The Effect of Acquisition Resolution and Magnetic Field Strength on Multivariate Decoding of fMRI

Dissertation

zur Erlangung des akademischen Grades
Doctor rerum naturalium

genehmigt durch die **Fakultät für Naturwissenschaften**
der Otto-von-Guericke-Universität

von : **Ayan Sengupta**

Gutachter
Jun.-Prof. Dr. Michael Hanke
Prof. Dr. David G. Norris

eingereicht am: 18.06.2016
verteidigt am: 23.11.2016

Magdeburg, November 2016

Abstract

A decade after it was shown that the orientation of visual grating stimuli can be decoded from human visual cortex activity by means of multivariate pattern classification of BOLD fMRI data, numerous studies have investigated which aspects of neuronal activity are reflected in BOLD response patterns and are accessible for decoding. However, it remains inconclusive what are the effects of acquisition resolution and MR field strength on BOLD fMRI decoding analyses. This thesis is the first to provide empirical ultra high-field (7 Tesla) fMRI data recorded at four spatial resolutions (0.8 mm, 1.4 mm, 2 mm, and 3 mm isotropic voxel size) on this topic — in order to test the hypotheses on the strength and spatial scale of orientation discriminating signals. Here I present detailed analysis, in line with predictions from previous simulation studies, about how the performance of orientation decoding varies with different acquisition resolutions. This study also for the first time investigates the effect of MR field strength on orientation decoding by comparing classification performance across field strengths (7T vs 3T) in 1.4 mm, 2 mm, and 3 mm resolutions. The interplay between acquisition resolution and the time series signal to noise ratio contributing to the effective decoding is also highlighted in this thesis. The potential of using multiband data acquistion in multivariate decoding studies to provide fast EPI acquisitions with relatively low signal losses as compared to parallel imaging techniques has been shown here. Moreover, I also examine different spatial filtering procedures and its effects on multivariate decoding across different resolutions, across field strengths and in different primary sensory regions of the brain (visual and auditory cortex). Here I show that higher-resolution scans with subsequent down-sampling or low-pass filtering yield no benefit over scans natively recorded in the corresponding lower resolution. The orientation-related signal in the BOLD fMRI data is spatially broadband in nature, includes both high spatial frequency components, as well as large-scale biases previously proposed in the literature.

Moreover, I found above chance-level contribution from large draining veins to orientation decoding. Multi-resolution raw EPI data acquired at the 7 Tesla were publicly released to facilitate further investigation.

M.S. Ayan Sengupta

The Effect of Acquisition Resolution and Magnetic Field Strength on Multivariate Decoding of fMRI

A decade after it was shown that the orientation of visual grating stimuli can be decoded from human visual cortex activity by means of multivariate pattern classification of BOLD fMRI data, numerous studies have investigated which aspects of neuronal activity are reflected in BOLD response patterns and are accessible for decoding. However, it remains inconclusive what are the effects of acquisition resolution and MR field strength on BOLD fMRI decoding analyses. This thesis is the first to provide empirical ultra high-field (7 Tesla) fMRI data recorded at four spatial resolutions (0.8 mm, 1.4 mm, 2 mm, and 3 mm isotropic voxel size) on this topic — in order to test the hypotheses on the strength and spatial scale of orientation discriminating signals. Here I present detailed analysis, in line with predictions from previous simulation studies, about how the performance of orientation decoding varies with different acquisition resolutions. This study also for the first time investigates the effect of MR field strength on orientation decoding by comparing classification performance across field strengths (7T vs 3T) in 1.4 mm, 2 mm, and 3 mm resolutions. The interplay between acquisition resolution and the time series signal to noise ratio contributing to the effective decoding is also highlighted in this thesis. The potential of using multiband data acquistion in multivariate decoding studies to provide fast EPI acquisitions with relatively low signal losses as compared to parallel imaging techniques has been shown here. Moreover, I also examine different spatial filtering procedures and its effects on multivariate decoding across different resolutions, across field strengths and in different primary sensory regions of the brain (visual and auditory cortex). Here I show that higher-resolution scans with subsequent down-sampling or low-pass filtering yield no benefit over scans natively recorded in the corresponding lower resolution. The orientation-related signal in the BOLD fMRI data is spatially broadband in nature, includes both high spatial fre-

quency components, as well as large-scale biases previously proposed in the literature. Moreover, I found above chance-level contribution from large draining veins to orientation decoding. Multi-resolution raw EPI data acquired at the 7 Tesla were publicly released to facilitate further investigation.

M.S. Ayan Sengupta

The Effect of Acquisition Resolution and Magnetic Field Strength on Multivariate Decoding of fMRI

Ein Jahrzehnt, nachdem gezeigt wurde, dass die Ausrichtung visueller Gitterreize durch multivariates Musterdecodieren ausgelesen werden kann, haben zahlreiche Studien mittels Klassifizierung von BOLD fMRI-Daten untersucht, welche Aspekte die neuronalen Aktivität in BOLD Antwortmuster reflektiert und zugänglich für die Decodierung sind. Es ist jedoch nicht eindeutig, was die Wirkung von MR Auflösung und MR Feldstärke auf die Dekodierung ist. Diese Arbeit ist die erste, die empirische Ultrahochfeld-7T fMRI-Daten zur Verfügung stellt, die in vier räumlichen Auflösungen aufgezeichnet (0,8 mm, 1,4 mm, 2 mm und 3 mm isotropen Voxelgröße) wurden, um die Hypothesen zur Stärke und räumlichen Skala der Orientierungsscheidungssignale zu testen. Hier stelle ich eine detaillierte Analyse im Einklang mit den Prognosen aus vorherigen Simulation Studien darüber vor, wie die Leistung der Orientierungsdekodierung mit verschiedenen Akquisitionsauflösungen variiert. In dieser Studie wurde auch zum ersten Mal die Wirkung von MR-Feldstärke untersucht, um die Orientierungsdecodierung durch Klassifikationsleistung über Feldstärken (7T vs 3T) in 1,4 mm, 2 mm und 3 mm Auflösungen zu vergleichen. Das Zusammenspiel zwischen der Auflösung der Erfassung und des Zeitreihe Signal-Rausch-Verhältnisses im Hinblick auf die wirksame Decodierung wird auch in dieser Arbeit hervorgehoben. Hier wurde das Potential der Verwendung von Multiband-Datenerfassung in multivariaten Dekodierungsstudien mit schnellen EPI Akquisitionen gezeigt, die relativ geringe Verluste im Vergleich zu parallelen Bildgebungsverfahren haben. Außerdem habe ich auch verschiedene räumliche Filterungsverfahren analysiert und seine Auswirkungen auf die multivariate Dekodierung in verschiedenen Auflösungen, über Feldstärken und in verschiedenen primären sensorischen Bereichen des Gehirns (Visueller und auditorischer Kortex) untersucht. Hier zeige ich, dass höhere auflösende Scans mit anschließender Abwärtsabtastung oder Tief-

passfilterung keinen Vorteil gegenüber Scans ergeben, die nativ in der entsprechenden geringeren Auflösung aufgezeichnet wurden. Die orientierungsbezogenen Signale der BOLD fMRI-Daten sind räumlich von der Art Breitband, das sowohl hohe räumliche Frequenzkomponenten enthält, als auch niedrig-frequenten Signale. Außerdem fand ich einen signifikanten Beitrag von großen drainierenden Venen zur Orientierungsdecodierung. Die rohen EPI Daten in multipler Auflösung, die am 7T erworben wurden, wurden für weitere Untersuchungen öffentlich zur Verfügung gestellt.

Acknowledgements

First of all, I would like to thank Prof. Dr. Stefan Pollmann for offering me the position of a Research Associate, and support me during my entire time as a doctoral student. This project would not have been possible without Jun.- Prof. Dr. Michael Hanke. Without his persistent guidance and his scientific, technical and philosophical inputs, I would not have been able to take this thesis to its conclusion. He introduced me to the idea of *Open Science* and it's importance in the modern scientific world. I have enjoyed so much in contributing to these projects, that I would like to keep on collaborating with him for a long time to come. Moreover, scientific discussions with Prof. Dr. Oliver Speck, Dr. Claus Tempelmann, Dr. Martin Kanowski, Dr. Jörg Stadler and Prof. Dr. Michael Hoffmann have provided very useful insights into the project. I would also like to thank Denise Sheermann for helping me with MR data acquisition and recruitment of participants.

I would also like to acknowledge the way in which I got enriched as a student of Science and also as a person by discussing a variety of topics over lunch and barbeque seesions with my colleagues. I consider myself lucky to have a vibrant social life in Magdeburg even though I don't speak the German language and I credit all my multi-cultural friends for that, especially the Skorwider family.

In the end, my whole journey of graduate studies would not have been possible without the constant support and encouragement from my family. My wife, Urmi and my parents have stood by me through thick and thin and given me the confidence to achieve whatever I could. Lastly, I would like to dedicate my entire work in this thesis to our little daughter Urja, whose cheerful smile makes my life rejuvinated with energy and hope.

CONTENTS

1. General Introduction

1.1 fMRI and Univariate Analysis

Functional Magnetic Resonance Imaging (fMRI) has been a fast and effective tool for understanding functioning of the brain *in-vivo*, by utilizing Blood Oxygen Level Dependent (BOLD) signals. BOLD signals are coupled to the haemodynamic activity of the underlying neural patterns [Ogawa et al., 1990]. Typically a magnetic resonance image consists of a three dimensional volumetric representation of the brain, where different regions of the brain are represented as different intentisites of multiple voxels. A voxel is a cuboid representing the BOLD signal of the brain at that particular location. Hence the neural activity of a brain performing a cognitive task is reflected in the local blood flow estimates (a proxy for local neural processing) and is represented as a modulation of a voxel intensity in a Magnetic Resonance image over a period of time. A fMRI data consist of several volumetric MR images of the brain acquired over a period of time, providing snapshots of the functioning of the brain (in terms of BOLD signal) every repetition time (TR). The modulation of image intensity of a particular voxel is referred to as the time series of that voxel and gives an estimate of the temporal haemodynamic activity in the brain at that particular location. Generally, fMRI analysis involves a mass-univariate approach to statistical analysis of the individual voxel time series. As a part of the analysis, the experimenter defines a reference model (experimental design) of the stimulus (or task) over the period of the experiment. A General Linear Model (GLM) is fitted to the time-series data of each voxel representing the linear estimation of the individual voxel activity in terms of the experimental design. It has to be noted that the BOLD signal is sluggish and temporaly smeared and in order to account for this, the experimental model is first convolved with an assumed Haemodynamic Response Function (HRF) before the linear corelation is performed with the fMRI time-series of

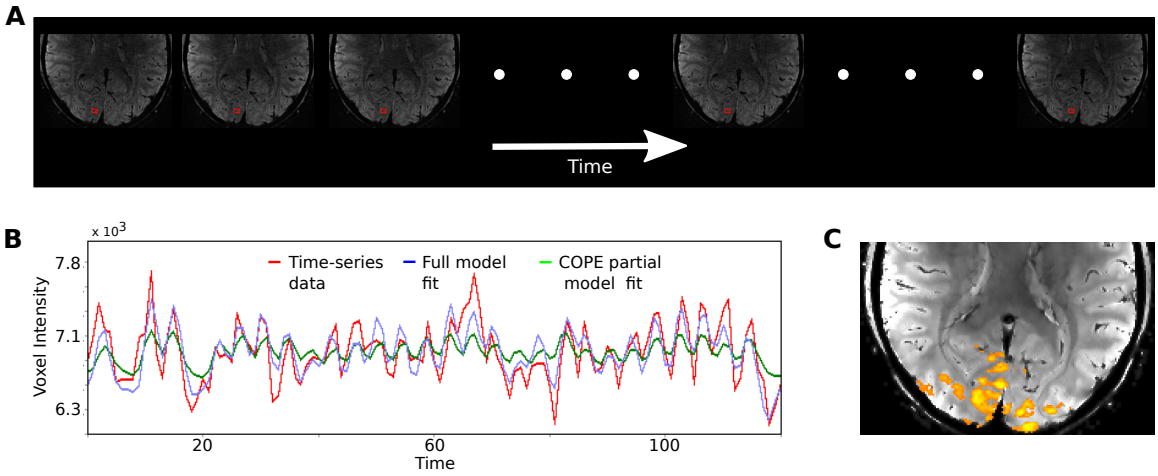

Figure 1.1: Mass-univariate General Linear Model analysis
(A) Voxel intensity variation with respect to time (also known as voxel time-series) **(B)** The full model fit (including all experimental conditions in the design) to the recorded time series data. Contrast of Parameter Estimates (COPE) partial model fit shows how the model fit to the data considering only the contrast of interest **(C)** Cluster of activated voxels in the visual cortex in response to a flickering orientation grating stimulus

individual voxels. A General Linear model in fMRI context can be expressed as:

$$Y = X\beta + \epsilon$$

where Y is the estimated voxel response, X is the pre-defined experimental design, β represents the model parameters and ϵ is the error term. The parameter estimates (PE) of the GLM fitted to the individual voxel data are defined as the set of β values which minimizes the sum of squared differences of the estimated BOLD response and the corresponding measured fMRI time-series. To find out whether a voxel is significantly more responsive to a particular experimental condition over the others, a t-statistic is performed with the β parameter estimates multiplied with a contrast vector between the two conditions [Mahmoudi et al., 2012, Poline and Brett, 2012]. Figure 1.1 shows the basic steps of a Univariate GLM approach of fMRI analysis.

1.2 Machine Learning Classifiers for analyzing fMRI data

In a recent approach to fMRI analysis, Machine Learning classification algorithms has been applied to patterns of brain activity across multiple voxels (instead of studying individual voxels) to differentiate between different experimental conditions. These analysis procedures are collectively known as Multivariate Pattern Analysis (MVPA). Though the General Linear Model approach has been the standard fMRI analysis procedure, it has several limitations for studying the voxels in isolation. The univariate model analysis ignores any covariance between neighboring voxel activities with respect to the cognitive task being performed. Hence as a standard step for noise reduction, spatial smoothing across voxels are performed with gaussian kernels, thus smearing out the fine-grained spatial patterns that might discriminate between experimental conditions [Norman et al., 2006]. Moreover, a tradition GLM approach tries to find the voxels which show statitically significant response to experimental conditions. But this ignores the contribution of the weaker voxel responses to a particular condition, which might carry important information. However, MVPA approach provides increased sensitivity by analyzing patterns of BOLD activity across voxels, irrespective of the fact that the voxel responses when studied individually may be non-significant. Thus MVPA provides a powerful tool to map a particular neural activity (represented as a BOLD activity patterns) with the corresponding cognitive state of the mind.

A standard MVPA procedure involves some basic steps as displayed in Figure 1.2. Firstly, a *feature selection* procedure is performed to choose, a set of voxels which are of relevance to the cognitive task being performed. For example, for a visual decoding study, voxels in the visual cortex will be selected. In some cases an univariate feature selection procedure may be performed to localize the voxels which are significantly more active than the others. A spatial searchlight algorithm [multivariate approach Kriegeskorte et al., 2006] is also used where adjacent sets of voxels are exhaustively tested over the brain for maximum informative content in their patterns of activity and a selection is made based on that criteria. There are several other methods of feature-

selection being to reduce the dimensioanality of the data [Norman et al., 2006]. In the second step, *pattern assembly* is performed. As shown in Figure 1.2A, the time-series of the voxels after feature selection are sorted in a manner that at a particular time point, the activation intensity of the selected voxels are considered to be a pattern and it is subsequently labeled with the corresponding cognitive task being performed by the brain at that time. The dataset thus created is partitioned into training and testing sections. The training dataset consists of labeled patterns of activation of the voxels (shown in Figure 1.2B) and is provided as an input to a machine learning classification algorithm, and is known as *classifier training*. The classifier learns to map the voxel activities to the provided labels. Then this trained classifier is applied on the unlabeled testing dataset, and the *classification* procedure assigns a predicted label for each time point based on patterns of voxel activities. In the final step of *cross-validation* the predicted labels are verified with the true labels and the classification accuracy is determined by the following formula:

$$Accuracy = \frac{TP + TN}{p + n}$$

where $p = TP + FN$ and $n = TN + FP$. The true positive count is TP and TN is the true negative count, FP is the count of total number of false positives and FN as false negatives. Generally the accuracy of a classification is represented as a mean of accuracies generated in a *Leave-one-run-out cross-validation* scheme. In this cross-validation procedure, the MVPA dataset is partitioned into chunks corresponding to each experimental run. Data from one chunk are treated as a testing dataset and the rest is used as a training dataset. The cross-validation procedure is repeated until all the runs (chunks) are invidually tested by the classifier.

Out of numerous machine learning classification algorithms, correlation based classifiers, linear discriminant analysis, linear support vector machines, Bayes classifiers, Radial basis function networks etc. have been used in the context of MVPA clas-

sification of cognitive states. Though each classification algorithm has its own share of advantages and limitations, linear classifiers are used very commonly in multivariate analysis of neural patterns. In Figure 1.2C it is shown the general working principle of a linear and a non-linear classification kernel. The MVPA dataset is represented in a multi-dimensional space with each of the voxels as one of the dimensions. Each time point is represented in this multi-dimensional space with each co-ordinate value as the intensity of each voxel at that time point. In general a trained classification algorithm determines decision boundary in the multi-dimensional space based on the labels provided in the training dataset. When the testing dataset is provided as an input to the trained classifier, the algorithm clusters the testing time-points into different labels based on the corresponding distance from the decision boundary. The implementation of different classification algorithms vary from each other, like for linear classifiers the decision boundary is a multidimensional plane but for non linear classifiers it is a multidimensional non-linear surface.

A machine learning classifer performance is determined by a set of its hyperparameters, for example the C value for a Linear SVM classifier [Burges, 1998, Chang and Lin, 2011]. As shown in Figure 1.2C, the C parameter represents the trade-off between width of the margin of the decision boundary and number of support vectors. Higher the value of C parameter, the more rigid is the margin of SVM. These hyperparameters could have substantial effect on decoding performance and needs to be tuned with the process of nested-cross validation. In this method, to optimize the value of the hyperparameter, the training dataset independently undergoes cross-validation estimations after being partitioned into nested training and testing datasets. Once the optimal nested-cross validation loop determines the best performing hyperparameter value the classification model is created with the entire training dataset and multivariate classification is performed on the testing dataset. Figure 1.3A shows the basic nested cross validation procedure.

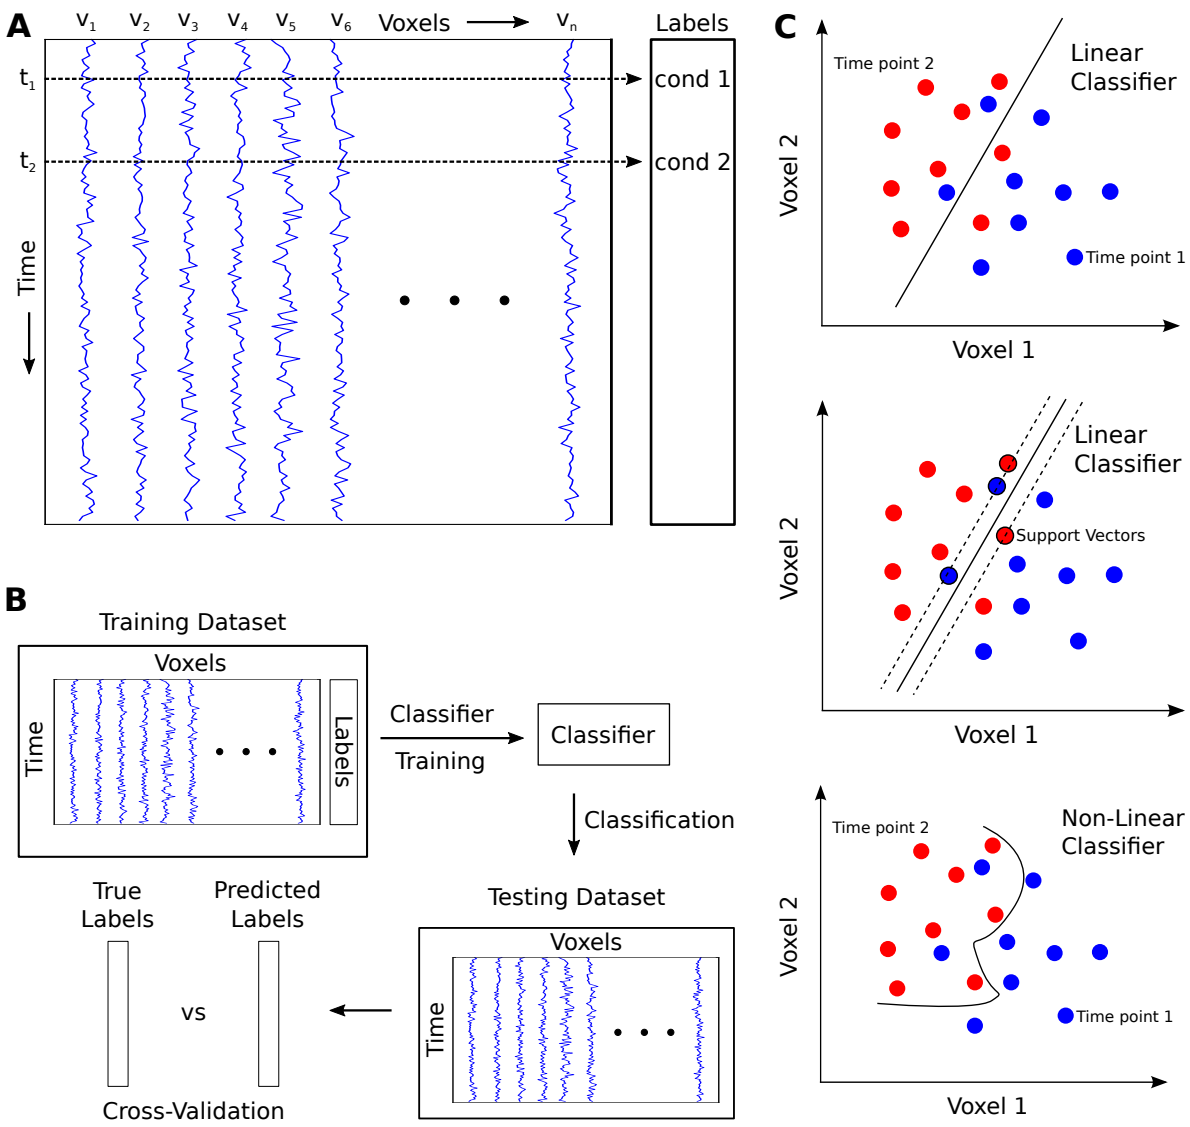

Figure 1.2: Multivariate Pattern Analysis

(A) BOLD time-series data represented as a multivariate dataset and corresponding labels provided to the pattern of voxel intensities at a particular time point. **(B)** The basic steps of a multivariate decoding study, partitioning dataset into training and testing parts, training the classifier with the training dataset, cross validation of the classifier predicted labels **(C)** Linear and Non-linear decision boundaries of a trained Linear and Non-linear classifier. The 'C' parameter represents the tradeoff between margin width and number of support vectors.

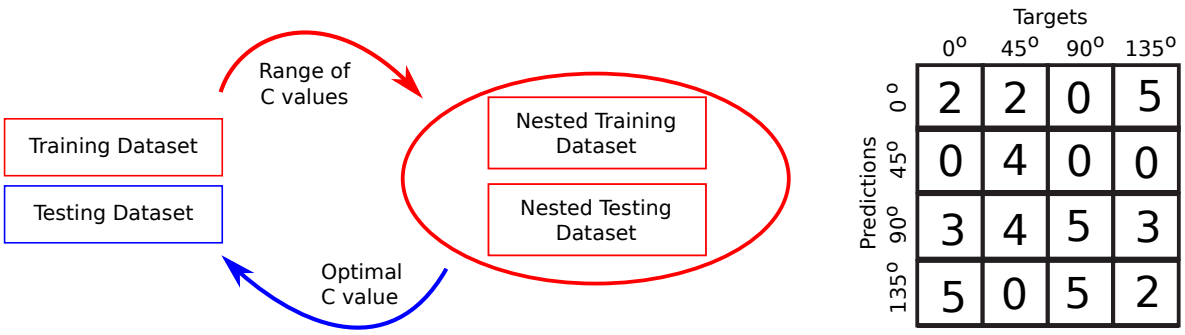

Figure 1.3: Nested Cross Validation

(A) Nested cross validation, a procedure to tune hyperparameters of the classifier. The training set is partitioned into nested training and testing dataset and an optimal hyperparameter value is determined in an iterative manner, which is then used for creating the classifier model for the entire dataset. **(B)** *confusion matrix* - Representation of the performance of a classifier. A sample confusion matrix of a 4 way classification of orientation gratings. The true labels (also known as 'Targets') are represented along the columns and corresponding predicted labels along the rows.

1.3 Current State of Research in Orientation Decoding

As described in the previous sections, multivariate pattern analysis approach has been successfully used in studying fMRI signal patterns from the brain and associate it with the corresponding cognitive task being performed by the participant at that time. *Orientation decoding* is one of the most extensively studied paradigm which uses multivariate analysis [Haynes and Rees, 2005, Kamitani and Tong, 2005]. In this paradigm a participant undergoes fMRI scanning while shown oriented gabor or sine-wave gratings in their visual field and machine learning classifiers are trained to decode the corresponding orientation of the gratings from the patterns of BOLD activity in the primary visual cortex.

From previous literature [Bartfeld and Grinvald, 1992], the presence of cortical columnar structures in the primary visual cortex is known, which are responsive to a particular orientation. There are some probable models of arrangement of the orientation columns in the striate cortex. Hubel and Wiesel [1972] proposed the 'ice cube' model (see Figure 1.4A), where the orientation columns and the ocular dominance columns are arranged orthogonal to each other forming a cuboid structure and the spa-

tial frequency of the orientation columns were higher than that of the ocular dominance columns. The 'ice cube' model by Hubel et al. [1995] was a speculative model and it was unlikely that an idealized model like this would be found in a variable biological system like the visual cortex [Dow, 2002]. From their electro-physiological studies on macaque striate cortex, Braitenberg and Braitenberg [1979] proposed the 'pinwheel' model of arrangement of the orientation columns in the visual cortex. In this model there are periodic orientation 'singularity' points around which various orientation columns were arrayed in a 'centric' fashion as they actually appear in the visual world [Dow, 2002](see Figure 1.4B). The same 'pinwheel' pattern was unveiled by ultra-high resolution fMRI in human primary visual cortex [Yacoub et al., 2008]. In Figure 1.4C the white and the black dots represent the 'singularity' points. Orientation pinwheels were arranged in a clockwise manner around the white and in a counter-clockwise manner around the black singularity points. These orientation columns are sub-millimeter in size and the orientation selectivity in the primary visual cortex cycle through all orientations approximately every millimeter. Thus signal from multiple orientation columns are averaged into a standard 3 mm isotropic voxel size due to partial volume effect. But orientation decoding has been successfully performed by several studies over the last decade [Alink et al., 2013, Boynton, 2005, Gardner, 2010, Haynes and Rees, 2005, Kamitani and Sawahata, 2010, Misaki et al., 2013, Swisher et al., 2010]. This has led to conflicting inferences about the true spatial scale of the orientation signals that the classifiers can use to learn to discriminate different orientations [Alink et al., 2013, Op de Beeck, 2010, Freeman et al., 2013, Swisher et al., 2010].

The mechanism by which signals from low-resolution voxels are being successfully decoded to predict information represented at a fine scale relative to the voxel size is a subject of ongoing debate. This paragraph describes different hypotheses and claims regarding this. Kamitani and Tong [2005] showed that the orientation decoding per-

[3]Adapted from Figure 3 of Dow [2002]
[3]Reproduced from Figure 4 of Dow [2002]
[3]Reproduced from Figure 2 of Yacoub et al. [2008]

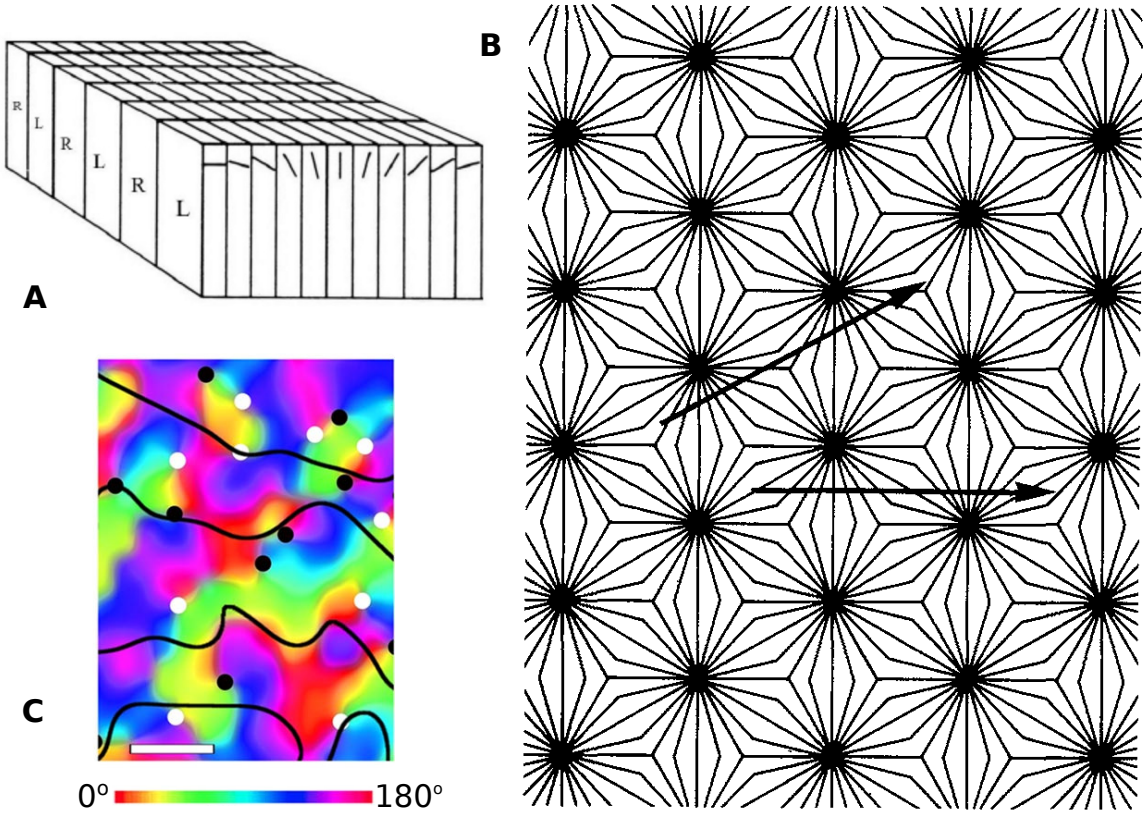

Figure 1.4: Orientation Columns in V1
(A) Arrangement of orientation columns and ocular dominance columns in the macaque striate cortex, as shown by an 'ice cube' model by Hubel et al. [1995].[1] **(B)** 'Pinwheel' model of orientation columns shown in macaque cortex found by an electro-physiological study by Braitenberg and Braitenberg [1979]. The arrows show the theoretical direction of the movement of the electrode in the visual cortex.[2] **(C)** 'Pinwheel' patterns shown by fMRI activity in human V1. Solid black lines represent the ocular dominance column borders.[3]

formance is the best in lower visual areas like V1 and V2 but higher visual regions like V4 and MT+ showed little or no orientation selectivity. This finding was in line with the previous electrophysiological and optical imaging studies of the primate visual cortex, which showed that the presence of orientation columns primarily in lower visual regions. From this, Kamitani and Tong [2005] inferred that in the early visual cortex, there are random variability in the spatial distribution of the orientation columns which lead to local biases in individual voxels leading to a robust decoding of sub-voxel structures. This view was supported by a simulation study by Chaimow et al. [2011], where he showed that on increasing the random irregularity of the simulated ocular

dominance columns enhanced the decoding performance in V1. However, Op de Beeck [2010] showed that spatial smoothing Gaussian kernels (upto 10mm FWHM) prior to orientation decoding did not affect the classification accuracies and hence claimed that the classifiers pick up orientation signals from large scale functional organizations in the visual cortex. These contribution of the large scale biases in the visual cortex include the oblique and radial effects as shown by Furmanski and Engel [2000], Sasaki et al. [2006]. This view is also supported by Freeman et al. [2011] where it has been shown that the large scale topographic radial maps in the V1 are necessary and sufficient for robust decoding of orientation gratings and globally coherent stimulus like spirals [Freeman et al., 2013]. The coarse-scale interpretation of orientation decoding has not gone unchallenged [see Alink et al., 2013]. According to Swisher et al. [2010] the orientation signals are spatially broadband in nature and present in the spatial range of 1 mm to 10 mm. In a recent paper Pratte et al. [2016] also showed how the fast temporal-encoding paradigm for spatial mapping [as implemented in Freeman et al., 2011], can lead to erroneous estimates of a voxel's orientation or retinotopic preference. They claimed that the radial bias is not the only source of orientation information in fMRI signal and hence not necessary for orientation decoding. Apart from the signal biases hypotheses, the ability to decode orientations from V1 at low frequencies is also attributed to the contribution of vasculature (large draining veins) [Gardner, 2010]. Kriegeskorte et al. [2010] introduced the concept of voxel sampling as a complex spatio-temporal filter due to the contributions of the veins, which make the fine-grained signals available at a much lower spatial frequency.

Due to technological advances in recent years, it is possible to measure very high resolution fMRI scans in ultra-high MR field strength (7 Tesla). A 7 Tesla scanner provides superior BOLD sensitivity than a conventional 3 Tesla fMRI acquisition. Keeping in mind the extensive debate about the spatial scale of orientation signals, it is highly motivating to acquire high resolution fMRI for performing multivariate analysis because it provides substantially low partial voluming effect [Weibull et al., 2008]. This is

especially relevant when Yacoub et al. [2008] has shown that modelling of the pinwheel patterns of the orientation columns in V1 is possible from 7 Tesla fMRI acquired with 0.5 mm in-plane resolution.

Moreover, in general, to study the spatial scale of orientation signals by comparing metrics like prediction accuracy, across a range of spatial frequencies, the authors typically acquired high-resolution fMRI and simulated a lower-resolution acquisition by applying spatial filters to the original data [see Swisher et al., 2010], or reconstruction of k-space data to different resolutions [Gardumi et al., 2016]. However, these approaches has not gone unchallenged as it is unclear to what degree particular filtering strategies [e.g. Gaussian filtering vs. low-pass filtering in the spatial frequency domain, see Misaki et al., 2013] can effectively simulate the properties of fMRI recorded at a lower physical resolution, where a change in slice thickness alone can significantly alter image contrast. Despite this criticism, I am not aware of any study that has compared the performance of orientation decoding in visual cortex across a range of physical acquisition resolutions. Hence, the effect data acquisition at different resolutions on multivariate decoding still remains inconclusive.

1.4 Thesis Outline

In this thesis, I have addressed these open questions in Experiment 1 (Chapter 3), by applying multivariate pattern analysis to fMRI data acquired with a standard orientation decoding paradigm [similar to the paradigm used in Swisher et al., 2010], in four different resolutions at 7 Tesla (0.8 mm, 1.4 mm, 2 mm and 3 mm isotropic voxel size). Moreover, to give a better perspective on the spatial scale of the orientation signals, spatial smoothing (including volumetric gaussian filtering and cortical surface based filtering) and spatial resampling were also performed. To study the contribution of veins to orientation decoding, separate susceptibility weighted images were analyzed to localize the venous voxels in V1 and subsequently orientation decoding was performed in the venous and non-venous voxels separately. In order to check whether orientation

decoding benefits from higher magnetic field strengths, in Experiment 2 (Chapter 4) I present a study to compare multivariate analysis on data acquired from 3 Tesla and 7 Tesla Siemens scanners, at three different resolutions (1.4 mm, 2 mm and 3 mm iso) with the same stimulation paradigm, almost identical scanning parameters and similar decoding analysis procedures. It is known that the data acquired from 7 Tesla show substantially better time-series signal to noise ratio than the 3 Tesla data. To enhance tSNR in the 3 Tesla acquisition keeping the stimulation paradigm unaltered, I also performed in Experiment 3 (Chapter 5) a multiband acquisition with no parallel imaging technique for 2 mm 3 Tesla acquisition and compared decoding accuracy with the normal 2 mm acquisition. The spatial scale of orientation signals as reflected in the 3 Tesla data, were studied with spatial filtering and compared with the corresponding 7 Tesla data. Additionally in this thesis, I analyzed whether the spatial scale of columnar structures as reflected by fMRI data are comparable across different sensory processing regions of the brain. Experiment 4 (Chapter 6) presents multivariate decoding of musical genres from fMRI signal patterns recorded in primary auditory cortex. The data were recorded in 7 Tesla with 1.4 mm resolution and subsequent multivariate analysis was performed on spatially filtered data at different levels of gaussian smoothing. Overall, this thesis provides a comprehensive analysis of the effect of acquisition resolution and MR field strength on multivariate decoding. To fascilitate future research, the multi-resolution data acquired in the 7 Tesla has been made publicly available (see Appendix A) in BIDS (Brain Imaging Data Structure) format as a part of the *studyforrest* project [Hanke et al., 2014] and in recent future, the rest of the data will also be published.

2. Retinotopic Mapping: Localization of V1

Experiment 1 (chapter 3) - Experiment 3 (chapter 5) used retinotopic mapping for localizing the primary visual cortex (V1) and here I include the detailed description of the experimental design, data acquisition protocols, and all computational processes required for the same [first published in Sengupta et al., 2016]. The data processing pipeline and the quality analysis of the retinotopic maps described in this chapter, were developed in collaboration with Falko R. Kaule and Professor Michael B. Hoffmann, who were co-authors of Sengupta et al. [2016].

2.1 Participants

All of the participants recruited in the following experiments previously volunteered for both studies of the *studyforrest* project [Hanke et al., 2014, 2015a]. The pool of the participants included a total of fifteen right-handed volunteers (mean age 29.4 years, range 21–39, 6 females). The demographics of the participants of the individual experiments are described in the specific methods sections. The integrity of their visual function was assessed at the Visual Processing Laboratory, Ophthalmic Department, Otto-von-Guericke University, Magdeburg, Germany (as specified in the following section), under the supervision of Professor Michael B. Hoffmann. Participants were fully instructed about the purpose of the study and received monetary compensation. They signed an informed consent for public sharing of all obtained data in anonymized form. All anonymized subject ID references made in this thesis are identical as in Hanke et al. [2014]. This study was approved by the Ethics Committee of the Otto-von-Guericke University.

2.2 Subjective measurements of visual function

To test whether the participants had normal visual function and to detect critical reductions of visual function, two important measures were determined: (1) visual acuity to identify dysfunction of high resolution vision and (2) visual field sensitivity to localize visual field defects. For each participant, these measurements were performed for each eye separately — if necessary with refractive correction. (1) Normal decimal visual acuity (>=1.0) was obtained for each eye of each participant. (2) Visual field sensitivities were determined with static threshold perimetry (standard static white-on-white perimetry, program: dG2, dynamic strategy; OCTOPUS Perimeter 101, Haag-Streit, Koeniz, Switzerland) at 59 visual field locations in the central visual field (30 ° radius) *i.e.*, covering the part of the visual field that was stimulated during the MRI scans. In all, except for two participants, visual field sensitivities were normal for each eye (MD (mean defect) dB<2.0 & >-2.0; LV (loss variance) $dB^2 < 6$) — indicating the absence of visual field defects. Visual field sensitivities for sub-04 (both eyes) were slightly lower than normal but not indicative of a distinct visual field defect.

2.3 Retinotopic Mapping

2.3.1 Stimulus

Similar to previous studies [Engel et al., 1997a, Sereno et al., 1995], traveling wave stimuli were designed to encode visual field representations in the brain using temporal activation patterns[Warnking et al., 2002]. Expanding/contracting rings and clockwise/counter-clockwise wedges (see Figure 2.1A) consisting of flickering radial checkerboards (flickering frequency of 5 Hz) were displayed on a gray background (mean luminance $\approx 100\, cd/m^2$) to map eccentricity and polar angle. The total run time for both eccentricity and polar angle stimuli was 180 s, comprising five seamless stimulus cycles of 32 s duration each along with 4 s and 12 s of task-only periods (no checkerboard stimuli) respectively at the start and the end.

The flickering checkerboard stimuli had adjacent patches of pseudo-randomly cho-

sen colors, with pairwise euclidean distances in the *Lab* color space (quantifying relative perceptual differences between any two colors) of at least 40. Each of these colored patches were plaided with a set of radially moving points. To improve the perceived contrast, the points were either black or white depending on the color of the patch on which the points were located. The lifetime of these points was set to 0.4 s, a new point at a random location was initialised after that. With every flicker, the color of the patches changed to its complementary luminance. Simultaneously, the color changed and the direction of movement of the plaided points also reversed.

Eccentricity encoding was implemented by a concentric flickering ring expanding and contracting across the visual field (0.95°of visual angle in width). The ring was not scaled with cortical magnification factor. The concentric ring traveled across the visual field in 16 equal steps, stimulating every location in the visual field for 2 s. After each cycle, the expanding or the contracting rings were replaced by new rings at the center or the periphery respectively.

Polar angle encoding was implemented by a single moving wedge (clockwise and counter-clockwise direction). The opening angle of the wedge was 22.5 degrees. Similar to the eccentricity stimuli, every location in the visual field was stimulated for 2 seconds before the wedge was moved to the next position.

2.3.2 Center letter reading task

In order to keep the participants' attention focused and to minimize eye-movements, they performed a reading task. A black circle (radius 0.4°) was presented as a fixation point at the center of the screen, superimposed on the main stimulus. Within this circle, a randomly selected excerpt of song lyrics was shown as a stream of single letters (0.5° height, letter frequency 1.5 Hz, 85% duty cycle) throughout the entire length of a run. Participants had to fixate, as they were unable to perform the reading task otherwise. After each acquisition run, participants were presented with a question related to the previously read text. They were given two probable answers, to which they replied by

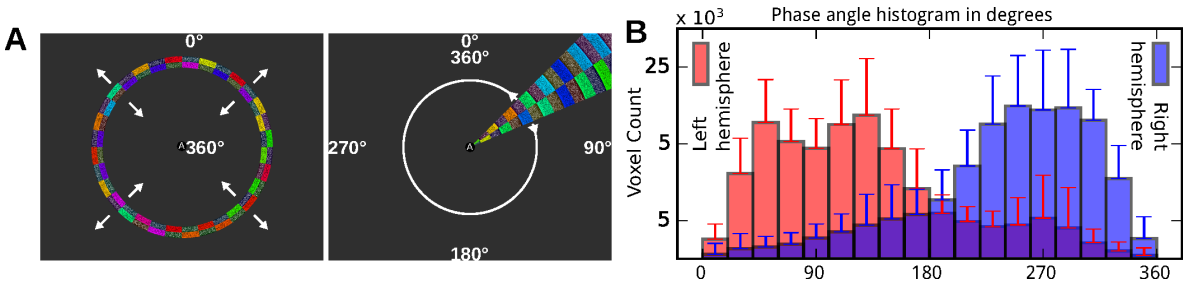

Figure 2.1: Retinotopic Mapping Stimuli
(A) Ring and wedge stimuli with continuous central letter reading task to encourage fixation. White numbers indicate the respective phase angle encoding. (B) Histogram of polar angles for all voxels in the MNI occipital lobe mask for the left and right hemisphere. Error bars indicate standard deviation across all subjects. [4]

corresponding button press (index or middle finger of their right hand). These question only served the purpose of keep participants attentive — and were otherwise irrelevant.

2.3.3 Stimulation setup

Visual stimuli were presented on a rear-projection screen inside the bore of the magnet using an LCD projector (JVC DLA RS66E, JVC Ltd., light transmission reduced to 13.7% with a gray filter) connected to the stimulus computer via a DVI extender system (Gefen EXT-DVI-142DLN with EXT-DVI-FM1000). The screen dimensions were 26.5 cm×21.2 cm at a resolution of 1280×1024 px with a 60 Hz video refresh rate. The binocular stimulation were presented to the participants through a front-reflective mirror mounted on top of the head coil at a viewing distance of 63 cm. Stimulation was implemented with PsychoPy v1.79 (with an early version of the `MovieStim2` component later to be publicly released with PsychoPy v1.81)[Peirce, 2007] on the (Neuro)Debian operating system [Halchenko and Hanke, 2012]. Participant responses were collected by a two-button keypad and was also logged on the stimulus computer.

2.3.4 Functional MRI acquisition

For all of the fMRI acquisitions for retinotopic mapping, the following parameters were used: T2*-weighted echo-planar images (gradient-echo, 2 s repetition time (TR), 30 ms

[4] Reproduced from Figure 3 of Sengupta et al. [2016]

echo time, 90° flip angle, 1943 Hz/px bandwidth, parallel acquisition with sensitivity encoding (SENSE) reduction factor 2) were acquired during stimulation using a whole-body 3 Tesla Philips Achieva dStream MRI scanner equipped with a 32 channel head coil. 35 axial slices (thickness 3.0 mm) with 80 × 80 voxels (3.0×3.0 mm) of in-plane resolution, 240 mm field-of-view (FoV), anterior-to-posterior phase encoding direction) with a 10% inter-slice gap were recorded in ascending order — practically covering the whole brain. Philips' "SmartExam" was used to automatically position slices in AC-PC orientation such that the topmost slice was located at the superior edge of the brain. This automatic slice positioning procedure was identical to the one used for scans reported in the companion article [Sengupta et al., 2016] and yielded a congruent geometry across all paradigms.

2.3.5 Structural MRI acquisition

Structural images were acquired for all participants in the same 3 Tesla Philips Achieva scanner. Individual T1-weighted images consisted of 274 sagittal slices (FoV = 191.8 × 256 × 256 mm). It was recorded using a 3D turbo field echo (TFE) sequence (TR 2500 ms, inversion time (TI) 900 ms, flip angle 8°, echo time (TE) 5.7 ms, bandwidth 144.4 Hz/px, SENSE reduction AP 1.2, RL 2.0) with an acquisition resolution of 0.7 mm. It was reconstructed using a 384 × 384 in-plane reconstruction matrix (0.67 mm isotropic resolution). A 3D turbo spin-echo (TSE) sequence (TR 2500 ms, TE eff 230 ms, strong SPIR fat suppression, TSE factor 105, bandwidth 744.8 Hz/px, SENSE reduction AP 2.0, RL 2.0, scan duration 7:40 min) was used to acquire a T2-weighted image whose geometric properties were identical to the T1-weighted image. All the anatomical images were recorded as a part of the *studyforrest* project [Hanke et al., 2014] and are publicly available from GitHub `https://github.com/psychoinformatics-de/studyforrest-data-structural`.

2.3.6 Experimental Design

Participants performed four acquisition runs in a single session with a total duration of 12 min, with short breaks in-between and without moving out of the scanner. In each run, participants performed the center reading task while passively watching the contracting, counter-clockwise rotating, expanding, and clockwise rotating stimuli in exactly this sequential order. For the retinotopic mapping experiment, 90 volumes of fMRI data were acquired for each run.

2.3.7 Retinotopic mapping analysis

Many regions of interest (ROI) in the human visual system follow a retinotopic organization [Engel et al., 1997a, 1994, Sereno et al., 1995]. The primary areas like V1 and V2 are also provided as labels with the Freesurfer segmentation using the `recon-all` pipeline [Dale et al., 1999]. But the higher visual areas (V3, VO, PHC, etc) need to be localized by retinotopic mapping [Arcaro et al., 2009, Sereno et al., 2012, Silver and Kastner, 2009a, Wandell et al., 2007] or probability maps [Van Essen et al., 2001, Wang et al., 2014].

An analysis pipeline was implemented for the acquired fMRI data based on standard algorithms publicly available in the software packages Freesurfer [Dale et al., 1999], FSL [Smith et al., 2004], and AFNI [Cox, 1996]. All analysis steps were performed on a computer running the (Neuro)Debian operating system [Halchenko and Hanke, 2012], and all necessary software packages (except for Freesurfer) were obtained from system software package repositories.

BOLD images time series for all scans of the retinotopic mapping paradigm were brain-extracted using FSL's BET and aligned (rigid-body transformation) to a participant-specific BOLD template image. All volumetric analysis was performed in this image space. An additional rigid-body transformation was computed to align the BOLD template image to the previously published cortical surface reconstructions based on T1 and T2-weighted structural images of the respective participants[Hanke et al.,

2014] for later delineation of visual areas on the cortical surface. Using AFNI tools, time series images were also "deobliqued" (`3dWarp`), slice time corrected (`3dTshift`), and temporally bandpass-filtered (`3dBandpass` cutoff frequencies set to 0.667/32 Hz and 2/32 Hz, where 32 s is the period of both the ring and the wedge stimulus).

For angle map estimation, AFNI's `waver` command was used to create an ideal response time series waveform based on the design of the stimulus. The bandpass filtered BOLD images were then processed by the `3dRetinoPhase` (`DELAY` phase estimation method was based on the response time series model). Expanding and contracting rings, as well as clockwise and counter-clockwise wedge stimuli, were jointly used to generate average volumetric phase maps representing eccentricity and polar angles for each participant. Polar angle maps were adjusted for a shift in the starting position of the wedge stimulus compared between the two rotation directions. The phase angle representations, relative to the visual field, are shown in Figure 2.1A. As an overall indicator of mapping quality, Figure 2.1B shows the distribution of the polar angle representations across all voxels in the MNI occipital lobe mask combined for all participants.

For visualization and subsequent delineation, all volumetric angle maps (after correction) were projected onto the cortical surface mesh of the respective participant using Freesurfer's `mri_vol2surf` command — separately for each hemisphere. In order to illustrate the quality of the angle maps, the subjectively best, average, and worst participants (respectively: participant 1, 10, and 9) have been selected on the basis of visual inspection. Figure 2.2C shows the eccentricity maps on the left panel and the polar angle maps for both hemispheres on the right panel. Table 2.1 summarizes the results of the manual inspections of all surface maps. Delineations of the visual areas depicted in Figure 2.2 were derived according to Kaule et al. [2014](page 4). Further details on the procedure can be found in Arcaro et al. [2009], Silver and Kastner [2009b], Wandell et al. [2007].

[5]Reproduced from Figure 3 of Sengupta et al. [2016]

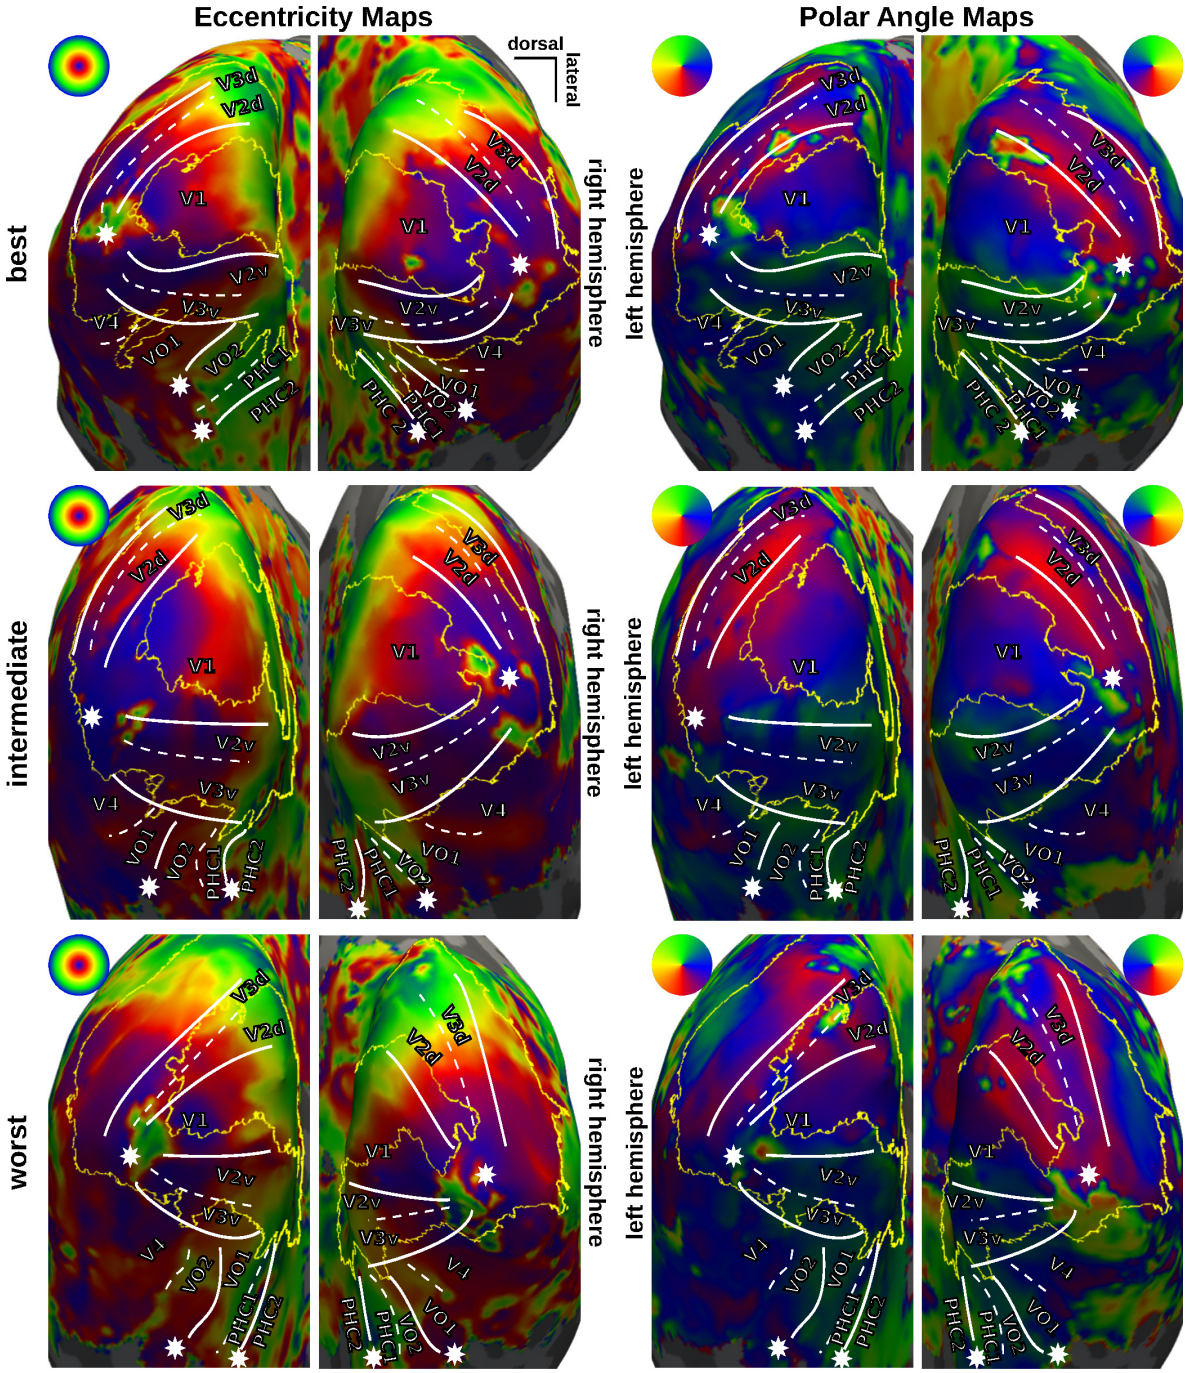

Figure 2.2: Retinotopic Mapping Quality Analysis

Inflated occipital cortex surface maps for eccentricity and polar angle for the best, intermediate, and worst participants: participants 1, 10, and 9 respectively. White lines indicate manually delineated visual area boundaries; stars mark the center of the visual field; yellow lines depict the outline of the autogenerated Freesurfer V2 label[Hanke et al., 2014] for comparison. All maps are constrained to the MNI occipital lobe mask.[5]

Participants	Phasemap Quality		Freesurfer Segmentation Fit			
	Left Hemifield	Right Hemifield	Left V1	Left V2/3	Right V1	Right V2/3
sub-04	4	5	3	3	1	2
sub-06	4	3	3	3	2	2
sub-09	2	2	1	1	1	1
sub-10	4	3	3	2	3	2
sub-16	3	4	3	1	3	2
sub-17	2	3	2	2	3	2
sub-18	3	2	3	2	2	3
sub-20	1	4	2	2	2	2
sub-21	2	3	2	3	3	2
Mean	2.78	3.22	2.44	2.11	2.22	2

Table 2.1: Quality analysis of the phasemaps generated by the retinotopic mapping processing pipeline.
The phasemaps of the participants recruited in the following experiments were checked for borders of V123 and Parahippocampal cortex (PHC) complex regions. The phasemap quality was graded from (1-5) with 5 being the best. The Freesurfer recon-all pipeline also provides segmentations of V1 and V2/3 regions as labels. The quality of the phasemaps generated here were also compared against Freesurfer segmentations. The fit of the Freesurfer segmentations to the phasemaps were graded on a scale of (1-3) with 3 being the best fit. [6]

[6]Generated in collaboration with Falko R. Kaule (available at `https://github.com/psychoinformatics-de/studyforrest-data-retinotopy`)

3. Experiment 1: The Effect of Acquistion Resolution on BOLD fMRI Decoding Analyses at 7 Tesla

3.1 Introduction

The term multivariate pattern (MVP) analysis summarizes a range of data analysis strategies that are highly suitable for studying neural representations encoded in distributed patterns of brain activity [see, for example, Bonte et al., 2014, Haxby, 2012, Haynes, 2009, Zhang et al., 2015]. While there is an ever increasing number of publications that demonstrate the power of MVP analysis for functional magnetic resonance imaging (fMRI) data [Alink et al., 2013, Op de Beeck, 2010, Freeman et al., 2011, 2013] with standard resolution (a voxel size of about 2-3 mm isotropic), MVP analysis is especially promising in the context of high-resolution fMRI. Ongoing technological improvements, such as ultra high-field MRI scanners (7 Tesla or higher) have pushed the resolution for fMRI to a level that is slowly approaching the spatial scale of the columnar organization of the brain [Heidemann et al., 2012, Yacoub et al., 2008]. Being able to use fMRI to sample brain activity patterns at a near-columnar level makes it feasible to employ MVP analysis with the aim to decode distributed neural representations of an entire cortical field at a level of detail that is presently only accessible to invasive recording techniques with limited spatial coverage. However, at this point, it is unclear which spatial resolution is most suitable for decoding neural representation from fMRI data with MVP analysis. While higher resolutions can improve the fidelity of the BOLD signal by, for example, reducing the partial volume effect [Weibull et al., 2008], the benefits can be counteracted by physiological noise (such as inevitable motion) and lower temporal signal-to-noise ratio (tSNR). This raises the question: does the decoding of neural representations continuously improve with increasing spatial resolution, or is there an optimal resolution for a given type of representation?

In this study, I provide empirical data on the effect of spatial acquisition resolution on the decoding of visual orientation from high field (7 Tesla) fMRI. I recorded BOLD fMRI data at 0.8 mm, 1.4 mm, 2 mm and 3 mm voxel size while participants were visually stimulated with oriented phase-flickering gratings using a uniform event-related paradigm. This is one of the most frequently employed MVP analysis technique: a cross-validated classification analysis, where a classifier is repeatedly trained to distinguish patterns of brain activation from fMRI data of a set of stimulus conditions, and its prediction accuracy is evaluated against a left-out data portion [Pereira et al., 2009]. Moreover, I focus on the decoding of the representation of oriented visual gratings in primary visual cortex. Chaimow et al. [2011] investigated the effect of acquisition resolution on decoding of the stimulated eye using simulated 3 Tesla fMRI data based on a model of ocular dominance columns. They found that a resolution of 3 mm was optimal for decoding and performance decreased with higher or lower resolution. It is known that the organization of orientation columns is characterized by higher spatial frequencies than ocular dominance columns [Obermayer and Blasdel, 1993] and the BOLD point-spread function (PSF) is considerably smaller than that at 3 Tesla [≈2.3 mm FWHM vs. ≈3.5 mm FWHM Engel et al., 1997b, Shmuel et al., 2007]. Considering that, I expect the maximum orientation decoding accuracy to be at a resolution higher than 3 mm. Though I are not trying to find or comment on a single *optimal resolution* for multivariate analysis, from this study I expect to get a better technical understanding of how information acquired at different resolutions at 7 Tesla contribute to orientation decoding performance. Multi-resolution data also allow for evaluating *filtering strategies* used in previous studies in terms of their validity regarding the simulation of lower-resolution fMRI acquisitions from high-resolution data. These data also enable the investigation of the *role of aliasing* of a high spatial-frequency signal (beyond the Nyquist frequency) into a lower frequency range sampled by fMRI voxels [sometimes referred to as "hyperacuity"; Op de Beeck, 2010, Swisher et al., 2010], as, in the case of spatial aliasing, the frequency bands carrying an orientation-selective signal would

vary with the sampling resolution of fMRI. Lastly, I collected high-resolution susceptibility weighted imaging data for blood-vessel localization in order to investigate the *role of large draining veins* that may carry orientation-selective signals reflected in low spatial frequency components when sampled by millimeter range voxels [Gardner, 2010, Kamitani and Tong, 2005, Kriegeskorte and Bandettini, 2007, Shmuel et al., 2010]. In combination with the multi-resolution fMRI data, I can investigate the effect of this potential signal source on the orientation decoding at various levels of spatial scale.

While my primary focus is on the technical aspect of acquisition resolution for decoding information from BOLD signal patterns using the representation of visual orientations as a well-researched example, I acknowledge that the data can be used to investigate a number of additional questions, such as the specific nature of the encoding of visual orientation in the BOLD signal pattern. It can also be a valuable resource in further optimization of the decoding procedure (classification algorithm, hyper-parameter optimization, etc.). In order to facilitate the required future analyses I have publicly released the data. It has been uploaded to OpenFMRI (accession number: ds000113c) and is also available without restrictions from GitHub `https://github.com/psychoinformatics-de/studyforrest-data-multires7t` and a description is available in Appendix A. This dataset will serve as starting point to a series of additional analysis that explore aspects beyond acquisition resolution.

3.2 Materials and methods

3.2.1 Participants

Seven healthy right-handed volunteers (age 21-38 years, 5 males) with normal or corrected to normal vision were recruited from the subject pool of the *studyforrest* project. Before every scanning session, they were provided with instructions for the experiment (approved by the Ethics Committee of the Otto-von-Guericke University) and signed an informed consent form.

3.2.2 Stimuli

A stimulus comprised two semi-annular patches of flickering sine-wave gratings left and right of a central fixation point on a medium gray background (0.8°-7.6° eccentricity, 160° width on each side with a 20° gap along the vertical meridian, above and below the fixation point, to aid separation of gratings between hemifields). Gratings on each side of the stimulus were independently oriented at either 0°, 45°, 90°, or 135°, with a constant spatial frequency of 1.4 cycles per degree of visual angle corresponding to the center of the screen, a contrast of 100%, and a flickering frequency of 4 Hz with 50% duty cycle [Swisher et al., 2010]. The phase of the gratings was changed at a frequency of 4 Hz and was chosen randomly from 0, $\frac{\pi}{2}$, π, or $\frac{3\pi}{2}$ degrees of phase angle (Figure 3.1).

Stimulus presentation and response logging were implemented using PsychoPy [v1.79; Peirce, 2008] running on a computer with the (Neuro)Debian operating system [Halchenko and Hanke, 2012]. Stimuli were displayed on a rear-projection screen (1280×1024 pixels resolution; 60 Hz video refresh rate; 25.5 cm wide) located behind the head coil. Participants viewed the screen via a mirror placed at a distance of ≈4 cm from their eyes. The total viewing distance was 100 cm.

3.2.3 Behavioral task

In order to keep the participants' attention focused and to minimize eye-movements, they performed the reading task (as described in section 2.3.2) throughout the entire length of a run. Each trial started with 3 s of stimulation with oriented gratings and continued for another 5 s of a task-only period (Figure 3.1). During task-only periods, a medium gray background was displayed in both hemifields. At the end of each run, the participant was asked a question related to the previously read text.

In a pilot experiment with in-scanner eye-movement recordings, the letter reading task was found to minimize eye-movements effectively; however, it was unsuitable to verify fixation accuracy on a trial-by-trial basis. In order to evaluate a potential impact of the reading task on the orientation decoding performance, the task was replaced for

one participant with a visual detection task. One participant was repeatedly presented with a Landolt C stimulus (radius 0.12°, left or right opening (0.048°) at random intervals in each run. The participants had to respond to the direction of the opening of the probe by pressing one of two buttons corresponding to a left or right opening (see Figure 4.1). Discrimination accuracy for this participant was 98.6%, while orientation decoding performance did not qualitatively differ from mean decoding accuracy of other participants. The performance of the subject with the Landolt C task was compared relative to the 95% binomial proportion confidence interval computed from the number of correct predictions (BOLD pattern classification), concatenated across hemispheres and cross-validation fold, and all subjects performing the reading task. For all resolutions (except 3 mm data) the performance of the subject performing the Landolt C task was within the confidence interval (for 3 mm the decoding accuracy was close to, but higher, than the upper boundary of the confidence interval). This suggests that the employed reading task was generally effective in keeping participants focused on the fixation point.

3.2.4 Procedures

Participants were scanned in four different sessions, one experiment session for each of the four acquisition resolutions (0.8 mm, 1.4 mm, 2.0 mm and 3.0 mm isotropic). These sessions took place on different days over the course of five weeks. The order of acquisition resolutions was randomized for each participant. In every experiment session, participants completed ten runs with short breaks in-between, without leaving the scanner. Each run comprised 30 trials (8 s duration; 4 min total run duration). Independent sequences were generated per hemifield with equal number of occurrences of each orientation. There were 4 different orientations (0°, 45°, 90°, or 135°) each occurring for exactly 5 times in the sequence, contributing to 20 trials in one run. The sequences were randomly shuffled per hemifield. This resulted in random pairings of orientations within trials. While analyzing, a single GLM was used to model the events in both

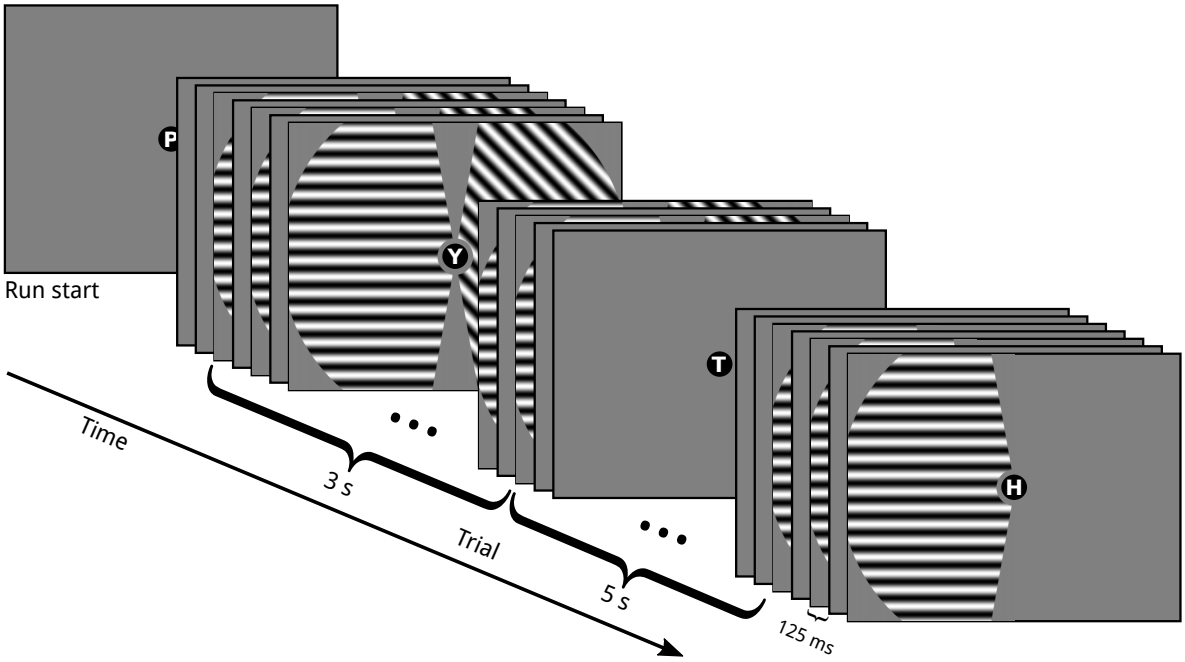

Figure 3.1: Stimulation paradigm
Independent flickering oriented grating stimuli on a medium gray background were presented in both hemifields for 3 s at the beginning of each trial. Stimulation was followed by a 5 s inter-trial interval. Throughout an entire experiment run, participants performed a continuous central letter reading task to maintain fixation. Interspersed trials where the previous stimulus was repeated in only one of the hemifields were used to decouple stimulation sequences.

the hemifields. This was done to account for potential inter-hemispheric cross-talk due to the simultaneous bilateral stimulation, and the correlation in this stimulus sequence between hemifields. Moreover, in order to minimize undesired attention shift effects, I opted for a simultaneous onset of the stimulation in both hemifields. Combined with the further constraint of the same number of stimulation trials per orientation in both hemifields, this would unavoidably lead to a singularity of the GLM design matrix, unless a further source of temporal variability is introduced. In order to decouple stimulation sequences between hemifields, 10 unilateral stimulation events (termed NULL events) were inserted into the trial sequence at pseudo-random positions (a run could not start with a NULL event and no two NULL events could occur in immediate succession). NULL events were identical to regular trials, except for the fact that in one hemifield the same oriented grating as in the previous trial was repeated while the

other hemifield remained empty. The side of the blank hemifield was chosen at random for each NULL event. For all participants, the actual generated trial sequences show a roughly equal count of NULL events for each hemifield and the NULL events were included in the GLM analysis.

3.2.5 Functional imaging

The objective for functional data acquisition was to obtain BOLD fMRI data from the V1 ROI at four different resolutions with an identical stimulation paradigm. MR acquisition parameters were chosen to be maximally similar across resolutions given two a priori constraints: 1) sufficient spatial coverage of the V1 ROI and 2) identical sampling frequency (TR) across resolutions.

T2*-weighted echo planar images (EPI) (TR/TE = 2000/22 ms, FA=90°) of the occipital lobe were acquired during visual stimulation using a 7 Tesla whole body scanner (Siemens, Erlangen, Germany) and a 32 receive channel head coil (Nova Medical, Wilmington, MA). Slices, oriented parallel to the calcarine sulcus (on a tilted axial plane), were acquired for 4 different spatial resolutions, i.e. 3 mm isotropic (FoV = 198 mm, matrix size 66 × 66, 37 slices, GRAPPA accel. factor 2), 2 mm isotropic (FoV = 200 mm, matrix size 100 × 100, 37 slices, GRAPPA accel. factor 3), 1.4 mm isotropic (FoV = 196 mm, matrix size 140 × 140, 32 slices, GRAPPA accel. factor 3) and 0.8 mm isotropic (FoV = 128 × 166.4 mm (AP × LR), matrix size 160 × 208, 32 slices, GRAPPA accel. factor 4). All EPI scans implemented ascending slice acquisition order and used a 10% inter-slice gap to minimize cross-slice excitation. The sequence for 0.8 mm isotropic resolution used a left-right phase encoding direction in order to avoid wrap-around artifacts, while all other sequences used anterior-posterior phase encoding. 121 volumes were acquired for each experiment run and 10 separate scans (one for each experimental run) were performed for each subject. An automatic positioning system (Siemens AutoAlign Head LS) was used to aid positioning of the field-of-view to the same volume in each scan for each subject similar to the procedure described in Dou

et al. [2014]. Online distortion correction [In and Speck, 2012] was applied to data from all the scans.

In order to aid co-registration of the small scan volume of the 0.8 mm acquisition with the structural image, an additional EPI acquisition was performed that used the same auto-alignment procedure, but with a 250×250 in-plane matrix and 57 slices (4 s TR). This setup increased the FoV in the axial plane to cover the full extent of the brain, while the 20 additional slices further increased the coverage along the inferior-superior direction. 60 volumes were acquired to improve image signal-to-noise ratio (SNR) by averaging across volumes. The resulting volume was used as an intermediate alignment target. Figure 3.2 illustrates the effect of distortion correction and the alignment quality of BOLD images to the respective structural images for two participants.

3.2.6 Structural imaging

T1 and T2-weighted structural images were acquired for all participants in a 3 Tesla Philips Achieva scanner equipped with a 32 channel head coil (refer to section 2.3.5).

3.2.7 Region of interest localization

As described in section 2.3, retinotopic measurements were performed using flickering checkerboard patterns. After retinotopic phase maps (polar angle and eccentricity) were generated, the V1 region was manually delineated on the cortical surface [following the procedure described in Warnking et al., 2002]. Surface reconstruction was performed using the default Freesurfer `recon-all` pipeline [Dale et al., 1999], using T1 and T2-weighted images as input. V1 delineations on the surface were projected back into a subject's individual volumetric space to generate a participant specific V1 ROI mask for the classification analyses.

3.2.8 Blood vessel localization

Susceptibility weighted (SW) imaging data [openly available from the *studyforrest* project Hanke et al., 2014] were used for localization of veins in V1 by utilizing the

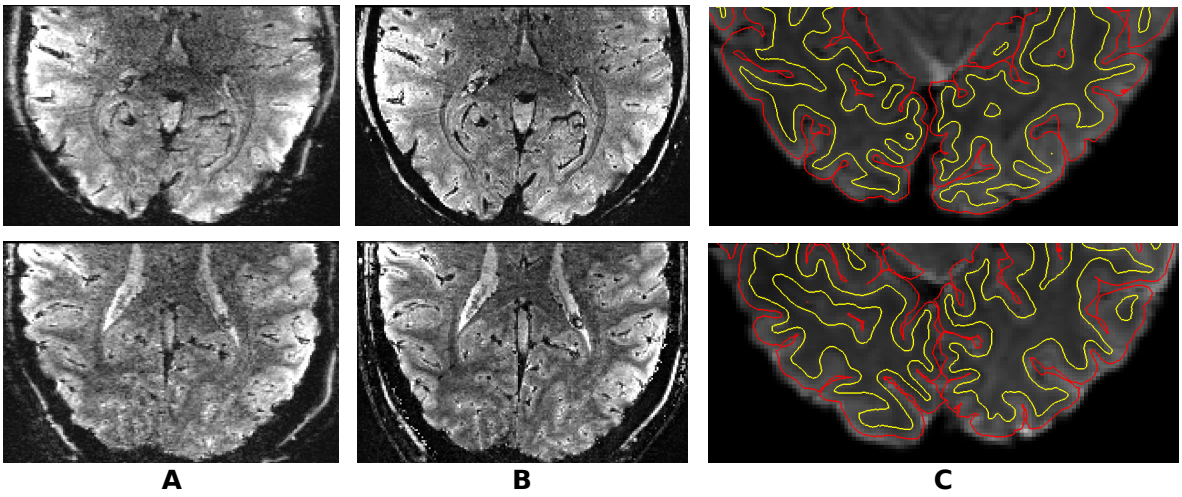

Figure 3.2: Alignment of EPI with structural data
The alignment of distortion corrected EPI functional data obtained at 7 Tesla to the structural data obtained at 3 Tesla from 2 subjects. (A) Uncorrected data from Siemens 7T Magnetom (B) Distortion corrected data [In and Speck, 2012] (C) Alignment of the EPI sequences acquired in 7T to the corresponding 3T Structural images. The white matter segmentation is shown with yellow lines and pial surface with red lines. The white matter and pial surface segmentations were performed on the structural data with Freesurfer and overlayed on the aligned EPI images to show the quality of the alignments.

difference in magnetic susceptibility of venous and neighboring non-venous tissues to improve contrast in venography [Liu et al., 2014]. These acquisitions were recorded in a 3 Tesla Philips Achieva scanner using a 3D Presto fast field echo (FFE) sequence (TR 19 ms, TE shifted 26 ms, flip angle 10°, bandwidth 217.2 Hz/px, NSA 2, SENSE reduction AP 2.5, FH 2.0). Susceptibility weighted images for every participant had 500 axial slices (thickness 0.35 mm, FoV 181×202×175 mm) and an in-plane acquisition voxel size of 0.7 mm reconstructed at 0.43 mm (512×512 matrix). The SW images of every participant consisted of seperate phase and magnitude components. The processing of the these components [similar to the procedure outlined in Haacke et al., 2004] are described in the following paragraphs.

Phase unwrapping Generally complex MR image acquisition can be expressed as $I = |I| * exp(\phi)$, where $|I|$ is the magnitude part and ϕ is the phase component of the

image. The phase image conveys several important information like field inhomogeniety, venous blood flow etc. But extracting the phase image from the measured complex image is non-trivial, because the any phase component beyond the range of $(-\pi, \pi]$ is wrapped back into the principal value range. So when the phase image is generated from the scanner it undergoes the *Phase Wrapping* process, as mentioned above. The phase image provided here is actually a phase wrapped image in which the wrapped phase is defined as $\Psi = W(\phi)$ where W is the wrapping operator. So for further processing, first, the phase components of the SWI scans were masked (using a brain mask derived from the magnitude component), and 3D phase unwrapped with PRELUDE [default settings; Jenkinson, 2003] from FSL [v5.0.9; Smith et al., 2004].

Contrast enhancement and localization of veins Similar to Haacke et al. [2004], the unwrapped phase image was spatially high-pass filtered using a mean 'box' filter kernel [65x65x65 voxels, as implemented in fslmaths; Smith et al., 2004]. The high pass filtered phase component $\varphi(x)$ was then transformed to a score $g(x)$ (value interval $[0, 1]$) using $g(x) = (\pi - \varphi(x))/\pi$ for $0 < \varphi(x) \leq \pi$ and 1 otherwise. These scores were multiplied 4 times with the original magnitude image, as suggested by Haacke et al. [2004], in order to enhance the contrast between venous and non-venous voxels. These contrast-enhanced images were suitably thresholded to perform segmentation of the venous voxels. For every participant, 2 different thresholds were chosen (60^{th} and 90^{th} percentile) and the blood vessel masks were constrained to individual V1 ROI. These were resliced into different acquisition resolutions using trilinear interpolation.

Separate MVP analyses were performed inside and outside the venous voxels (with variable threshold) in V1 to investigate their individual contributions at different acquisition resolutions across different threshold levels.

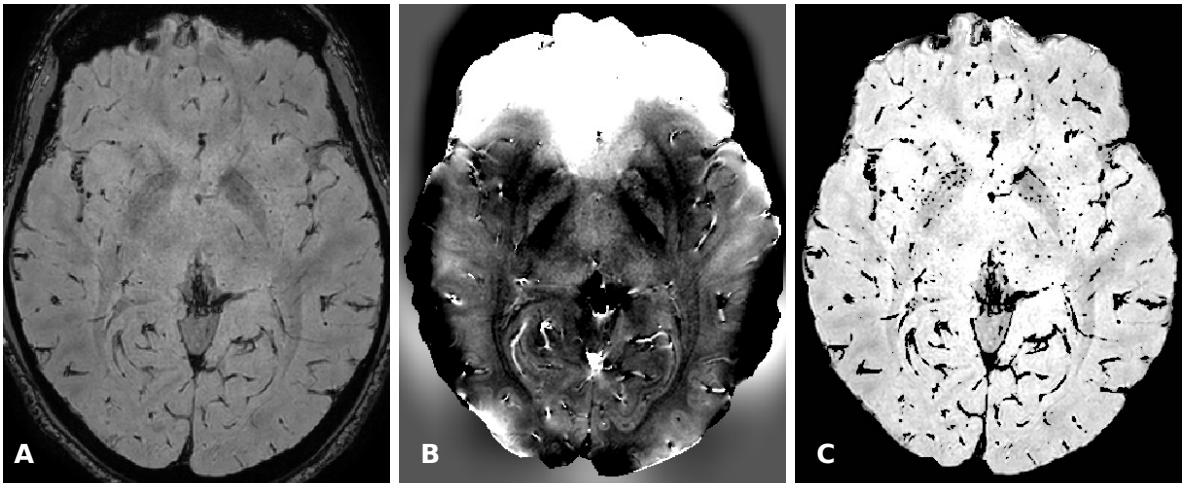

Figure 3.3: Localization of veins with SWI
(A) Magnitude image of SWI (B) High-pass filtered phase unwrapped image (C) Contrast enhanced masked image which is thresholded for vein localization

3.2.9 Orientation decoding analysis

MVP analysis for orientation decoding was performed with PyMVPA [v2.4.1; Hanke et al., 2009] on a compute cluster running (Neuro)Debian [v8.0; Halchenko and Hanke, 2012]. For feature extraction, BOLD fMRI time series from an individual experimental run were voxel-wise fitted to hemodynamic response (HR) regressors (boxcar function convolved with the canonical Glover HRF kernel [Glover, 1999] for each experimental condition using a general linear model (GLM). Additionally, the GLM design matrix included temporal derivatives of HR regressors, six nuisance regressors for motion (translation and rotation), and polynomial regressors (up to 2nd-order) modeling temporal signal drift as regressors of no-interest. GLM β weights were computed using the GLM implementation in NiPy [v0.3; Millman and Brett, 2007] while accounting for serial correlation with an autoregressive term (AR1). Lastly, separately for every run β scores were Z-scored per voxel. The resulting dataset for MVP analysis contained 40 samples (one normalized β score per condition per run) for each participant.

Linear support vector machines [SVM; PyMVPA's `LinearCSVMC` implementation of the LIBSVM classification algorithm; Chang and Lin, 2011] were used to perform a within-subject leave-one-run-out cross-validation of 4-way multi-class orientation classi-

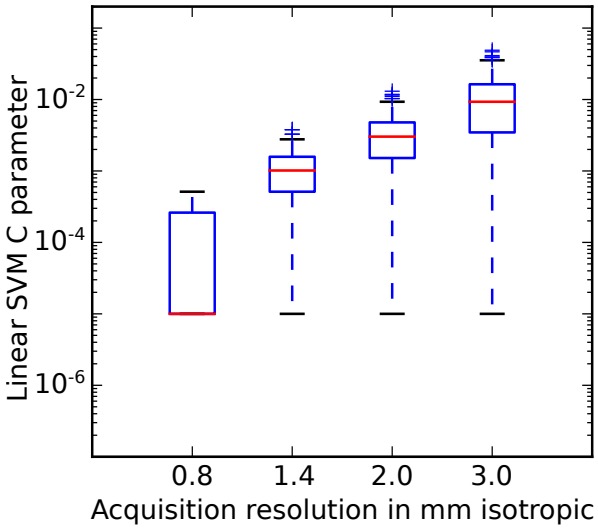

Figure 3.4: Range of tuned Linear SVM C parameters in the orientation decoding analysis across different resolutions.

fication. This method was selected based on its prevalence in the literature, not because of an assumed optimal performance in this context. This linear SVM algorithm has one critical hyper-parameter C that indicates the trade-off between width of the margin of the classifying hyperplane and number of correctly classified training data points. While it seems uncommon for neuroimaging studies to optimize this parameter for a particular application, I observed substantial variability in performance with varying number of input features. Consequently, I decided to tune this parameter using a nested cross-validation approach, where the training portion within each cross-validation fold was subjected to a series of leave-another-run-out cross-validation analyses in order to perform a grid search for the optimal C value (search interval $[10^{-5}, 5 \times 10^{-2}]$ in 200 equal steps). The "optimal" C value was then used to train a classifier on the full training dataset, which was subsequently evaluated on the data from the left out run. Reported accuracies always refer to the performance on the test dataset using the tuned C setting. Tuning of the C parameter was performed independently for each participant, resolution, and hemisphere. The ranges of tuned C parameters for all resolutions are illustrated in Figure 3.4.

3.2.10 Spatial filtering strategies

In order to investigate how signal for orientation decoding is distributed across the spatial frequency spectrum, two different strategies for volumetric spatial filtering of the functional imaging data were implemented.

Gaussian smoothing Similar to Swisher et al. [2010], I used Gaussian filtering prior feature extraction for MVP analysis to estimate the spatial scale of the orientation specific signal. In the following, the size of the Gaussian filter kernel is described by its full width at half maximum (FWHM) in mm. Individual filters were implemented using the following procedure: *Low-pass* (LP) 3D Gaussian spatial filtering was performed with the `image_smooth()` function in the nilearn package [Pedregosa et al., 2011]. *High-pass* (HP) filtered images for a particular filter size were computed by subtracting the respective LP filtered image from the original, unfiltered image. *Bandpass* (BP) filtering was implemented by a Difference-of-Gaussians (DoG) filter [Alink et al., 2013]. Filtered images were computed by subtracting the LP filtered images for two filter sizes from each other. For example, an image for the "4-5 mm" band was computed by subtracting the 5 mm LP filtered image from the 4 mm LP filtered image. It is important to note that, due to the nature of the filter, the pass-band of a DoG filter is not as narrow as the filter-size label might suggest. Figure 3.5 illustrates the attenuation profile of an exemplary 4-5 mm DoG filter. However, for compactness and compatibility for previous studies [*e.g.*, Alink et al., 2013] I are characterizing DoG BP filters by the FWHM size of the underlying LP filters. The respective *band-stop* (BS) filtered image were computed by subtracting the corresponding BP filtered image from the original, unfiltered image.

Because of its prevalence in standard fMRI analysis pipelines, spatial filtering was always applied to the whole volume, prior to any masking. However, as this procedure can potentially introduce signal from outside an ROI, particularly with large-sized LP filters, I also performed a supplementary analysis where filtering was restricted to the V1 ROIs in each hemisphere to prevent information propagation by smoothing (see

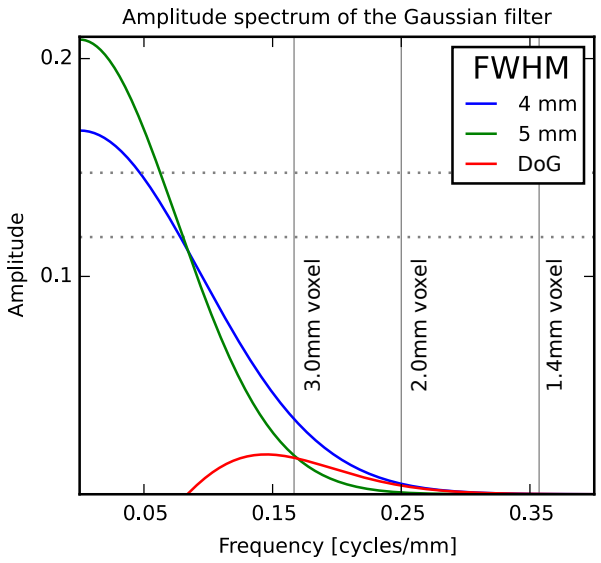

Figure 3.5: Illustration of the attenuation profile of a Difference-of-Gaussian (DoG) band-pass filter
The blue and green curve represent the profiles of Gaussian low-pass filters (4 mm and 5 mm respectively) in the frequency domain. Horizontal lines represent the -3 db points of the Gaussians. Band-pass filtering is implemented by subtracting the two low-pass filter outputs from each other. The profile of the resulting DoG band-pass filter is shown in red. Vertical lines show the Nyquist-frequencies for the three lowest resolutions in the study. The pass-band of this exemplary DoG filter (corresponding to an axis label "5 mm" in Figure 3.10 contains frequencies higher than what can be appropriately measured with a 3 mm acquisition.

supplementary material).

All spatial filtering procedures described above were volumetric, using 3D Gaussian kernels and ROI voxel selection was performed after spatial filtering with different Gaussian kernel widths on the entire volume. Though this 3D filtering procedure was being extensively used in previous studies like [Op de Beeck, 2010, Swisher et al., 2010], this approach can lead to information propagation from adjacent parts of the cortex, white matter and superficial vessels. Moreover, unconstrained 3D filtering does not respect the cortical folding pattern and, given a large enough filter, can smooth across sulcal boundaries, such as the two banks of the calcarine sulcus. This confounds filter width with the extent of the cortical region from which information is drawn. To avoid this problem, two additional spatial filtering approaches were implemented, namely

volumetric filtering restricted to the V1 ROI, and surface-based smoothing.

Volumetric filtering restricted to the V1 ROI Similar to the spatial filtering procedure performed in Alink et al. [2013], the voxel values outside the V1 ROI were considered to be missing values (NaN) instead of applying spatial filtering on the whole volume, prior to any masking. To eliminate a potential effect of smoothing across hemispheres with large Gaussian kernels, filtering was restricted to individual hemispheres. First, voxel values outside the left V1 ROI was considered to be NaNs and spatial smoothing was applied. The same procedure was applied to the right ROI, and then the smoothed left and right V1 ROI were combined to form the smoothed BOLD volume. The same nested cross validation approach was performed on the smoothed data.

Surface-based smoothing Freesurfer's `mri_vol2surf` function [Dale et al., 1999] was used for smoothing gray matter BOLD data on the cortical surface, while specifying the filter size with the `surf-fwhm` parameter. In the next step surface-projected data was mapped back into the BOLD volume using Freesurfer's `mri_surf2vol` function (trilinear interpolation, `fill-projfrac` parameter with range 0-1 in steps of 0.01). This procedure was performed for each hemisphere separately. Back projection into the volume was performed to maintain an equal number of input features for the decoding analysis. Subsequently, the same nested cross validation approach was performed on the smoothed data.

Spatial resampling to other resolutions, with and without Gaussian filtering A frequently expressed concern in the literature with respect to Gaussian smoothing is that a linear transformation does not actually remove high spatial frequency information [Alink et al., 2013, Kamitani and Sawahata, 2010]; instead, it merely implements a relative scaling of frequency components [see Misaki et al., 2013]. In order to investigate the potential impact of an irreversible frequency-domain transformation, I performed a

Fourier (FFT) based spatial frequency resampling, which destructively removes high-frequency components. Resampling BOLD fMRI data from one resolution to the other was implemented as a two-step procedure. In the following paragraphs, I describe the procedure using resampling from 0.8 mm to 3.0 mm resolution as an example, but the procedure was analogous for all resolution pairs.

First FFT-based spatial filtering was performed on the distortion corrected 0.8 mm data (see Figure 3.6A) using the scipy function `signal.resample()`. This removed the higher frequency components, but the voxel grid remained unchanged (in-plane matrix size (208, 160) with 32 slices). In the next step, linear resampling/reslicing was performed with nilearn function `resample_img()` to convert the FFT filtered image to the corresponding 3.0 mm voxel grid (see Fig. 3.6B for an example). Importantly, other than changing the voxel size, no further transformation, for example, to align a resampled image to the orientation of the corresponding native acquisition, were applied.

FFT resampling was also combined with subsequent Gaussian low-pass filtering in order to evaluate a suggestion by Freeman et al. [2013] that one way of testing the contribution of fine scale signals to orientation decoding is to compare high-resolution BOLD fMRI data down-sampled to conventional resolutions, with or without first removing high spatial frequency signals. For all spatial resampling analysis, with or without Gaussian filtering, all voxels in the respective V1 ROI masks were considered for multivariate decoding.

3.3 Results

3.3.1 Maximum orientation decoding accuracy

Effect of acquisition resolution and number of input voxels In order to determine the effect of acquisition resolution, I performed orientation decoding at all resolutions. Figure 3.7A shows the mean classification accuracy across participants and hemispheres as a function of acquisition resolution in the V1 ROI. In the set of tested

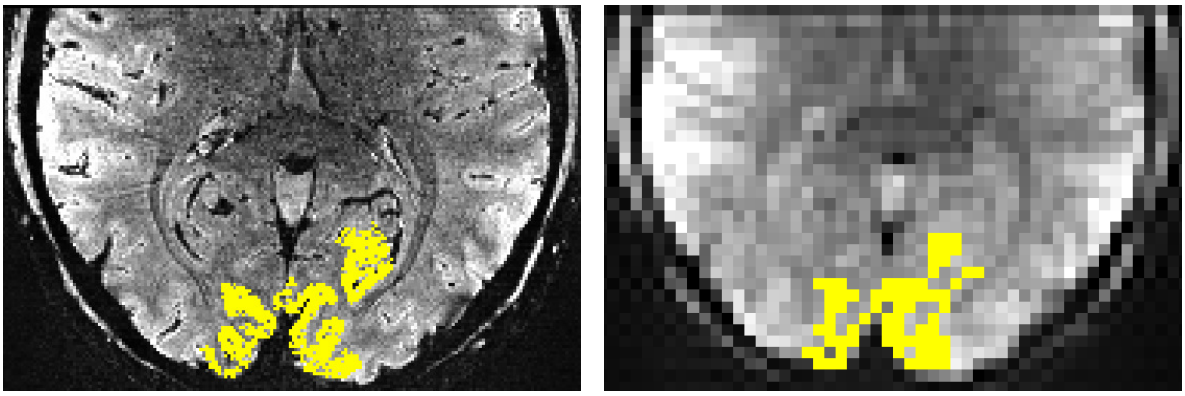

Figure 3.6: Resampling from 0.8mm iso to 3.0mm iso resolution
(A) Distortion corrected 0.8mm isotropic BOLD image overlayed with V1 ROI mask. (B) Removal of high-frequency components using scipy function signal.resample() overlayed with resampled V1 ROI mask (linear interpolation using scipy function ndimage.interpolation.zoom())

acquisition resolutions, I found the peak classification performance of 40.89% at 2 mm isotropic resolution.

In this analysis, the NULL events (unilateral stimulation events) were included in the GLM. Additionally I analyzed the data using two separate models for both hemifields, while excluding NULL events from the modeling. This resulted in an overall improved classification performance, but did not impact the structure of the relative performance differences between resolutions (0.8 mm: 32.32%, 1.4 mm:41.78%, 2.0 mm: 46.42%, and 3.0 mm: 40.17%). The orientation decoding performance in the ipsilateral V1 ROI gives an idea about the combined impact of potential interhemispheric cross-talk and random correlations of the stimulus sequence between hemispheres. The ipsilateral accuracies show similar trend as the contralateral accuracies but are substantiaally lower. The ipsilateral accuracies for 1.4 mm and 2 mm resolution show poor decoding performance (¡ 30%) and the 0.8 mm and 3 mm decoding accuracies are at chance level.

For the above analysis, all voxels in the respective V1 ROIs were used. As the number of voxels in a 0.8 mm V1 mask was substantially higher than those in a 3.0 mm

	V1 Region of Interest				Venous voxels in V1 for two thresholds			
	Left hemisphere		Right hemisphere		$>60^{th}$ percentile		$>90^{th}$ percentile	
Resolution	#voxels	std	#voxels	std	#voxels	std	#voxels	std
0.8 mm	7312	1912	7683	2556	1148	446	287	111
1.4 mm	2084	626	2169	710	518	186	130	47
2.0 mm	883	273	898	311	231	84	58	21
3.0 mm	324	94	327	104	105	36	26	9

Table 3.1: V1 ROI size
Average number of voxels for both hemispheres with standard deviation across participants. The four rightmost columns indicate the number of voxels within the ROI that are considered to be intersecting veins for two different thresholds (the 40% of voxels with the highest volume fraction of blood vessels; and the same for the top 10% voxels; see Figure 3.13 for an illustration).

V1 mask (Table 3.1) and the number of input features/voxels can impact the classification performance, I repeated the analysis, but held the number of voxels constant across participants and resolutions (50, 100, 125, and 150 voxels). Voxel sub-selection was done randomly, and the analysis was repeated 100 times with a new random selection of voxels. Figure 3.7B shows that a constant and smaller number of input voxels had a negative effect on classification performance. Classification performance was better with 2.0 mm and 3.0 mm data as compared to 0.8 mm and 1.4 mm data.

Time-series signal-to-noise ratio (tSNR) It has been shown that overall contrast-to-noise ratio (OCNR) is a factor that impacts classification performance [Chaimow et al., 2011]. According to Chaimow et al. [2011] OCNR is a measure is proportional to contrast range and the square root of the number of voxels and is inversely proportional to the noise level. The noise level was calculated as the inverse of time course signal-to-noise ratio, which in turn depends on voxel size [Triantafyllou et al., 2005]. In this study, tSNR is modulated across acquisition resolutions due to differential impact of technical/thermal and physiological noise components. In order to characterize this impact, I computed tSNR for each voxel as the ratio of mean signal intensity across

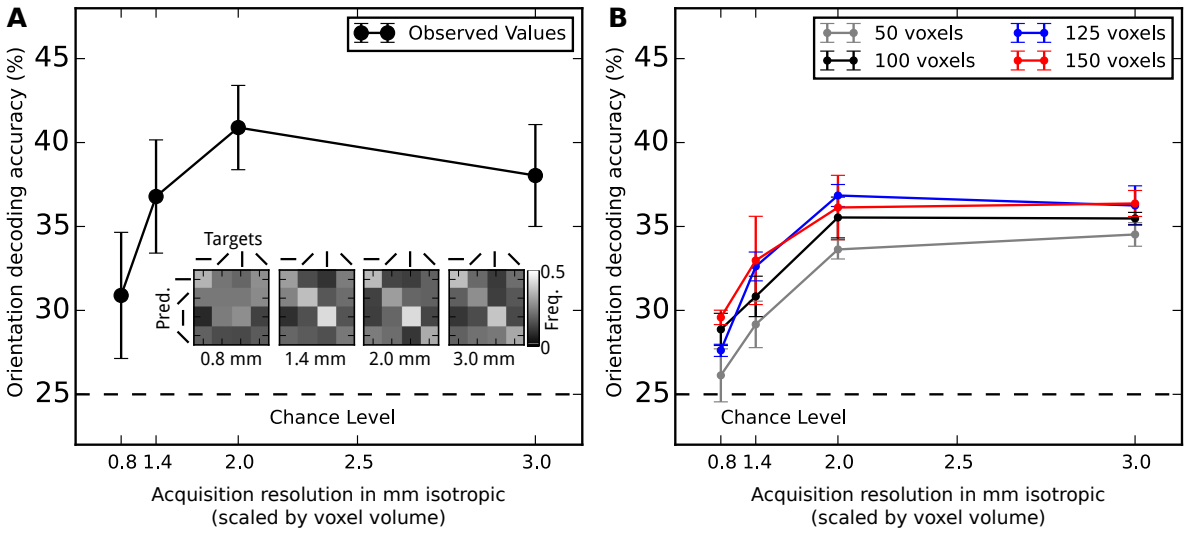

Figure 3.7: Orientation decoding accuracy on spatially unfiltered data **(A)** Orientation decoding accuracy on spatially unfiltered data as a function of acquisition resolution in the whole contra-lateral V1 ROI. Error bars show the standard error of the mean (SEM) across 7 participants averaged across hemispheres. Chance level accuracy (25%) is indicated as a horizontal dashed line. Classification performance is detailed in confusion matrices for each resolution depicting the frequency of correct classification for each combination of prediction and target values. **(B)** Analog to (A), but with a constant number of input voxels across resolutions. 50, 100, 125, or 150 voxels were selected at random from the the whole contra-lateral V1 ROI for the classification analysis. Selection was repeated 100 times. Error bars show SEM across repetitions. Upper range limit of 150 voxels was determined by the ROI with the least number of voxels at 3 mm resolution.

all time points after polynomial detrending (1st and 2nd order; analog to preprocessing for MVP analysis) of scanner drift noise and the corresponding standard deviation. Voxel-wise tSNR was averaged across all experiment runs. For a tSNR estimate of the whole ROI, I averaged this score across all voxels. The relationship of voxel volume and tSNR in the empirical data can be well explained by the following model [Triantafyllou et al., 2005]:

$$\text{tSNR} = \kappa V / \sqrt{1 + \lambda^2 \kappa^2 V^2}$$

where V is the voxel volume, κ is the proportionality constant, and λ is the magnetic field strength independent constant parameter with λ=0.0117, κ=22.74 (R^2=0.95) The estimated asymptotic tSNR limit of ≈85 ($\frac{1}{\lambda}$) is similar to the report of Triantafyllou

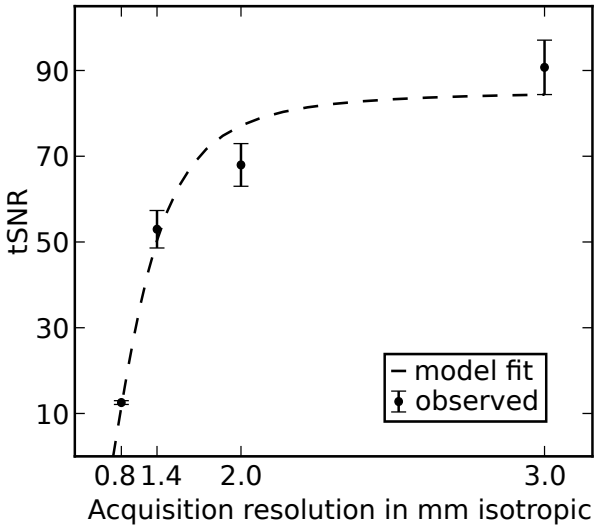

Figure 3.8: Temporal signal-to-noise ratio (tSNR) as a function of voxel volume
The observed data are represented by dots and the error bars represent the SEM across subjects. The dashed line shows the fit to the following model $\text{tSNR} = \kappa V/\sqrt{1+\lambda^2\kappa^2V^2}$ similar to the report of Triantafyllou et al. [2005]

et al. [2005] for 7 Tesla acquisitions and is reached around 2.5 mm acquisition resolution (see supplementary Figure 3.8).

Figure 3.9A illustrates the non-linear relation of tSNR and orientation decoding accuracy. I observe a substantial drop in accuracy when decreasing resolution from 2 mm to 3 mm, despite a further increase in tSNR. This non-linearity was not observed by Gardumi et al. [2016], who only reported a positive trend for the correlation between decoding accuracy and tSNR, based on a single acquisition (1.1 mm resolution with comparable tSNR of ≈32, and other resolutions being generated by reconstructing k-space data to lower resolutions).

BOLD signal change Another potential source of differences in orientation decoding accuracy across resolutions are BOLD signal amplitude differences due to, for example, differential impact of a partial voluming effect [see Alink et al., 2013, Tong et al., 2012]. In order to quantify this effect, I calculated mean percentage BOLD signal change in response to any flickering orientation stimulus across resolutions using FeatQuery in

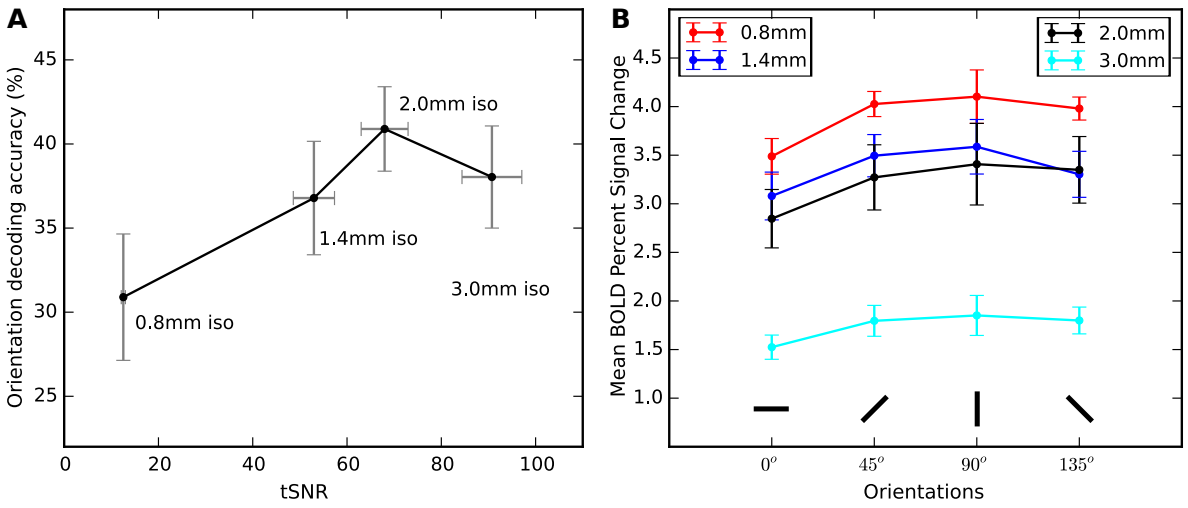

Figure 3.9: Temporal signal-to-noise ratio and Percentage BOLD signal change
(A) Temporal signal-to-noise ratio (tSNR) as a function of resolution (voxel volume). Error bars show the SEM for tSNR and accuracy across subjects and hemispheres. (B) Estimated BOLD signal change by orientation for all resolutions. Maximum pairwise signal change difference is observed for the cardinal directions 0°and 90°. This pattern is congruent with the confusion plots in Figure 3.7A.

FSL [v5.0.9; Smith et al., 2004]. Similar to preprocessing in MVP analysis, no spatial smoothing was performed before calculating the percentage signal change. In order to obtain comparable percentage signal change across resolutions, I obtained a mask of all responsive V1 voxels ($z > 2.3$ with $p < 0.05$ default parameters of FSL FEAT) in 0.8 mm data for every subjects [Swisher et al., 2010, similar to Figure 3]. The responsive V1 voxel mask obtained at 0.8 mm was resliced into 1.4 mm, 2.0 mm and 3.0 mm resolutions. Percentage signal change was calculated with FeatQuery within these masks. I found that the mean percentage BOLD signal change was the highest for 0.8 mm resolution (0.8 mm: 4.51%, 1.4 mm:3.92%, 2.0 mm: 3.73%, and 3.0 mm: 2.05%).

In addition, it may also be that particular orientation stimuli elicit stronger BOLD responses than others [e.g., a grating along the cardinal orientations; Furmanski and Engel, 2000]. In order to test for a differential effect and a possible interaction between orientation and acquisition resolution, I computed a 2-factor (orientation and resolution) within-subject ANOVA for the estimated BOLD signal change from all 7

subjects (Figure 3.9B). There was a significant main effect of acquisition resolution ($F(3, 18)$=32.99, p=1.58e-07) and orientation, $F(1, 6)$=12.31, p=0.01), and significant interaction between the factors, resolution, and orientation ($F(3, 18)$=4.27, p=0.01). However, pairwise t-test (Bonferroni-corrected) did not reveal significant differences between any two particular orientations.

Impact of head motion on decoding accuracy Head motion is a likely factor to impact decoding accuracy. In order to evaluate this effect, I calculated a head motion index suggested by Alink et al. [2013] for every participant and acquisition resolution. Inline with the findings of Gardumi et al. [2016], I found a consistent, but non-significant trend towards a negative correlation between head motion and decoding accuracy across acquisition resolutions. (0.8 mm: r=-0.45, p=0.3; 1.4 mm:r=-0.64, p=0.11; 2.0 mm: r=-0.68, p=0.09; 3.0 mm: r=-0.23, p=0.6).

3.3.2 Filtering Strategies

Impact of volumetric Gaussian smoothing Figure 3.10 A-D show the impact of Gaussian filtering on the classification performance for data from all four acquisition resolutions. LP spatial filtering is most commonly performed as a noise reduction step in fMRI data pre-processing. The classification performance achieved on HP filtered data of the same filter size is an indication of the amount of usable information removed by LP filtering. Classification performance on BP filtered data indicates whether usable information is present in a particular band of spatial frequencies. Likewise, band-stop performance indicates the presence of usable information anywhere, except in a particular band.

Except for 0.8 mm and 1.4 mm data, LP filtering did not aid classification performance, relative to the performance on unfiltered data. For all resolutions, except for 0.8 mm, the best performance was achieved by LP filtering with kernel sizes no larger than 3 mm FWHM. Peak performance on HP filtered data was achieved for filter sizes

larger than 9 mm FWHM, except for the 0.8 mm acquisition resolution. BP filtering yielded peak performance for all acquisition resolutions in the in the range of ≈5-8 mm, using DoG BP filters with a 1 mm difference in the FWHM size of the underlying LP filters. Classification performance of BS filtered data remained above-chance for all spatial frequency bands. The BS performance curve initially follows the LP performance for small filter sizes, but resembles the HP performance for larger filter sizes.

Impact of alternative spatial filtering procedures Figure 3.11 E-H shows the performance of orientation decoding following low-pass, high-pass, band-pass, and band-stop surface based spatial filtering. The results of surface based smoothing were similar to those of the volumetric Gaussian filter, but the decoding accuracy did not decrease rapidly with greater filtering. The band pass filtering peak was present at ≈5-8 mm but less pronounced more evenly sloped than what was obtained from volumetric filtering. This result is in congruence Swisher et al. [2010]

The results of the volumetric filtering analysis restricted to the V1 ROI are highly similar to the results of the unconstrained filtering prior masking (Fig. 3.11 A-D).

Impact of spatial resampling to other resolutions, with and without Gaussian smoothing As an alternative approach to Gaussian LP filtering for simulating a resolution reduction, data acquired in a particular resolution were resampled (FFT-based transformation) to all other resolutions and classification analysis was performed with and without additional prior Gaussian LP filtering, as suggested by Freeman et al. [2013].

Decoding performance on down-sampled data was lower than the accuracy obtained from data recorded in the respective native resolution. Gaussian LP filtering prior to down-sampling generally did not make the decoding accuracies better than that of the native resolution data.

Data acquired at 2.0 mm and 3.0 mm resolutions, showed a general trend towards

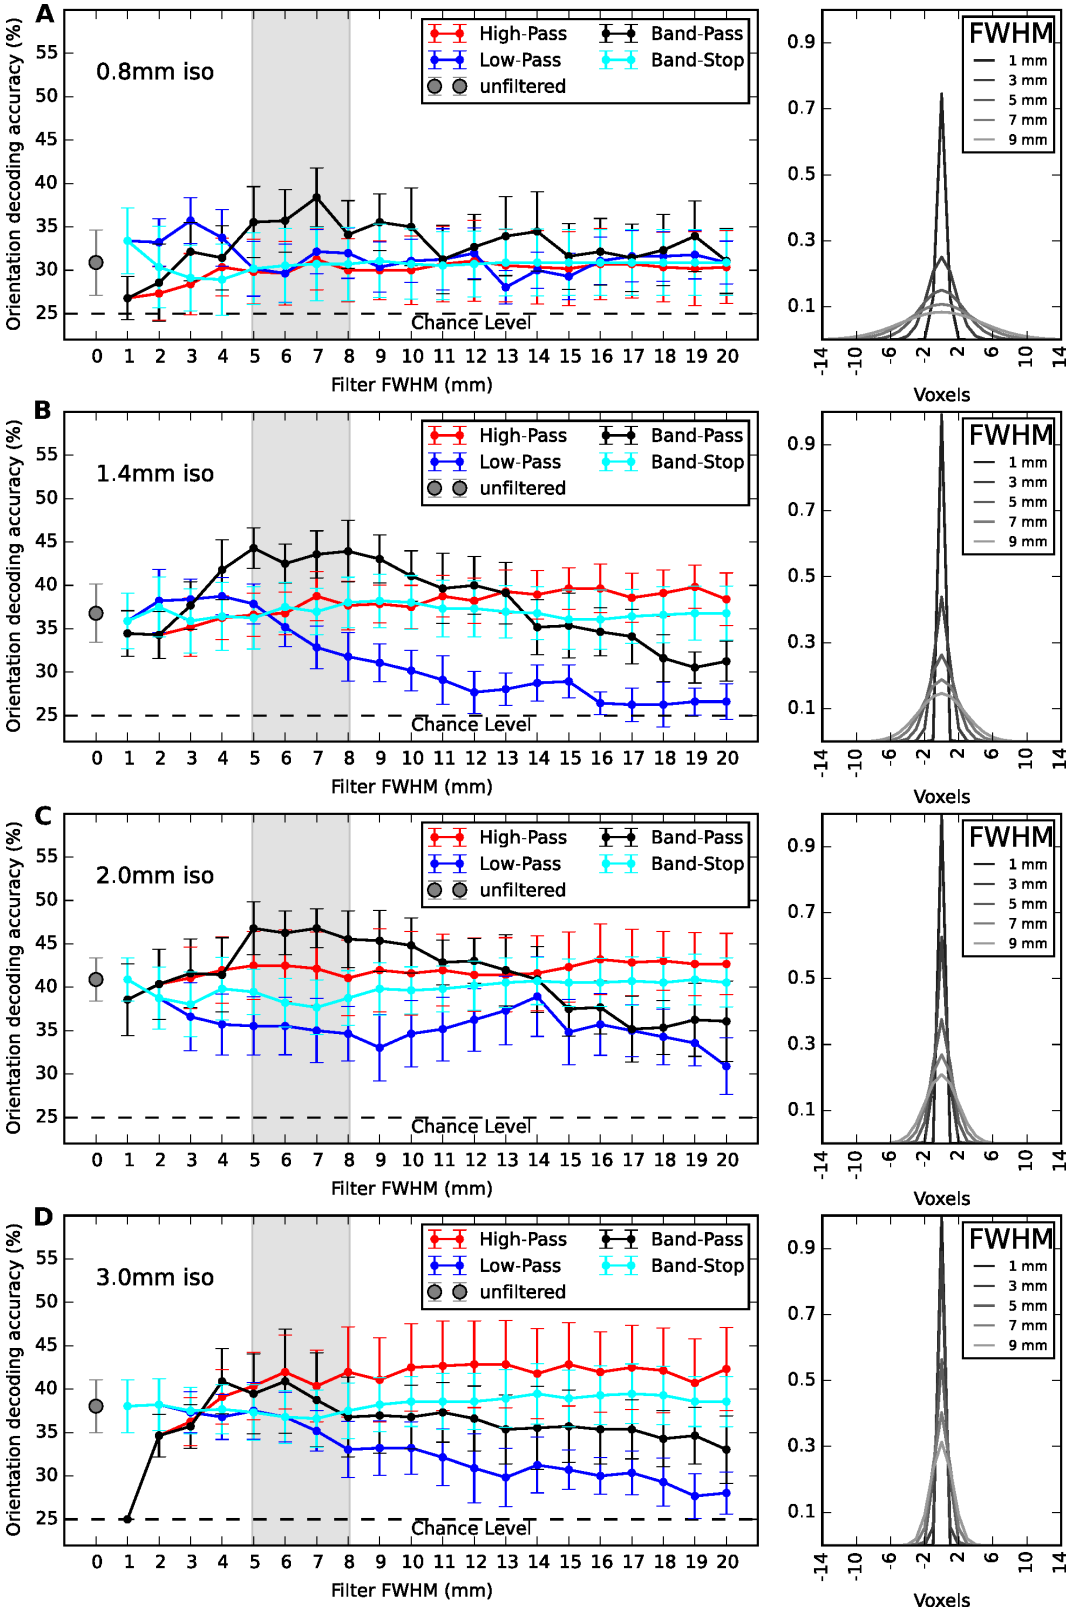

Figure 3.10: Effect of volumetric spatial filtering on orientation decoding Orientation decoding accuracies for all acquisition resolutions (increasing acquisition voxel size from top to bottom) and levels of spatial high-pass, low-pass, band-pass, and band-stop Gaussian filtering. Panels on the right visualize the size of selected Gaussian filter kernels with respect to the voxel size at each resolution. FWHM values for band-pass and band-stop filters refer to the corresponding 1 mm band to the closest smaller filter size (e.g., 5 mm refers to the 4-5 mm band).

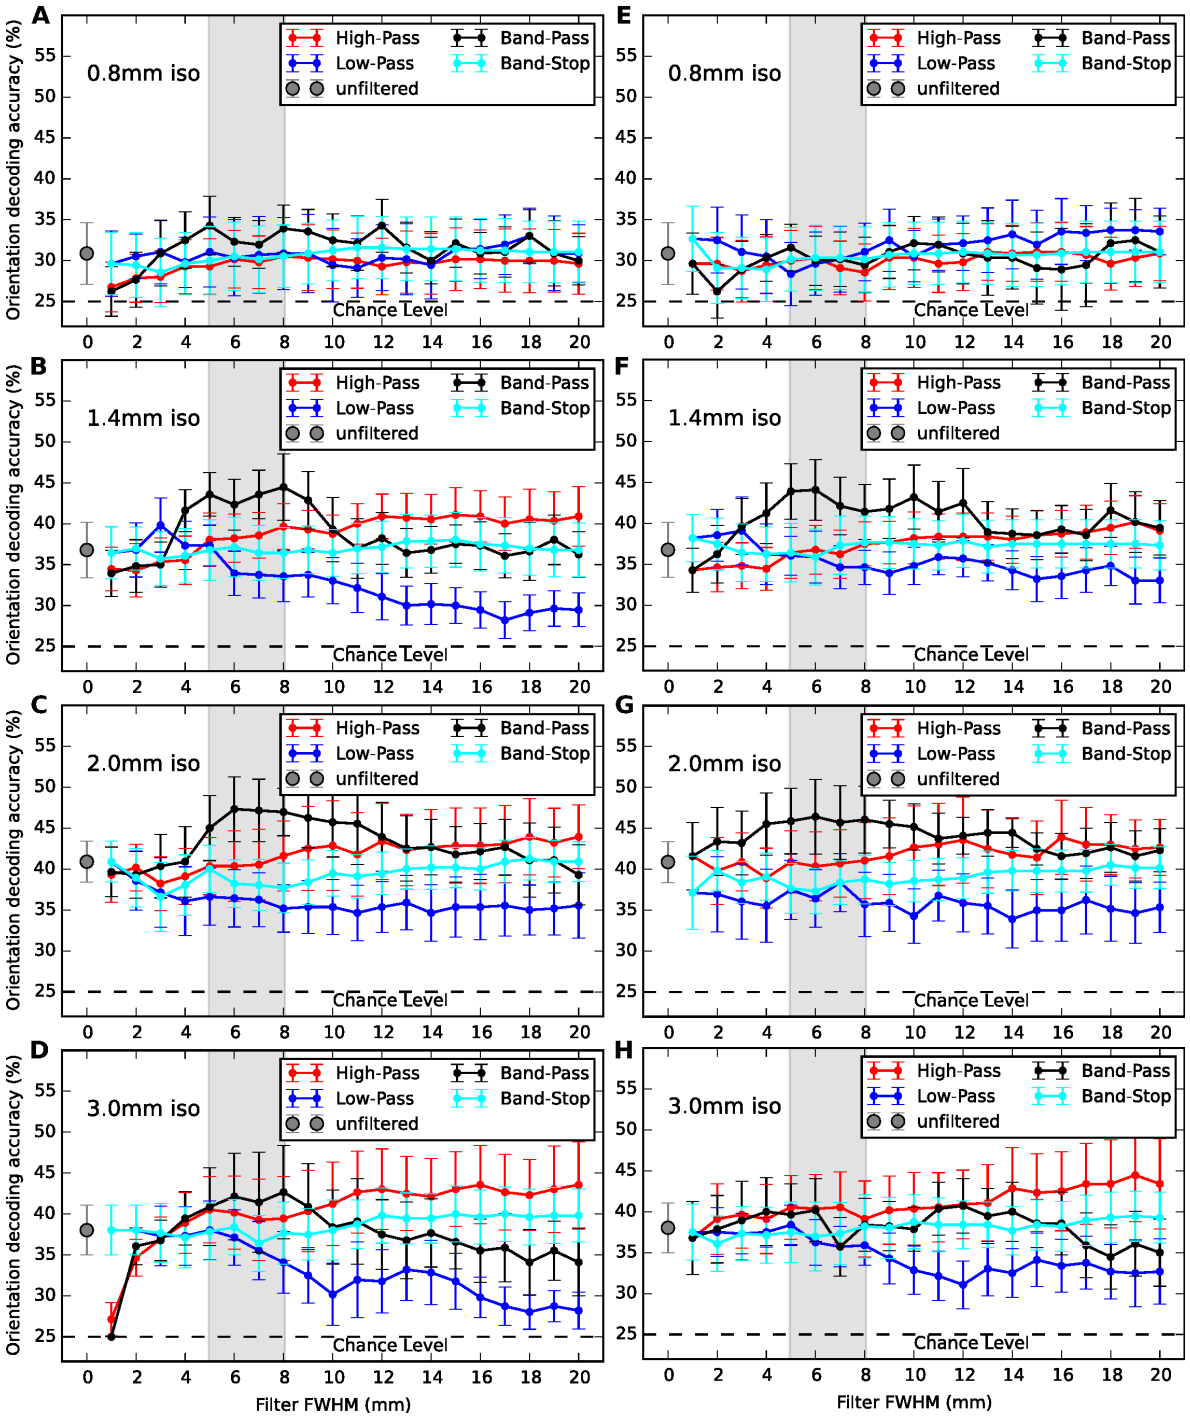

Figure 3.11: Results of alternative spatial filtering procedures
Volumetric spatial filtering restricted to V1 ROI (A-D), cortical surface-based smoothing (E-H).

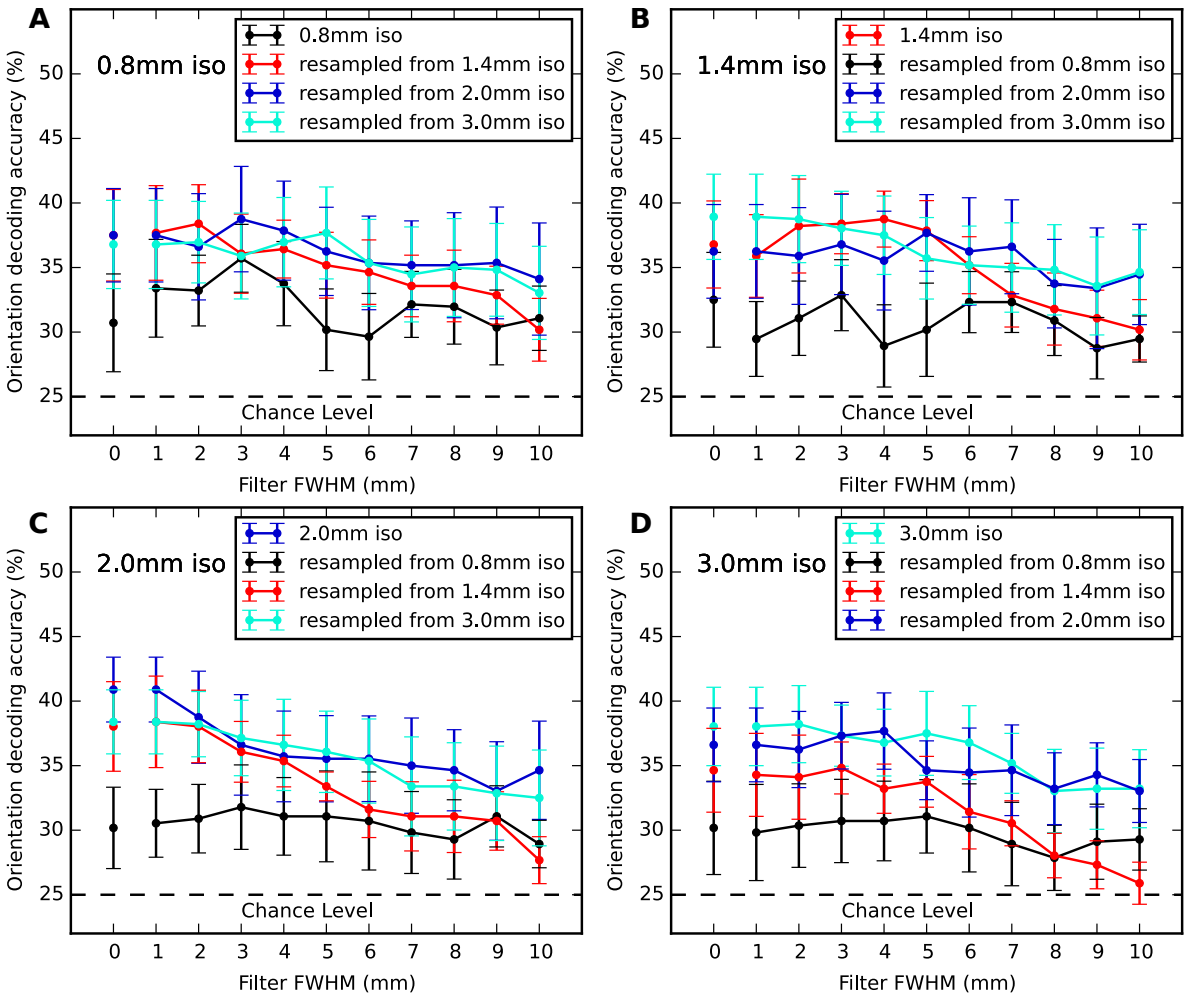

Figure 3.12: Orientation decoding performance on fMRI data resampled to other spatial resolutions
Resampling operation was performed with and without different levels of prior low-pass Gaussian spatial filtering. Recording high-resolution data with subsequent spatial down-sampling showed a consistent trend of lower classification accuracy compared to the native resolution acquisition, with or without prior Gaussian low-pass filtering of any tested kernel size.

better performance after resampling (up-sampling or down-sampling) compared to the corresponding native acquisition resolution, even with prior Gaussian LP filtering of different kernel sizes. 0.8 mm data consistently showed low decoding accuracy when resampled to any other resolution with or without Gaussian filtering.

3.3.3 Aliasing

In the case of frequency aliasing, a (spatial) source frequency is aliased into a lower frequency when the sampling frequency is too low (Nyquist-Shannon sampling theorem). If aliasing occurs, the apparent frequency is dependent on the sampling frequency. I investigated via BP filtering which frequency bands were most informative for orientation decoding across all acquisition resolutions using Gaussian BP filtered data (Figure 3.10 A-D; orange curves). Peak accuracy was consistently located in the ≈5-8 mm bands (highlighted range).

3.3.4 Vascular contribution to orientation decoding

Orientation decoding was performed inside and outside the vein localizer mask in order to evaluate the availability of orientation discriminating signal in the vascular system. Two different, arbitrary thresholds were used to classify voxels as intersecting vs. non-intersecting with veins, based on the co-registered and re-sliced vein mask: the top 40% and top 10% of voxels with the highest value after realignment and reslicing to the target resolution with trilinear interpolation. The resulting number of voxels is presented in Table 3.1.

Decoding accuracy was computed inside and outside the vein mask within the V1 ROI. Analyses outside the vein mask were performed twice: once for the entire region and again for a subset of voxels that was constrained to the number of voxels inside the vein mask for the corresponding resolution. In the latter case, the analysis was repeated with a new random voxel selection 100 times.

Figure 3.13A (right panel) shows that voxels with the highest venous content in their volume still yield above change decoding performance. The performance drop for the two lowest resolutions between the two vein mask thresholds may be explained by the low number of input features going into the classification at high threshold (compare Figure 3.13A, left panel). At 0.8 mm, the 10% most venous voxels yield the same decoding performance as the rest of the V1 ROI combined (Fig. 3.13, middle

panel), and noticeably more than a corresponding number of randomly samples non-venous voxels (Fig. 3.13A, middle panel). Similar results can be observed for the 1.4 mm resolution.

3.4 Discussion

In order to investigate the effect of acquisition resolution and spatial filtering on the decoding of visual orientations from primary visual cortex, I measured ultra-high field 7 Tesla fMRI data in four different resolutions from seven participants. Linear SVM classifiers were trained to classify voxel patterns of regression weights of hemodynamic response models for the visual stimulation with four different oriented gratings. Cross-validated classification accuracy was used as quality metric.

The overall classification accuracies reported here are deceptively low (peaking at 40-50% with a theoretical chance-level performance of 25% for the 4-way classification analyses employed in this study). Other decoding studies in the literature have often used binary classification paradigms [for example, Alink et al., 2013, Chaimow et al., 2011] or reported average pairwise accuracy for classification performance results like [e.g., Op de Beeck, 2010, Kamitani and Tong, 2005]. Converted into average pairwise binary accuracies, the results reported here range from 55% to 70% (for 0.8 mm and 2 mm respectively, each accuracy corresponding to an analysis of the full V1 ROI and with no additional smoothing; see Figure 3.7A; theoretical chance-performance: 50%), hence accuracies are of the same magnitude as in other studies [see, for example, Alink et al., 2013, Haynes and Rees, 2005]. In addition, some studies like Swisher et al. [2010] also reported similar unfiltered accuracy results (≈50%) in a 4-way classification analysis with 0°, 45°, 90°, and 135°gratings with much longer stimulation time (a block design of 18s of block duration and 8 blocks/run). Therefore, I conclude that the overall quality of the present data is comparable to that of previous studies, and that the results presented here can be used to address open questions regarding the impact of data acquisition and spatial filtering parameters on the decoding of orientation from

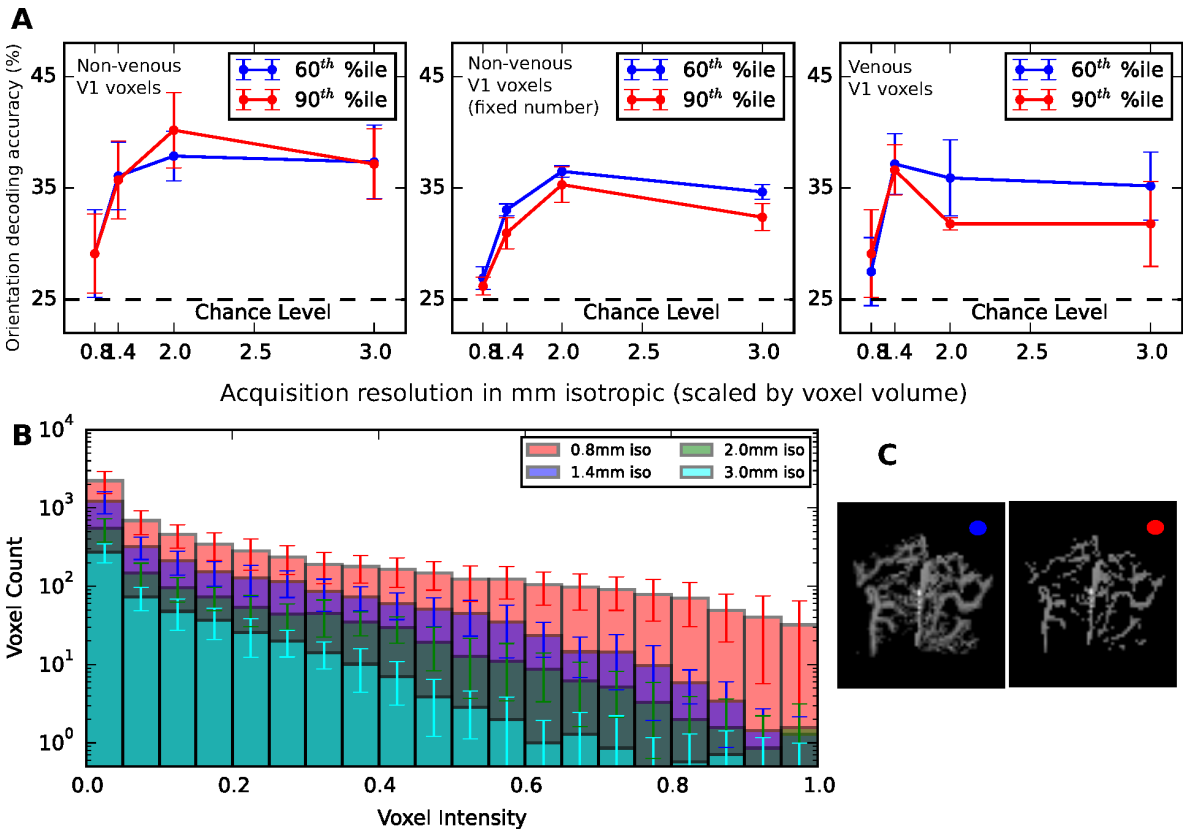

Figure 3.13: Vascular contribution in orientation decoding

(A) Decoding accuracy was computed inside and outside the vein mask within the V1 ROI. The vein masks obtained from susceptibility weighted imaging were thresholded at two different levels i.e. 60 percentile(%ile) and 90 percentile(%ile). The panel on the left shows the performance of the entire V1 ROI outside the vein mask (non-venous voxels) for the two different thresholds. Orientation decoding accuracy on V1 voxels restricted to the veins mask (venous voxels) is shown on the right panel. The middle panel depicts the decoding performance of a fixed number of non-venous voxels, the number of voxels being equal to the number of venous voxels in the right panel corresponding to every resolutions and thresholds. The dashed horizontal lines indicate the chance performance. (B) Trilinear interpolation was used to reslice the vein mask to all four target resolutions. The histogram shows the distribution of mask voxel intensities corresponding to the volumetric fraction of "vein voxels" in the high-resolution vein mask (voxel count axis in log-scale). (C) Axial maximum intensity projection of the vein mask of one participant resliced to the 0.8 mm resolution; illustrates the two chosen thresholds. The color indicator correspond to the curves depicted in panel A.

the early visual cortex. Moreover, in this experiment I did not use a univariate feature selection approach to define the "visually responsive" voxels in V1 (for example a GLM contrast) in an attempt to improve the decoding accuracy. Studying the potential impact of such an approach is left to a future study.

Optimal acquisition resolution Among the four tested acquisition resolutions, the highest decoding accuracy was achieved with a 2 mm resolution (Figure 3.7A). This result is congruent with a simulation study by Chaimow et al. [2011] that analyzed the impact of anatomical and physiological properties of primary visual cortex, as well as technical parameters of BOLD fMRI acquisition on the accuracy of decoding the stimulated hemifield from signal sampled from ocular dominance columns. The aforementioned study included a number of predictions for choosing optimal voxel size and number of input voxels to maximize decoding accuracy for 3 Tesla fMRI [see Figure 6 in Chaimow et al., 2011] that show a striking similarity to the results presented here (Figure 3.7). For 3 Tesla fMRI, Chaimow et al. [2011] showed that peak decoding accuracy is achieved between around 3 mm in-plane voxel size for ocular dominance. Given that the profile of orientation columns has higher spatial frequency compared to ocular dominance columns [Obermayer and Blasdel, 1993] and the BOLD PSF at 7 Tesla is considerably smaller compared to 3 Tesla [Engel et al., 1997b, Shmuel et al., 2007] a higher optimal resolution was to be expected for this study, and this hypothesis is supported by my results. This finding is also inline with a recent study by Gardumi et al. [2016] showing that optimal decoding accuracy of speaker identity, or phonemes, from auditory cortex BOLD patterns could be achieved with an effective voxel size of 2.2 mm (acquisition resolution was 1.1 mm and target resolution was achieved by reconstructing k-space data to a lower resolution).

Superior decoding performance at 2 mm could still be observed even when the number of input voxels for classification was held constant across resolutions, although the performance differences between resolutions are reduced (Figure 3.7B). The ratio of

input features (voxels) and the number of observations (fixed in this study) is a critical factor for the training of a classification model, as with increasing dimensionality the sampling of the feature space becomes sparser, and, consequently, the estimated decision surface suffers from increased uncertainty [curse of dimensionality, Bellman, 1961, after Friedman et al. 2001] In this study, the number of voxels in the ROI varies by a factor of >20 from the lowest to the highest resolution (Table 3.1). The pattern of decoding accuracy differences when using the full ROI vs. a constant number of voxels across all resolutions could indicate that ≈700 input voxels (size of the ROI at 2 mm) represents the optimal trade-off between the number of observations and input voxels, given the noise in the data and the fixed number of observations in this study.

Moreover, the present data suggest, in line with Chaimow et al. [2011], that temporal signal-to-noise-ratio, an indicator of temporal signal stability, is a critical factor for optimal decoding accuracy (Figure 3.9A). I also checked the effect of mean percentage BOLD signal change on decoding accuracy across different acquisition resolutions. Though the overall BOLD signal change amplitude in 0.8 mm data (4.51%) was higher than that in 2.0 mm data (3.73%), the decoding performance was better in the 2.0 mm data. In fact, 0.8 mm data had the highest percentage of BOLD signal change but showed the lowest decoding accuracy among all resolutions. An 'oblique effect' has been described in the literature in that cardinal orientations elicited higher activation changes than oblique orientations of circular gratings Furmanski and Engel [2000]. The reverse, higher activation for oblique than cardinal orientations, was found by Swisher et al. [2010], who used the same kind of hemifield gratings as in the present study. my pattern diverges from both previous results showing activation that was lowest for 0° and highest for 90° orientations, with oblique orientations in between. A possible explanation may be collateral summation of iso-orientation neurons, because the oriented lines were longest in the 90° stimuli, shortest in the 0° stimuli and of in-between length in the oblique orientations. However, this cannot explain the differences between the Swisher et al. [2010] data and ours, so this remains a speculative interpretation.

Optimal low-pass filter size Gaussian spatial LP filtering is one of the most common preprocessing steps for fMRI data analyses. However, the present findings indicate that explicit spatial LP filtering, in addition to the implicit spatial filtering due to inherent motion, and the effect of head movement correction algorithms is generally not beneficial for orientation decoding (Figure 3.10). Only for resolutions higher than 2 mm does additional spatial smoothing with 2-3 mm FWHM show a tendency for improved decoding accuracy. This suggests that, given a resolution, a spatial smoothness equivalent to a Gaussian kernel size of ≈2 mm FWHM is optimal. This is congruent with the observation of overall lower decoding accuracies for 3 mm scans and is in line with the prediction of optimal acquisition resolution between 2 mm and 3 mm as presented above.

Moreover, spatial down-sampling is not beneficial for orientation decoding either. As shown in Figure 3.12 (0 mm data points, corresponding to no Gaussian smoothing), orientation decoding on down-sampled data never outperforms the decoding on data natively recorded in the corresponding resolution (as for example, in the 2.0 mm panel of Figure 3.12, the 0.8 mm and 1.4 mm downsampled data performed lower than native 2.0 mm data).

Spatial characteristics of orientation specific signals The analysis of individual spatial frequency bands via BP filtering (Fig. 3.10) revealed that orientation-related signal is present in a wide range of spatial frequencies as indicated by above-chance decoding performance for nearly all tested bands. However, a drop in decoding accuracy can be observed across all resolutions for bands with a 12 mm FWHM (or larger) Gaussian kernel as the smaller kernel in the LP filter pair used for BP filtering.

Freeman et al. [2013], states that it is still an open question whether fMRI can reflect signals originating from sampling random irregularities in the fine-scale columnar architecture (spatial scale ≈1 mm). This study also suggests that given a columnar architecture in the human visual cortex [Adams et al., 2007], BOLD fMRI measurements

at conventional resolution ≈2 mm iso might reflect a combination of fine-scale and coarse-scale (spatial scale ≈10 mm) contributions. Similarly, I can interpret the present results such that the orientation-related signal in the BOLD fMRI data is spatially broadband in nature, includes both high spatial frequency components, as well as large-scale biases. On one hand the highest decoding accuracy was recorded at 2 mm iso resolution, and low pass filtered components generated above chance accuracies beyond 10 mm FWHM Gaussian smoothing [similar to Op de Beeck, 2010]. These observations point to the fact that the low frequency components provide orientation specific signals. On the other hand I found that for DoG BP filters (Gaussian kernel sizes of 4 and 5 mm FWHM and larger, decoding performance on BP filtered data was better than the LP filtered components in all acquisition resolutions. This phenomena shows that low spatial frequency fMRI components also contribute to noise.

According to Freeman et al. [2013], a convincing proof of fine-scale signals (≈1 mm, according to the definition by Freeman et al.) underlying the ability to decode orientations would be a comparison between decoding accuracies after down-sampling high-resolution measurements to conventional scanning resolutions, with and without prior removal of the columnar-scale contributions. To test this hypothesis, I did FFT based resampling of the BOLD fMRI data from their native resolution into all three alternative resolutions with or without low frequency components (Fig. 3.12). I generally observe a drop in accuracy after down-sampling data from my two highest resolutions (0.8 mm and 1.4 mm), regardless of the presents of prior LP Gaussian filtering (except for a singular slight increase in performance when resampling 0.8 mm to 1.4 mm data without prior LP filtering). From these findings I conclude that the orientation-related signal used for decoding is unlikely to comprise of low-frequency components alone. This conclusion is in line with Swisher et al. [2010] who also reported that "majority of orientation information in high resolution fMRI activity patterns can be found at spatial scales ranging from the size of individual columns to about a centimeter".

Carlson [2014] identified neuronal activity patterns related to stimulus edges that

mimic a radial bias as a potential source of a global signal bias. The stimuli employed in this study had clearly visible, unsmoothed edges, hence edge-related activity is a valid explanation for the observed orientation-related large-scale signals. It can be argued that the V1 ROI could be adjusted by a "safety margin" to the representation of the edge of the stimuli to reduce edge related signals. I have tested various criteria for ROI definition and sizes. I have found very little variation of the results with respect to the particular shape and size of the ROI. The reported results are based on a V1 ROI generated by retinotopic mapping that used a stimulus that was larger than my visual orientation stimulus, hence I are likely to sample voxels representing edge-related signals. In other words, my ROI should contain a maximum amount of stimulus-related information present in V1. I leave an analysis exploring aspects of the relationship of individual stimulus properties and ROI shapes with the BOLD signal and decoding to a future study.

Overall, BP filtering yielded peak performances for all resolutions (except for the 3 mm acquisition). Consistent with Alink et al. [2013], the present results suggest that a band matching a DoG BP filter consisting of a 5 mm and an 8 mm FWHM Gaussian LP filter) carries most (but not all) orientation-related signal. This band covers wavelength from about 4.5 mm to 1.6 cm. The Nyquist-Shannon Sampling Theorem dictates that, in order to measure a particular signal appropriately, the sampling frequency has to be at least twice the critical frequency of that signal. Hence, a 3 mm acquisition can only sample frequencies with a wavelengths of 6 mm or larger, and consequently misses some part of this most informative band (Fig. 3.5).

This is consistent with my finding that optimal decoding accuracy required a resolution higher than 3 mm. The nearly identical peak performance on 1.4 mm and 2 mm data is also compatible with this minimum frequency rule. However, the markedly lower decoding performance on 0.8 mm data is likely evidence that a minimum sampling resolution is necessary but not sufficient for optimal decoding performance. In this study, an optimal balance of scanning resolution and temporal signal-to-noise-ratio is

reached at 2 mm resolution. Higher resolution reduce tSNR and lower resolutions do not provide sufficient sampling of higher frequency signals.

The present data does not support the hypothesis that the high spatial frequency signal of orientation columns in early visual cortex could be reflected in (much larger) fMRI voxels by means of spatial aliasing. In the case of frequency aliasing due to an insufficient sampling frequency by the voxel grid, the frequency of the aliased signal would vary depending on the actual sampling frequency (size of the voxel). However, the peak decoding performance is always located in the same band across all four resolutions. my findings are in line with Kamitani and Tong [2005] and Chaimow et al. [2011] which show that the spatial frequencies of columnar structures (0.5 cycles/mm) do not contribute signal for decoding, due to several technical limitations like inherent head motion and reduced SNR proportional to reduction in voxel volume. Moreover, Shmuel et al. [2007] state that the PSF — that captures blurring factors due to eye movements, neuronal response, BOLD response PSF in gray matter, as well as the PSF of the data acquisition process — makes fMRI data inherently LP filtered and, as such, poses a physical limitation on the spatial frequency scale from which fMRI signal can be obtained. Kamitani and Tong [2005] and Chaimow et al. [2011] identify contributions from random variations and irregularities in the columnar structures captured by larger voxels as the main source of information for decoding. These are of considerably lower frequency than the primary spatial frequency characteristics of the columnar organization and are lower than the Nyquist criterion of the BOLD fMRI sampling frequencies.

It could be speculated that the spatial scale of the orientation signal as estimated by volumetric spatial filtering is, to some degree, determined by the representation of the cortical folding pattern in the scan volume. As volumetric filtering procedures using 3D Gaussian kernels inherently mixes signals from gray matter, white matter, and superficial vessels. It might be that a volumetric BP filter corresponding to the most informative spatial frequency band is beneficial because it is of sufficient size to average

signal across the entire diameter of the folded calcarine sulcus, whereas a smaller filter is not, and a bigger filter includes a substantial fraction of the surrounding white matter and adjacent cortical fields. If the above speculation is correct, I could expect lower decoding accuracy in the most informative band band when replacing the employed spatial filtering procedure with a cortical surface-based smoothing or a spatial filtering that is restricted to V1 ROIs in each hemisphere. I performed these two alternative analyses and found only minor differences in the results (see supplementary material Fig. 3.11). Similar to the report of Swisher et al. [2010], the band-pass, high-pass, low-pass components based on these alternative spatial smoothing schemes perform very similar, but more evenly sloped than what was obtained from the unconstrained volumetric filtering. Except for the 0.8 mm data, where the insufficient signal is even more evident, the BP performance is extremely similar. Consequently, I find little evidence for an impact of using standard, unmasked, volumetric spatial filtering for this decoding analysis.

Venous voxels in V1 ROI contribute above chance classification of orientation gratings Several authors have cited an orientation-related BOLD signal originating from the vascular system (draining veins) as a potential information source for decoding that may introduce spatial biases in the representation of orientation as measured with fMRI [Chaimow et al., 2011, Kriegeskorte et al., 2010, Shmuel et al., 2010]. The present results confirm the presence of such a signal. Particularly for the two highest resolutions tested here the decoding accuracy obtained from voxels sampling veins is equal to the performance obtained from the non-venous rest of the V1 ROI, or even outperforms it when controlling for the number of input voxels for the classification model (Fig. 3.13A).

A BOLD signal originating in the blood vessels has the potential to introduce complex transformations of the spatial representation of orientation in the BOLD response patterns. Due to the structural properties of the vascular system this signal

is likely to be of lower spatial frequency, compared to the underlying neuronal activation pattern, and is superimposed on a potential high-frequency pattern reflecting the columnar structure of V1. This explanation has been put forth by Kriegeskorte et al. [2010] who describe voxels as "complex spatio-temporal filters" and my results are compatible with this model.

It should also be mentioned that previous studies found a substantial reduction of intra-vascular BOLD signals at higher magnetic field strength [Yacoub et al., 2001], and enhanced signal contributions from microvascular structures at 7T [Shmuel et al., 2007]. Consequently, the particular composition of the compound signal captured with BOLD fMRI will vary with the magnetic field strength. A future study should compare the present results with data acquisitions at a different field strength to shed more light on nature of the underlying signal and the implications for decoding analysis.

Limitations The focus of the present study was to investigate the effect of acquisition resolution and spatial filtering on the decoding of visual orientations from primary visual cortex. In order to yield comparable results, the acquisition parameters were constrained to guarantee a certain minimum coverage of the V1 ROI even at the highest resolutions and to have an identical temporal sampling frequency (TR) to yield the same number of observations across all resolutions. This choice implied that the GRAPPA acceleration factor had to be increased with increasing resolution, hence leading to an increased under-sampling of the k-space with higher resolutions. This could impact the sensitivity of the scan to high-frequency spatial signals. A future study will have to test whether the present findings hold when constraints on coverage and sampling frequency are relaxed. For example, a study by De Martino et al. [2013] using a 3D gradient and spin echo (GRASE) sequence suggests that such a sequence outperforms a gradient echo sequence, such as the one employed in this study, for high-resolution imaging at 0.8 mm isotropic resolution — at the expense of a vastly reduced scan volume.

The present study is exclusively based on 7 Tesla fMRI data, hence it remains

unclear in which way the characteristics of the relation of decoding performance and acquisition resolution are dependent on MR field-strength. The differences in the sizes of the BOLD point-spread functions [Engel et al., 1997b, Shmuel et al., 2007] suggest a lower resolution limit for 3 Tesla scans. However, the reported optimal resolution is within the range of conventional acquisition resolutions of today's 3 Tesla scanners. A future study should address the question of how the decoding performance varies with field-strength for identical resolutions.

While this study focused on the optimal acquisition parameters for decoding of visual orientation from fMRI BOLD response patterns in early visual cortex, I acknowledge other possibilities of further optimization of the decoding procedure (classification algorithm, hyper-parameter optimization, etc.) and their potential impacts on results and interpretations. To facilitate the required future analyses I have publicly released the data (available without restrictions from GitHub `https://github.com/psychoinformatics-de/studyforrest-data-multires7t`) and a "Data in brief" manuscript along with this. In this study I have found that given a neural signal with known fine-scale spatial characteristics, there are technical and physiological factors that place the acquisition resolution optimal for decoding at a substantially coarser scale. Future studies should investigate whether the optimal settings for other decoding paradigms and different cortical areas, beyond the findings for visual orientation in visual cortex presented here, and the congruent results for auditory representations reported by Gardumi et al. [2016], are similar in nature.

4. Experiment 2: The Effect of MR field strengths (7 Tesla vs. 3 Tesla) on Orientation Decoding: A Comparison Study

4.1 Introduction

As multivariate pattern analysis (MVPA) approaches are being increasingly used to analyze fMRI data in decoding cognitive states represented in the distributed patterns of brain activity [Bonte et al., 2014, Haxby, 2012, Haynes and Rees, 2005, Haynes, 2009, Kamitani and Tong, 2005, Zhang et al., 2015], the true origin and spatial scale of the decoding signals picked by the classifiers, is being strongly debated [Alink et al., 2013, Op de Beeck, 2010, Freeman et al., 2013, Swisher et al., 2010]. It is interesting to note that the data acquisitions in these previous studies were performed not only in different acquisition resolutions but also in scanners of different field strengths. For example, Swisher et al. [2010] reached a conclusion of the broadband nature of orientation decoding signals from a high resolution dataset acquired in a high-field 7 Tesla scanner, whereas, Op de Beeck [2010] and Freeman et al. [2011, 2013] concluded that orientation decoding is driven by a much larger coarse scale map in V1, by performing MVPA on a dataset acquired on a conventional 3 Tesla scanner. Until now, it hasn't been investigated whether these studies reached conflicting conclusions due to the differential effect of the magnetic field strengths on the acquired data with respect to BOLD signal sensitivity.

Experiment 1 (chapter 3) investigated the effect of acquisition resolution and spatial filtering on the decoding of visual orientations from V1 by performing multivariate cross-validated decoding analysis on ultra-high field 7 Tesla fMRI data in four different resolutions (0.8 mm, 1.4 mm, 2 mm and 3 mm iso) from 7 participants. Among the four tested acquisition resolutions, the highest decoding accuracy was achieved with a 2 mm resolution. The results shown in Experiment 1 (chapter 3) were congruent with the

simulation study by Chaimow et al. [2011] which reported that the optimal accuracy of decoding ocular dominance is reached between 2-3 mm voxel size at 3 Tesla. As signal to noise ratio (SNR) increases with magnetic field strength, fMRI acquisition at ultra-high field 7 Tesla is expected to strongly benefit as compared to 3 Tesla. Though 7 Tesla fMRI offers much higher blood-oxygen level dependent (BOLD) sensitivity and substantial reduction in partial volume effect enabling higher acquisition resolution, [Weibull et al., 2008], it suffers from higher distortion artifacts and increased physiological noise (such as inevitable motion). Though relative advantages and shortcomings of 7 Tesla vs. 3 Tesla have been extensively studied with respect to clinical MR, it remains inconclusive what effect MR field strength has on the performance of multivariate decoding. This becomes even more relevant when in Experiment 1 the best performance at 7 Tesla was obtained with 2 mm iso data (see Figure 3.7), and at 3 Tesla, fMRI data can be recorded at 2 mm iso with subtantially less distortion and comparable signal to noise ratio. In this study, I aim to address this question by performing orientation decoding analysis similar to that described in Experiment 1 on data acquired from 7 participants at 3 Tesla in three different acquisition resolutions (1.4 mm, 2 mm and 3 mm iso). Moreover, in decoding studies in the visual cortex [like Chaimow et al., 2011], it has been shown that though neural signals emanate from columnar structures, there are technical and physiological factors that place the optimal acquisition resolution for decoding at a substantially coarser scale. Whether the pattern of decoding performance with respect to different acquisition resolutions (as shown in Experiment 1) remains similar across a different magnetic field strength (3 Tesla) with an identical orientation decoding paradigm, is investigated in this study.

Though the main focus of this experiment is the comparison of orientation decoding performance in different MR field strengths, it was also examined how post-acquisition volumetric gaussian filtering affects the decoding performance in 3 Tesla. This spatial filtering approach before performing multivariate decoding enables to understanding the spatial scale of the orientation discriminating signals from the visual

cortex.

4.2 Materials and methods

4.2.1 Participants

Seven healthy right-handed participants (5 males) were recruited from the subject-pool of the *studyforrest* project [Hanke et al., 2014, 2015a]. They all had normal or corrected to normal vision and were paid for their participation. There were 5 subjects who previously participated in the 7 Tesla experiment described in Experiment 1 (due to unavailability of sub-04 and sub-18, sub-09 and sub-10 were recruited instead). Before every scanning session, they were provided with instructions for the experiment and signed an informed consent form. The study was approved by the Ethics Committee of the Otto-von-Guericke University.

4.2.2 Stimulus and Experimental Design

In order to keep parity between orientation decoding experiments performed in the 7 Tesla and 3 Tesla experiments, the stimulus was kept identical to that described in Experiment 1 (chapter 3). Flickering sine-wave orientation gratings (flicker frequency = 4 Hz, constant spatial frequency 1.4 cycles per degree of visual angle with 100%contrast) were displayed in both hemifields on medium gray background in form of semi-annular patches (0.8°-7.6° eccentricity, 160° width on each side with a 20° gap along the vertical meridian). Orientation gratings (0°, 45°, 90°, or 135°) were displayed with random phase (0, $\frac{\pi}{2}$, π, or $\frac{3\pi}{2}$ degrees) changed at a frequency of 4 Hz.

The presentation computer ((Neuro)Debian operating system [Halchenko and Hanke, 2012]) performed stimulus presentation and response logging using PsychoPy [v1.79; Peirce, 2008]. The stimulus was displayed on a high definition rear-projection screen (1140×780 pixels, 18 cm wide), 60 Hz video refresh rate) placed at a total viewing distance of 35 cm. In order to keep the participants' attention focused and to minimize eye-movements, they performed a center fixation task that was unrelated to the stim-

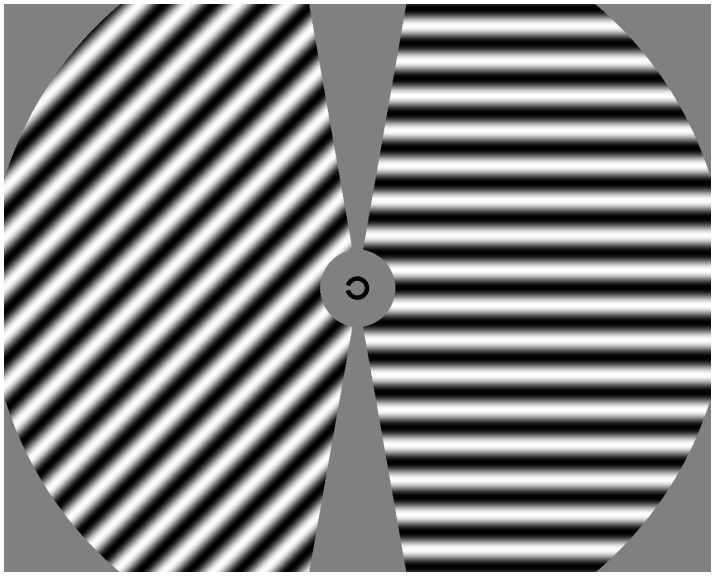

Figure 4.1: Oriented grating stimulus with Landolt C fixation task
Event-related stimulation paradigm with 3 s of flickering stimulus at the beginning of each trial followed by a 5 s inter-trial interval. The participants had to fixate at the center of the screen performing Landolt C task throughout the entire length of the experiment. The sequence of oriented gratings displayed in each hemifield were independently randomized. The paradigm also included ten trials when the previous stimulus was repeated in only one of the hemifields and were used to decouple stimulation sequences (refer to section 3.2.4).

ulation with oriented gratings. Participants were asked to fixate on the Landolt-Ring (radius 0.12°) presented at the center of the screen. At random intervals in each run the Landolt-C stimulus was shown (left or right opening of 0.048°) and the participants had to respond to the direction of the opening of the probe by pressing one of two buttons corresponding to a left or right opening. The mean accuracy for this task was 90.1% correct across all participants.

FMRI acquisitions (1.4 mm, 2.0 mm and 3.0 mm isotropic) were performed in three separate sessions, acquisition order being randomized for each participant. Each run comprised of 30 trials (8 s duration; 4 min total run duration) and each participant performed 10 experimental runs in every session. The pseudo-random sequence of displaying orientation gratings in both the hemifields were identical to that of the 7 Tesla experiment and are explained in details in section 3.2.4.

4.2.3 MR acquisition

BOLD fMRI data from the V1 ROI were acquired at 3 Tesla at three different resolutions with a stimulation paradigm which was almost identical to that used in the 7 Tesla experiment. For direct comparison of multivariate accuracies, the MR acquisition parameters were carefully chosen so that they could be maximally similar across multiple resolutions and across field strengths.

A Siemens Prisma 3 Tesla scanner with a 64 receive channel head coil (Siemens, Erlangen, Germany) was used to acquire T2*-weighted echo planar images (EPI) (TR/TE = 2000/30 ms, FA = 90°). The fMRI data were acquired on a tilted axial plane for 3 different spatial resolutions, i.e. 3 mm isotropic (FoV = 216 mm, matrix size 72 × 72, 36 slices, GRAPPA accel. factor 2), 2 mm isotropic (FoV = 216 mm, matrix size 108 × 108, 32 slices, GRAPPA accel. factor 2) and 1.4 mm isotropic (FoV = 210 mm, matrix size 150 × 150, 28 slices, GRAPPA accel. factor 3) with slices parallel to the calcarine sulcus. Slices were acquired in an ascending order with a 10% inter-slice gap to minimize cross-slice excitation. All acquisition sequences used anterior-posterior phase encoding. Each experimental run consists of 120 volumes and 10 separate scans (one for each experimental run) were performed for each subject.

To be able to directly compare the multivariate decoding performance with respect to MR field strength, the distortion corrected [In and Speck, 2012] 7 Tesla data (1.4 mm, 2.0 mm and 3.0 mm isotropic) acquired in Experiment 1 were re-analysed in this study with the same pre-processing steps (like feature selection) and orientation classification procedure as was done on the 3 Tesla dataset in this study. The analysis procedures performed on both 7 Tesla and 3 Tesla data in this experiment were substantially different from the previous one and are described in details in the following sections.

Structural images for all participants acquired as a part of the *studyforrest* project [Hanke et al., 2014] in a 3 Tesla Philips Achieva scanner, were reused in this study (refer to section 2.3.5).

4.2.4 Region of interest localization

As described in section 2.3 and 3.2.7, all participants underwent retinotopic mapping for localization of V1 and a participant specific V1 ROI mask was created in respective acquisition resolutions. Feature selection is an important pre-processing step in the MVPA pipeline applied to reduce excessive noise and also for dimensionality reduction [Haxby, 2012]. Unlike Experiment 1, in this study, univariate feature selection process was performed in the V1 ROI mask for localizing the visually responsive voxels using a second-level fixed-effects GLM analysis in FSL FEAT (default parameters - $z > 2.3$ with $p < 0.05$) [v5.0.9; Smith et al., 2004]. The visually responsive V1 ROI masks thus created, were used for all further analyses.

4.2.5 tSNR calculation

The voxelwise time-series signal to noise ratio calculation method was identical to the method described in section 3.3.1. The overall tSNR value for every acquisition was obtained by averaging voxelwise tSNR across only the visually responsive voxels in the V1 ROI to get a better estimate of tSNR, in response to the flickering visual stimulus.

4.2.6 Orientation decoding analysis

Multivariate orientation decoding was performed with PyMVPA [v2.4.1; Hanke et al., 2009] on a compute cluster running (Neuro)Debian [v8.0; Halchenko and Hanke, 2012]. All functional scans acquired at the 3 Tesla underwent motion correction using MCFLIRT in FSL [v5.0.9; Smith et al., 2004] and the 7 Tesla data which were re-analysed in this experiment were already motion corrected as part of the distortion correction procedure [In and Speck, 2012]. Similar to section 3.2.9, univariate GLM analysis was performed on BOLD fMRI time series from an individual experimental run using GLM implementation in NiPy [v0.3; Millman and Brett, 2007] while accounting for serial correlation with an autoregressive term (AR1). The GLM design matrix included hemodynamic response regressors and corresponding temporal deriva-

tives, six nuisance regressors for motion (translation and rotation), and polynomial regressors (up to 2nd-order) modeling temporal signal drift as regressors of no-interest. The β weights thus computed for every run were Z-scored per voxel. The resulting dataset for MVP analysis contained 40 samples (one normalized β score per condition per run) for each participant.

Similar to previous literature, Linear support vector machines [SVM; PyMVPA's `LinearCSVMC` implementation of the LIBSVM classification algorithm; Chang and Lin, 2011] were used to perform a within-subject leave-one-run-out cross-validation of 4-way multi-class orientation classification. This linear SVM algorithm has a critical hyper-parameter C that indicates the trade-off between the width of the margin of the classifying hyperplane and number of correctly classified training data points. In this experiment the value of the C parameter was scaled according to the norm of the data (default operation for a linear kernel in PyMVPA).

4.2.7 Spatial filtering

In order to reveal the spatial scale of orientation specific information distributed across the spatial frequency spectrum, the decoding procedure was repeated after spatial filtering of the functional imaging data with volumetric Gaussian kernels. As it has been shown in the previous studies like Swisher et al. [2010] and in Experiment 1 (chapter 3) that there was no substantial difference between decoding performance after a 3D Gaussian filtering and cortical surface-based smoothing approach, in this study I implemented only the 3D Gaussian filtering scheme. Similar to Swisher et al. [2010], I used Gaussian filtering prior to any feature extraction for MVPA analysis in the particular ROI to estimate the spatial scale of the orientation specific signal. To be able to compare with results across previous literature, the level of spatial filtering used in this method is expressed in terms of the size of the Gaussian filter kernel described by its full width at half maximum (FWHM) in mm. All spatial smoothing procedures were implemented with the `image_smooth()` function in the nilearn package [Pedregosa et al.,

2011]. The respective implementations of *Low-pass* (LP), *High-pass* (HP), *Bandpass* (BP) and *band-stop* (BS) filtering kernels were identical to the procedure described in Experiment 1 (see 3.2.10).

4.3 Results

4.3.1 Decoding accuracy and tSNR comparison

In order to determine the effect of MR field strength on orientation decoding, I compared decoding performance averaged across 7 participants in three different acquistion resolutions (1.4 mm, 2.0 mm and 3.0 mm isotropic) in both 7 Tesla and 3 Tesla. Figure 4.2A shows the comparison of mean classification accuracy across participants from 7 Tesla and 3 Tesla. For the above analysis, only visually reponsive voxels in the respective V1 ROIs (as described in 4.2.4) were used. For 7 Tesla acquisition the peak classification performance of 38.92% was recorded at 2 mm isotropic resolution. Mean accuracy at 1.4 mm isotropic and 3.0 mm isotropic performed lower than 2 mm data. The pattern of decoding accuracy across resolutions were similar to the study shown in Experiment 1 (chapter 3), though the ROI and the optimization of SVM classifier hyperparameters done in this experiment were different from the previous study. For the 3 Tesla data, the above pattern of decoding accuracy across resolutions was not observed. The lowest mean decoding accuracy at 3 Tesla was recorded with the 1.4 mm data. For 2.0 mm data acquired at 3 Tesla decoding performance was lower than that at 7 Tesla data (the mean decoding accuracy for 2.0 mm data at 3 Tesla was 32.14% and it was lower than the 95% binomial proportion confidence interval of the 2.0 mm acquisition at 7 Tesla [33.73% - 44.31%] computed from the number of correct predictions concatenated across hemispheres and cross-validation folds for all subjects). The 3 Tesla data of 3.0 mm resolution performed better (35.89%) than the other 2 resolutions and was almost identical to the decoding accuracy 3 mm data at 7 Tesla.

According to Chaimow et al. [2011], the classification performance is impacted by the overall contrast-to-noise ratio (OCNR), which is directly proportional to the

time course signal to noise ratio (tSNR) (depends on voxel size [Triantafyllou et al., 2005]). In previous literature, it has been shown that the noise factor in fMRI time-series are dominated by physiological and thermal noise. As described in Triantafyllou et al. [2005], these noise factors vary primarily on signal intensity, echo time of the EPI sequence and magnetic field strength. Figure 4.2B shows how time-series signal to noise ratio calculated for only the visually responsive voxels in V1 varies across different resolutions and across field strengths. It has been observed that the tSNR calculated in this study is very similar to the tSNR values reported by Triantafyllou et al. [2005] in figure 6. Similar to Experiment 1, the asymptotic relation of tSNR as a function of voxel volume and MR field strength was modelled as $\text{tSNR} = \kappa V/\sqrt{1+\lambda^2\kappa^2V^2}$, where V is the voxel volume, κ is the proportionality constant, and λ is the magnetic field strength independent constant parameter. The empirical data obtained in this experiment from both 7 Tesla and 3 Tesla were fitted to the above model. The goodness of fit of the data to the model was obtained by calculating R^2 score. For 7 Tesla, the model parameters were calculated to be λ=0.013 and κ=22.7 (R^2=0.85) and for 3 Tesla, they were λ=0.014, κ=6.6 (R^2=0.99). Similar to the observation in Figure 4.2A, that 3 mm iso data acquired at 3 Tesla and at 7 Tesla showed nearly identical orientation decoding performance, it was observed that the tSNR calculations of these 2 acquisitions were similar.

4.3.2 BOLD signal change

In order to quantify whether orientation decoding at 3 Tesla is driven by differential impact in fMRI BOLD signal in response to different orientation gratings [see Alink et al., 2013, Furmanski and Engel, 2000, Tong et al., 2012], I calculated mean percentage BOLD signal change in response to orientation stimulus across resolutions using Feat-Query in FSL [v5.0.9; Smith et al., 2004]. Similar to preprocessing in MVP analysis, no spatial smoothing was performed before calculating the percentage signal change. Percentage signal change was calculated with Featquery considering only the responsive

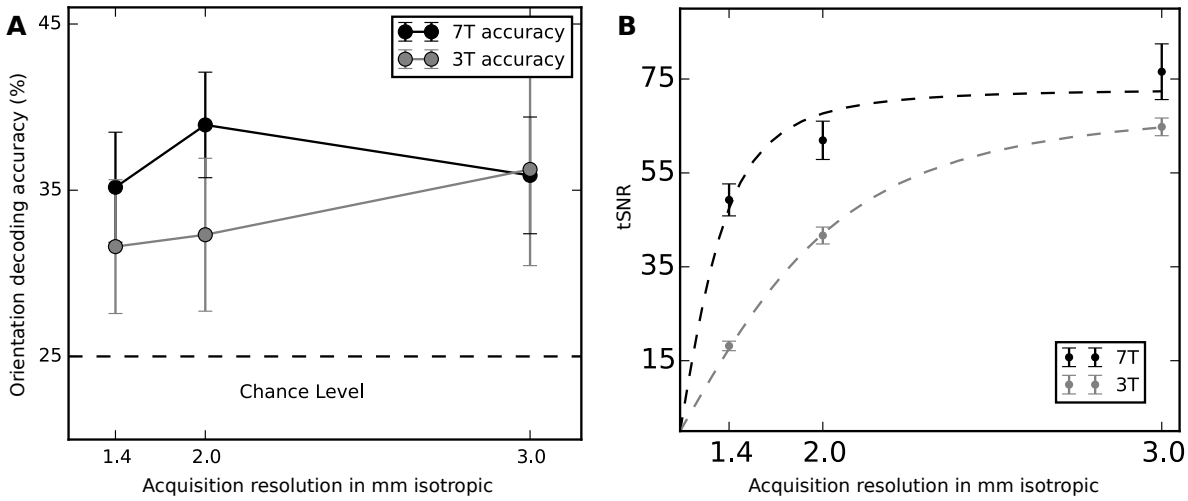

Figure 4.2: Comparison of Magnetic field strengths
(A) Orientation Decoding accuracies across field strengths. For 7 Tesla acquisitions, the peak decoding accuracy found at 2 mm iso resolution. The 3 Tesla data showed monotonic increase in accuracy with increase in voxel size. Decoding performance for 3 Tesla data were lower than the 7 Tesla data for 1.4 mm and 2.0 mm resolutions, but the 3.0 mm acquisition showed identical performance as the 7 Tesla data. **(B)** Time-series signal to noise ratio is lower in the 3 Tesla data than 7 Tesla for all acquisition resolutions. However tSNR level reached an asymptotic level as shown in Triantafyllou et al. [2005]

voxels in V1. As shown in Figure 4.3B that the mean percentage BOLD signal change was the highest for 1.4 mm resolution and reduced with increase in voxel dimensions, similar to Experiment 1 (chapter 3). It was observed that the 0° orientation produced lowest percentage signal change and only for 3 mm data the percentage signal change for the oblique (45° and 135°) orientations were more than the cardinal (0° and 90°) orientations, similar to the 'oblique effect' [shown in Furmanski and Engel, 2000]. A 2-factor (orientation and resolution) within-subject ANOVA for the estimated BOLD signal change from all 7 participants were performed to test the significance of effect of orientation and acquisition resolution and a possible interaction between them. It was found that there is a significant effect of acquisition resolution on BOLD % signal change ($F(2, 12)$=65.5, p=3.49e-07). But there was no significant effect of orientation ($F(1, 6)$=5.645, p=0.055) and no significant interaction of acquisition resolution and orientation (p=0.609). It was also observed that for every acquisition resolutions 0°

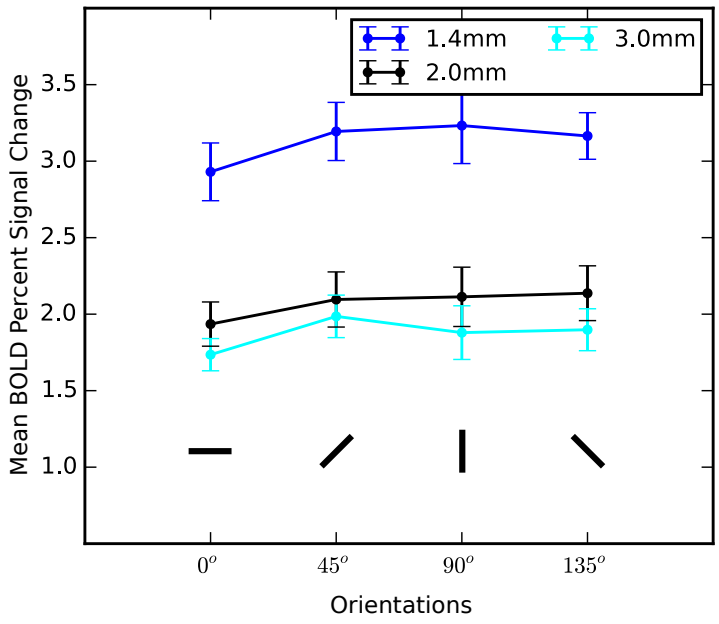

Figure 4.3: Percent signal change in response to different orientations across acquistion resolutions in 3 Tesla
Mean BOLD Percentage signal change calculated for every orientation across resolutions. For all resolutions, 0° orientations showed lowest signal change. Though the effect on orientation on signal change was non-significant, 'oblique effect' [shown in Furmanski and Engel, 2000] was shown by the 3 mm data

orientation gratings showed the lowest percentage signal change.

4.3.3 Impact of Gaussian smoothing

Figure 4.4 A-C show how spatial smoothing with Gaussian kernels affects the performance of multivariate orientation decoding using linear SVM classifier across data acquired in 3 Tesla from three different acquisition resolutions (1.4 mm, 2.0 mm and 3.0 mm). LP volumetric gaussian filtering shows monotonic decrease in decoding accuracy with increase in kernel size in all resolutions. But the decoding performance remained above chance-level even after spatial smoothing of 10 mm. The complimentary HP filtered images indicate the amount of high-frequency spatial information is left after low frequency components are removed from the image by LP filtering. HP components showed above chance decoding accuracy and performed better than the LP filtered image beyond the 9-10 mm filtering kernel in all resolutions. Band pass

filtered images show how much decoding information is present at a particular spatial frequency band of 1 mm FWHM bandwidth. The overall best decoding accuracy was again obtained by the BP filtered images. The BP components showed a prominent peak in accuracy in the range of ≈5-8 mm for all resolutions. This observation is similar to the findings in Experiment 1 and in Alink et al. [2013]. Classification performance of BS filtered data remained above-chance for all spatial frequency bands. The BS performance curve initially follows the LP performance for small filter sizes, but resembles the HP performance for larger filter sizes.

4.4 Discussion

The primary focus of this study was to investigate the effect of MR field strength on orientation decoding from primary visual cortex (V1). With that objective, a cross-validated orientation classification analyses with Linear Support Vector Machine classifier were performed on fMRI data acquired from 7 right handed participants in both 7 Tesla and 3 Tesla Siemens scanners in 3 different acquisition resolutions (1.4 mm, 2.0 mm and [mm3.0] isotropic). The classification model was trained to decode four different orientation gratings (0°, 45°, 90°, or 135°) from the patterns of β weights obtained by fitting BOLD fMRI data from V1 to a GLM. Classification accuracy averaged across both hemifields and across participants was used as a metric for comparison of decoding performance between MR field strengths. Similar to Experiment 1, the decoding accuracies of the 4-way classification in this experiment were lower as compared to pairwise classification accuracies reported in other studies like Haynes and Rees [2005], Kamitani and Tong [2005].

In Experiment 1, it was found that the optimal C parameter for a Linear SVM classifier might vary upto a large extent between resolutions (Figure 3.4), the search range of the hyper-parameter being manually determined. It is to be explicitly mentioned that the main focus of this study was to compare orientation decoding performance across MR field strengths and not a comparison across acquisition resolutions.

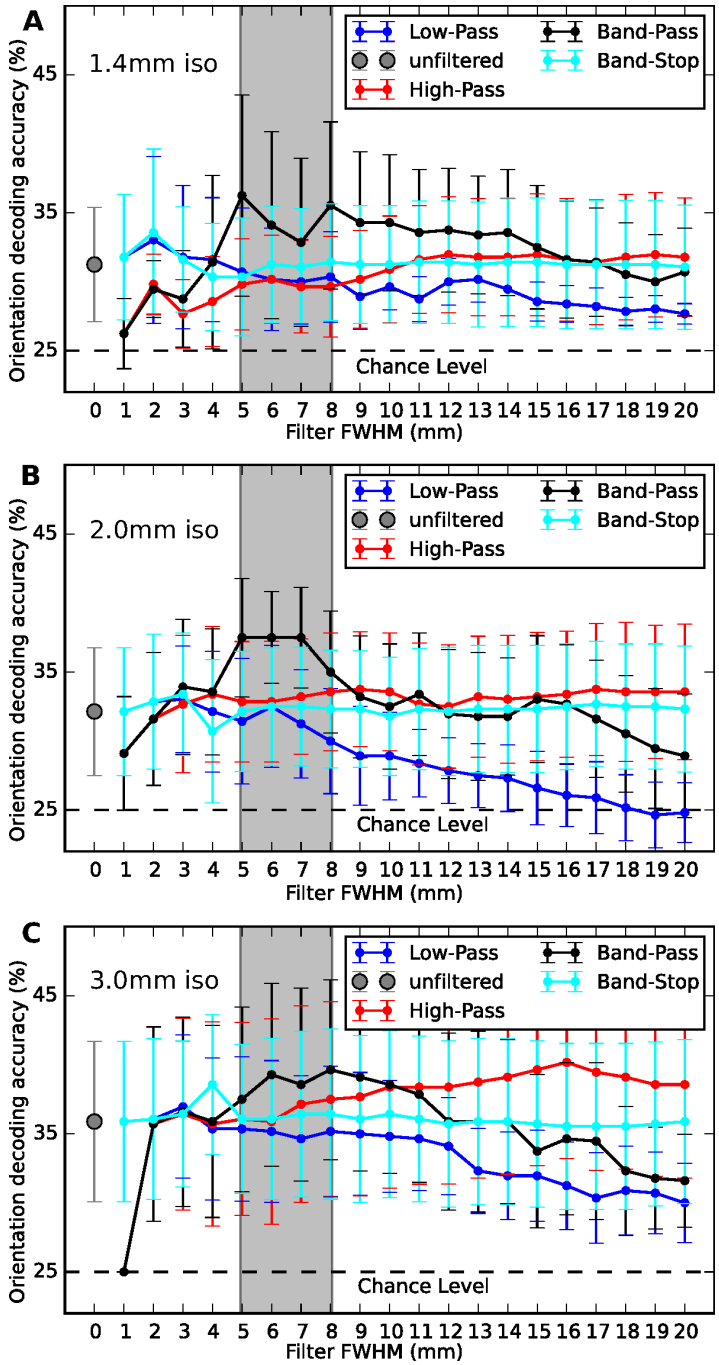

Figure 4.4: Spatial smoothing with volumetric Gaussian filter
Above chance level orientation decoding accuracy was shown in all resolutions even with spatial smoothing beyond 10 mm FWHM of gaussian smoothing. Band pass filtering showed the highest decoding accuracy peak in the ≈5-8 mm FWHM in all resolutions of 3 Tesla showing the absence of aliasing of fine scale signals from V1 by the voxel grid. Band stop components show decoding accuracies similar to low pass components until ≈5-8 mm band and then shows similar decoding accuracies like the highpass components. This shows the broadband nature of the orientation signals.

Hence, the value of the C parameter was intentionally scaled according to the norm of the data and not fine-tuned within a given range by nested cross validation.

Orientation classification analysis at 7 Tesla shows a distinct peak in decoding performance at 2 mm isotropic resolution. Both 1.4 mm data (higher acquisition resolution) and 3 mm data (lower acquisition resolution) showed similar decoding performance, lower than that of the 2 mm data. This pattern of variation of decoding accuarcy with respect to the acquisition resolution is 7 Tesla is inline with the study in Experiment 1 (chapter 3). In case of 3 Tesla data, the decoding accuracy monotonically increased with increase in voxel size. in fact, the decoding accuracy at 3 mm resolution acquired at 3 Tesla was almost identical to the 3 mm data acquired at 7 Tesla. This brings me to the discussion, whether I am benefitting from acquiring fMRI data at ultra high field strengths. Though the choice of MR field strength (7 Tesla vs. 3 Tesla) should depend on the particular study and the hypothesis being tested, it is a well known fact that higher field strengths like 7 Tesla provides superior signal to noise ratio and improved BOLD signal sensitivity over 3 Tesla. But as shown in previous literature, fMRI time-series is dominated by physiological and thermal noise under some conditions. Specifically, Triantafyllou et al. [2005] showed that unlike image SNR, the fMRI time-series signal to noise ratio (tSNR) reached its asymptotic limit with moderate spatial resolution ($\approx$ between 2-3 mm isotropic) at 3 Tesla and it shows only marginal gains from data acquisition at higher field strengths (7 Tesla). This asymptotic trend of tSNR with respect to voxel dimensions in different field strengths was found in the empirical data analyzed in this study (refer to Figure 4.2B). Moreover, it has been shown in the simulation study of Chaimow et al. [2011], how multivariate decoding performance at a particular resolution varies proportionally with the tSNR factor. This explains the reason why 3 mm data of 7 Tesla and 3 mm data of 3 Tesla showing very similar tSNR values performed almost identically in orientation decoding.

From the discussion in the previous paragraph, it is evident that tSNR is an important factor in orientation decoding, but it is not the exclusive driving factor. The

spatial scale of the orientation specific signal is also a primary key which determines the performance of orientation decoding at a particular combination of acquisition resolution and MR field strength. As shown in Figure 3.9A, at 7 Tesla, though the tSNR value of the 3 mm data was more than the 2 mm acquisition, the decoding accuracy for the 3 mm data was lower, because it did not have the sufficient spatial sampling frequency to be able to capture the orientation specific signals. On the other hand, 0.8 mm and 1.4 mm data acquired in 7 Tesla provided good spatial sampling frequency but lacked in tSNR due to smaller voxel size. So better orientation decoding performance requires an optimal balance of acquisition resolution and tSNR factor, which was obtained in 2 mm data at the 7 Tesla. But in the 3 Tesla acquisition, due to reduced field strength, 2 mm data did not have adequate tSNR (41.66) as compared to the 7 Tesla 2 mm acquisition (61.94). This was also reflected at the substantially lower decoding accuracy of 3 Tesla 2 mm data as compared to that of the 7 Tesla.

Spatial scale of orientation specific signals is a widely debated area of research [Alink et al., 2013, Op de Beeck, 2010, Freeman et al., 2011, 2013, Gardner, 2010]. There are several competing ideas, one group of scientists have shown that the orientation decoding classifiers reflect these underlying fine scale organizations (spatial scale ¡ 1 mm) of orientation columns in the primary visual cortex [Adams et al., 2007, Hubel and Wiesel, 1972]. Boynton [2005], Kamitani and Tong [2005] showed that random irregularites in the columnar organizations as sampled by the acquisition voxel grid provides a reliable bias of orientation decoding. Conversely, the concept of orientation decoding depending entirely on low frequency coarse-scale orientations, including biases for radial and cardinal orientations has been proposed by Furmanski and Engel [2000], Mannion et al. [2010], Sasaki et al. [2006]. In this study, I performed a 2-factor (orientation and resolution) within-subject ANOVA for the mean BOLD signal change from all 7 participants in the visually responsive voxels of V1 and could not find significant effect of orientation. Only in the 3 mm data it was observed that the percentage signal change for the cardinal orientations were lower than the radial orientations, similar to 'oblique

effect' reported by Swisher et al. [2010]. Spatial filtering with volumetric Gaussian kernel also indicate the spatial scale of orientation specific signals. In all resolutions I find that the Band pass filtered component showed highest decoding accuracy in the ≈5-8 mm. This is identical to the results shown in the gaussian filtering approach on 7 Tesla data (Figure 3.10). It indicates that orientation specific signals which are present in a particular band of frequencies, are preferred and beneficially sampled by the classifiers to perform orientation decoding across MR field strengths. But the spatial scale of the orientation signals were not confined to that particular band. LP filtered components showed gradual decline in decoding accuracy but performance was above chance beyond 10 mm smoothing showing the contribution of coarse scale signals. On the other hand the HP filtered components beyond 10 mm show high decoding performance indicating high frequency signals. Interestingly, the band stop components always perform above chance level with decoding accuraies similar to the low pass filtered components until the ≈5-8 mmand then resembling the accuracies of the high pass filtered component. This shows how the low frequency and the high frequency components of the orientation specific signals contribute to decoding in different ranges of the spectrum. Overall I again conclude that the orientation signals are spatially broadband in nature, starting from millimeter range columnar signals to large-scale orientation biases (beyond 10 mm), in line with the findings of Experiment 1 and Swisher et al. [2010].

5. Experiment 3: Dependence of Orientation Decoding on tSNR

5.1 Background and Motivation

In Experiment 1 and Experiment 3, it was shown how the balance of BOLD acquisition resolution and the corresponding tSNR plays an important role in optimal decoding performance. The primary difference between the pattern of decoding accuracy across acquisition resolutions in 7 Tesla vs. 3 Tesla is at 2 mm resolution. The substantially lower decoding accuracy with 2 mm data at 3 Tesla, can be attributed to the fact that the 7 Tesla acquisition provides better BOLD sensitivity and hence the tSNR figures were substantially higher (tSNR: 41.66 for 2 mm data in 3 Tesla and 61.94 for 2 mm data in 7 Tesla). This explanation of lower decoding performance of 2 mm data in 3 Tesla is supported by the findings of Tong et al. [2012], where it is shown how orientation classification performance was highly correlated to the amplitude of the stimulus-driven fMRI response. In Experiment 1, evidences of this correlation can be found, where it is shown that for all resolutions the highest difference in mean percentage signal change between 0° and 90° (see Figure 3.9B) orientations is also reflected in the maximum decoding performance obtained for discriminating between 0° and 90° orientations (see confusion plots in Figure 3.7A). Keeping in mind these findings, it can be hypothesized that by modifying the scanning protocol parameters, if the tSNR factor of the 2 mm data in 3 Tesla could be enhanced, then a better decoding accuracy should be obtained.

5.2 Methods

5.2.1 Data Acquisition

In the 2 mm data acquisition protocol at 3 Tesla parallel image acquisition technique [GRAPPA accel. factor 2 Griswold et al., 2002] was used for keeping repetition time

(TR) unchanged across resolutions. Though parallel imaging techniques are used in MR imaging to reduce acquisition time, it has a limitation of decreased SNR than the MR images obtained from Fourier transform of the k-space in the traditional manner. For parallel imaging acceleration factor of R, theoretically SNR is reduced by a factor of $\sqrt{R}$ [Glockner et al., 2005]. To check the effect of enhanced tSNR on decoding accuracy without modifying the repetition time (TR), a separate dataset was acquired with Simultaneous Multi-slice EPI (multiband) acquisition after switching off GRAPPA acceleration. The multiband technique was chosen because this method provides faster fMRI acquisition without significant loss of SNR. FoV in this acquisition setup was kept exactly same as the single band acquistion (FoV=216 mm, matrix size 108×108, 32 slices). In order to keep identical repetition time (TR) across all resolutions multiband acquisition factor of 2 was introduced and the 7/8 partial fourier transform of the k-space was performed (SMS=2, TR = 2000 ms, FA = 90°, 7/8 k-space, 120 volumes). The EPI data acquisition was performed with a paradigm identical to Experiment 2. The same participants who previously volunteered for Experiment 2 also participated in this experiment.

5.2.2 Data Analysis

The Region of Interest localization, the tSNR calculation procedure and the orientation decoding analysis were identical to that of Experiment 2.

5.3 Results and Conclusion

In this experiment, I compared the tSNR calculations and corresponding decoding accuracies of the normal 2 mm iso 3 Tesla acquisition, the multiband 2 mm iso 3 Tesla acquisition and the 2 mm iso 7 Tesla acquisition. As shown in Figure 5.1, 2 mm data with GRAPPA acceleration factor 2, mean tSNR across all V1 responsive voxels across participants was 41.66. On the other hand, mean tSNR of the 2 mm multiband acquisition without GRAPPA acceleration was calculated as 57.68 showing a 38.45% increase.

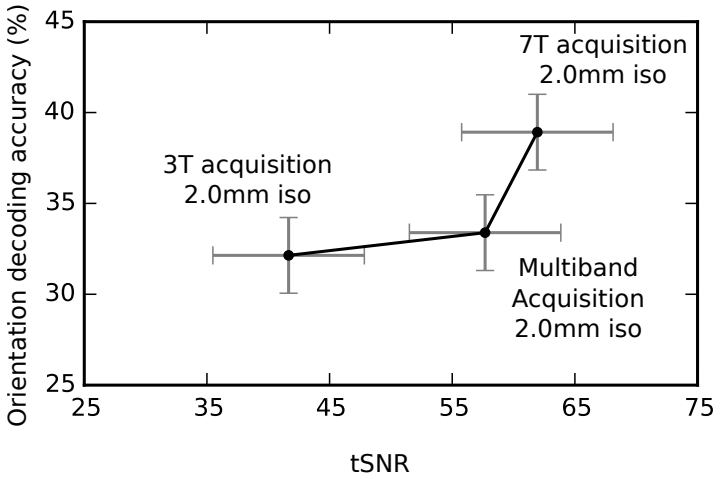

Figure 5.1: Dependence of orientation decoding on tSNR
With multiband 2 mm iso acquisition 38.45% increase of tSNR was achieved over normal 2 mm data at 3 Tesla, but the decoding accuracy showed marginal enhancement (≈3%). However, the tSNR of 2 mm iso 7T acquisition was 48.67% more than mm2 data at 3 Teslaand the decoding accuracy was ≈21% better.

It was observed in Figure 5.1 that, with this enhancement in tSNR, the mean orientation decoding performance at 2 mm iso increased from 32.1% for normal EPI acquisition to 33.4% for multiband acquisition. The 2 mm acquisition at 7 Tesla had tSNR of 61.94 and it showed 38.92% decoding accuracy.

Only a very moderate improvement in decoding accuracy was observed as shown in Figure 5.1. This phenomena can be attributed to the much larger BOLD point spread function in 3 Tesla as compared to that of 7 Tesla. According to Shmuel et al. [2007], at 7 Tesla the BOLD PSF has an upper bound of 2.34 ± 0.20 mm, whereas, at 3 Tesla the obtained PSF is much wider, ≈ 3.4 mm. This is because, at 3 Tesla stimulus-evoked BOLD signals from intravascular contributions of large draining veins constitute $\approx 50\%$ of the fMRI signals [Shmuel et al., 2007]. On the other hand the intravascular BOLD contributions become negligible only at MR field strength of 7 Tesla or above, because it provides reduced susceptibility due to lowered echo time (TE) [Duong et al., 2003, Jochimsen et al., 2004].

Inline with Tong et al. [2012], from these findings it can be concluded that there is a positive correlation between the Orientation decoding accuracy and the tSNR

obtained for that acquisition with respect to the particular stimulus. But the correlation is not linear and it depends on other factors like the acquisition resolution, BOLD PSF of that magnetic field strength, contribution of veins etc. Though it is out of scope of this experiment, Tong et al. [2012] also showed how decoding accuracy is varying with the spatial frequency and the contrast of the orientation stimulus. In a future study, it can be investigated, whether there is an optimal combination of stimulus parameters, acquisition resolution, scanning protocol parameters and magnetic field strength which will provide improved decoding accuracies and more detailed insight into the true spatial scale of orientation specific signals.

6. Experiment 4: Multivariate Decoding in Auditory Cortex

6.1 Background and Motivation

Though recently, multivariate decoding is studied in auditory cortex [Gardumi et al., 2016, Vetter et al., 2014], the most extendively studied paradigm over the last decade has been decoding orientation gratings from the BOLD fMRI patterns from the visual cortex. Like most previous literature, the Experiment 1 (Chapter 3) and Experiment 2 (Chapter 4) used an event-related orientation decoding paradigm for investigating the effect of acquistion resolution and MR field strength on multivariate classification analysis. Upon investigating the spatial scale of orientation specific signals with volumetric Gaussian spatial filtering, emerged a particular pattern. In line with Swisher et al. [2010], in both the experiments I found evidences to conclude that orientation specific signals were broadband in nature ranging from $\approx$ 1 mm-10 mm, representing columnar scale contributions to large scale signal biases. Especially, the $\approx$5-8 mm FWHM band showed maximum orientation decoding accuracy across acquistition resolutions and MR field strengths. I wanted to investigate whether the same patterns of decoding were displayed for a different decoding paradigm in a different region of the brain. The subjects who volunteered for the Experiment 1, also participated in the perception of musical genres study [Hanke et al., 2015a], as a part of the *studyforrest* project [Hanke et al., 2014]. Here I re-analyzed the auditory fMRI data for decoding musical genres from activation patterns from auditory cortex and investigate the spatial scale of decoding signals with spatial smoothing.

6.2 Methods

The seven participants passively listened to five natural, stereo, high-quality music stimuli (6 s duration; 44.1 kHz sampling rate) for each of five different musical genres: 1) Ambient, 2) Roots Country 3) Heavy Metal, 4) 50s Rock'n'Roll, and 5) Symphonic

while fMRI was recorded in 7 Tesla Siemens scanner [please see procedural details in Hanke et al., 2015a]. The EPI scans were recorded with 1.4 mm isotropic voxel size (TR=2.00 sec, matrix size 160×160, 36 slices, 10% interslice gap) and were distortion corrected [In and Speck, 2012] and aligned to a per-subject BOLD template image (voxel size 1.2 mm). The multivariate decoding of musical genres from the suditory cortex fMRI signal patterns was performed on the template aligned functional images. The Region of interest (ROI) in the primary auditory cortex (Broadmann Area 41 and 42) included Freesurfer cortical parcellation of superiortemporal and transversetemporal gyri in both left and right hemispheres. Leave-one-run-out nested cross validation was performed with Linear CSVM classifier (similar to Section 3.2.9) and mean classification accuracy across 7 participants are reported in Figure 6.1. The spatial scale of decoding signals were studied with MVP analysis on fMRI images after 3D volumetric Gaussian spatial filtering, with Low pass, High pass, Band Pass and Band stop implementations (implementation details in Section 3.2.10).

6.3 Results and Conclusion

Figure 6.1 shows the multivariate decoding performance of 5 different types of musical genres from primary auditory cortex signal patterns. The theoretical chance level of 20% is shown by dashed line. The unfiltered decoding performance and the decoding accuracies after spatial filtering were higher than chance level. But the pattern of decoding accuracies at different levels of spatial filtering with volumetric Gaussian kernels, resembled the results of orientation decoding from V1 (as shown in Figure 3.10 and Figure 4.4). The bandpass component also show maximum decoding accuracy in the ≈5-8 mm. The lowpass components showed steady decline of performance with increasing kernel width after the 3 mm FWHM mark but does not reach chance level even after 20 mm of smoothing. The highpass components showed better decoding accuracies beyond 9 mm FWHM. The Bandstop component's performance was similar to the lowpass components until the 5 mm FWHM band and then started following highpass

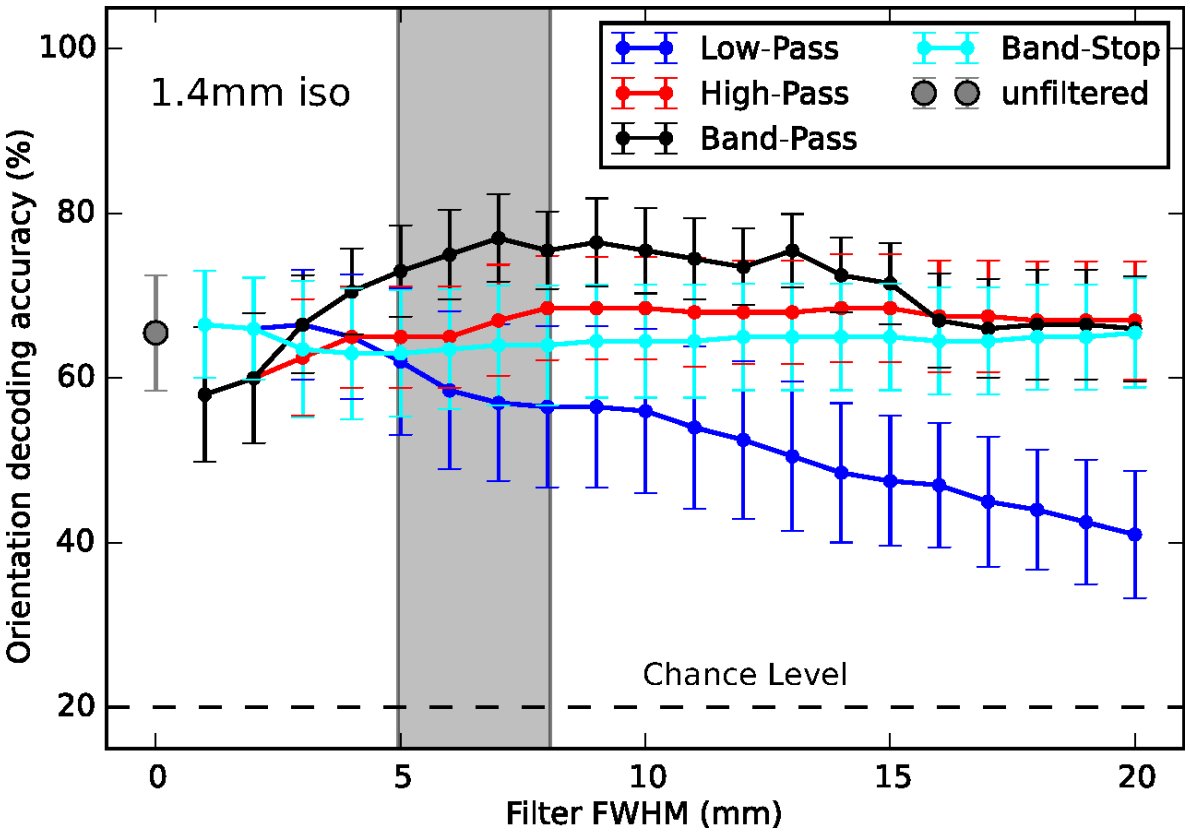

Figure 6.1: Decoding accuracy in auditory cortex after spatial smoothing Classification accuracy of decoding musical genres from BOLD signal patterns from the primary auditory cortex. The unfiltered data showed much higher decoding accuracy than orientation classification in V1. After low pass spatial filtering by gaussian kernel the decoding accuracy stayed above chance level even at 20 mm FWHM. High pass, band pass and band stop filters showed very similar trend as found in orientation decoding across the smoothing spectrum. Peak decoding accuracy was reached at ≈5-8 mm band in the band pass filtered components. The bandpass components showed better accuracy than the high pass components until 15mm FWHM smoothing. This hints at larger contribution of low frequency decoding signal unlike orientation decoding.

components.

In line with Gardumi et al. [2016], the results suggest that decoding information for musical genre classification from the primary auditory cortex are spatially distributed and are represented at different spatial scales. However, similar to V1, though the decoding signals originating from the primary auditory cortex are also spatially broadband in nature, ≈5-8 mm band contributes more than other frequency bands. Keeping in mind, the broadband nature of the decoding signals and the results of the vowel decoding shown in Gardumi et al. [2016], it can be concluded that for multivariate

decoding in the auditory cortex, there is no particular optimal spatial resolution but the acquisition resolution and the corresponding spatial smoothing parameters need to be tailored to the particular decoding task which is being investigated.

It is to be noted that the findings of this experiment show striking similarity with the orientation decoding accuracy patterns obtained from the visual cortex. Until now one of the most debated points in multivariate pattern analysis literature has been the origin and the scale of the orientation decoding signals and most of the previous studies have related it to some anatomical or topographical structure, which is specific to the visual cortex, namely, submillimeter scale orientation columns or large scale orientation maps etc. But it was shown in Linden and Schreiner [2003] how the anatomy and the synaptic physiology of the auditory cortex have important differences from the columnar organization of the visual cortex and these differences have corresponding functional significances too. Irrespective of these anatomical differences and the differences in the experimental stimulations and the psychophysical task performed by the participants, the consistency of the pattern of decoding accuracy (along different levels of spatial smoothing) across two different sensory cortices is highly significant. Does that mean that multivariate decoding is actually dependant on much more basic physical parameters which are consistent across different sensory cortices? Finding an answer to this question is out of the scope of this experiment, but it is extremely relevant for better understanding of the process of multivariate decoding and needs to be investigated in details in future studies.

7. Summary and General Conclusions

Over the last decade, there has been a multitude of studies which used Multivariate Pattern Analysis approach to successfully decode viewed orientations from the BOLD signal patterns in the primary visual cortex [Alink et al., 2013, Op de Beeck, 2010, Boynton, 2005, Freeman et al., 2011, 2013, Gardner, 2010, Haynes and Rees, 2005, Kamitani and Sawahata, 2010, Kamitani and Tong, 2005, Misaki et al., 2013, Pratte et al., 2016, Swisher et al., 2010]. In contrast to the conventional univariate fMRI analysis approach of studying individual voxel response in isolation, multivariate techniques apply machine learning approaches to analyze the patterns of activity of multiple voxels together even if the activity in individual voxel might not be significant. The ability to decode orientations from BOLD patterns recorded by conventional 3 mm isotropic voxels at 3 Tesla(though it has been shown by previous electrophysiological studies that orientation selectivity in V1 originate from sub-millimeter range columnar structures) has led to an ongoing debate about the true spatial scale of orientation signals picked up by the classifiers. There are several hypotheses put forward to explain this mechanism, for example, random irregularities in V1 columnar structures introducing local bias effects in voxel signals [Kamitani and Tong, 2005], exclusive coarse-scale radial biases [Freeman et al., 2011, 2013], functionally organized cortical vasculature providing a complex spatio-temporal filter [Gardner, 2010, Kriegeskorte et al., 2010] etc. In these reports the spatial scale of the orientation signals has been studied by running decoding analysis on either spatially filtered images (by gaussian filtering kernel of variable width) or on spatially resampled images (k-space resampling to other lower resolutions). However it remains inconclusive, what effect data acquisition in different resolutions have on orientation decoding and whether higher magnetic field strength (7 Tesla) has any benefit over the conventional 3 Tesla. These open questions were addresed in this thesis.

7.1 Summary of the experimental procedures and results

In Experiment 1, fMRI data was acquired in four different resolutions at 7 Tesla (0.8 mm, 1.4 mm, 2 mm and 3 mm iso) from 7 participants while they passively watched independent semi-annular orientation gratings (0°, 45°, 90°, or 135°) in both hemifields. Linear Support Vector machine classifier was used with nested cross validation to decode the viewed orientation from the BOLD activity pattern in the contralateral hemisphere. The highest recorded decoding accuracy in the 4-way classification was found in the 2 mm data. From the results, I could conclude that at 2 mm resolution there was an optimal balance of tSNR and sampling frequency, If I went to higher resolution (0.8 mm, 1.4 mm) I could not get enough tSNR due to reduced voxel size, and 3.0 mm data had more tSNR but lacked sampling frequency. Spatial filtering prior to decoding showed that the ≈5-8 mm is the most informative band in orientation decoding from V1. Moreover I showed that the orientation signal is highly broadband in nature, including columnar scale contributions to signals of spatial scale of 10 mm. In this study I also localized the veins in V1 with susceptibility weighted imaging and did decoding analysis both inside and outside the venous voxels. It showed that the vasculature in V1 contributes to above-chance decoding accuracies which is inline with the concept of complex spatio-temporal filters.

In Experiment 1, the 2 mm iso data, which provided maximum decoding accuracy at 7 Tesla, can be acquired with the conventional 3 Tesla scanner with much less distortion than 7 Tesla.Hence the comparison study described in Experiment 2 was performed where orientation decoding was compared across MR field strengths (7 Tesla vs 3 Tesla) in multiple acquisiition resolutions(1.4 mm, 2 mm and 3 mm) to get a better understanding whether multivariate decoding profits from high magnetic field. It was reported that in 3 Tesla decoding accuracies monotonically increased with increase in voxel size and the decoding performance in 2 mm 3 Tesla data was lower than 2 mm 7 Tesla data. It was revealed that the 3 Tesla acquisitions had lower tSNR values than the corresponding 7 Tesla acquisitions. Spatial filtering prior to decoding analysis in the

3 Tesla data showed almost identical patterns of accuracy across the entire spectrum substantiating the claim of orientation signals being broadband.

To determine the dependence of orientation decoding on time series signal to noise ratio in Experiment 3 I performed a separate multiband acquisition at 3 Tesla with 2 mm resolution, with the same experimental paradigm, with no parallel imaging technique to enhance the tSNR. In the multiband 3 Tesla acquisition though substantial increase in tSNR was achieved, the decoding accuracy marginally increased, showing the blurring effect of the much larger BOLD point spread function of 3 Tesla acquisitiion as compared to the 7 Tesla.

As shown in Experiment 1 and Experiment 2, a very consistent pattern of orientation decoding accuracy in V1 emerged with Gaussian spatial filtering across magnetic field strengths and acquisition resolutions. In Experiment 4 I also found that this same pattern is reflected in decoding musical genres from fMRI activation patterns in the primary auditory cortex.

7.2 Conclusion and Future Research

In conclusion, the experiments included in this thesis present an in-depth analysis of some technical aspects in multivariate orientation decoding from V1. This is the first time that multi resolution fMRI data has been compared against each other and also across magnetic field strengths. The results give us a clearer idea about the spatial broadband nature of the orientation specific signals, with particularly high contributions from the ≈5-8 mm band. In line with the simulation studies by Chaimow et al. [2011], it can be concluded from this thesis, that data acquisition for orientation decoding approximately around 2 mm iso provide better decoding performance. The consistency of the pattern of decoding accuracy across different levels of Gaussian smoothing in different experimental paradigms in different regions of the brain, hints at rather uniform factors underlying the ability to decode. These factors might include the technical parameters of scanning, uniformity of anatomical structures across primary sensory areas of the

cortex, decoding analysis procedures etc. and is a subject of further research. Though there has been previous simulation studies modelling the functional organization of vasculature in V1, this thesis for the first time provides empirical evidence about the contribution of veins in orientation decoding across different resolutions. The venous voxels in V1 providing above chance decoding accuracy, support the claim of complex spatio-temporal filter [Kriegeskorte et al., 2010] where it plays an important role in bringing fine scale functional signals to a much lower resolution which can be sampled by the voxel grid. The comparison of MR field strength on orientation decoding shows that 7 Tesla provides better decoding performance than 3 Tesla at resolutions where the higher magnetic field strength can provide better tSNR due to improved BOLD sensitivity. Though decoding accuracy is shown to be positively correlated with tSNR, it is not the sole driving factor. From the results of comparison between normal EPI acquisition and multiband acquisition in 3 Tesla, it can be concluded that the balance of acquisition resolution and corresponding tSNR along with the BOLD PSF, plays an important role in decoding accuracy.

It is to be understood that multivariate decoding of fMRI data is a relatively new field in Neuroscience and that there are a multitude of other factors which are still to be explored and are out of scope of this thesis. One of the most important factors that needs to be investigated is the effect of the classification algorithm on decoding performance. This thesis uses the most widely used Linear support vector machine classifier for easy comparison with previous literature. Keeping in mind the intrinsic non linearity of the decision surface in the multidimensional classification space, non-linear kernels may be better suited for this purpose. Not only different classification algorithms, the optimization of the hyperparameters of these kernels is a subject of extensive research. For example, in this thesis it was shown that C parameter optimization of the Linear SVM kernel by nested cross validation improved decoding accuracy. Regarding feature selection before classification, it is known to be an important step but it has to be investigated what is the optimal way of performing it. In this thesis, univariate

feature selection was performed, but the effect of classifiers like Sparse Multinomial Linear Regressor, which provides in-built feature selection functionality, needs to be studied further. This study shows the immense potential of multiband imaging in case of orientation decoding, because it provides fast acquisition time without substantial loss in BOLD signal which is not possible with parallel imaging techniques. Whether at 7 Tesla multiband imaging can be used to provide better SNR at high resolution EPI acquisitions and the correponding effects on decoding are left for a future study. With the increase of the application of the powerful multivariate methods in fMRI analysis, this thesis is expected to provide valuable technical resources for future studies. I understand that the ultra-high field multi-resolution data acquired as part of this thesis could prove very useful to other researchers and might provide a widely used resource for years to come. Even to facilitate future research, the data was made publicly available in the widely accepted BIDS format and the comparison with 3 Tesla dataset will also be published soon.

BIBLIOGRAPHY

Adams, D.L., Sincich, L.C., Horton, J.C., 2007. Complete pattern of ocular dominance columns in human primary visual cortex. The Journal of Neuroscience 27, 10391–10403.

Alink, A., Krugliak, A., Walther, A., Kriegeskorte, N., 2013. fMRI orientation decoding in V1 does not require global maps or globally coherent orientation stimuli. Frontiers in Psychology 4, 493.

Arcaro, M.J., McMains, S.A., Singer, B.D., Kastner, S., 2009. Retinotopic organization of human ventral visual cortex. The Journal of Neuroscience 29, 10638–10652.

Bartfeld, E., Grinvald, A., 1992. Relationships between orientation-preference pinwheels, cytochrome oxidase blobs, and ocular-dominance columns in primate striate cortex. Proceedings of the National Academy of Sciences 89, 11905–11909.

Op de Beeck, H., 2010. Probing the mysterious underpinnings of multi-voxel fMRI analyses. Neuroimage 50, 567–71.

Bellman, R.E., 1961. Adaptive control processes: a guided tour. volume 4. Princeton University Press: Princeton.

Bonte, M., Hausfeld, L., Scharke, W., Valente, G., Formisano, E., 2014. Task-dependent decoding of speaker and vowel identity from auditory cortical response patterns. The Journal of Neuroscience 34, 4548–4557.

Boynton, G., 2005. Imaging orientation selectivity: decoding conscious perception in V1. Nature Neuroscience 8, 541–2.

Braitenberg, V., Braitenberg, C., 1979. Geometry of orientation columns in the visual cortex. Biological cybernetics 33, 179–186.

Burges, C.J., 1998. A tutorial on support vector machines for pattern recognition. Data mining and knowledge discovery 2, 121–167.

Carlson, T.A., 2014. Orientation decoding in human visual cortex: new insights from an unbiased perspective. The Journal of Neuroscience 34, 8373–8383.

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, 627–42.

Chang, C.C., Lin, C.J., 2011. LIBSVM: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology 2, 27:1–27:27. Software available at `http://www.csie.ntu.edu.tw/~cjlin/libsvm`.

Cox, R.W., 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Computers and Biomedical research 29, 162–173.

Dale, A.M., Fischl, B., Sereno, M.I., 1999. Cortical surface-based analysis: I. Segmentation and surface reconstruction. Neuroimage 9, 179–194.

De Martino, F., Zimmermann, J., Muckli, L., Ugurbil, K., Yacoub, E., Goebel, R., 2013. Cortical depth dependent functional responses in humans at 7t: improved specificity with 3d grase. PloS one 8, e60514.

Dou, W., Speck, O., Benner, T., Kaufmann, J., Li, M., Zhong, K., Walter, M., 2014. Automatic voxel positioning for MRS at 7 t. Magnetic Resonance Materials in Physics, Biology and Medicine 28, 259–270.

Dow, B.M., 2002. Orientation and color columns in monkey visual cortex. Cerebral Cortex 12, 1005–1015.

Duong, T.Q., Yacoub, E., Adriany, G., Hu, X., Uĝurbil, 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. Magnetic resonance in medicine 49, 1019–1027.

Engel, S.A., Glover, G.H., Wandell, B.A., 1997a. Retinotopic organization in human visual cortex and the spatial precision of functional mri. Cerebral cortex 7, 181–192.

Engel, S.A., Glover, G.H., Wandell, B.A., 1997b. Retinotopic organization in human visual cortex and the spatial precision of functional MRI. Cerebral cortex 7, 181–192.

Engel, S.A., Rumelhart, D.E., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.J., Shadlen, M.N., 1994. fMRI of human visual cortex. Nature .

Freeman, J., Brouwer, G., Heeger, D., Merriam, E., 2011. Orientation decoding depends on maps, not columns. The Journal of Neuroscience 31, 4792–804.

Freeman, J., Heeger, D., Merriam, E., 2013. Coarse-scale biases for spirals and orientation in human visual cortex. The Journal of Neuroscience 33, 19695–703.

Friedman, J., Hastie, T., Tibshirani, R., 2001. The elements of statistical learning. volume 1. Springer series in statistics Springer, Berlin.

Furmanski, C.S., Engel, S.A., 2000. An oblique effect in human primary visual cortex. Nature Neuroscience 3, 535–536.

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

Gardumi, A., Ivanov, D., Hausfeld, L., Valente, G., Formisano, E., Uludağ, K., 2016. The effect of spatial resolution on decoding accuracy in fmri multivariate pattern analysis. NeuroImage 132, 32–42.

Glockner, J.F., Hu, H.H., Stanley, D.W., Angelos, L., King, K., 2005. Parallel mr imaging: A user's guide 1. Radiographics 25, 1279–1297.

Glover, G.H., 1999. Deconvolution of impulse response in event-related BOLD fMRI. Neuroimage 9, 416–429.

Gorgolewski, K.J., Poline, J.B., Keator, D.B., Nichols, B.N., Auer, T., Craddock, R.C., Flandin, G., Ghosh, S., Sochat, V., Rokem, A., et al., 2015. Brain imaging data structure-a new standard for describing and organizing human neuroimaging data. INCF Neuroinformatics .

Griswold, M.A., Jakob, P.M., Heidemann, R.M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A., 2002. Generalized autocalibrating partially parallel acquisitions (grappa). Magnetic resonance in medicine 47, 1202–1210.

Haacke, E.M., Xu, Y., Cheng, Y.C.N., Reichenbach, J.R., 2004. Susceptibility weighted imaging (swi). Magnetic Resonance in Medicine 52, 612–618.

Halchenko, Y.O., Hanke, M., 2012. Open is not enough. Let's take the next step: An integrated, community-driven computing platform for neuroscience. Frontiers in Neuroinformatics 6.

Hanke, M., Adelhöfer, N., Kottke, D., Iacovella, V., Sengupta, A., Kaule, F.R., Nigbur, R., Waite, A.Q., Baumgartner, F.J., Stadler, J., 2016. Simultaneous fMRI and eye gaze recordings during prolonged natural stimulation – a studyforrest extension. Scientific Data , Under reviewURL: `http://biorxiv.org/content/early/2016/03/31/046581`.

Hanke, M., Baumgartner, F.J., Ibe, P., Kaule, F.R., Pollmann, S., Speck, O., Zinke, W., Stadler, J., 2014. A high-resolution 7-Tesla fMRI dataset from complex natural stimulation with an audio movie. Scientific Data 1.

Hanke, M., Dinga, R., Häusler, C., Guntupalli, J.S., Casey, M., Kaule, F.R., Stadler, J., 2015a. High-resolution 7-Tesla fMRI data on the perception of musical genres – an extension to the studyforrest dataset. F1000Research 4:174.

Hanke, M., Dinga, R., Häusler, C., Guntupalli, J.S., Casey, M., Kaule, F.R., Stadler, J., 2015b. High-resolution 7-tesla fmri data on the perception of musical genres–an extension to the studyforrest dataset. F1000Research 4.

Hanke, M., Halchenko, Y., Sederberg, P., Olivetti, E., Fründ, I., Rieger, J., Herrmann, C., Haxby, J., Hanson, S., Pollmann, S., 2009. PyMVPA: A Unifying Approach to the Analysis of Neuroscientific Data. Frontiers in Neuroinformatics 3, 3.

Haxby, J.V., 2012. Multivariate pattern analysis of fMRI: the early beginnings. Neuroimage 62, 852–855.

Haynes, J., Rees, G., 2005. Predicting the orientation of invisible stimuli from activity in human primary visual cortex. Nature Neuroscience 8, 686–91.

Haynes, J.D., 2009. Decoding visual consciousness from human brain signals. Trends in cognitive sciences 13, 194–202.

Heidemann, R.M., Ivanov, D., Trampel, R., Fasano, F., Meyer, H., Pfeuffer, J., Turner, R., 2012. Isotropic submillimeter fMRI in the human brain at 7T: Combining reduced field-of-view imaging and partially parallel acquisitions. Magnetic Resonance in Medicine 68, 1506–1516.

Hubel, D.H., Wensveen, J., Wick, B., 1995. Eye, brain, and vision. Scientific American Library New York.

Hubel, D.H., Wiesel, T.N., 1972. Laminar and columnar distribution of geniculo-cortical fibers in the macaque monkey. Journal of Comparative Neurology 146, 421–450.

In, M.H., Speck, O., 2012. Highly accelerated PSF-mapping for EPI distortion correction with improved fidelity. Magnetic Resonance Materials in Physics, Biology and Medicine 25, 183–192.

Jenkinson, M., 2003. Fast, automated, N-dimensional phase-unwrapping algorithm. Magnetic resonance in medicine 49, 193–197.

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. Magnetic resonance in medicine 52, 724–732.

Kamitani, Y., Sawahata, Y., 2010. Spatial smoothing hurts localization but not information: pitfalls for brain mappers. Neuroimage 49, 1949–52.

Kamitani, Y., Tong, F., 2005. Decoding the visual and subjective contents of the human brain. Nature Neuroscience 8, 679–85.

Kaule, F.R., Wolynski, B., Gottlob, I., Stadler, J., Speck, O., Kanowski, M., Meltendorf, S., Behrens-Baumann, W., Hoffmann, M.B., 2014. Impact of chiasma opticum malformations on the organization of the human ventral visual cortex. Hum. Brain Mapp. , 5093–5105.

Kriegeskorte, N., Bandettini, P., 2007. Analyzing for information, not activation, to exploit high-resolution fMRI. Neuroimage 38, 649–62.

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, 1965–76.

Kriegeskorte, N., Goebel, R., Bandettini, P., 2006. Information-based functional brain mapping. P. Natl. Acad. Sci. USA 103, 3863–3868.

Linden, J.F., Schreiner, C.E., 2003. Columnar transformations in auditory cortex? a comparison to visual and somatosensory cortices. Cerebral Cortex 13, 83–89.

Liu, S., Mok, K., Neelavalli, J., Cheng, Y.C.N., Tang, J., Ye, Y., Haacke, E.M., 2014. Improved mr venography using quantitative susceptibility-weighted imaging. Journal of Magnetic Resonance Imaging 40, 698–708.

Mahmoudi, A., Takerkart, S., Regragui, F., Boussaoud, D., Brovelli, A., 2012. Multivoxel pattern analysis for fmri data: A review. Computational and mathematical methods in medicine 2012.

Mannion, D.J., McDonald, J.S., Clifford, C.W., 2010. Orientation anisotropies in human visual cortex. Journal of Neurophysiology 103, 3465–3471.

Millman, K.J., Brett, M., 2007. Analysis of functional magnetic resonance imaging in Python. Computing in Science & Engineering 9, 52–55. doi:10.1109/MCSE.2007.46.

Misaki, M., Luh, W., Bandettini, P., 2013. The effect of spatial smoothing on fMRI decoding of columnar-level organization with linear support vector machine. Journal of Neuroscience Methods 212, 355–61.

Norman, K., Polyn, S., Detre, G., Haxby, J., 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends Cogn Sci 10, 424–30.

Obermayer, K., Blasdel, G.G., 1993. Geometry of orientation and ocular dominance columns in monkey striate cortex. The Journal of Neuroscience 13, 4114–4129.

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.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al., 2011. Scikit-learn: Machine learning in Python. The Journal of Machine Learning Research 12, 2825–2830.

Peirce, J., 2007. PsychoPy–Psychophysics software in Python. Journal of Neuroscience Methods 162, 8–13.

Peirce, J.W., 2008. Generating stimuli for neuroscience using PsychoPy. Frontiers in Neuroinformatics 2.

Pereira, F., Mitchell, T., Botvinick, M., 2009. Machine learning classifiers and fMRI: a tutorial overview. Neuroimage 45, S199–S209.

Poline, J.B., Brett, M., 2012. The general linear model and fmri: does love last forever? Neuroimage 62, 871–880.

Pratte, M.S., Sy, J.L., Swisher, J.D., Tong, F., 2016. Radial bias is not necessary for orientation decoding. NeuroImage 127, 23–33.

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

Sengupta, A., Kaule, F., Guntupalli, J.S., HHoffmann, M., Häusler, C., Stadler, J., Hanke, M., 2016. An extension of the studyforrest dataset for vision research. Scientific Data , In PressURL: http://biorxiv.org/content/early/2016/03/31/046573.

Sereno, M., Dale, A., Reppas, J., Kwong, K., Belliveau, J., Brady, T., Rosen, B., Tootell, R., 1995. Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. Science 268, 889–893.

Sereno, M.I., Lutti, A., Weiskopf, N., Dick, F., 2012. Mapping the human cortical surface by combining quantitative T1 with retinotopy. Cerebral Cortex 23, 2261–2268.

Shmuel, A., Chaimow, D., Raddatz, G., Ugurbil, K., Yacoub, E., 2010. Mechanisms underlying decoding at 7T: 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., Ugurbil, K., 2007. Spatio-

temporal point-spread function of fMRI signal in human gray matter at 7 Tesla. Neuroimage 35, 539–52.

Silver, M.A., Kastner, S., 2009a. Topographic maps in human frontal and parietal cortex. Trends in Cognitive Sciences 13, 488–495.

Silver, M.A., Kastner, S., 2009b. Topographic maps in human frontal and parietal cortex. Trends in Cognitive Sciences 13, 488–495.

Smith, S., Jenkinson, M., Woolrich, M., Beckmann, C., Behrens, T., Johansen-Berg, H., Bannister, P., De, L.M., Drobnjak, I., Flitney, D., Niazy, R., Saunders, J., Vickers, J., Zhang, Y., De, S.N., Brady, J., Matthews, P., 2004. Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage 23 Suppl 1, S208–19.

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

Tong, F., Harrison, S., Dewey, J., Kamitani, Y., 2012. Relationship between BOLD amplitude and pattern classification of orientation-selective activity in the human visual cortex. Neuroimage 63, 1212–22.

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.

Van Essen, D.C., Lewis, J.W., Drury, H.A., Hadjikhani, N., Tootell, R.B., Bakircioglu, M., Miller, M.I., 2001. Mapping visual cortex in monkeys and humans using surface-based atlases. Vision research 41, 1359–1378.

Vetter, P., Smith, F.W., Muckli, L., 2014. Decoding sound and imagery content in early visual cortex. Current Biology 24, 1256–1262.

Wandell, B.A., Dumoulin, S.O., Brewer, A.A., 2007. Visual field maps in human cortex. Neuron 56, 366–383.

Wang, L., Mruczek, R.E., Arcaro, M.J., Kastner, S., 2014. Probabilistic maps of visual topography in human cortex. Cerebral Cortex , bhu277.

Warnking, J., Dojat, M., Guérin-Dugué, A., Delon-Martin, C., Olympieff, S., Richard, N., Chéhikian, A., Segebarth, C., 2002. fMRI retinotopic mapping–step by step. Neuroimage 17, 1665–83.

Weibull, A., Gustavsson, H., Mattsson, S., Svensson, J., 2008. Investigation of spatial resolution, partial volume effects and smoothing in functional MRI using artificial 3d time series. NeuroImage 41, 346–353.

Yacoub, E., Harel, N., Ugurbil, K., 2008. High-field fMRI unveils orientation columns in humans. Proc Natl Acad Sci U S A 105, 10607–12.

Yacoub, E., Shmuel, A., Pfeuffer, J., De Moortele, V., Adriany, G., Andersen, P., Vaughan, J.T., Merkle, H., Ugurbil, K., Hu, X., et al., 2001. Imaging brain function in humans at 7 tesla. Magnetic Resonance in Medicine 45, 588–594.

Zhang, F., Wang, J.P., Kim, J., Parrish, T., Wong, P.C., 2015. Decoding multiple sound categories in the human temporal cortex using high resolution fmri. PloS one 10, e0117303.

Appendices

A. Data Availability

The empirical ultra high-field fMRI dataset recorded at four spatial resolutions (0.8 mm, 1.4 mm, 2 mm, and 3 mm isotropic voxel size) for orientation decoding in visual cortex — in order to test hypotheses on the strength and spatial scale of orientation discriminating signals are openly accessible from the *OpenfMRI* portal (dataset accession number: ds000113c) in *BIDS* (Brain Imaging Data Structure) format.

A.1 Data Specifications Table

Subject area	Neuroimaging
More specific subject area	Early visual system
Type of data	Ultra High Field (7 Tesla) BOLD fMRI
Data format	Raw and distortion corrected BOLD fMRI data stored in compressed NIFTI format; BIDS-compliant
Experimental factors	Acquisition resolution (within-subject factor; 0.8 mm, 1.4 mm, 2 mm, and 3 mm isotropic voxel size)
Data source location	Magdeburg, Germany
Data accessibility	Data available at *OpenfMRI* portal (dataset accession number: ds000113c), as well as Github/ZENODO (DOI: 10.5281/zenodo.46756).

A.2 Value of the data

- first publicly available dataset to provide ultra high-field, multi-resolution BOLD fMRI data for a uniform stimulation paradigm targeting the representation of visual orientations in early visual cortex

- compliant with the brain imaging data structure (BIDS) standard, hence highly suitable for automated processing
- potent dataset for optimization and benchmarking of algorithms, such as pattern classification and feature extraction
- flexible and unrestricted data access down to the level of individual files facilitate cloud-based analysis and utilization in (web-based) demonstrations

A.3 Data structure and usage information

This dataset is compliant with the Brain Imaging Data Structure (BIDS) specification[Gorgolewski et al., 2015], which is a new standard to organize and describe neuroimaging and behavioral data in an intuitive and common manner. Data are shared in documented standard formats, such as NIfTI or plain text files, to enable further processing in arbitrary analysis environments with no imposed dependencies on proprietary tools. Extensive documentation of this standard is available at `http://bids.neuroimaging.io`. This section provides information about the released data, but limits its description to aspects that extends the BIDS specifications. For a general description of the dataset layout and file naming conventions, the reader is referred to the BIDS documentation. In summary, all files related to the data acquisitions for a particular participant described in this manuscript can be located in a `sub-<ID>/ses-r<RES>/` directory, where `ID` is the numeric subject code, and `RES` is a two-digit acquisition resolution identifier.

In order to de-identify data, information on center-specific study and subject codes have been removed using an automated procedure. All human participants were given integer IDs that are consistent across all other data releases of the *studyforrest* project [Hanke et al., 2016, 2014, 2015b, Sengupta et al., 2016].

All data are made available under the terms of the Public Domain Dedication and License (PDDL; `http://opendatacommons.org/licenses/pddl/1.0/`). All

source code is released under the terms of the MIT license (`http://www.opensource.org/licenses/MIT`). In short, this means that anybody is free to download and use this dataset for any purpose as well as to produce and re-share derived data artifacts. While not legally required, we hope that all users of the data will acknowledge the original authors by citing this publication and follow good scientific practise as laid out in the ODC Attribution/Share-Alike Community Norms (`http://opendatacommons.org/norms/odc-by-sa/`).

Participant demographics

A plain text table (`participants.tsv`) contains basic demographics for each participant: gender, age group (five-year bin size), and self-reported handedness.

fMRI data

fMRI data are provided in two flavors: raw (`*run-??_bold.nii.gz`) and distortion-corrected (`*rec-dico_run-??_bold.nii.gz`). While raw BOLD data are suitable for further analysis, they suffer from severe geometric distortions. Distortion correction was applied using an online procedure [In and Speck, 2012] and the resulting data represents the primary data type for further analysis.

Motion estimates

Data motion correction was performed scanner-side as part of the distortion correction procedure, and the associated motion estimates are provided in a whitespace-delimited 6-column text file (`*motion_physio.tsv.gz`; translation X, Y, Z in mm, rotation around X, Y, Z in deg) with one row per fMRI volume for each acquisition run separately.

Stimulus timing

Stimulation timing information for each acquisition run is provided in corresponding `*_events.tsv` files. These four-column text files describe the `onset` and `duration` of

a stimulus trial (in seconds from the acquisition run start) and identify the associated stimulus orientation (in deg) presented in the left (`lh_orientation`), and in the right hemifield (`rh_orientation`). A stimulus orientation label of `none` indicates that no stimulus was present in the respective trial (unilateral stimulation).

Auxilliary scans to facilitate alignment

Data for the additional fMRI acquisition with enhanced spatial coverage at 0.8 mm resolution is provided in `*task-coverage*` files. These images can be used to aid alignment of high-resolution BOLD images with limited coverage to other functional or structural images.

LIST OF FIGURES

LIST OF TABLES

Declaration of Originality

I hereby declare that I have authored this thesis titled *The Effect of Acquisition Resolution and Magnetic Field Strength on Multivariate Decoding of fMRI* independently, that I have not used other than the declared sources/resources, and that I have explicitly marked all material which has been quoted either literally or by content from the used sources. Additionally, this work has neither been used by myself nor by anybody else to attain any academic degree.

Name _____

Magdeburg, _____ _____

Date Signature

Eigenständigkeitserklärung

Hiermit erkläre ich, dass ich die von mir eingereichte Dissertation zum dem Thema *The Effect of Acquisition Resolution and Magnetic Field Strength on Multivariate Decoding of fMRI* selbständig verfasst, nicht schon als Dissertation verwendet habe und die benutzen Hilfsmittel und Quellen vollständig angegeben wurden.

Weiterhin erkläre ich, dass ich weder diese noch eine andere Arbeit zur Erlangung des akademischen Grades doctor rerum naturalium (Dr. rer. nat.) an anderen Einrichtungen eingereicht habe.

Name _____

Magdeburg, am _____ _____

Datum Unterschrift