb $f(k)$ into the modulation-transfer function, now defined as $MTF(k) = \Delta x \Delta k f(k) \sum_{p=-n}^{n-1} \delta(k - p\Delta k)$.

Appendix B. Linear approximation of the MR imaging process

Let $r: [-L/2, L/2] \rightarrow \mathbb{R}$ be a spatial pattern of relative responses of BOLD contrast and let

$$y[r](x) = \begin{cases} 1 + r(x) & \text{if } x \in [-L/2, L/2] \\ 0 & \text{otherwise} \end{cases},$$

be the associated pattern of absolute BOLD signal.

First we compute the result of convolving such a signal pattern with a point-spread function defined according to Appendix A.

$$\begin{aligned}
(y[r] * psf)(x) &= \int_{-\infty}^{\infty} y[r](x')psf(x-x')dx' \\
&= \int_{-L/2}^{L/2} (1+r(x'))psf(x-x')dx' \\
&= \int_{-L/2}^{L/2} psf(x-x')dx' + \int_{-L/2}^{L/2} r(x')psf(x-x')dx' \\
&= \int_{-L/2}^{L/2} \int_{-\infty}^{\infty} MTF(k)e^{i2\pi k(x-x')}dkdx' + \int_{-L/2}^{L/2} r(x')psf(x-x')dx' \\
&= \int_{-L/2}^{L/2} \int_{-\infty}^{\infty} \Delta x \Delta k f(k) \sum_{p=-n}^{n-1} \delta(k-p\Delta k) e^{i2\pi k(x-x')} dk dx' + \int_{-L/2}^{L/2} r(x')psf(x-x')dx' \\
&= \Delta x \Delta k \sum_{p=-n}^{n-1} f(p\Delta k) \int_{-L/2}^{L/2} e^{i2\pi p(x-x')/L} dx' + \int_{-L/2}^{L/2} r(x')psf(x-x')dx' \\
&= \Delta x f(0) + \int_{-L/2}^{L/2} r(x')psf(x-x')dx'. \tag{B.1}
\end{aligned}$$

We define an operator MRI that models the MRI acquisition process by mapping the spatial pattern $r(x)$ onto a measured MRI pattern $(MRI[r]_q)_{-n \leq q \leq n-1}$ according to A.3, such that

$$\begin{aligned}
MRI[r]_q &= |(y[r] * psf)(q\Delta x)| \\
&= \sqrt{\text{Re}((y[r] * psf)(q\Delta x))^2 + \text{Im}((y[r] * psf)(q\Delta x))^2} \\
&= \sqrt{\left(\Delta x f(0) + \int_{-L/2}^{L/2} r(x')psf_{\text{Re}}(x-x')dx'\right)^2 + \left(\int_{-L/2}^{L/2} r(x')psf_{\text{Im}}(x-x')dx'\right)^2}, \tag{B.2}
\end{aligned}$$

where $psf_{\text{Re}}(x) = \text{Re}(psf(x))$ and $psf_{\text{Im}}(x) = \text{Im}(psf(x))$.

We also note that

$$\begin{aligned}
MRI[0]_q &= \sqrt{\left(\Delta x f(0) + \int_{-L/2}^{L/2} 0 \cdot psf_{\text{Re}}(x-x')dx'\right)^2 + \left(\int_{-L/2}^{L/2} 0 \cdot psf_{\text{Im}}(x-x')dx'\right)^2} \\
&= \text{Re}((y[0] * psf)(q\Delta x)) \\
&= \int_{-L/2}^{L/2} psf_{\text{Re}}(x-x')dx'. \tag{B.3}
\end{aligned}$$

In order to linearly approximate the MR imaging process, we compute $MRI'[0][r]_q$, the functional derivative of $MRI[r]_q$ with respect to r , evaluated at $r_0 = 0$ (representing no response, only baseline). This functional derivative is a linear operator that maps a response pattern $r(x)$ onto an MRI response pattern $(MRI'[0][r]_q)_{-n \leq q \leq n-1}$, resulting in a linear approximation of the true MRI measurement according to $MRI[r]_q \approx MRI[0]_q + MRI'[0][r]_q$.

We calculate $MRI'[0][r]_q$ from B.2 using the chain rule and the fact that $\int_{-L/2}^{L/2} r(x')psf_{\text{Re}}(x-x')dx$ is already a linear operator on r :

$$\begin{aligned}
MRI'[0][r]_q &= \frac{1}{2\sqrt{(\Delta x f(0))^2}} 2(\Delta x f(0)) \int_{-L/2}^{L/2} r(x') psf_{Re}(x-x') dx' \\
&= \int_{-L/2}^{L/2} r(x') psf_{Re}(x-x') dx'
\end{aligned} \tag{B.4}$$

Finally, using B.3 and B.4 the linear approximation of $MRI[r]_q$ is

$$\begin{aligned}
MRI[r]_q &\approx MRI[0]_q + MRI'[0][r]_q \\
&= \int_{-L/2}^{L/2} psf_{Re}(x-x') dx + \int_{-L/2}^{L/2} r(x') psf_{Re}(x-x') dx' \\
&= y[r] * psf_{Re}.
\end{aligned} \tag{B.5}$$

STUDY 4

Optimization of functional MRI for detection, decoding and imaging the response patterns of cortical columns

Denis Chaimow, Kâmil Uğurbil, Amir Shmuel

The manuscript is under revision for publication in NeuroImage.

Author contributions

Denis Chaimow developed and implemented the model, and wrote the manuscript. Kâmil Uğurbil contributed to discussions. Amir Shmuel advised on the development of the model and edited the manuscript.

Acknowledgements

We thank Jiachen Liang and Dante Louie for editing the manuscript. This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada (AS, NSERC Discovery grants RGPIN 375457-09 and RGPIN- 2015-05103).

Abstract

The capacity of functional MRI (fMRI) to resolve cortical columnar organizations depends on several factors, e.g. the spatial scale of the columnar pattern, the point-spread of the fMRI response, the voxel size, and the SNR considering thermal and physiological noise. How these factors combine, and what is the voxel size that optimizes fMRI of cortical columns remain unknown.

Here we combine current knowledge into a quantitative model of fMRI of patterns of cortical columns. We compare different approaches for imaging patterns of cortical columns, including univariate and multivariate based detection, multi-voxel pattern analysis (MVPA) based decoding, and reconstruction of the pattern of cortical columns. We present the dependence of their performance on the parameters of the imaged pattern and the data acquisition, and predict voxel sizes that optimize fMRI under various scenarios.

To this end, we modeled differential imaging of realistic patterns of cortical columns with different spatial scales and degrees of irregularity. We quantified the capacity to detect and decode stimulus-specific responses by analyzing the distribution of voxel-wise differential responses relative to noise. We quantified the accuracy with which the spatial pattern of cortical columns can be reconstructed as the correlation between the underlying columnar pattern and the imaged pattern.

For regular patterns, optimal voxel widths for detection, decoding and reconstruction were close to half the main cycle length of the organization. Optimal voxel widths for irregular patterns were less dependent on the main cycle length, and differed between univariate detection, multivariate detection and decoding, and reconstruction. We compared the effects of different factors of Gradient Echo fMRI at 3 Tesla (T), Gradient Echo fMRI at 7T and Spin-Echo fMRI at 7T, and found that for all measures (detection, decoding, and reconstruction), the width of the fMRI point-spread has the most significant effect. In contrast, different response amplitudes and noise characteristics played a comparatively minor role. We recommend specific voxel widths for optimal univariate detection, for multivariate detection and decoding, and for reconstruction under these three data-acquisition scenarios. Our study supports the planning, optimization, and interpretation of fMRI of cortical columns and the decoding of information conveyed by these columns.

Introduction

With the advent of high-field functional MRI (fMRI), it has become possible to image organizations of cortical columns (Cheng et al., 2001; Goodyear and Menon, 2001; Nasr et al., 2016; Shmuel et al., 2010; Yacoub et al., 2008; 2007; Zimmermann et al., 2011). However, fMRI of cortical columns is still challenging. The ability to resolve the true underlying columnar organization depends on several factors, such as the spatial scale of the columnar pattern, the point-spread of the measured blood oxygenation level-dependent (BOLD) response, the voxel size and the signal-to-noise ratio (SNR). How the combination of these factors influences our capacity to image cortical columns is unclear. As a

result, which voxel size optimizes imaging of cortical columns remains unresolved.

Anatomical studies in humans (Adams et al., 2007; Horton et al., 1990) and invasive optical imaging studies in animals (Bonhoeffer and Grinvald, 1991; Shmuel and Grinvald, 2000; 1996) revealed characteristic organizations of columnar patterns. The structure and development of such organizations were modeled (Erwin et al., 1995; Niebur and Wörgötter, 1994; Rojer and Schwartz, 1990). In addition, we have good knowledge of the spatial and temporal characteristics of the BOLD response. The BOLD response to prolonged stimuli can be approximated as a convolution of a hemodynamic response function with stimulus-induced neuronal activity (Boynton et al., 1996; Logothetis et al., 2001). The BOLD point-spread function (Chaimow et al., 2016; Engel et al., 1997; Parkes et al., 2005; Shmuel et al., 2007) is a measure of the spatial spread of the BOLD response to a localized neuronal activation. The spread and the amplitude of the response depend on the relative contributions of the micro- and macro-vasculature to the total signal (Uludağ et al., 2009). These contributions depend on the magnetic field strength of the MR scanner and on the fMRI contrast. Lastly, models of noise in MRI (Krueger and Glover, 2001; Triantafyllou et al., 2005) explain how SNR and temporal SNR depend on the field strength and the voxel volume. All the aforementioned factors contribute to and interact in fMRI of cortical columns.

In a review article, Formisano et al. (2012) discussed these factors and their potential relevance to pattern information fMRI applied to responses at the resolution scale of cortical columns. In their Figure 2 they show in a schematic way how pattern information may depend on voxel size and field strength. They pose a number of “...important questions [that] remain unanswered”, such as what spatial resolution maximizes pattern information at 3T and 7T, whether 7T is better than 3T for revealing pattern information and whether the optimal field strength for pattern information depends on the resolution.

To answer these and similar questions, we combine currently available knowledge in a quantitative model of functional imaging of cortical columns. We quantify the dependence of measures associated with imaging cortical columns on properties of the columnar pattern, the BOLD point-spread function, the voxel size, and the field strength-specific noise characteristics. To this end, we build on our previous study (Chaimow et al., 2011), where we developed a model of imaging ocular dominance columns in order to investigate the basis of multivariate-pattern decoding analysis of cortical columns at 3T.

In the current study, we model generic isotropic columnar patterns of varying spatial scale and degree of irregularity. We demonstrate how contrast-to-noise ratio (CNR) can be determined from the spatial frequency spectrum of the columnar pattern, the BOLD point-spread function, the voxel width, and noise. We consider physiological noise correlations and explore a variety of noise and BOLD point-spread parameters including those typical for 3T GE, 7T GE, and 7T SE imaging. We examine different aspects of imaging cortical columns by quantifying the capacity to detect responses and decode the stimulus, and the accuracy with which we can reconstruct the spatial pattern of the columns. In addition, we demonstrate how response detection, decoding, and im-

aging the pattern of cortical columns depend on parameters of the columnar pattern and the imaging process. Lastly, we determine optimal voxel widths for various patterns of columns imaged using 3T GE, 7T GE, and 7T SE fMRI.

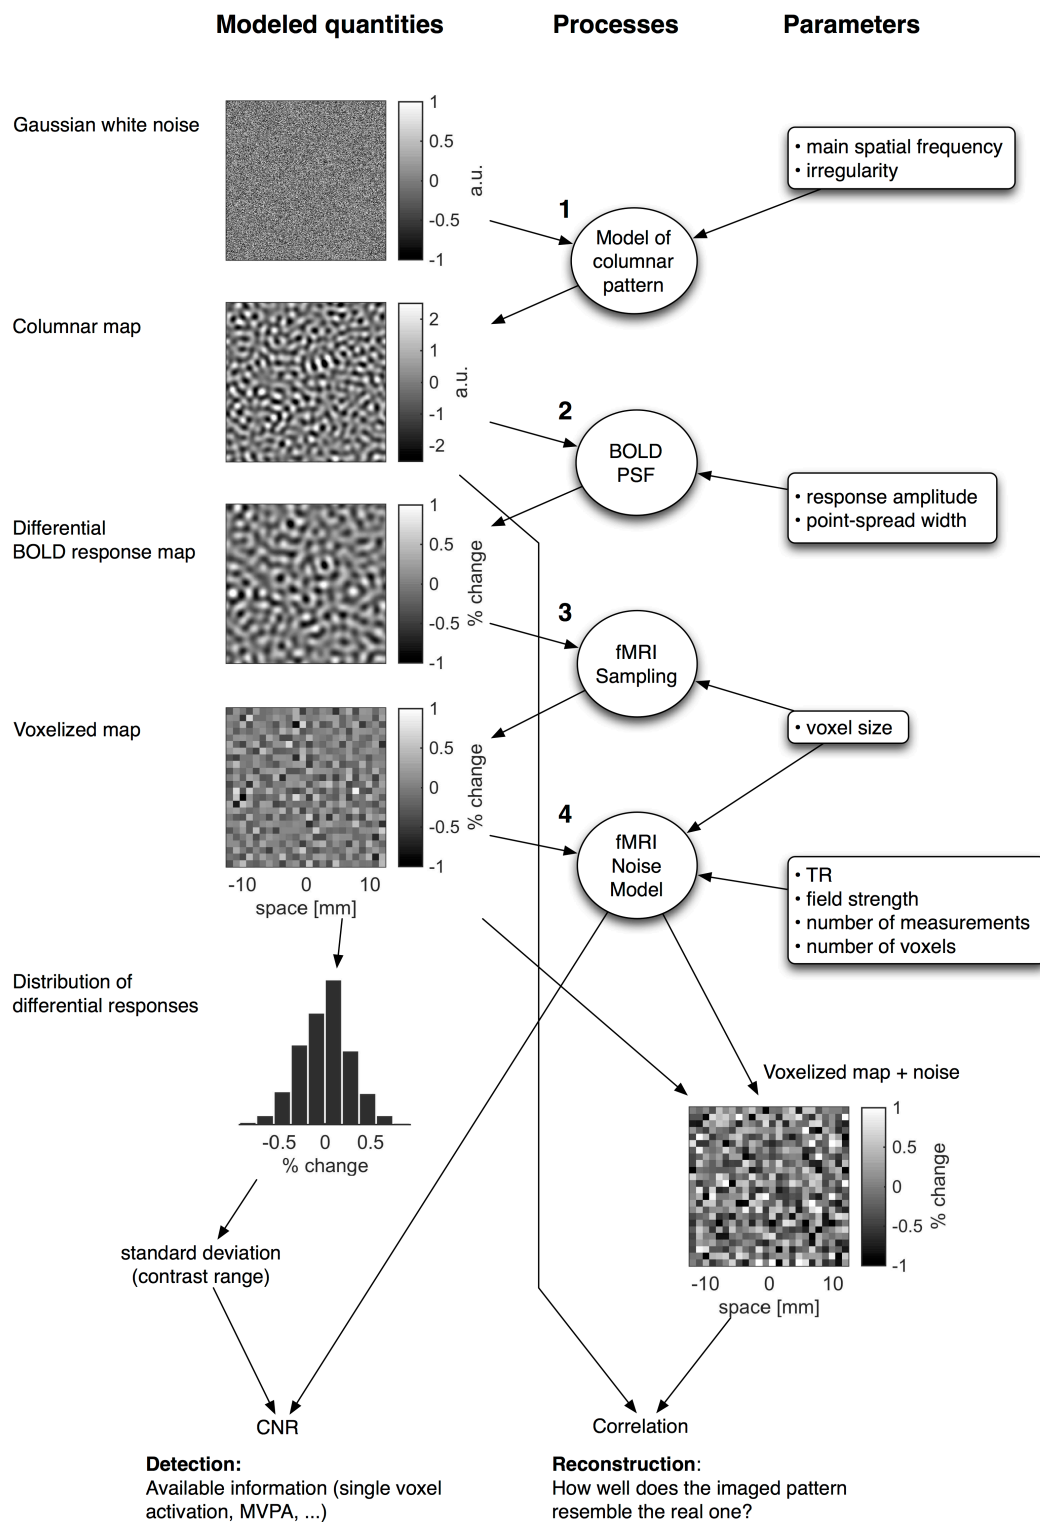


Fig. 1 (previous page) Model overview. In stage 1, columnar maps representing neuronal responses to two orthogonal stimulation conditions were generated by spatial filtering of white noise. In stage 2, the neuronal response was convolved with a BOLD-fMRI point-spread function. In stage 3, this BOLD response map was further transformed into a voxel pattern by simulating k-space sampling. All voxels in this pattern create the distribution of differential responses. This distribution is characterized by its

standard deviation, which reflects the range of contrasts in the set of imaged voxels (contrast range). In stage 4, realistic fMRI noise amplitudes were computed as a function of voxel size. The contrast range was divided by the noise level resulting in contrast to noise ratio (CNR). CNR (and the number of voxels) determine our likelihood to detect information about the stimulus. In order to evaluate our capacity to reconstruct a pattern of cortical columns, we added noise to the voxelized image, then computed the correlation between the outcome (voxelized map + noise) and the real pattern (bottom right).

Methods

Modeling fMRI of a columnar organization

We build on a model that we previously developed to study the mechanisms underlying fMRI-based decoding of information conveyed in cortical columns (Chaimow et al., 2011). Here, the purpose of the model was to analyze the distribution of voxel-wise differential responses relative to noise in order to optimize data acquisition parameters for detection of differential responses, for decoding and for pattern correlation.

Details of the underlying model have been described previously. All mathematical details and derivations with modifications and extensions to the model (as described below) are presented in the Appendix. Table 1 lists all parameters and their values. Here we briefly describe the structure of the model, and the different stages it involves.

The model (Figure 1) consisted of a stage for creating a columnar pattern, a neuronal response stage, a spatiotemporal BOLD response stage and an MRI sampling stage. In addition, noise was modeled according to Triantafyllou et al. (2005), while taking into account estimation of differential responses obtained from multiple measurements in time (Appendix C). The model followed Chaimow et al. (2011), differing only in the following aspects.

Isotropic patterns of cortical columns

In Chaimow et al. (2011) we simulated ocular dominance maps. Except for the occasional appearance of linear zones in other cortical maps (Shmuel and Grinvald, 2000), the anisotropic pattern of ocular dominance columns is unique. In order to study generic columnar patterns, we modeled isotropic patterns by spatial filtering of Gaussian white noise (Rojer and Schwartz, 1990) using an isotropic band-pass filter. The filter was parameterized by the main pattern frequency ρ , which determined the width of the columns, and by an irregularity parameter δ (Fig. 2). The irregularity parameter controlled the distribution width of contributing spatial frequency components. The inverse of the main pattern frequency is the pattern's main cycle length, which is twice as long as the width of a typical column (defined as the local set of neurons that respond preferentially to the same stimulation condition in differential analysis of responses to two orthogonal stimuli).

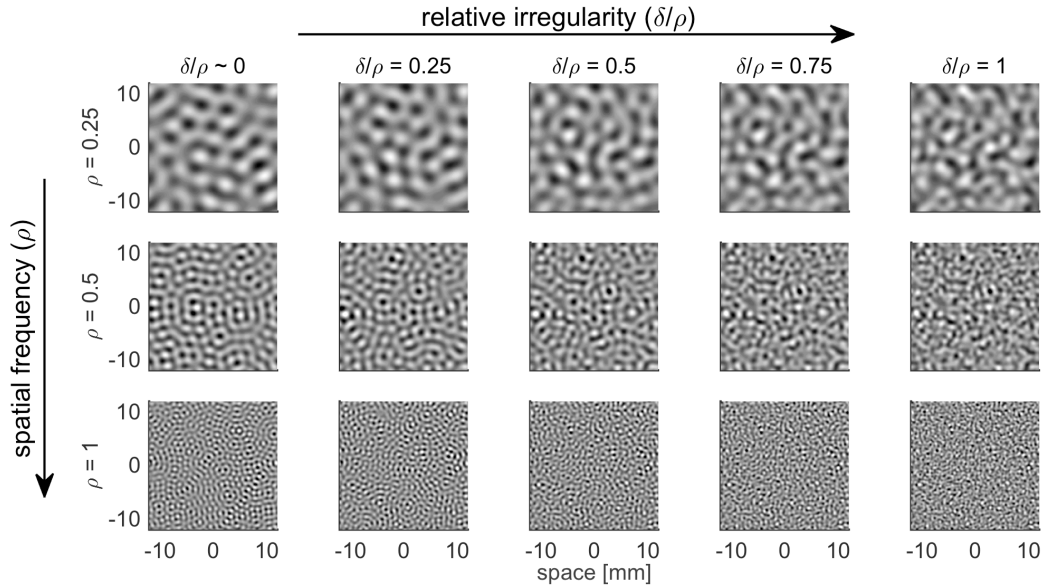


Fig. 2 Parameterization of patterns of cortical columns. Patterns of cortical columns were modeled as spatially band-pass filtered white noise. The filter was parameterized by a main spatial frequency (ρ) that controlled the typical diameter of columns, and by the relative width of the band-pass (δ/ρ) controlling their irregularity. This figure illustrates the effects of the two parameters on the modeled pattern. It shows different patterns resulting from combinations of $\rho \in \{0.25, 0.5, 1.0\}$ and $\delta/\rho \in \{0, 0.25, 0.5, 0.75, 1.0\}$.

BOLD point-spread function widths and response amplitudes

BOLD point-spread function widths of 2.8 mm, 1.02 mm and 0.82 mm full-width-at-half maximum (FWHM) (Chaimow et al., 2016) and BOLD response amplitudes of 5%, 6% and 4% were used for gradient echo (GE) BOLD imaging at 3T, GE imaging at 7T and spin echo (SE) BOLD imaging at 7T, respectively (see Discussion for rationale of parameter choices). When not specified otherwise, parameters of GE fMRI at 7T were implemented.

Noise model

Voxel volume dependent noise was modeled by extending a model from Triantafyllou et al. (2005). We fitted the model to the voxel volume dependent measurements of time-course signal-to-noise ratio (tSNR) that Triantafyllou et al. (2005) reported for 3T and 7T using a repetition time (TR) of 5.4 s (their Figure 5). The fitting was done using a non-linear least squares algorithm (Trust-Region, MATLAB, The MathWorks Inc., Natick, MA, USA). It yielded estimated parameter values of $\kappa = 6.657$ at 3T, $\kappa = 9.9632$ at 7T (κ is a field-dependent constant governing the relationship between voxel volume and image SNR), $\lambda = 0.0129$ at 3T and $\lambda = 0.0113$ at 7T (λ is a mostly field-independent constant determining the asymptotic level of physiological noise). In order to model noise using different TRs, we implemented a modification of the tSNR equation (Chaimow et al., 2011). Gray matter longitudinal relaxation constants (T1), necessary for converting between different TRs, were taken from Wright et al. (2008) (1.607 s at 3T, 1.939 s at 7T). For all simulations, we used a TR of 2 s.

Table 1 List of parameters. This table shows a list of all parameters employed by the model together with their assigned values.

Parameter	Description	Value
L	side length of simulation grid	24 mm
N	number of grid points per dimension	512
ρ	main pattern spatial frequency	1/1.6 mm (default)
δ/ρ	relative degree of irregularity	0.5 (default)
$\text{fwhm}_{\text{psf}}^{3\text{TGE}}$	point-spread function FWHM, GE at 3T	2.8 mm
$\text{fwhm}_{\text{psf}}^{7\text{TGE}}$	point-spread function FWHM, GE at 7T	1.02 mm (default)
$\text{fwhm}_{\text{psf}}^{7\text{TSE}}$	point-spread function FWHM, SE at 7T	0.82 mm
$\beta^{3\text{TGE}}$	amplitude of point-spread function, GE at 3T	5%
$\beta^{7\text{TGE}}$	amplitude of point-spread function, GE at 7T	6% (default)
$\beta^{7\text{TSE}}$	amplitude of point-spread function, SE at 7T	4%
TR	repetition time	2 s
N_T	number of measurements (sum of the two conditions)	1000 (default)
w	voxel width	no specific value
d_{slice}	slice thickness	2.5 mm
$\kappa^{3\text{T}}$	image SNR relative to voxel volume at 3T	6.6567
$\kappa^{7\text{T}}$	image SNR relative to voxel volume at 7T	9.9632
$\lambda^{3\text{T}}$	asymptotic level of physiological noise at 3T	0.0129
$\lambda^{7\text{T}}$	asymptotic level of physiological noise at 7T	0.0113
$T_1^{3\text{T}}$	longitudinal relaxation time constant of gray matter at 3T	1.607 s
$T_1^{7\text{T}}$	longitudinal relaxation time constant of gray matter at 7T	1.939 s
τ	time constant of temporal correlations of physiological noise	15 s
$p_{\text{detect}}^{\text{min}}$	minimal detection probability	80%
α	maximal type I error rate	5%
R^{min}	minimal pattern correlation	$0.7071 = \sqrt{0.5}$

tSNR was also adapted to estimate the signal-to-noise ratio of a differential response obtained from multiple measurements (Appendix C). Temporal correlations were taken into account while considering thermal and physiological noise components separately. Thermal noise was assumed to be independent over time. Physiological noise correlations were modeled using an autoregressive model (Purdon and Weisskoff, 1998). The adapted noise model was then used to compute differential multi-measurement SNR. Data for each condition was assumed to consist of 500 measurements.

Derived quantities

Contrast range and contrast to noise ratio

We use the term *contrast range* to characterize the expected differential responses distributed over voxels. We defined contrast range as the standard deviation of the distribution of differential responses obtained from all voxels.

Contrast range was multiplied with expected differential multi-measurement SNR, resulting in contrast to noise ratio (CNR).

In our model, contrast range completely specifies the expected distribution of differential responses, which is a normal distribution with zero mean. The reason is that modeled differential fMRI responses are the result of multiple linear transformations of Gaussian white noise, which were used for creating the columnar pattern. Similarly, the expected distribution of differential con-

trasts relative to noise is a normal distribution with zero mean and standard deviation defined by contrast to noise ratio.

Detection probability

The differential response of a voxel relative to noise determines the probability of univariate detection of this response using the time-course of that voxel (Appendix D).

Percentiles of the distribution of detection probabilities across voxels were computed by transforming percentiles of the CNR dependent distribution of differential responses relative to noise. Average detection probability across voxels was computed by numerical integration over the CNR dependent distribution of differential responses relative to noise.

Together with the number of voxels, CNR also determines the probability of multivariate detection of a condition-specific response from the time-courses of a number of voxels (Appendix D). Unless stated otherwise, detection probabilities were computed analytically using formulas derived in Appendix D. Minimally required detection probability was set to 0.8.

Simulation of multivariate decoding

Multivariate detection and decoding was simulated as a function of differential multi-measurement CNR (assuming 500 volumes for each condition) and number of voxels. Single condition response trials of 10 volumes were assumed.

Each trial's response was modeled as the sum of a condition specific response pattern and noise. The response pattern for condition one was a random vector. It was drawn once for each simulation consisting of 100 response trials. The vector's components followed independent normal distributions with zero means and standard deviations equal to one half of the differential multi-measurement CNR. The response pattern for condition two was the negative of the response pattern for condition one. This is not a restriction because it can be achieved for any data set by mean-centering the data. The noise was a random vector, drawn for each response trial. Its components followed independent normal distributions with zero means and standard deviations equal to one half of the square root of the number of response trials. This assignment of response and noise amplitudes ensured that the CNR that can be obtained from the difference between the average of all condition one and the average of all condition two responses was equal to the specified multi-measurement differential CNR.

As part of a leave-one-out procedure, all simulated responses except for one were averaged according to their condition. The left out test trial's response was projected onto the difference between the separately averaged responses of condition one and condition two. This classification was counted as successful if the projection was positive and the test trial was associated with condition one or if it was negative and the test trial was associated with condition two. This procedure was repeated such that each trial served as the test trial once and the overall rate of successful classification was computed.

The entire simulation was repeated 100,000 times, resulting in a distribution of successful classification rates. In addition, the same simulation was run

100,000 times, with the difference that the response pattern was always set to zero. These latter simulations served to provide a distribution of successful classification rates under the null-hypothesis of no condition specific response.

Finally, the probability of statistically significant decoding was computed as the fraction of the distribution of successful classification rates that was higher than the 95 percentile of successful classification rates under the null hypothesis.

Simulation of multivariate detection

Multivariate detection was simulated within the same simulations as described for simulation of decoding. All simulated responses were averaged according to their conditions. The differential response was calculated as the differences of these averages. For each simulation, the vector-norm of this response was calculated resulting in a distribution of response vector-norms. The same distribution was calculated with response patterns set to zero, resulting in a distribution of response vector-norms under the null-hypothesis of no condition specific response.

The probability of detection was computed as the fraction of the distribution of response vector-norms that were higher than the 95 percentile of vector-norms under the null hypothesis.

Pattern correlation

Pattern correlation describes how similar the imaged pattern is to the real pattern. It was computed as the correlation coefficient between the modeled real columnar pattern and the corresponding imaged pattern including noise. We added noise to each simulated voxel response by drawing a random number from a normal distribution with standard deviation equal to the voxel size dependent noise level (using single-voxel multi-measurement SNR). Before computing the correlation coefficient, the voxel pattern with the added noise was interpolated to match the resolution of the modeled real pattern by zero padding in the spatial frequency domain.

Minimally required pattern correlation was set to $\sqrt{0.5}$, so that the coefficient of determination (R^2) was at least 0.5.

Model Implementation

The model was implemented using numerical simulations in MATLAB (The MathWorks Inc., Natick, MA, USA) on a grid of 512×512 evenly spaced points, representing a field of view of 24×24 mm². For each simulation, the contrast range estimates were averaged over 32 individual simulations in order to increase the accuracy of the results.

Estimation of optimal voxel size

Simulating k-space sampling required restricting the spatial frequency representation to a voxel size dependent subset of frequencies. Due to the discrete nature of the simulation grid, only certain voxel sizes could be simulated. In order to estimate the optimal voxel size, 40 different voxel sizes were tested. The voxel sizes were spaced less than 0.03 mm apart for voxels of up to 0.6

mm in width, less than 0.09 mm apart for voxels of up to 1 mm in width and less than 0.3 mm apart for voxels of up to 2 mm in width. The three largest tested voxels were 2.4 mm, 3 mm and 4 mm wide.

Linear functions were fitted to the estimated optimal voxel sizes as a function of the main cycle length of the columnar organization using MATLAB Statistics Toolbox (fitlm; The MathWorks Inc., Natick, MA, USA). For the fitting, we considered only the range in which detection probability was higher than 0.8 and pattern correlation was higher than 0.7071.

Results

Modeling functional MRI of a columnar organization

In order to evaluate how well a cortical columnar pattern can be imaged using fMRI and to analyze the contributions of different mechanisms, we adapted a model of imaging cortical columns (Chaimow et al., 2011). More specifically, we modeled the differential responses obtained under two opposing stimulation conditions in a set of voxels. We assumed these voxels to be part of a single 2.5 mm thick slice, parallel to and overlapping with a flat region of cortical gray matter.

Figure 1 illustrates the model and presents one instance of its numerical realization. The model made it possible to vary the spatial frequency and the regularity of the columnar pattern (see also Figure 2), the response amplitude, the width of point-spread function of the BOLD response and the voxel size of the MR measurement. In Figure 1 we used a point-spread function with a full-width-at-half-maximum (FWHM) of 1.5 mm, a voxel width of 1 mm, an average columnar pattern cycle length of 1.6 mm (typical for orientation columns in human V1, Yacoub et al., 2008) and a moderate degree of irregularity. The results of both the BOLD response stage and the subsequent voxel sampling showed condition-specific patterns. Nonetheless, these patterns show some loss of fine-scale detail compared to the original columnar patterns. In addition, following the BOLD response and voxel sampling stages, the functional contrasts are relatively small.

Measures of the quality of imaging a pattern of cortical columns

We quantified how well a columnar pattern could be imaged using two different approaches.

The first approach was to quantify the successful *detection* of differences between responses to two stimulus conditions. In practice, this detection could be carried out as an explicit statistical test or, alternatively, by means of demonstrating successful decoding of the stimulus condition from the data using a classification algorithm. To this end, we estimated the standard deviation of the distribution of differential responses as it reflects the dispersion of condition-specific contrasts (e.g., contrast between responses to stimuli with orthogonal orientation) in a set of imaged voxels (Figure 1, bottom left). The larger this standard deviation, the larger the contrast values that exist in the distribution. Throughout the rest of the manuscript, the standard deviation of

the distribution of differential functional contrast will be referred to as the *contrast range*.

The second approach was to quantify the capacity to *reconstruct* the pattern of cortical columns. We measured this capacity by calculating the similarity (correlation coefficient) between the underlying columnar pattern and the imaged pattern with added noise (pattern correlation; Figure 1, bottom right).

Measures influenced by voxel size

We analyzed how measures related to detection, decoding and pattern reconstruction depend on voxels size through analytical calculations and numerical simulations of our model. All simulations in this section assume a columnar pattern resembling orientation columns in humans (cycle length = 1.6 mm, moderate degree of regularity). Furthermore, we considered noise and BOLD point-spread parameters associated with 7T GE imaging (point-spread function FWHM = 1.02 mm).

Contrast range as voxel size dependent integration of power in the frequency space

First, we studied the factors that determine the distribution of differential contrast values and contrast range. Using a mathematical formulation of the model, we derived a relationship between the power spectrum of a columnar pattern and contrast range (Appendix A). When given the power spectrum $|B(k_1, k_2)|^2$ of a columnar fMRI response pattern of size $L_1 \times L_2$, the contrast range c can be calculated as:

$$c = \sqrt{\frac{1}{L_1 L_2} \int_{-1/2w}^{+1/2w} \int_{-1/2w}^{+1/2w} |B(k_1, k_2)|^2 dk_1 dk_2},$$

where w is the voxel width, and k_1 and k_2 are coordinates in k-space.

This relationship gives rise to a simple interpretation of contrast range (Fig. 3). The power spectrum of a neuronal columnar response pattern characterizes the contrast available at each spatial frequency in the original pattern. The fMRI response modulates this information by suppressing higher frequencies relative to all other frequencies. In the case of a unimodal spatial frequency distribution, it transforms the original pattern into a BOLD-fMRI pattern with a lower apparent main pattern frequency ($0.54 = 1/1.84$ cycles/mm relative to $0.63 = 1/1.6$ cycles/mm main pattern frequency; compare vertical lines in Figure 3A) and an apparent lower degree of irregularity (not shown). Finally, the MRI voxel sampling performs a limited integration in the spatial frequency domain, cutting off all components with spatial frequencies higher than what the voxel can sample (highest fully sampled frequency equals one over twice the voxel width).

Therefore, when the distribution of spatial frequency components is unimodal, the largest increase in this integral with decreasing voxel width is expected to occur when the integral includes the peak of that distribution (vertical red line in Figure 3B left and right). Accordingly, contrast range, which is proportional to the square root of this integral, increases most when reducing voxel width to approximately half the cycle length of the apparent main pattern frequency ($1.84 \text{ mm}/2 = 0.92 \text{ mm}$).

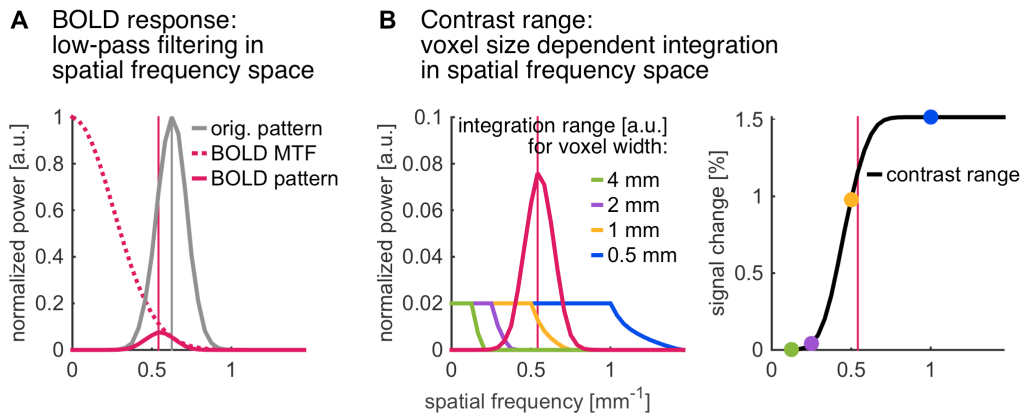


Fig. 3 Contrast range as integration of spatially filtered pattern components. A The spatial power spectrum (in gray; averaged over all directions) of a simulated columnar pattern with a cycle length of 1.6 mm ($\rho=1/1.6$) and a moderate degree of irregularity ($\delta/\rho=0.5$). The dotted red curve shows the squared BOLD modulation transfer function (MTF; the spatial frequency representation of the BOLD point spread function). The BOLD response acts as a spatial filter that modulates the spectrum, resulting in the spatial power spectrum of the BOLD response pattern (red curve). The peak of the spectrum (vertical red line) is shifted towards lower frequencies (the vertical gray line shows the peak of the original spectrum for comparison). B Contrast range obtained by MRI voxel sampling can be modeled as integration of contrast in the spatial frequency space. The red curve shows the spatial power spectrum of the BOLD response pattern (same as in A). The remaining curves represent the weighting with which voxels with width 4 mm (green), 2 mm (purple), 1 mm (orange) and 0.5 mm (blue) integrate the pattern's power spectrum. The gradual fall-off in these weighting curves results from the averaging over all directions of the two-dimensional rectangular integration area. C Contrast range as a function of spatial frequencies associated with different voxel widths (marked by dots colored as in B). Contrast range is the square root of the voxel-size-dependent integral of the BOLD-filtered pattern's power spectrum.

Voxel size that optimizes CNR and univariate detection of differential response

In order to understand the processes involved in optimizing univariate detection of information represented in columnar patterns, we combined contrast range and SNR as a function of voxel size. Time-course SNR characterizes the noise level in single measurements from single voxels (Triantafyllou et al., 2005). We adapted this SNR measure by estimating effective noise increases due to differential analysis and decreases due to averaging multiple measurements within each condition, taking temporal correlations into account. We calculated differential multi-measurement SNR (using 500 measurements for each condition, $TR = 2$ s) and contrast range as a function of voxel width.

Contrast range decreased with increasing voxel width (Figure 4A, black). It sharply fell off with increasing voxel size for voxel width close to 1 mm, which is approximately equal to half the apparent columnar cycle length (0.92 mm; true cycle length = 1.6 mm). SNR, however, increased with increasing voxel width (Figure 4A, orange). The increase in SNR was more gradual than the drop of contrast range.

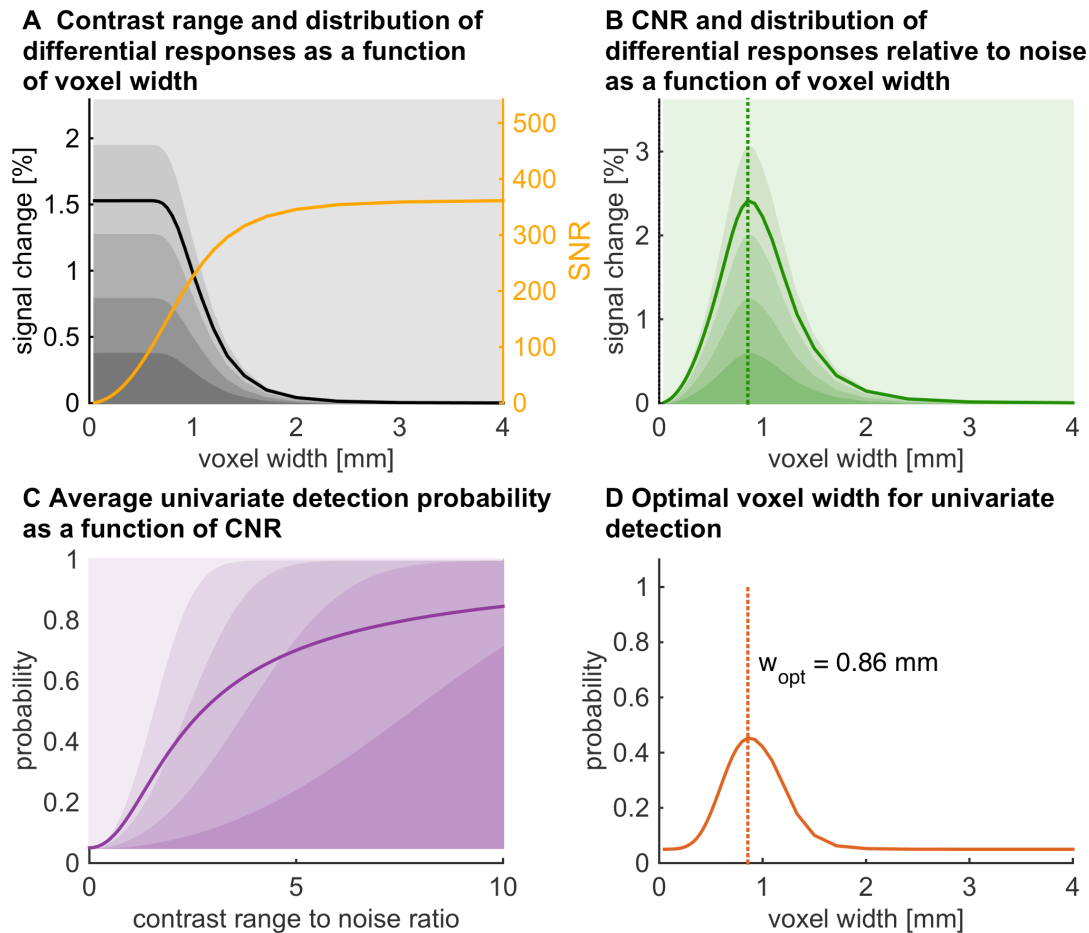


Fig. 4 CNR and univariate detection: dependence on voxel size. We simulated fMRI of a columnar pattern with a main cycle length of 1.6 mm ($\rho=1/1.6$) and intermediate degree of irregularity ($\delta/\rho=0.5$), using parameters of 7T GE fMRI. A Contrast range (black curve) and differential multi-measurement SNR (in orange) are presented as a function of voxel width. Contrast range is the standard deviation of the distribution of differential contrasts sampled by all voxels. Individual voxels differ in contrast. The gray bands make it possible to visualize the distributions of contrasts as a function of voxel-width. For each voxel width, each band presents the differential responses of 20% of voxels. With increasing voxel width, the contrast range decreased while the SNR increased. B Contrast to noise ratio (CNR; the product of contrast range and single-voxel SNR) is presented as a function of voxel width (green curve). It reached its maximum (at 0.86 mm, dotted line) close to the point where contrast range started to drop (A). The distributions of differential responses relative to the noise level from all voxels are presented in green bands as a function of voxel width. Each band represents 20% of the voxels. C The probability of detecting a differential response in single voxels (univariate analysis) is shown as a function of CNR. For each CNR there is a distribution of differential responses relative to the noise level, from we computed the corresponding distribution of detection probabilities. This latter distribution is presented in purple bands. Each purple band represents 20% of all voxels. The purple curve shows the average detection probability computed over the entire distribution of voxels for each CNR value. D CNR as a function of voxel width (B) and average univariate detection probability as a function of CNR (C) were combined, resulting in average univariate detection probability as a function of voxel width (solid red curve). Similar to CNR, it reached its maximum for a voxel width of 0.86 (dotted line).

Contrast range characterizes the distribution of differential responses of a set of voxels. Figure 4A also presents the expected distributions of actual differential responses (gray bands; each band represents 20% of voxels).

Next, we multiplied voxel dependent contrast range and differential multi-measurement SNR in order to obtain contrast to noise ratio (CNR). CNR relates the range of differential contrasts that can be expected in a set of voxels to the level of noise in which they are immersed. It also characterized the dis-

tribution of differential responses relative to noise across voxels. The higher the differential response relative to the expected noise in a single voxel, the more likely it is that this differential response can be detected using the time courses of that voxel. Figure 4B shows CNR as a function of voxel width. It assumed its maximum at a voxel width of 0.86 mm, which is close to half the apparent columnar cycle length (0.92 mm, true cycle length = 1.6 mm). Again, we visualized the expected distributions of actual differential responses relative to noise (Figure 4B; green bands; each band represents 20% of voxels).

The ratio of the differential response in a voxel to the expected noise determines how likely it is to detect this response in the time-course of that voxel. We calculated this univariate detection probability assuming a type-I error rate smaller than 5% (Appendix D). Figure 4C shows the distribution of detection probabilities over voxels as a function of CNR (purple bands; each band contains 20% of voxels). For low CNR, most of the voxels' univariate detection probabilities were not higher than the type-I error level of 5%. With increasing CNR the detection probability of more and more voxels increased and approached 1 eventually. However, even with CNR as high as 10 there was still a significant fraction of voxels with low or intermediate detection probabilities.

In order to evaluate univariate detection across voxels, we calculated the probability of detecting a differential response in any voxel given the CNR of the entire pattern (not knowing where this voxel's response lies in the distribution of responses; Figure 4C; purple curve). This was done by calculating the expected value of detection probability with respect to its distribution over voxels. As expected, low CNR resulted in low average detection probability. For our employed parameters (7T GE imaging, moderately irregular pattern with 1.6 mm cycle length) a differential multi-measurement CNR of 7.67 was needed in order to achieve an average detection probability of 80%.

Having studied how CNR depends on voxel width and how univariate detection probability depends on CNR, we were able to combine both stages in order to analyze the dependence of univariate detection probability on voxel width. Figure 4D shows univariate detection probability (across voxels) as a function of voxel width. Like CNR, it assumed its maximum at a voxel width of 0.86 mm. The resulting detection probability was 45%.

Voxel size that optimizes multivariate detection and decoding

In multivariate analysis of fMRI data, information from individual responses in multiple voxels is combined. As a result, the probability to detect a multivariate response depends on the differential responses of individual voxels (whose distribution depends on CNR) as well as the number of available voxels (Appendix D). Figure 5A shows multivariate detection probability as a function of CNR and number of voxels.

An alternative approach to multivariate analysis is to train an algorithm on a subset of the data with the goal to classify (decode) the stimulation condition of an unseen separate subset of the data. A high rate of successful decoding would indicate the presence of stimulus specific responses in the data.

We simulated linear classification of responses in a pattern of cortical columns as a function of CNR and number of voxels. We then calculated the probability of obtaining a rate of successful decoding higher than that obtained by chance, assuming a type-I error rate smaller than 5%. The probability of significant decoding (Figure 5A right, purple circles) was very similar to the probability of multivariate detection (Figure 5A right, green curves). This similarity increased with increasing number of voxels. Therefore, for the remainder of the study, we used multivariate detection probability as a single measure for optimizing multivariate analysis of fMRI data.

In analogy to univariate detection, we combined our findings on CNR as a function of voxel width, our analytical formula of multivariate detection probability as a function of CNR and the number of voxels to calculate multivariate detection as a function of voxel width. Note that we assumed a fixed area of cortex to image from. Therefore, the number of available voxels decreased with increased voxel width.

Figure 5B shows multivariate detection probability as a function of voxel width. Multivariate detection probability was numerically indistinguishable from 1 across a range of voxel widths (0.35 mm – 1.2 mm). We defined the optimal voxel width to be the average over this range, which was 0.78 mm.

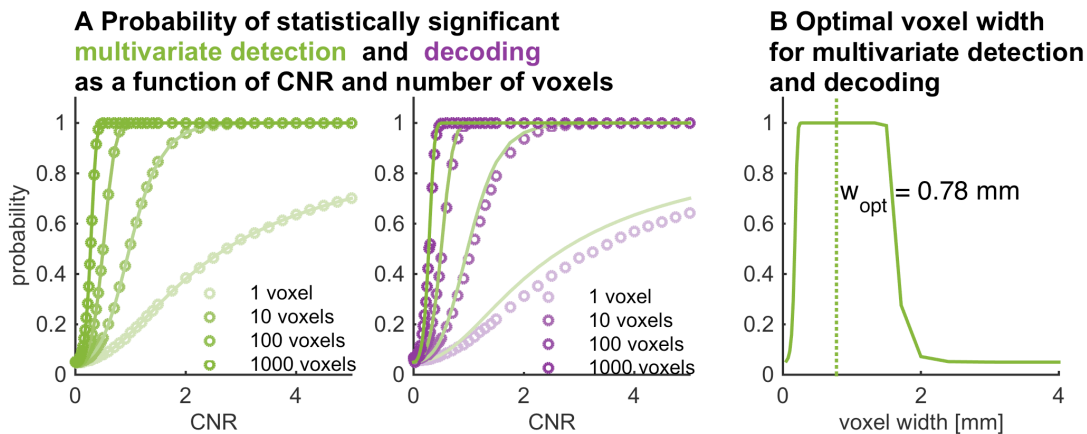


Fig. 5 Multivariate detection and decoding: dependence on voxel size. A Multivariate detection and decoding as a function of CNR and number of voxels. We simulated multivariate detection and decoding from 1,000 volumes, assuming response patterns with a given differential multi-measurement CNR and a given number of voxels. From these simulations we calculated the probability of statistically significant multivariate detection (green circles, left) and decoding (purple circles, right) assuming a type 1 error not higher than 5% ($p < 0.05$). In addition, we used analytical formulas to calculate multivariate detection probability directly (green curves in the two panels of A). Multivariate detection probability was very similar to the probability of decoding. B CNR as a function of voxel width (Figure 4B) and multivariate detection probability as a function of CNR and number of voxels (A) were combined, resulting in multivariate detection probability as a function of voxel width (green curve). The multivariate detection probability had a broad peak for a range of voxel widths (0.35 mm – 1.2 mm) whose mean was 0.78 mm (vertical line).

Voxel size that optimizes reconstruction of a pattern of cortical columns

In order to understand the factor underlying the emergence of optimal voxel size for reconstruction of a columnar pattern, we calculated the correlation between an imaged pattern and its true underlying pattern as a function of voxel width.

Figure 6A shows pattern correlation for an idealized noiseless measurement together with differential multi-measurement SNR as a function of voxel width. Similar to the contrast range in Figure 4A, the noiseless pattern correlation decreased with increasing voxel width and fell of very fast when the voxel width became larger than half the apparent columnar cycle length (0.92 mm; true cycle length = 1.6 mm).

Unlike CNR, pattern correlation cannot be decomposed into a product of SNR and a noiseless measure of contrast. Nonetheless, the noiseless pattern correlation and the SNR curves illustrate the two antagonistic factors for optimal voxel size for pattern correlation of imaged patterns containing realistic noise.

Figure 6B presents pattern correlation incorporating realistic noise, as a function of voxel width for varying numbers of measurements (nT). The number of measurements had a small influence on the optimal voxel width by reducing the level of noise. This caused pattern correlation to become more similar to the ideal noiseless pattern correlation curve, leading to decreases in optimal voxel width. However, the differences of optimal voxel width for different numbers of measurements were too small to be practically relevant. The optimal voxel width when using 1000 measurements was 0.67 mm, considerably smaller than half the apparent cycle length (0.92 mm) and smaller than half the true cycle length (1.6 mm/2 = 0.8 mm).

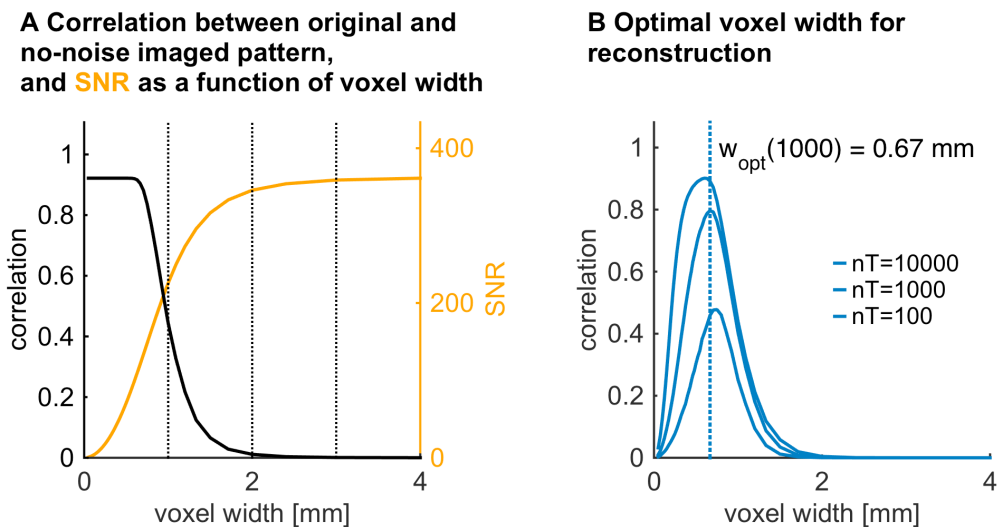


Fig. 6 Reconstruction of a pattern: dependence on voxel size. We simulated fMRI of a columnar pattern with a main cycle length of 1.6 mm ($\rho=1/1.6$) and intermediate degree of irregularity ($\delta/\rho=0.5$), using parameters of 7T GE fMRI. A The correlation between imaged patterns (simulated without measurement noise) and the real pattern (pattern correlation; black) is shown as a function of voxel width. The orange curve shows the differential multi-measurement SNR as a function of voxel width. With increasing voxel width, the noise-less pattern correlation decreased while the SNR increased. B The correlation between imaged patterns (with added measurement noise) and the real pattern is shown as a function of voxel width for different number of measurements. Optimal voxel width decreased with increasing number of measurements. The voxel width that optimized pattern correlation for 1,000 measurements was 0.67 mm (vertical dotted line).

Summary: optimal voxel widths for detection, decoding and reconstruction

In summary, for a columnar pattern resembling orientation columns in humans (cycle length = 1.6 mm, with moderate degree of regularity) and imaging parameters associated with 7T GE imaging, the voxel width that optimized the reconstruction of a columnar pattern (0.67 mm) was smaller than the voxel width that optimized univariate detection of responses from that pattern (0.86 mm). The voxel width that optimized multivariate detection of condition-specific information was in-between (0.78 mm), although a wide range of voxel widths resulted in very similar multivariate detection probabilities.

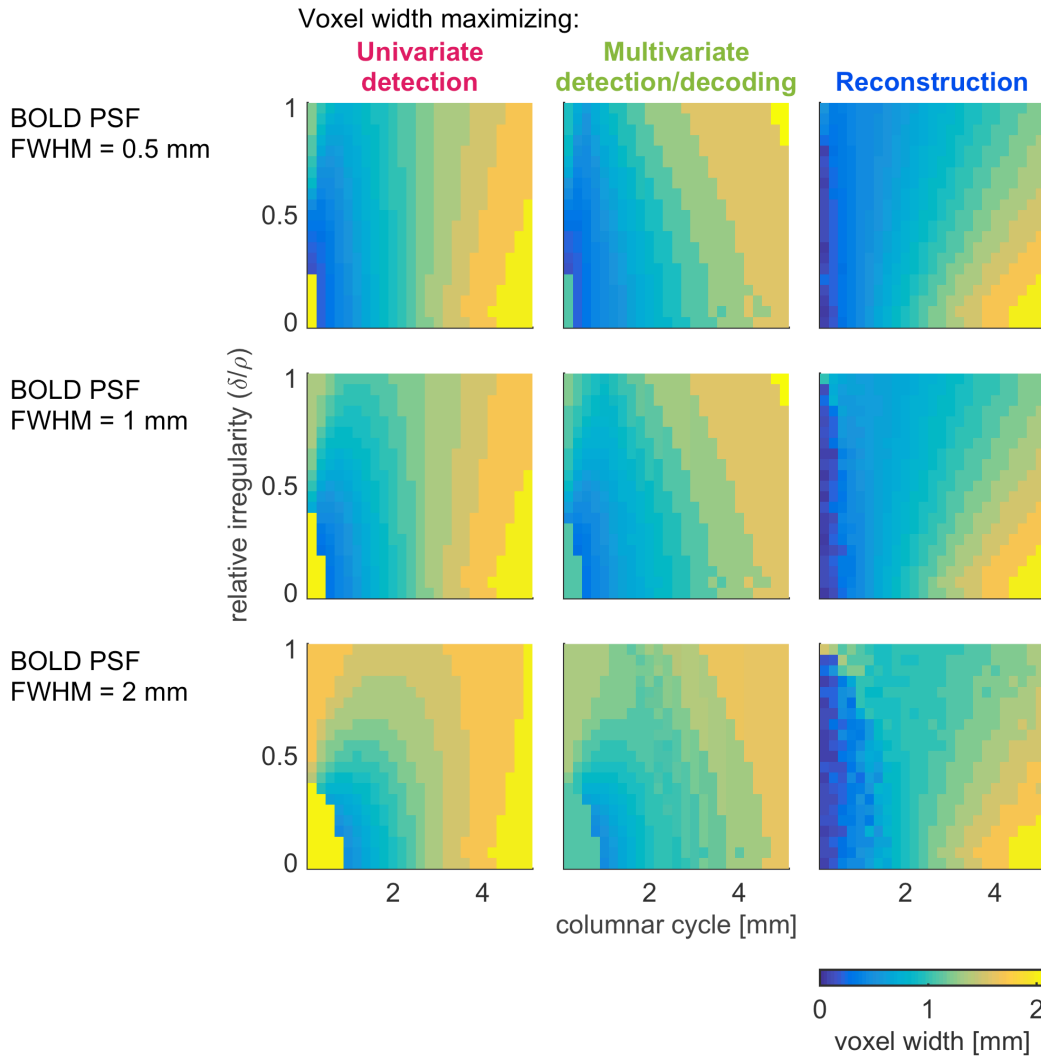


Fig. 7 Optimal voxel size as a function of the pattern's main cycle, regularity and BOLD point-spread. The plots show the voxel widths that maximized univariate detection, multivariate detection and decoding, and pattern correlation as a function of the main cycle length of the columnar pattern ($1/\rho$; horizontal axes) and relative pattern irregularity (δ/ρ ; vertical axes) for three different BOLD point-spread function widths. The simulated BOLD response and the noise corresponded to 7T GE imaging.

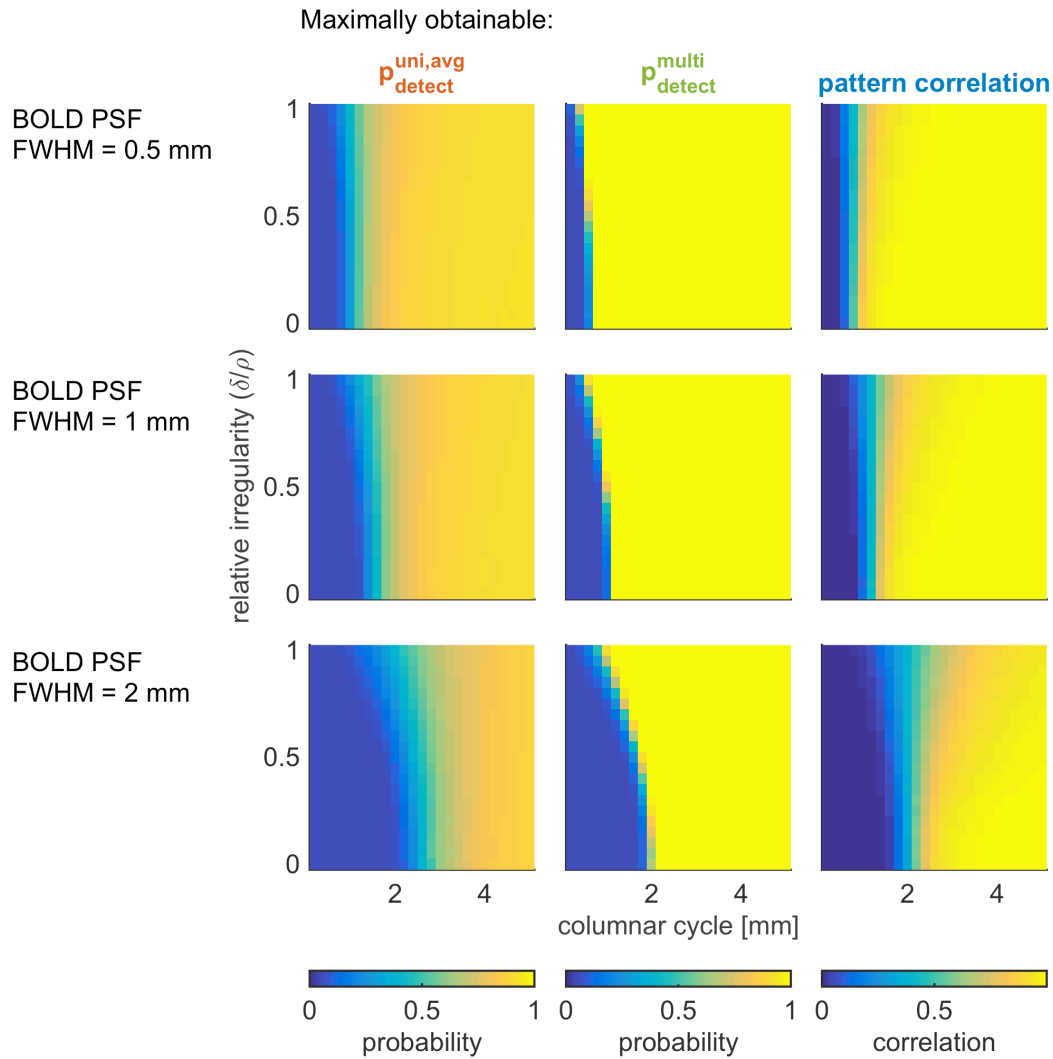


Fig. 8 Maximally obtainable detection probabilities and pattern correlations using optimal voxel sizes as a function of features of the pattern’s main cycle, regularity and BOLD point-spread. The plots show the average univariate detection probability ($p_{\text{detect}}^{\text{uni,avg}}$), the multivariate detection and decoding probability ($p_{\text{detect}}^{\text{multi,avg}}$), and pattern correlation that could be obtained (using the optimal voxel sizes from Figure 7) as a function of cycle length of the columnar pattern ($1/\rho$; horizontal axes) and relative pattern irregularity (δ/ρ ; vertical axes) for 3 different BOLD point-spread function widths. The BOLD response and the noise corresponded to 7T GE imaging.

Dependence of optimal voxel width on columnar pattern organization and BOLD point-spread width

So far we have shown how functional contrast, the probability to detect a differential response and the capacity to reconstruct a columnar pattern depend on voxel size and how the optimal voxel width emerges. Next, we estimated how optimal voxel width changes as a function of the spatial organization of the pattern and as a function of the BOLD point-spread function.

We calculated univariate and multivariate detection probability and pattern correlation as a function of voxel width by repeating the simulations presented in the previous section (Figures 4D, 5B and 6B) while varying the main pattern frequency, the regularity of the pattern and the width of the BOLD point-

spread function. We considered BOLD response amplitude and noise that corresponded to GE fMRI at 7T. From these simulation results, we determined the voxel widths that optimized univariate and multivariate detection probability and pattern correlation for each combination of pattern parameters and BOLD point-spread width (Figure 7). We also computed the univariate and multivariate detection probability and pattern correlation that can be obtained by these optimal voxel widths (Figure 8). However, in order to better understand the effect that pattern frequency, pattern regularity and BOLD point-spread width had, we considered regular, moderately irregular and irregular patterns separately (Figure 9). In addition, we estimated minimum cycle lengths below which average detection probability would be smaller than 0.8 or pattern correlation would be smaller than 0.7071 (corresponding to $R^2=0.5$). For regular patterns, optimal voxel widths for univariate detection as well as for reconstruction were close to half the columnar cycle length, while optimal voxel widths for multivariate detection were slightly smaller (Figure 9; row 1). Furthermore, optimal voxel widths for regular patterns were relatively independent of the BOLD point-spread function width.

Moving to moderately irregular (Figure 9; row 3) and irregular patterns (Figure 9; row 5) we observed multiple effects. First, compared to regular patterns, optimal voxels for univariate and multivariate detection and reconstruction were less dependent on the columnar cycle length.

As a result, very short cycle lengths resulted mostly in wider optimal voxel widths (more similar to optimal voxel widths for longer cycle lengths) compared to regular patterns. This effect increased with increasing BOLD point-spread width. Note that the shortest cycle lengths of very regular patterns do not follow this trend (Figure 9; row 1, column 1). However, their optimal voxel widths are not practically relevant as their associated detection probabilities are virtually zero.

In contrast to short cycle lengths, optimal voxel widths for long cycle lengths in univariate detection and reconstruction decreased when going to more irregular patterns. For multivariate detection, optimal voxel widths increased for both long and short cycle lengths compared to regular patterns.

Maximally obtainable detection probabilities and pattern correlations increased continuously with increasing cycle length (Figure 9; rows 2, 4 and 6). In the case of univariate detection and reconstruction, these increases became slightly more gradual with increasing irregularity. Multivariate detection probability transitioned rather abruptly as a function of cycle length for all levels of irregularity.

Lastly, with increasing irregularity, the cycle length associated with univariate detection probability of 0.8 and pattern correlation of 0.7017 increased slightly, while the cycle length associated with multivariate detection probability of 0.8 decreased. This shows that the presence of irregularities facilitates the multivariate detection of information represented by the pattern while at the same time it makes it more difficult to reconstruct that pattern accurately.

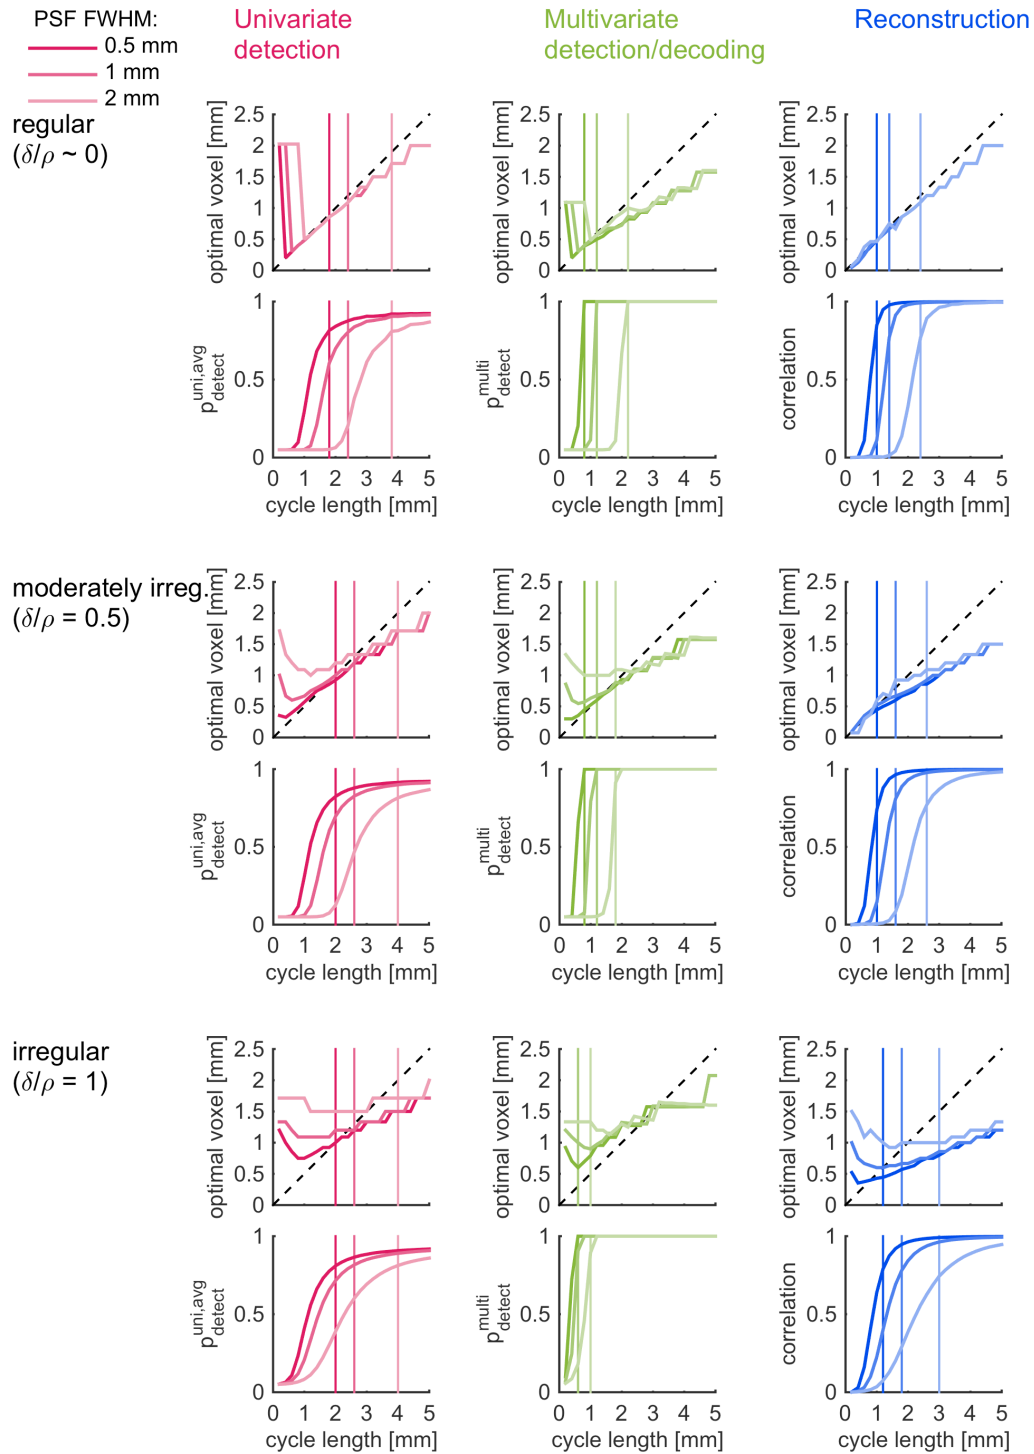


Fig. 9 Dependence of the optimal voxel size on the pattern's main cycle, regularity and BOLD point-spread. This figure highlights a subset of the simulation results presented in Figures 7 and 8. Optimal voxel widths and maximally obtainable univariate and multivariate detection probability and pattern correlation are shown as a function of the main cycle length for regular (top), moderately irregular (middle) and irregular (bottom) patterns. Simulation results obtained with different point-spread widths are color coded. The optimal voxel width for regular patterns is close to half the cycle length (compare to diagonal black dashed lines). Color coded vertical dashed lines show the minimal main cycle length for which detection probability of at least 0.8 or pattern correlation of at least 0.7071 could be achieved. Results for shorter cycles are plotted as dotted lines.

Magnetic field-strength and pulse sequence dependent differences

Thus far, we have analyzed the dependence of optimal voxel sizes and the resulting detection probabilities and pattern correlations on parameters related to the pattern and the MRI measurement process. fMRI at different magnetic field strengths differs with respect to several of these parameters. To assess the implications of these differences between fMRI at ultra-high magnetic field (7T) and a more standard magnetic field (3T), we varied field strength specific parameter settings individually and collectively. For 7T we also compared spin echo BOLD imaging (SE) to gradient echo BOLD imaging (GE).

There are three main magnetic field strength dependent factors with respect to fMRI of cortical columns. These are the amplitude of the BOLD response, the spatial width of the BOLD response, and the voxel size-dependent characteristics of noise. For each factor we calculated univariate and multivariate detection probability and pattern correlation as a function of voxel width. To this end, we repeated the simulations used for presenting Figures 4D, 5B and 6B while varying the parameters associated with each of the three factors.


Varying the response amplitude and BOLD point spread width

The first factor, response amplitude, is the relative signal change evoked by a specific neuronal event. It has been shown that GE response amplitude slightly increases at higher field strength and that it is higher than SE response amplitude (Hennig et al., 2003; Yacoub et al., 2001; 2005). Our response amplitudes characterize the BOLD response one would expect from a hypothetical spatially extended neuronal response, maximally selective for one stimulus condition. Here we used response amplitudes of 4% (7T SE), 5% (3T GE) and 6% (7T GE) (see Discussion for rationale of parameter choices). The upper row of Figure 10 shows univariate and multivariate detection probability and pattern correlation as a function of voxel width for all three amplitudes. Noise characteristics and point spread width were held constant (7T noise and 1.02 mm respectively). All measures benefited from higher responses. The largest increase could be observed for univariate detection (0.30 for 4%, 0.45 for 6% response). The optimal voxel size remained constant for univariate and multivariate detection. Higher responses did not change the voxel size that optimized univariate or multivariate detection, which could be expected since the underlying CNR scales with response amplitude independent of the voxel size. Optimal voxel size for pattern correlation was slightly lower for the highest response amplitude (0.67 mm for 6% response) compared to the other two response amplitudes (0.71 mm for 4% and 5% response).


The second factor is the width of the BOLD point-spread function. The BOLD point-spread function acts as a spatial low pass filter. Therefore, the effect of this factor will depend on the spatial properties of the columnar pattern. Here, we again used a pattern that resembles the properties of orientation columns in V1 (columnar cycle length = 1.6 mm).

Effect of:


A BOLD response amplitude



B BOLD response point-spread width



C Field strength dependent noise



D All field strength and pulse sequence specific factors combined




Fig. 10 Magnetic field strength and pulse sequence dependent effects. The BOLD response amplitude, BOLD point spread width and noise are different for fMRI at different field strength and different pulse sequences. Univariate and multivariate detection probability and pattern correlation are presented as a function of voxel width while varying these contributing factors. Vertical dashed lines show the optimal voxel widths. **A** Effect of response amplitude. Univariate and multivariate detection probability and pattern correlation are shown as a function of voxel width using BOLD response amplitudes of 4% (7T SE; blue), 5% (3T GE; red) and 6% (7T GE; violet); all other factors were held constant (BOLD point-spread function FWHM=0.8 mm, 7T noise). Higher response amplitude obtained with GE fMRI at 7T led to increases in univariate detection probability. It increased multivariate detection probability and pattern correlation only moderately. **B** Effect of point-spread width. Univariate and multivariate detection probability and pattern correlation are shown as a function of voxel width using a BOLD point-spread FWHM of 0.82 mm (7T SE; blue), 1.02 mm (7T GE; violet) and 2.8 mm (3T GE; red); all other factors were held constant (response amplitude of 4%, 7T noise). The wide point-spread of 3T GE led to dramatic decreases in all 3 measures, and to increases in optimal voxel width (dashed lines). **C** Effect of noise. Univariate and multivariate detection probability and pattern correlation are shown as a function of voxel width for 7T noise (blue) and 3T noise (red); all other factors were held constant (response amplitude of 4%, BOLD point-spread function FWHM=0.8 mm). The noise characteristics at 7T relative to 3T caused increases in univariate detection probability, and moderate increases in multivariate detection probability and pattern correlation. **D** All effects combined. Univariate and multivariate detection probability and pattern correlation are presented as a function of voxel width for imaging with 7T SE (blue), 7T GE (violet) and 3T GE (red). Univariate and multivariate detection probability and pattern correlation were dramatically lower at 3T GE relative to 7T SE and 7T GE.

The second row of Figure 10 shows univariate and multivariate detection probability and pattern correlation as a function of voxel width for narrow BOLD point-spread function width (FWHM = 0.82 mm, for 7T SE imaging), intermediate BOLD point-spread function width (FWHM = 1.02 mm, for 7T GE imaging) and wide BOLD point-spread function width (FWHM = 2.8 mm, for 3T GE imaging). Noise characteristics and response amplitude were held constant (7T noise and 4% respectively). The widest point spread (3T GE, FWHM = 2.8 mm) resulted in dramatic reductions in all three measures, making detection or reconstruction virtually impossible. For the other two point spreads, the slightly narrower point spread of 7T SE (0.82 mm) compared to that of 7T GE (1.02 mm) resulted in an increase in univariate detection probability from 0.45 to 0.59 with no change in optimal voxel width. It had no effect on multivariate detection probability which was 1 in either case. Pattern correlation increased from 0.80 (7T GE point spread) using a voxel width of 0.67 mm to 0.90 (7T SE point spread) using a voxel width of 0.63 mm.

The third field-dependent factor is the noise characteristics associated with different magnetic field strengths. The third row of Figure 10 shows univariate and multivariate detection probability and pattern correlation as a function of voxel width for 3T noise compared to 7T noise (noise model from Triantafyllou et al., 2005). Compared to 3T noise, 7T noise resulted in increases of all three measures, in a manner comparable to those obtained by increasing the response amplitude. The largest increase could be observed for univariate detection probability (0.33 to 0.45 at voxel widths of 0.92 and 0.86, respectively). Pattern correlation increased from 0.73 (optimal voxel width 0.71 mm) to 0.80 (optimal voxel width 0.67 mm). Optimal achievable multivariate detection probability remained 1.

In a last step, we combined all three factors in order to estimate the overall differences between different field strengths. The fourth row of Figure 10 shows univariate and multivariate detection probability and pattern correlation as a function of voxel width for 3T (5% response amplitude, 2.8 mm BOLD point-spread function width, 3T noise), 7T GE (6% response amplitude, 1.02 mm BOLD point-spread function with, 7T noise) and 7T SE (4% response amplitude, 0.82 mm BOLD point-spread function with, 7T noise).

Combining all factors, all measures were considerably reduced at 3T GE compared to 7T GE and 7T SE. While response amplitude and noise contributed to this difference, the dominating factor was the width of the BOLD point-spread function. 7T GE and 7T SE showed very similar detection probabilities and pattern correlations. It appeared that SE's advantages due to a narrower point-spread width and disadvantages due to a lower response amplitude cancelled each other out. 7T GE had a slight advantage in univariate detection (detection probability 0.45 for 7T GE, 0.44 for 7T SE), while 7T SE was better at reconstruction (pattern correlation 0.84 for 7T SE, 0.80 for 7T GE). For our employed pattern parameters (intermediate irregularity and 1.6 mm cycle length) there was no detectable difference in optimal voxel widths between 7T GE and 7T SE imaging for any of the three measures.

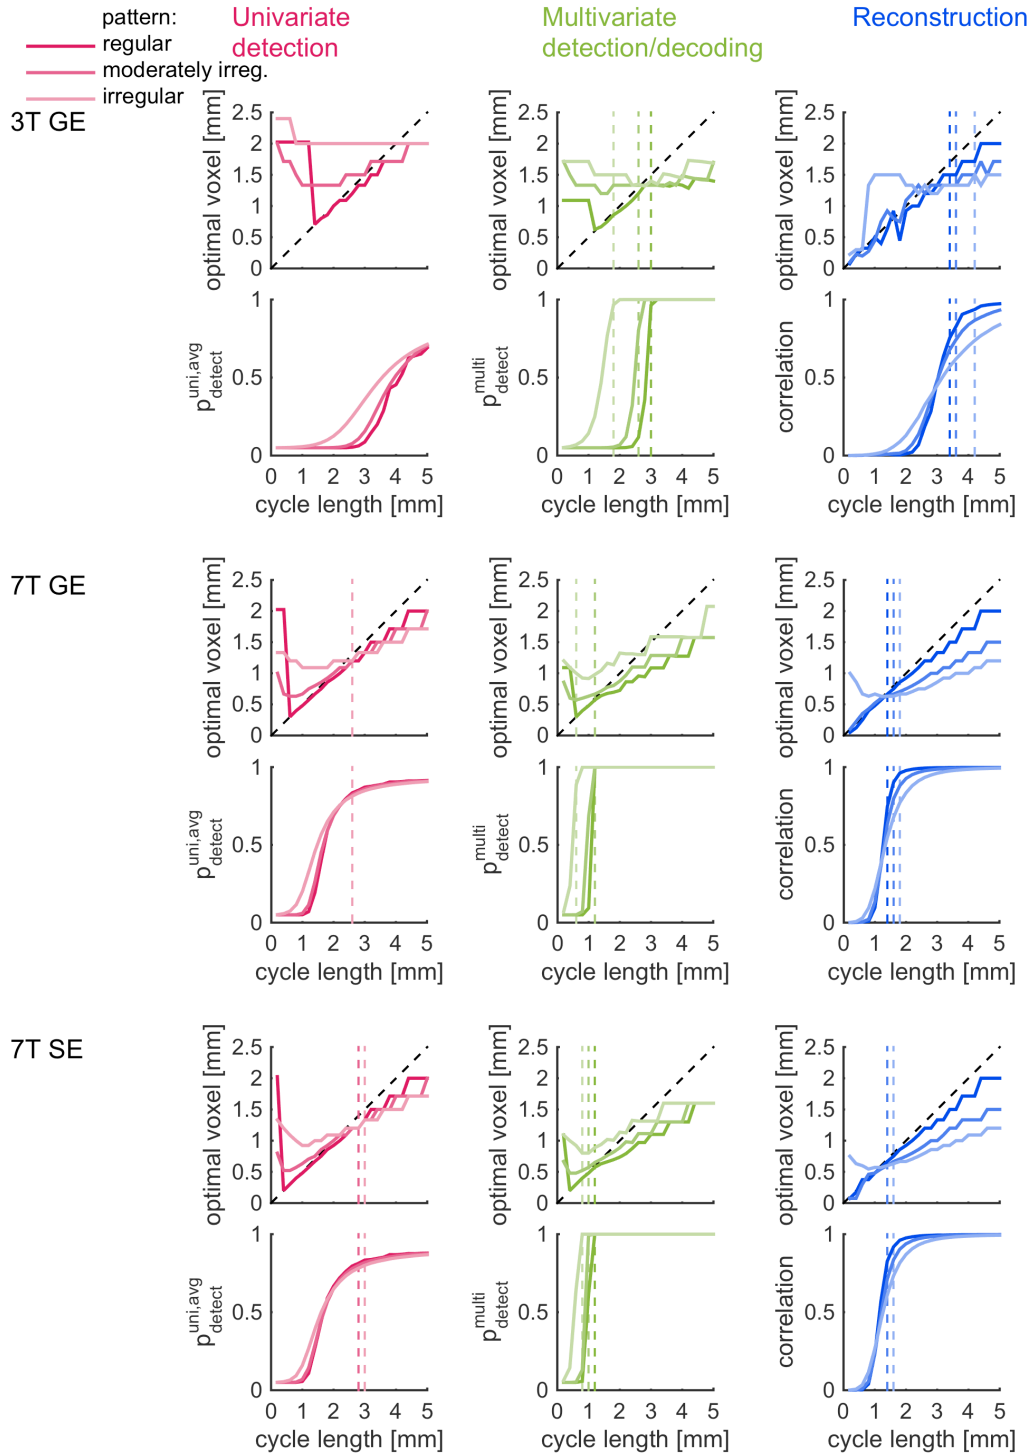


Fig. 11 Optimal voxel size as a function of the main cycle of the columnar pattern for 3T GE, 7T GE and 7T SE BOLD imaging. Optimal voxel widths and maximally obtainable univariate and multivariate detection probabilities and pattern correlation are shown as a function of the main cycle length of the columnar pattern for 3T GE (top), 7T GE (middle) and 7T SE (bottom) BOLD imaging. Simulation results for different pattern levels of irregularity are presented in color code. Color coded vertical dashed lines mark the shortest cycle lengths associated with detection probabilities of at least 0.8 or pattern correlation of at least 0.7071. Results for shorter cycle lengths are plotted as dotted lines. For main cycle length longer than these thresholds, the optimal voxel widths vary approximately linearly with cycle length. The optimal voxel width for regular patterns is close to half the cycle length (compare to the diagonal dashed black lines). The parameters obtained from linear fits to these simulation results are presented in Table 2 as recommendations for optimal voxel size in different scenarios.

Table 2 Optimal voxel widths. This table presents optimal voxel widths for specific imaging scenarios, together with minimal pattern cycle lengths below which detection, decoding, and reconstruction are not feasible. Optimal voxel widths are reported as the sum of a minimal voxel width and a pattern cycle length-dependent term. These linear functions were estimated by fitting lines to optimal voxel width curves in Figure 11. The cycle length range was limited to where univariate and multivariate detection probability and pattern correlation exceeded predefined feasibility thresholds ($P_{\text{detect}} > 0.8$; $R > 0.7071$).

	Univariate detection				Multivariate detection and decoding				Reconstruction				
	very regular	moderately irreg.	irregular	outside of simulated range	very regular	moderately irreg.	irregular	outside of simulated range	very regular	moderately irreg.	very regular	moderately irreg.	irregular
min. cycle length	2.6 mm + 0.2 mm	2.6 mm + 0.5 mm	2.6 mm + 0.7 mm	2.6 mm + 0.7 mm	3 mm + 1.2 mm	2.6 mm + 1.0 mm	1.8 mm + 1.2 mm	1.8 mm + 1.2 mm	3.4 mm + 0.3 mm	3.6 mm + 0.5 mm	3.4 mm + 0.3 mm	3.6 mm + 0.5 mm	4.2 mm + 1.1 mm
optimal voxel width	0.4 × cycle length	0.3 × cycle length	0.2 × cycle length	0.2 × cycle length	0.1 × cycle length	0.1 × cycle length	0.1 × cycle length	0.1 × cycle length	0.4 × cycle length	0.2 × cycle length	0.4 × cycle length	0.2 × cycle length	0.1 × cycle length
3T													
min. cycle length	2.6 mm + 0.2 mm	2.6 mm + 0.5 mm	2.6 mm + 0.7 mm	2.6 mm + 0.7 mm	1.2 mm + 0.2 mm	1.2 mm + 0.4 mm	0.6 mm + 0.8 mm	0.6 mm + 0.8 mm	1.4 mm + 0.1 mm	1.6 mm + 0.2 mm	1.4 mm + 0.1 mm	1.6 mm + 0.2 mm	1.8 mm + 0.3 mm
optimal voxel width	0.4 × cycle length	0.3 × cycle length	0.2 × cycle length	0.2 × cycle length	0.3 × cycle length	0.3 × cycle length	0.2 × cycle length	0.2 × cycle length	0.4 × cycle length	0.3 × cycle length	0.4 × cycle length	0.3 × cycle length	0.2 × cycle length
7T													
min. cycle length	2.8 mm + 0.2 mm	2.8 mm + 0.4 mm	3.0 mm + 0.6 mm	3.0 mm + 0.6 mm	1.2 mm + 0.2 mm	1 mm + 0.3 mm	0.8 mm + 0.7 mm	0.8 mm + 0.7 mm	1.4 mm + 0.1 mm	1.4 mm + 0.2 mm	1.4 mm + 0.1 mm	1.4 mm + 0.2 mm	1.6 mm + 0.3 mm
optimal voxel width	0.4 × cycle length	0.3 × cycle length	0.2 × cycle length	0.2 × cycle length	0.3 × cycle length	0.3 × cycle length	0.2 × cycle length	0.2 × cycle length	0.4 × cycle length	0.3 × cycle length	0.4 × cycle length	0.3 × cycle length	0.2 × cycle length
7T													
SE													

The dependence of optimal voxel widths using 3T GE, 7T GE and 7T SE on the pattern of cortical columns

Lastly, we generated a number of scanner and pulse sequence dependent simulations. We estimated optimal voxel widths, univariate and multivariate detection probabilities and pattern correlations that resulted from these voxel widths for each combination of pattern parameters (cycle length and irregularity) under three different scenarios: 3T GE, 7T GE and 7T SE imaging. The imaging scenarios determined BOLD point-spread function width, BOLD amplitude and noise characteristics.

We present results for three different levels of irregularity (Figure 11). Optimal voxel widths as a function of cycle length appeared approximately linear for cycle lengths longer than that associated with the defined detection thresholds (Figure 11; first column).

We fitted linear functions to these optimal voxel curves. The estimated coefficients made it possible to predict optimal voxel widths as the sum of a minimum voxel width and a cycle length dependent part (Table 2).

Discussion

Validity of our model and results

Model of cortical columns

We modeled columnar patterns as band-pass filtered spatial white noise. This results in spatial frequency spectra that are directly determined by the band-pass filter shape (Appendix B). The resulting frequency spectra consist of a main spatial frequency and a range of normally distributed frequency components that contribute to the irregularity of the pattern. This model is based on the assumption that on average, columns have a specific size and distance between each other, and that their structure shows some degree of random variation. (Rojer and Schwartz, 1990) demonstrated that such a model accounts for realistic cortical columns by comparing modeled patterns and their power spectra to patterns of orientation columns in monkey V1 (see also Obermayer and Blasdel, 1993).

In theory, patterns of columns may have a multimodal frequency distribution. Multimodal distributions will be associated with columnar structures of multiple scales. However additional components with spatial frequency significantly different than that of the main frequency of the organization may not be relevant to the local characterization of cortical columns. For example, the radial bias (overrepresentation or enhanced responses to an oriented grating in cortical sites that are retinotopically radial to the center of the visual field along that orientation; Sasaki et al., 2006) or the oblique effect (overrepresentation of cardinal orientations; Furmanski and Engel, 2000) manifest themselves as an additional low spatial frequency component of the columnar organization. Although various frequency distributions are conceivable as models for columnar patterns, a normal distribution similar to the one used in our model appears to

be a good approximation to the distribution around the main frequency of the organization (Yao et al., 2016).

We have shown that contrast range is completely determined by the voxel width dependent integral of the spatial power spectrum of the BOLD response pattern (Figure 3; Appendix A). Furthermore, the BOLD response, modeled as a convolution of the pattern of neuronal responses with a BOLD PSF in the image space, is equivalent to a spatial frequency dependent multiplication in the spatial frequency space. As a result, if we consider the features of the neuronal columnar organization, contrast range depends only on the spatial power spectrum of the columnar pattern, not on its phase.

Detection and decoding probabilities depend on contrast range and the number of voxels (Figures 4-5; Appendix D). Their derivation assumes that MR measured differential responses follow a normal distribution. In our model, the distribution of neuronal responses follows a normal distribution, because the pattern of neuronal response is a convolution of Gaussian white noise with a linear filter. The distribution of BOLD responses follows a normal distribution too, because it is modeled as a convolution (a linear transformation) of the neuronal response with a BOLD PSF.

These conditions are likely to apply not only in our model, but also in measured responses. First, differential neuronal responses follow an approximate symmetric distribution, with a mode near zero, and density that commonly decreases with increasing distance from zero. Then in practice, the convolution with a BOLD PSF is local averaging over space, which according to the central limit theorem, is expected to result in BOLD responses that by default approach a normal distribution even more than the neuronal responses do.

All this implies that very different columnar patterns will result in the same detection and decoding probabilities as long as their spatial power spectra are identical. This makes our detection and decoding results depend only on the spatial power spectrum of columnar patterns independent of the specific model that simulated them or whether they were obtained empirically.

2D functional MRI of a flat cortical region

Our model assumes that fMRI data were acquired from a slice that overlapped tangentially with a flat region of cortical gray matter whose thickness was similar to the thickness of the visual cortex (2.5 mm). This approach has been successfully implemented in several studies of columnar imaging from V1 (Cheng et al., 2001; Goodyear and Menon, 2001; Shmuel et al., 2010; Yacoub et al., 2008; 2007). Still, it poses limitations on the area of interest and on the cortical anatomy of subjects.

A more general approach would consider curved cortex as well as arbitrary positions and orientations of voxels relative to the cortex. As a result, partial volume effects could cause voxels to partially overlap with tissue outside of the gray matter, including the white matter and cerebrospinal fluid. This would lead to reductions in the response amplitude and proportional reductions in the CNR. Such effects can be minimized when voxels are sufficiently small so that one can select voxels whose volume overlaps largely with gray matter.

In addition, arbitrarily oriented voxels influence the alignment of the voxel integration functions relative to the cortical manifold, affecting the sampling of the columnar pattern. In three dimensions, the orientation of individual voxels can vary parallel to the plane of the cortical manifold (i.e. rotations around the normal vector of the cortical manifold). Such variations are likely to have no substantial effect when considering isotropic columnar patterns. Additional variations of orientation relative to arbitrary axes may affect local sampling, but there is no reason to assume that such variations would systematically result in either increases or decreases of spatial sampling frequencies relative to the columnar pattern.

Lastly, a curved cortex would cause differences in column-to-column distances along the pial surface of the cortex compared to distances along the gray matter/white matter boundary. In general, a curved cortex can be viewed as local compressions and dilations of columnar map representations within the gray matter. Assuming that compressions are as likely as dilations, we can expect the resulting local changes of columnar cycle length to cancel out on average, but to also introduce a certain blurring in voxels that integrate over layers with different degrees of dilation and compression.

In summary, we believe that our results are approximate estimates for the scenario of arbitrary voxel orientations and sampling of curved cortex. It is however possible, that under these circumstances, detection, decoding performance, and reconstruction quality are somewhat reduced, and that optimal voxels will be somewhat smaller in order to reduce effects related to partial volume and curved cortex.

Modeling the BOLD response as a convolution with a Gaussian independent of cortical site

We assumed that the BOLD response is linear with respect to the underlying neuronal activity and that its spatial profile can be modeled as a convolution of the neuronal response with a Gaussian point-spread function (PSF) which is invariant to the cortical site. This approach is supported by several empirical results. Boynton et al. (1996) showed that the temporal BOLD response to temporally long stimuli is linear in time. Logothetis et al. (2001) confirmed this linearity in time, and demonstrated a linear relationship between neuronal responses and the BOLD response. Hansen et al. (2004) reported spatial linearity by using a set of different visual stimuli.

It has been proposed (Kriegeskorte et al., 2010) and demonstrated (Polimeni et al., 2010) that the BOLD response depends on the cortical site, suggesting that it is more complex than a convolution with a single prototypical Gaussian. The spatial extent and the magnitude of the response may vary due to local variations in vascular geometry. As a consequence, a convolutional model with a single Gaussian function could only be an approximating simplification. However, for patterns of cortical columns with a relatively high spatial frequency, such as ocular dominance and orientation columns in V1, and columns for axis of motion in area MT, the only currently feasible approach is that of differential imaging. Differential imaging substantially reduces blood vessel responses, making the BOLD response more homogeneous parallel to the cor-

tical manifold. Indeed, our own results (Chaimow et al., 2016) suggest that differential fMRI of cortical columns can be modeled as a convolution of neuronal activity with a Gaussian invariant to the cortical site. Therefore, modeling the spatial BOLD response as a Gaussian PSF is a useful approximation for comparing fMRI contrasts, quantitative modeling, interpretation, and planning of high-resolution fMRI studies.

We expect random variations in local vascular geometry to be more substantial for veins and venules rather than capillaries, because of their respective densities. Consequently, responses measured with GE fMRI at ultra-high field or at lower magnetic field strengths are more sensitive to responses from larger pial veins, and will be more affected by these local variations. As a result, GE BOLD imaging at 7T does not only suffer from a slightly wider point-spread function than SE BOLD at the same field strength, but it is also subject to responses from draining veins if these are present.

Here, we did not model the influence of larger blood vessels. We have previously demonstrated that they can carry informative functional contrast (Shmuel et al., 2010). Therefore, when the region of interest contains voxels affected by larger blood vessels, CNR and detection of differential responses using GE at 7T does not necessarily need to be compromised. However, the measured BOLD pattern may be substantially distorted relative to the underlying neuronal organization.

Choice of point-spread function parameters and applicability to non-BOLD fMRI

Engel et al. (1997) estimated the BOLD point spread width as 3.5 mm (FWHM) at 1.5T. (Parkes et al., 2005) reported a point spread of 3.9 mm for GE fMRI at 3T. Using GE imaging at 7T, (Shmuel et al., 2007) measured 2.34 mm wide PSF, and due to confounding effects, estimated that the PSF is smaller than 2 mm. These previous estimates did not account for the spread of the neuronal response due to the receptive field size and scatter in area V1. Therefore, they estimated the capacity of the BOLD response to resolve retinotopic representations, but not the capacity to resolve responses of cortical columns. To address the latter, we have estimated point-spread function widths by fitting a model of imaging columns to ocular dominance responses from V1 (Chaimow et al., 2016). The FWHM for GE and SE fMRI at 7T is 1.02 mm and 0.82 mm, respectively. Based on these results we also derived an estimate for the FWHM of the PSF of GE BOLD fMRI at 3T GE (2.8 mm).

Response amplitudes can vary between subjects, stimuli and imaged area. According to a review by Hennig et al. (2003), BOLD response amplitudes depend only moderately on magnetic field strength. Yacoub et al. (2001) reported average response to a visual stimulus of 2.86% at 4T (GE) and 3.73% at 7T (GE). The ratio of GE and SE BOLD fMRI amplitude obtained at 7T from the gray matter is 1.29 (Yacoub et al., 2005). We recently estimated the median response from ocular dominance column responses averaged across left and right eye stimulation as 3.1% for GE and 2.0% for SE (7T; average from two subjects; Chaimow et al., 2016). According to our model, this averaged response is expected to be one half of the maximum response amplitude, defined as the maximally selective response for one stimulus condition. Guided by

these findings, we set the maximal response amplitudes of our model to 5% for 3T GE, 6% for 7T GE and 4% for 7T SE.

Our choice of parameters is motivated by the current common use of high-field BOLD fMRI for imaging columns. Nonetheless, our model can be applied to any fMRI contrast whose spatial responses can be approximated by a point-spread function of known shape and width.

Measures for evaluating the quality of imaging a pattern of cortical columns

In order to quantify the quality of imaging a pattern of cortical columns, we implemented four different measures: the probability of univariate detection of a response, the probability of multivariate detection of a response, the probability of statistically significant decoding of a stimulus condition from a response pattern and the correlation between the true columnar pattern and the measured response pattern. These four measures can be grouped into two classes, according to the possible objectives which one might have when imaging and analyzing a columnar structure.

The first three measures (univariate detection probability, multivariate detection probability and probability of statistically significant multivariate decoding) evaluate whether the features of the columnar pattern and the data-acquisition parameters make it possible to detect differences between responses to two different stimuli that represent a specific dimension (such as ocular dominance or preferred orientation). These responses can either be considered voxel by voxel (univariate detection) or as multi-voxel patterns (multivariate detection and multivariate decoding).

Historically, single voxel detection by means of the univariate approach has been the dominant method of fMRI analysis. An early study (Menon et al., 1997) identified single voxels within human V1 that responded with a higher amplitude to stimulating one eye relative to the other. Based on previous studies in primates (e.g. Ts'o et al., 1990), this feature should be represented in ocular dominance columns. Therefore, stimulus specific responses of a columnar organization are not necessarily confined to single voxels. Instead, we can expect a distribution of single-voxel differential responses: some may respond positively, some negatively and others may not respond at all.

Multivariate techniques consider the entire pattern as a single high-dimensional response, combining all available information. Furthermore, their multivariate response estimates can benefit from spatially uncorrelated noise components. For these reasons, multivariate techniques outperform univariate analysis and should be the method of choice for detecting differences between responses to stimuli organized in a columnar pattern.

Furthermore, in the multivariate case, detection is not necessarily part of an explicit statistical test. It is more common to train a classification algorithm on a subset of the data in order to “decode” the stimulus condition from the rest of the data. However, successful decoding is evidence that the data shows condition specific responses.

Note that in both cases, detection and decoding, the objective is to test whether the responses carry information on the studied dimension of the stim-

ulus (e.g., stimulated eye, orientation etc.). In other words, the objective is to test whether there is a statistical dependence between the specific stimuli and the corresponding responses (Kriegeskorte and Bandettini, 2007). Such statistical dependence can be evaluated in two directions, e.g. different stimuli cause different responses (= classical statistical analysis, e.g. GLM) and different responses are associated with different stimuli (= ‘decoding’) (Kriegeskorte and Bandettini, 2007). Indeed, in our simulations, the probability to detect a multivariate response and the probability to achieve a statistically significant classification/decoding performance were virtually identical.

The forth measure (pattern correlation) evaluates to which extent the features of the columnar pattern and the data-acquisition parameters make it possible to accurately reconstruct the spatial structure of the columnar pattern. Under this objective, i.e. to reconstruct the pattern of cortical columns, it is not sufficient to show that information is encoded in a set of voxels; instead, the spatial structure of this information (the pattern of cortical columns) is of interest. Examples of this approach include successful imaging of ocular dominance columns in human V1 (Cheng et al., 2001; Goodyear and Menon, 2001; Yacoub et al., 2007), maps of temporal frequency in V1 (Sun et al., 2007), orientation columns in human V1 (Yacoub et al., 2008) and columns in area MT showing preference to an axis of direction of motion in MT (Zimmermann et al., 2011). Even when CNR is high, the reconstructed pattern may possibly not resemble the true neuronal organization. This can be the case either because of voxels that are too large (thus reducing contrast but also reducing noise, and too large to capture the pattern’s frequencies), wide fMRI point-spread function or additional distortions of the reconstructed pattern due to differential responses of large blood vessels.

Thus far, we used the terms ‘*objective*’ to define the goal of a study, and ‘*measure*’ to describe quantities we use in our current study for evaluating different ‘*approaches*’ to achieve an objective. As an example, we consider detection of a difference between responses to two stimuli as an objective. This objective can be *addressed* by 3 different approaches, e.g. univariate detection, multivariate detection or multivariate decoding. The 3 measures we use to evaluate these 3 different approaches are univariate detection probability, multivariate detection probability and probability of statistically significant decoding.

Making an explicit distinction between objectives (e.g., in our case, (1) detection of information on a specific dimension of the stimuli and (2) reconstruction of the organization) and between associated methodological approaches (e.g. univariate vs. multivariate analysis) made it possible to define appropriate measures that reflect the quality of imaging a pattern of columns. This in turn allowed us to estimate optimal imaging parameters for each of the approaches, and to demonstrate that optimal parameters are not necessarily identical between approaches.

We have shown that in general, reconstruction requires smaller voxel size than detection. This is especially true for irregular patterns, where detection can benefit from low-frequency components using relatively large voxels with high SNR, whereas accurate reconstruction requires high spatial-frequency

content to be captured, which can only be achieved with smaller voxels. In addition, optimal voxel size for reconstruction decreases further with increasing SNR. Optimal voxel size for detection does not decrease if this SNR increase is largely voxel size independent (e.g. due to increased data volume or technological improvements).

Comparison to previous results

A number of studies estimated optimal voxel sizes for detection of responses of non-columnar organizations. Bodurka et al. (2007) proposed an optimal voxel width of 1.8 mm (isotropic) at 3T, based on equal contributions of physiological and thermal noise. They argued that further increases in voxel size led to diminished returns in SNR while at the same time increasing the chance of partial volume effects.

Hyde et al. (2001) showed that 1.5 mm wide isotropic voxels were optimal for detecting activation in somatosensory cortex evoked by finger tapping at 3T. They also explain this as the optimal tradeoff between noise and partial volume effects, based on the fact that this voxel size matches the activation volume, which they equate with the thickness of layers 3 and 4.

Similarly, Glover and Krueger (2002) developed a theoretical model of imaging a Gaussian-shaped activated region. They then showed empirically using a task of bilateral finger apposition at 1.5 T that CNR as a function of voxel size agreed with their model and that voxels of 4 x 4 mm and a slice thickness that was left constant at 5 mm were optimal.

Our study differs from those mentioned above in that we consider fine scale columnar patterns. We have shown that optimal voxel widths for detection, decoding, and reconstruction depend on the spatial properties (scale and irregularity) of the pattern and on the fMRI point-spread function. Partial volume effects are minimized when imaging a slice that overlaps tangentially with a flat region of cortical gray matter and whose thickness does not exceed the thickness of cortex. In this situation, for imaging sufficiently coarse patterns, even small improvements in SNR by using larger voxel widths can be beneficial for detection, decoding, and reconstruction. Furthermore, voxels much smaller than 1.5 mm and 1.8 mm may be optimal for imaging fine-grained patterns using fMRI contrast with relatively narrow PSF.

In Chaimow et al. (2011), we considered performance of a multivariate pattern analysis algorithm on simulated data from ocular dominance columns acquired at 3T. In that study we chose a fixed number of voxels rather than a fixed volume, and we did not model noise correlations (since we were interested in modeling a best-case scenario for decoding based on cortical columns). Furthermore we used a point-spread function that was confounded by receptive-field effects. For these reasons voxel sizes that optimized decoding were larger than those we report here.

Implications for fMRI of cortical columns

Which voxel size should we use?

In Table 2 we provide practical recommendations for optimal voxel sizes according to our model, under a number of scenarios. In general, if reconstruc-

tion is the goal, voxel size needs to be at least smaller than half of the columnar cycle length, and even smaller when taking irregularities into account.

If, however, the goal is to detect information present in the columnar pattern, the question will be whether the main pattern components are expected to survive the BOLD point-spread. If the answer is positive, the columnar cycle length determines optimal voxel size. However, if the point spread is too wide, not all is lost. Depending on the level of irregularities, sets of voxels can be used to detect low spatial frequency information about the stimulus with optimal widths that exceed the typical widths of columns.

The advantage of high-field fMRI

When only considering increases in response amplitude and SNR, the effects of high field are rather moderate. The differences in noise at 7T compared to 3T led to CNR improvements of 35% while the larger response amplitude of 7T GE contributed a 20% increase. Taken together, this makes for an increase of at most 62% in CNR. Although these are significant increases, similar results could be achieved by acquiring 2-3 times larger amount of data. However, for imaging cortical columns we found that the most important difference between 3T and 7T is the narrower point-spread function at 7T, even more so when using SE imaging. This factor caused a massive difference in detection probability and pattern reconstruction.

Only the narrow point spreads at 7T make detection and reconstruction of a columnar pattern with the scale of orientation columns possible. Yet, GE at 7T may be less suited for reconstruction than SE, given that it is prone to distortions of the pattern of cortical columns due to residual differential responses of large pial blood vessels.

These results suggest that developments that focus on improving the spatial specificity of the functional contrast would be the most promising strategy for improving fMRI of cortical columns.

Imaging of unknown patterns

Let us assume a semi-regular pattern of columns with a spatial frequency distribution that can be approximated by a Gaussian. However, we have no a priori knowledge about the main cycle length or degree of irregularities. How can we determine whether the imaged pattern resembles the underlying pattern of columns and whether the point-spread function width of our imaging technique and our employed voxel size were sufficiently small to resolve this pattern?

First, our results have illustrated that reconstruction is generally more difficult than detection. Therefore, the ability to detect a stimulus specific multivariate response should be a prerequisite for accurate reconstruction.

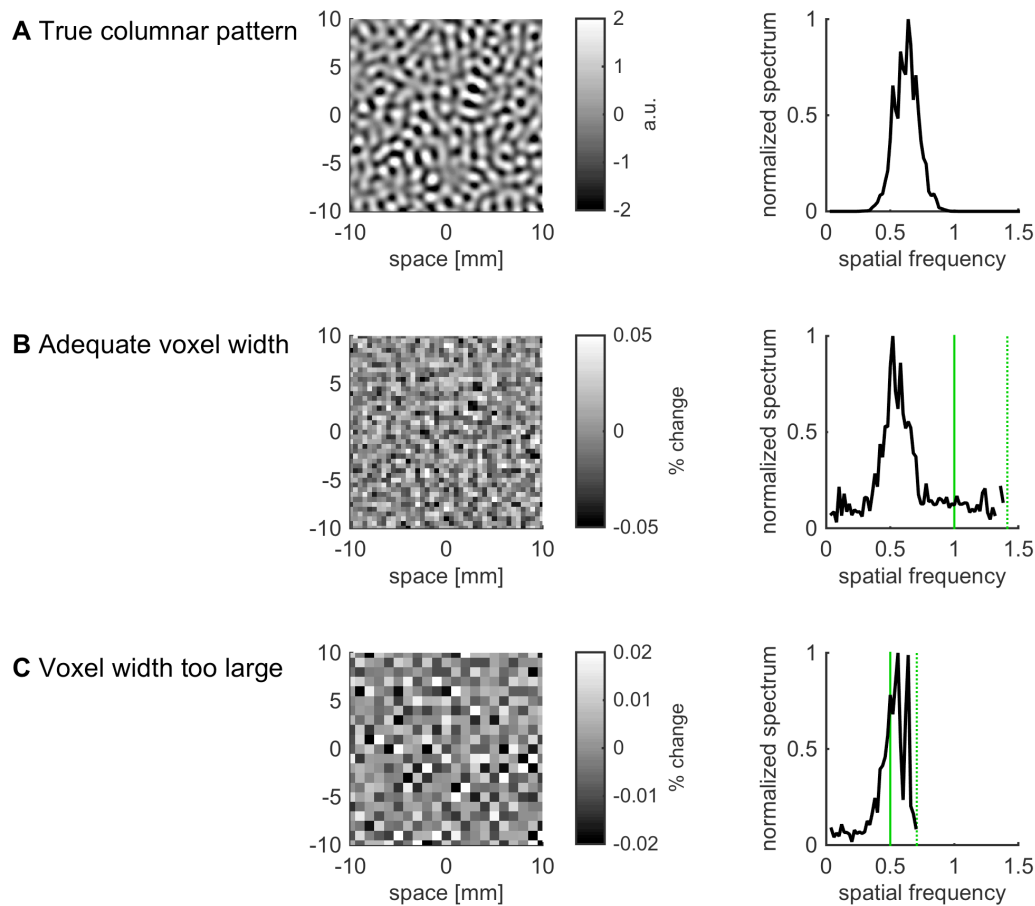


Fig. 12 Imaging of unknown patterns. A columnar pattern and the corresponding fMRI maps are shown for adequate voxel width and voxel width that is too large. The first row shows the true columnar pattern and its spatial frequency spectrum (main pattern cycle length = 1.6 mm, moderate irregularity). The second row shows the results of fMRI (7T GE, including noise) when voxel width is adequate. The spatial frequency spectrum averaged over all orientations shows a clear peak and returns to low values well below the highest frequency that the voxels sample (vertical green lines; the dotted line represent highest spatial frequency along k-space diagonal). When voxel width is too large (third row), the spatial frequency spectrum is being cut off abruptly.

The next question then is whether the voxel is small enough to resolve the BOLD response pattern. Here the spectrum of the response can help (Figure 12). A well resolved pattern will result in a gradual decrease of spatial frequency components higher than the main pattern frequency towards the highest frequency that the voxel can sample (Figure 12B). If the voxel size is too small, we can expect an abrupt cutoff close to the highest sampled frequencies with no significant decreases when approaching this point (Figure 12C). However in practice, high measurement noise may make it difficult to differentiate between spectra resulting from adequate voxel widths and those resulting from voxels that are too wide.

Finally, how can we tell if the point-spread function width is too wide, and does not allow resolving the pattern? Unfortunately, this is impossible to tell from the spatial structure of the imaged pattern alone, as the BOLD response results in similar types of spectra compared to the neuronal response, only with different apparent cycle lengths and degrees of irregularity. Another factor to consider is that wide point-spreads result in low contrast. In order to

quantify whether the contrast is too low given the spatial pattern, quantitative modeling may be necessary. In our study (Chaimow et al., 2016) we used a probabilistic extension of the present model together with Markov Chain Monte Carlo sampling in order to estimate the posterior distribution of point-spread widths given prior knowledge about parameters of ocular dominance columns. A similar approach can be taken here, using prior information about the point-spread width in order to infer the unknown pattern parameters.

Table 3 In a review on pattern information fMRI, Formisano and Kriegeskorte (Formisano and Kriegeskorte, 2012) show in a schematic way (their Figure 2) how pattern information may depend on voxel size and field strength. They pose a number of “...important questions [that] remain unanswered”. This table provides answers to these questions within the context of our modeling results.

Which spatial resolution maximizes pattern information at 3T?	It depends on the structure of the pattern and on the available volume. For a moderately irregular isotropic pattern with a cycle length of 1.6 mm and a volume of 24 mm x 24 mm x 2.5 mm we estimate the optimal voxel width as 1.2 mm.
Which spatial resolution maximizes pattern information at 7T?	For a pattern similar to that described for 3T, the optimal voxel size for gradient echo at 7T is 0.78 mm; for spin echo at 7T it is 0.77 mm.
Is 7T better than 3T for revealing pattern information?	Yes it is. This is mainly due to the narrower point spread function of the BOLD response at 7T.
Does the optimal field strength for pattern information studies depend on the resolution (e.g. does 3T yield greater pattern information than 7T at lower resolutions)?	There is no effect that would result in 3T having an advantage over 7T. However, when comparing 7T GE to 7T SE, we see that the gradient echo gains a slight advantage at voxel sizes larger than 2-3 mm, due to a higher response amplitude and a lesser relevance of the point spread of the BOLD response.

Conclusion

The success of imaging cortical columns depends on a number of factors and their interrelationships. To quantify these dependencies, here we have combined current knowledge into an integrated, quantitative model of imaging cortical columns. Our model makes it possible to evaluate how detection and decoding of a stimulus-specific response and the reconstruction of a pattern of cortical columns depend on the parameters of the pattern, voxel size, fMRI point-spread, and noise characteristics.

Our model addresses open questions that are of interest to the neuroimaging community (Formisano and Kriegeskorte, 2012; see our Table 3). While our findings need to be tested empirically, they can already guide the modeling, planning, and interpretation of imaging cortical columns.

References

- Adams, D.L., Sincich, L.C., Horton, J.C., 2007. Complete pattern of ocular dominance columns in human primary visual cortex. *J Neurosci* 27, 10391–10403. doi:10.1523/JNEUROSCI.2923-07.2007
- Bodurka, J., Ye, F.Q., Petridou, N., Murphy, K., Bandettini, P.A., 2007. Mapping the MRI voxel volume in which thermal noise matches physiological noise--implications for fMRI. *Neuroimage* 34, 542–549. doi:10.1016/j.neuroimage.2006.09.039
- Bonhoeffer, T., Grinvald, A., 1991. Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* 353, 429–431. doi:10.1038/353429a0
- Boynton, G.M., Engel, S., Glover, G., Heeger, D.J., 1996. Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* 16, 4207–4221.
- Chaimow, D., Yacoub, E., Uğurbil, K., Shmuel, A., 2016. Spatial specificity of the functional MRI blood oxygenation response relative to metabolic activity, in: Presented at the 22nd Annual Meeting of the Organization for Human Brain Mapping, Geneva.
- Chaimow, D., Yacoub, E., Uğurbil, K., Shmuel, A., 2011. Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns. *Neuroimage* 56, 627–642. doi:10.1016/j.neuroimage.2010.09.037
- Cheng, K., Waggoner, R.A., Tanaka, K., 2001. Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron* 32, 359–374.
- Engel, S.A., Glover, G.H., Wandell, B.A., 1997. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7, 181–192. doi:10.1093/cercor/7.2.181
- Erwin, E., Obermayer, K., Schulten, K., 1995. Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. *Neural Comp* 7, 425–468.
- Formisano, E., Kriegeskorte, N., 2012. Seeing patterns through the hemodynamic veil--the future of pattern-information fMRI. *Neuroimage* 62, 1249–1256. doi:10.1016/j.neuroimage.2012.02.078
- Furmanski, C.S., Engel, S.A., 2000. An oblique effect in human primary visual cortex. *Nat Neurosci* 3, 535–536. doi:10.1038/75702
- Glover, G.H., Krüger, G., 2002. Optimum voxel size in BOLD fMRI. *Proc. Intl. Soc. Mag. Reson. Med.*
- Goodyear, B.G., Menon, R.S., 2001. Brief visual stimulation allows mapping of ocular dominance in visual cortex using fMRI. *Hum Brain Mapp* 14, 210–217.
- Haacke, M.E., Brown, R.W., Thompson, M.R., 1999. *Magnetic resonance imaging: physical principles and sequence design.*
- Hansen, K.A., David, S.V., Gallant, J.L., 2004. Parametric reverse correlation reveals spatial linearity of retinotopic human V1 BOLD response. *Neuroimage* 23, 233–241.
- Hennig, J., Speck, O., Koch, M.A., Weiller, C., 2003. Functional magnetic resonance imaging: a review of methodological aspects and clinical applications. *J Magn Reson Imaging* 18, 1–15. doi:10.1002/jmri.10330
- Holla, M.S., 1970. On a Noncentral Chi-Square Distribution in the Analysis of Weapon Systems Effectiveness. *Metrika.*
- Horton, J.C., Dagi, L., McCrane, E., de Monasterio, F., 1990. Arrangement of ocular dominance columns in human visual cortex. *Arch Ophthalmol* 108, 1025–1031.

- Hyde, J.S., Biswal, B.B., Jesmanowicz, A., 2001. High-resolution fMRI using multislice partial k-space GR-EPI with cubic voxels. *Magn Reson Med* 46, 114–125.
- Kriegeskorte, N., Bandettini, P.A., 2007. Analyzing for information, not activation, to exploit high-resolution fMRI. *Neuroimage* 38, 649–662.
- 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. doi:10.1016/j.neuroimage.2009.09.059
- Krueger, G., Glover, G., 2001. Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magn Reson Med* 46, 631–637.
- Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412, 150–157. doi:10.1038/35084005
- Menon, R.S., Ogawa, S., Strupp, J., Uğurbil, K., 1997. Ocular dominance in human V1 demonstrated by functional magnetic resonance imaging. *J Neurophysiol* 77, 2780–2787.
- Nasr, S., Polimeni, J.R., Tootell, R.B.H., 2016. Interdigitated Color- and Disparity-Selective Columns within Human Visual Cortical Areas V2 and V3. *J Neurosci* 36, 1841–1857. doi:10.1523/JNEUROSCI.3518-15.2016
- Niebur, E., Wörgötter, F., 1994. Design principles of columnar organization in visual cortex. *Neural Comp* 6, 602–614. doi:10.1162/neco.1994.6.4.602
- Obermayer, K., Blasdel, G.G., 1993. Geometry of orientation and ocular dominance columns in monkey striate cortex. *J Neurosci* 13, 4114–4129.
- Parkes, L.M., Schwarzbach, J.V., Bouts, A.A., Deckers, R.H.R., Pullens, P., Kerskens, C.M., Norris, D.G., 2005. Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. *Magn Reson Med* 54, 1465–1472. doi:10.1002/mrm.20712
- Polimeni, J.R., Fischl, B., Greve, D.N., Wald, L.L., 2010. Laminar analysis of 7T BOLD using an imposed spatial activation pattern in human V1. *Neuroimage* 52, 1334–1346. doi:10.1016/j.neuroimage.2010.05.005
- Purdon, P.L., Weisskoff, R.M., 1998. Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. *Hum Brain Mapp* 6, 239–249.
- Roberts, M.J., 2012. *Signals and Systems*. McGraw-Hill Higher Education.
- Roger, A., Schwartz, E., 1990. Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biol Cybern* 62, 381–391.
- Sasaki, Y., Rajimehr, R., Kim, B.W., Ekstrom, L.B., Vanduffel, W., Tootell, R.B.H., 2006. The radial bias: a different slant on visual orientation sensitivity in human and nonhuman primates. *Neuron* 51, 661–670. doi:10.1016/j.neuron.2006.07.021
- Shmuel, A., Chaimow, D., Raddatz, G., Uğurbil, 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. doi:10.1016/j.neuroimage.2009.08.040
- Shmuel, A., Grinvald, A., 2000. Coexistence of linear zones and pinwheels within orientation maps in cat visual cortex. *Proceedings of the National Academy of Sciences* 97, 5568–5573. doi:10.1073/pnas.97.10.5568
- Shmuel, A., Grinvald, A., 1996. Functional organization for direction of motion and its relationship to orientation maps in cat area 18. *J Neurosci* 16, 6945–6964.
- Shmuel, A., Yacoub, E., Chaimow, D., Logothetis, N.K., Uğurbil, K., 2007. Spatio-temporal point-spread function of fMRI signal in human gray matter at 7 Tesla. *Neuroimage* 35, 539–552. doi:10.1016/j.neuroimage.2006.12.030
- Sun, P., Ueno, K., Waggoner, R.A., Gardner, J.L., Tanaka, K., Cheng, K., 2007. A tem-

- poral frequency-dependent functional architecture in human V1 revealed by high-resolution fMRI. *Nat Neurosci* 10, 1404–1406.
- Triantafyllou, C., Hoge, R.D., Krueger, G., Wiggins, C.J., Potthast, A., Wiggins, G.C., L, W.L., 2005. Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage* 26, 243–250. doi:10.1016/j.neuroimage.2005.01.007
- Ts'o, D.Y., Frostig, R.D., Lieke, E.E., Grinvald, A., 1990. Functional organization of primate visual cortex revealed by high resolution optical imaging. *Science* 249, 417–420. doi:10.1126/science.2165630
- Uludağ, K., Müller-Bierl, B., Uğurbil, K., 2009. An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. *Neuroimage* 48, 150–165. doi:10.1016/j.neuroimage.2009.05.051
- Wright, P.J., Mougín, O.E., Totman, J.J., Peters, A.M., Brookes, M.J., Coxon, R., Morris, P.E., Clemence, M., Francis, S.T., Bowtell, R.W., Gowland, P.A., 2008. Water proton T1 measurements in brain tissue at 7, 3, and 1.5 T using IR-EPI, IR-TSE, and MPRAGE: results and optimization. *MAGMA* 21, 121–130. doi:10.1007/s10334-008-0104-8
- Yacoub, E., Harel, N., Uğurbil, K., 2008. High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences* 105, 10607–10612. doi:10.1073/pnas.0804110105
- Yacoub, E., Shmuel, A., Logothetis, N.K., Uğurbil, 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., Shmuel, A., Pfeuffer, J., van de Moortele, P.F., Adriany, G., Andersen, P., Vaughan, J.T., Merkle, H., Uğurbil, K., Hu, X., 2001. Imaging brain function in humans at 7 Tesla. *Magn Reson Med* 45, 588–594. doi:10.1002/mrm.1080
- Yacoub, E., van de Moortele, P.F., Shmuel, A., Uğurbil, K., 2005. Signal and noise characteristics of Hahn SE and GE BOLD fMRI at 7 T in humans. *Neuroimage* 24, 738–750.
- Yao, Z., Fukuda, M., Moon, C.-H., Kim, S.-G.G., Shmuel, A., 2016. CBV fMRI with no large vessel signals rules out decoding of orientation through local irregularities, in: Presented at the 22nd Annual Meeting of the Organization for Human Brain Mapping, Geneva.
- Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P.-F., Feinberg, D., Adriany, G., Chaimow, D., Shmuel, A., Uğurbil, 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. doi:10.1371/journal.pone.0028716

Appendix A. Contrast range can be computed from the columnar pattern power spectrum

Let $(y_{m,n})$ be a $2N_1 \times 2N_2$ sized array of MR measured differential responses relative to baseline representing an area of size $L_1 \times L_2$ and let $b(x_1, x_2)$ be the underlying true response pattern which is only non-zero for $(x_1, x_2) \in [-\frac{L_1}{2}, \frac{L_1}{2}] \times [-\frac{L_2}{2}, \frac{L_2}{2}]$. The MR measurement can be modeled (Chaimow et al. 2011, Haacke et al. 1999) as

$$\begin{aligned} y_{m,n} &= \frac{1}{L_1 L_2} \sum_{p=-N_1}^{N_1-1} \sum_{q=-N_2}^{N_2-1} B(p/L_1, q/L_2) e^{i\pi \frac{mp}{N_1} \frac{nq}{N_2}} \\ &= \sum_{p=1}^{2N_1} \sum_{q=1}^{2N_2} B_{p,q} e^{i\pi \frac{mp}{N_1} \frac{nq}{N_2}} = 2N_1 2N_2 \left(\text{idft2} \left[(B_{p,q}) \right] \right)_{m,n}, \end{aligned} \quad (\text{A.1})$$

where $B(k_1, k_2)$ is the Fourier transform of $b(x_1, x_2)$ and where we have defined $B_{p,q} := \frac{1}{L_1 L_2} B \left(\frac{p-N_1-1}{L_1}, \frac{q-N_2-1}{L_2} \right)$. Contrast range c of the measured responses $(y_{m,n})$ is defined as

$$c = \sqrt{\frac{1}{2N_1 2N_2} \sum_{m=1}^{2N_1} \sum_{n=1}^{2N_2} y_{m,n}^2}, \quad (\text{A.2})$$

assuming that the mean of differential responses is known to be zero.

Inserting A.1 into A.2 and applying Parseval's theorem yields

$$c = \sqrt{\sum_{p=1}^{2N_1} \sum_{q=1}^{2N_2} |(B_{p,q})|^2}.$$

Finally we approximate the sum by an integral with $w = \frac{L_1}{2N_1} = \frac{L_2}{2N_2}$ as the voxel width,

$$c \approx \sqrt{\frac{1}{L_1 L_2} \int_{-\frac{1}{2w}}^{+\frac{1}{2w}} \int_{-\frac{1}{2w}}^{+\frac{1}{2w}} |(B(k_1, k_2))|^2 dk_1 dk_2}. \quad (\text{A.3})$$

When simulating or sampling a differential response pattern we will work with a discretized approximation $(b_{m,n})$ of size $2N_1^{sim} \times 2N_2^{sim}$ such that $b_{m,n} := b \left((m - N_1^{sim} - 1)w_1, (n - N_2^{sim} - 1)w_2 \right)$. The discrete Fourier transform of $(b_{m,n})$ will in turn serve as a model for $B(k_1, k_2)$ with $\text{dft2} \left[(b_{m,n}) \right] \approx \frac{1}{L_1 L_2} B \left(\frac{p-N_1-1}{L_1}, \frac{q-N_2-1}{L_2} \right) = B_{p,q}$ (Roberts, 2012). This allows us to compute contrast range from $(b_{m,n})$ by means of its power spectrum,

$$c \approx \sqrt{\sum_{p=1}^{2N_1} \sum_{q=1}^{2N_2} \left| \left(\text{dft2} \left[(b_{m,n}) \right] \right)_{p,q} \right|^2}. \quad (\text{A.4})$$

Appendix B. Modeling realistic isotropic columnar patterns

Let $(g_{m,n})$ be a $2N_1^{sim} \times 2N_2^{sim}$ sized Gaussian white noise array of independent sample from $\mathcal{N}(0, 1)$. Its discrete Fourier transform is $(G_{p,q})$, with expected values $E \left[|G_{p,q}|^2 \right] = 2N_1^{sim} 2N_2^{sim}$.

An isotropic columnar pattern $(a_{m,n})$ can be modeled by filtering the white noise array $(g_{m,n})$ with an isotropic filter (Rojer and Schwartz, 1990). The filtering can be implemented as a multiplication in spatial frequency space such that

$$(a_{m,n}) = \text{idft2} \left[G_{p,q} \cdot F_{p,q} \right].$$

We can express the filter ($F_{p,q}$) as

$$F_{p,q} = \frac{s(k(p,q))}{C_F},$$

where $s(k)$ is a one-dimensional and direction-independent filter-shape function and $k(p,q) = \sqrt{\frac{(p-N_1-1)^2}{L_1^2} + \frac{(q-N_1)^2}{L_1}}$ is absolute spatial frequency. The normalization constant C_F (see Appendix B.1) ensures that the filter output has an expected variance of one.

In particular we define the filter shape function as

$$s(k) = e^{-\frac{(k-\rho)^2}{2\delta^2}} + e^{-\frac{(k+\rho)^2}{2\delta^2}},$$

where parameters ρ and δ control the main spatial frequency and the degree of irregularity, respectively. The filter shape function defines the distribution of absolute spatial frequencies in the modeled pattern as can be shown by calculating the expectation of $|A_{p,q}|^2$. Using $F_{p,q} \in \mathbb{R}$ and $E[|G_{p,q}|^2] = 2N_1^{sim}2N_2^{sim}$ we get

$$\begin{aligned} E[|A_{p,q}|^2] &= E[|G_{p,q} \cdot F_{p,q}|^2] \\ &= E[|G_{p,q}|^2] \cdot F_{p,q}^2 \\ &= \frac{2N_1^{sim}2N_2^{sim}}{C_F} s(k(p,q))^2. \end{aligned}$$

Appendix B.1. Filter normalization

We derive the normalization constant C_F by requiring the expected variance of the modeled pattern ($a_{m,n}$) to be 1. $a_{m,n}$ are weighted sums of zero-mean random variables and therefore $E[a_{m,n}] = 0$. It follows that

$$1 =: E \left[\frac{1}{2N_1^{sim}2N_2^{sim}} \sum_{m=1}^{2N_1^{sim}} \sum_{n=1}^{2N_2^{sim}} |a_{m,n}|^2 \right].$$

We apply Parseval's theorem:

$$= E \left[\frac{1}{2N_1^{sim}2N_2^{sim}} \sum_{p=1}^{2N_1^{sim}} \sum_{q=1}^{2N_2^{sim}} |A_{p,q}|^2 \right],$$

and use $A_{p,q} = G_{p,q} \cdot F_{p,q}$:

$$= E \left[\frac{1}{2N_1^{sim}2N_2^{sim}} \sum_{p=1}^{2N_1^{sim}} \sum_{q=1}^{2N_2^{sim}} |G_{p,q} \cdot F_{p,q}|^2 \right].$$

Bringing the expectation inside and using $F_{p,q} \in \mathbb{R}$ and $E[|G_{p,q}|^2] = 2N_1^{sim}2N_2^{sim}$ we further get:

$$\begin{aligned}
&= \frac{1}{2N_1^{sim}2N_2^{sim}} \sum_{p=1}^{2N_1^{sim}} \sum_{q=1}^{2N_2^{sim}} E[|G_{p,q}|^2] \cdot F_{p,q}^2 \\
&= \frac{1}{2N_1^{sim}2N_2^{sim}} \sum_{p=1}^{2N_1^{sim}} \sum_{q=1}^{2N_2^{sim}} F_{p,q}^2 \\
&= \frac{1}{2N_1^{sim}2N_2^{sim}} \sum_{p=1}^{2N_1^{sim}} \sum_{q=1}^{2N_2^{sim}} \frac{s(k(p,q))^2}{C_F^2}.
\end{aligned}$$

It follows that the normalization constant C_F is given by

$$C_F = \sqrt{\frac{1}{2N_1^{sim}2N_2^{sim}} \sum_{p=1}^{2N_1^{sim}} \sum_{q=1}^{2N_2^{sim}} s(k(p,q))^2} \quad (\text{B.1})$$

Appendix C. Differential multi-measurement SNR

Time-course signal-to-noise ratio (tSNR) can be written as:

$$tSNR = \frac{S}{\sqrt{\sigma_0^2 + \sigma_p^2}}, \quad (\text{C.1})$$

where S is the baseline signal intensity and σ_0 and σ_p are the standard deviations of thermal and physiological noise components, respectively. These independent noise components can be modeled as $\sigma_0 = \frac{S}{\kappa V}$ and $\sigma_p = S\lambda$, where V is the voxel volume and κ and λ are parameters that can be estimated by fitting C.1 to real tSNR measurements (Trinatafyllou, 2005).

The effect of arbitrary repetition times (TR) can be modeled using a correction factor for κ (Chaimow et al. 2011).

In order to obtain SNR estimates for the difference between two condition averages over N_t measurements each we consider σ_p^2 and σ_0^2 separately. Let ϵ_t^0 and ϵ_t^p be thermal and physiological noise values on individual measurements. Thermal noise is independent between measurements. Therefore, the variance of the difference of averaged thermal noise measurements $\bar{\sigma}_0^2(N_t)$ is

$$\begin{aligned}
\bar{\sigma}_0^2(N_t) &= E \left[\left(\frac{1}{N_t} \sum_{t=1}^{N_t} \epsilon_t^0 - \frac{1}{N_t} \sum_{t=N_t+1}^{2N_t} \epsilon_t^0 \right)^2 \right] \\
&= \frac{1}{N_t^2} \left[\left(\sum_{t_1=1}^{N_t} \sum_{t_2=1}^{N_t} E[\epsilon_{t_1}^0 \epsilon_{t_2}^0] - 2 \sum_{t_1=1}^{N_t} \sum_{t_2=N_t+1}^{2N_t} E[\epsilon_{t_1}^0 \epsilon_{t_2}^0] + \sum_{t_1=N_t+1}^{2N_t} \sum_{t_2=N_t+1}^{2N_t} E[\epsilon_{t_1}^0 \epsilon_{t_2}^0] \right) \right] \\
&= \frac{2\sigma_0^2}{N_t}.
\end{aligned}$$

Physiological noise is correlated in time. We model these correlations by means of an autoregressive model. Purdon and Weiskoff (1998) found that the physiological noise correlations between subsequent measurements q can be modeled as $q = \exp(-\text{TR}/15)$, where TR is the repetition time in seconds. The correlation between measurements at times t_1 and t_2 is then given by

$E[\epsilon_{t_1}^p \epsilon_{t_2}^p] = q^{|t_1-t_2|} = \exp\left(\frac{-\text{TR} \cdot |t_1-t_2|}{15}\right)$. It follows that the variance of the difference of averaged physiological noise measurements $\bar{\sigma}_p^2(N_t)$ is

$$\begin{aligned}\bar{\sigma}_p^2(N_t) &= E \left[\left(\frac{1}{N_t} \sum_{t=1}^{N_t} \epsilon_t^p - \frac{1}{N_t} \sum_{t=N_t+1}^{2N_t} \epsilon_t^p \right)^2 \right] \\ &= \frac{1}{N_t^2} \left[\left(\sum_{t_1=1}^{N_t} \sum_{t_2=1}^{N_t} E[\epsilon_{t_1}^p \epsilon_{t_2}^p] - 2 \sum_{t_1=1}^{N_t} \sum_{t_2=N_t+1}^{2N_t} E[\epsilon_{t_1}^p \epsilon_{t_2}^p] + \sum_{t_1=N_t+1}^{2N_t} \sum_{t_2=N_t+1}^{2N_t} E[\epsilon_{t_1}^p \epsilon_{t_2}^p] \right) \right] \\ &\approx 2 \frac{\sigma_p^2}{N_t^2} \sum_{t_1=1}^{N_t} \sum_{t_2=1}^{N_t} \exp\left(\frac{-\text{TR} \cdot |t_1 - t_2|}{15}\right).\end{aligned}$$

Here we assumed that block size is large enough that temporal correlations between conditions can be neglected.

Taken together the resulting differential multi-measurement SNR is

$$d\text{SNR}(N_t) = \frac{1}{\sqrt{\frac{1}{\kappa^2 V^2 N_t} + \frac{\lambda^2}{N_t^2} \sum_{t_1=1}^{N_t} \sum_{t_2=1}^{N_t} \exp\left(\frac{-\text{TR} \cdot |t_1-t_2|}{15}\right)}}.$$

Appendix D. Detection probability

Appendix D.1. Multivariate detection

Let us consider a set of N voxels and let the measured differential response of voxel with index $m \in 1, \dots, N$ be

$$\hat{y}_m = y_m + \epsilon_m,$$

where y_m is the true differential response without noise and

$$\epsilon_m \sim \mathcal{N}(0, \sigma)$$

is the normally distributed differential multi-measurement noise with standard deviation σ .

We define noise-level-normalized differential responses

$$\hat{u}_m := \frac{\hat{y}_m}{\sigma} = \frac{y_m}{\sigma} + \frac{\epsilon_m}{\sigma}.$$

Given a true differential response y_m , the measured noise-level-normalized differential response \hat{u}_m is distributed according to

$$\hat{u}_m | u_m \sim \mathcal{N}(y_m/\sigma, 1).$$

When trying to detect a multivariate response, we consider (u_m) and equivalently (y_m) to be non-zero if the test statistic $\|(\hat{u}_m)\|^2 = \sum_{m=1}^N \hat{u}_m^2$ is greater or equal to some critical value χ_{crit}^2 .

We choose χ_{crit}^2 , such that the probability to falsely detect a response is smaller than some constant α (e.g. $\alpha = 0.05$). Such false positives occur when $\|(\hat{u}_m)\|^2 \geq \chi_{\text{crit}}^2$ and $u_m = 0$ for all m (null-hypothesis). The latter condition implies that under the null hypothesis, $\|(\hat{u}_m)\|^2$ is χ^2 distributed and therefore

$$\chi_{\text{crit}}^2(\alpha) = F_{\chi_N^2}^{-1}(1 - \alpha/2),$$

where $F_{\chi^2_N}^{-1}$ is the inverse cumulative distribution function of the χ^2 distribution with N degrees of freedom.

The distribution of $\|(\hat{u}_m)\|^2$ given an arbitrary noise-level-normalized true differential response (u_m) is a non-central χ^2 distribution with non-centrality parameter $\lambda = \sum_{m=1}^N u_m^2$ and N degrees of freedom.

Given a certain contrast-to-noise ration $cnr = c/\sigma$, where c is the contrast range, and assuming normally distributed differential responses, individual u_m can be considered random variables distributed according to

$$u_m \sim \mathcal{N}(0, cnr).$$

Therefore, if we express the non-centrality parameter λ divided by the squared contrast-to-noise ratio as:

$$\frac{\lambda}{cnr^2} = \sum_{m=1}^N \left(\frac{u_m}{cnr} \right)^2,$$

we see that λ divided by the squared contrast-to-noise ratio is itself distributed according to a (central) χ^2 distribution with N degrees of freedom, or in other words, λ is distributed according to a scaled χ^2 distribution with N degrees of freedom and scale parameter $1/cnr^2$.

Let $f(x; a, \lambda, N)$ be the probability density function of the scaled non-central χ^2 distribution with scale parameter a , non-centrality parameter λ and N degrees of freedom (setting $a = 1$ results in a non-scaled distribution and setting $\lambda = 0$ results in a central χ^2 distribution). The probability density function of $g(x; cnr, N)$ of the unconditioned distribution $\|(\hat{u}_m)\|^2$ (independent of the specific realization of differential responses) can be obtained by integrating over λ and was derived by Holla (1970):

$$\begin{aligned} g(x; cnr, N) &= \int_0^\infty f(x; 1, \lambda, N) f(\lambda; cnr^{-2}, 0, N) d\lambda \\ &= \frac{x^{N/2-1} e^{-\frac{x}{2(1+cnr^2)}}}{2(1+cnr^2)\Gamma(N/2)} \end{aligned}$$

The resulting distribution is a gamma distribution with shape parameter $k = N/2$ and scale parameter $\theta = 2(1 + cnr^2)$.

The probability of detecting a multivariate response can then be calculated as the probability of $\|(\hat{u}_m)\|^2$ being at least χ_{crit}^2 :

$$\begin{aligned} p_{\text{detect}}^{\text{multi}}(cnr, N) &= P \left[\|(\hat{u}_m)\|^2 \geq \chi_{\text{crit}}^2 \right] \\ &= \int_{\chi_{\text{crit}}^2}^{\infty} g(x; cnr, N) dx \\ &= 1 - G(\chi_{\text{crit}}^2; N/2, 2(1 + cnr^2)). \end{aligned} \tag{D.1}$$

where $G(\chi_{\text{crit}}^2; N/2, 2(1 + cnr^2))$ is the Gamma cumulative distribution function with shape parameter $N/2$ and scale parameter $2(1 + cnr^2)$.

Appendix D.2. Univariate detection

Univariate detection probability is a special case of multivariate detection probability. However, in this study we used two different probabilities that need to be distinguished.

The first probability is the detection probability given a specific true differential response y . Referring to the multivariate derivation, in this case λ is given as y^2/σ^2 (the squared differential response relative to noise) and no integration over λ is necessary:

$$\begin{aligned}
 p_{\text{detect}}^{\text{uni,specific}}(cnr) &= P \left[|u|^2 \geq \chi_{\text{crit}}^2 \right] \\
 &= \int_{\chi_{\text{crit}}^2}^{\infty} f(x; 1, y^2/\sigma^2, 1) dx \\
 &= 1 - F(\chi_{\text{crit}}^2; y^2/\sigma^2, 1).
 \end{aligned} \tag{D.2}$$

where $F(\chi_{\text{crit}}^2; y^2/\sigma^2, 1)$ is the non-central χ^2 cumulative distribution function with non-centrality parameter y^2/σ^2 and one degree of freedom, and χ_{crit}^2 is given by D.1 with $N = 1$.

The second probability is the average univariate detection probability over voxels. The multivariate detection probability takes into account that the specific true differential response pattern given the contrast-to-noise ration is not known. In the case of $N = 1$ the multivariate detection probability averages over the distribution of differential responses. This can also be interpreted as averaging over a distribution of voxels. As a consequence average univariate detection probability is a special case of multivariate detection probability as derived above for $N = 1$:

$$p_{\text{detect}}^{\text{uni,avg}}(cnr) = p_{\text{detect}}^{\text{multi}}(cnr, 1). \tag{D.3}$$

Statement of contributions

Study 1: Modeling and analysis of mechanisms underlying fMRI-based decoding of information conveyed in cortical columns.

The manuscript has been published in *NeuroImage* 56 (2011), pp. 627—642.

Denis Chaimow developed and implemented the model and wrote the manuscript. Essa Yacoub contributed to discussions. Kâmil Uğurbil contributed to discussions. Amir Shmuel advised on the developed of the model and edited the manuscript.

Study 2: Spatial specificity of the functional MRI blood oxygenation response relative to neuronal activity

The manuscript is under review for publication in *NeuroImage*.

Denis Chaimow developed and implemented the model, analyzed the data, and wrote the manuscript. Essa Yacoub acquired the data, contributed to discussions and commented on the manuscript. Kâmil Uğurbil contributed to discussions and commented on the manuscript. Amir Shmuel acquired the data, advised on the development of the model, and edited the manuscript.

Study 3: A more accurate account of the effect of k-space sampling and signal decay on the effective spatial resolution in functional MRI

The manuscript is about to be submitted for publication.

Denis Chaimow developed and implemented the model and wrote the manuscript. Amir Shmuel advised on the development of the model and edited the manuscript.

Study 4: Optimization of functional MRI for detection, decoding and imaging the response patterns of cortical columns

The manuscript is under review for publication in *NeuroImage*.

Denis Chaimow developed and implemented the model and wrote the manuscript. Kâmil Uğurbil contributed to discussions. Amir Shmuel advised on the development of the model and edited the manuscript.