

Representational content in human cortical systems and visual consciousness

Citation for published version (APA):

Schneider, M. (2019). *Representational content in human cortical systems and visual consciousness*. [Doctoral Thesis, Maastricht University]. ProefschriftMaken Maastricht. <https://doi.org/10.26481/dis.20191009ms>

Document status and date:

Published: 01/01/2019

DOI:

[10.26481/dis.20191009ms](https://doi.org/10.26481/dis.20191009ms)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Doctoral thesis

**REPRESENTATIONAL CONTENT
IN HUMAN CORTICAL SYSTEMS
AND VISUAL CONSCIOUSNESS**

Marian Schneider

2019

© Marian Schneider, Maastricht 2019.

This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). You are free to share and adapt the material for any purpose, even commercially, under the following terms: Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use. No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits. This is a human-readable summary of (and not a substitute for) the license. For the full license text, see: <https://creativecommons.org/licenses/by/4.0/legalcode>

Cover	Marian Schneider, 2019
Production	proefschriftmaken.nl
ISBN	978-94-6380-550-6

**REPRESENTATIONAL CONTENT
IN HUMAN CORTICAL SYSTEMS
AND VISUAL CONSCIOUSNESS**

Dissertation

To obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus, Prof. Dr. Rianne M. Letschert,
in accordance with the decision of the Board of Deans,
to be defended in public on
Wednesday 9th of October 2019, at 10.00 hours

by

Marian Schneider

Promotors

Prof. Dr. Rainer Goebel

Prof. Dr. Nikolaus Weiskopf (Max Planck Institute Leipzig)

Copromotor

Dr. Federico De Martino

Assessment Committee

Prof. Dr. Peter De Weerd (Chair)

Prof. Dr. Serge O. Dumoulin (Spinoza Centre for Neuroimaging)

Dr. Laurentius Huber

Prof. Dr. Lars Muckli (University of Glasgow)

The research presented in this thesis was supported by the Netherlands Organisation for Scientific Research (NWO) Research Talent grant 406-14-108 awarded to Marian Schneider and Rainer Goebel.

Contents

1	GENERAL INTRODUCTION	1
1.1	Overview	3
1.2	Neural correlates of consciousness	3
1.3	Neuroimaging at uhf: new opportunities	10
1.4	Neuroimaging at uhf: new challenges	17
1.5	Outline of the thesis	22
1.6	References	24
2	A SCALABLE METHOD TO IMPROVE GRAY MATTER SEGMENTATION AT ULTRA-HIGH FIELD MRI	31
2.1	Abstract	33
2.2	Introduction	33
2.3	Transfer functions and histograms	37
2.4	Input data requirements and preparation	43
2.5	Validation methods	44
2.6	Validation results	50
2.7	Discussion	52
2.8	Supplementary material	58
2.9	Acknowledgments	67
2.10	References	68
3	COLUMNAR CLUSTERS IN THE HUMAN MOTION COMPLEX REFLECT CONSCIOUSLY PERCEIVED MOTION AXIS	77
3.1	Abstract	79
3.2	Introduction	79
3.3	Methods	82
3.4	Results	102
3.5	Discussion	107
3.6	Supplementary material	111
3.7	Acknowledgments	134
3.8	References	135

4	MOTION DISPLACES POPULATION RECEPTIVE FIELDS IN THE DIRECTION OPPOSITE TO MOTION	143
4.1	Abstract	145
4.2	Introduction	145
4.3	Methods	149
4.4	Results	168
4.5	Discussion	176
4.6	Supplementary material	185
4.7	Acknowledgments	194
4.8	References	195
5	GENERAL DISCUSSION	201
5.1	Overview	203
5.2	Summary	203
5.3	Past and current context	205
5.4	Limitations and future directions	214
5.5	Conclusion: Chalmer's fallen caveat	224
5.6	References	225
	SUMMARY	233
	KNOWLEDGE VALORIZATION	235
	ACKNOWLEDGMENTS	241
	ABOUT THE AUTHOR	245

1

General Introduction

1.1 OVERVIEW

The work presented in this thesis can be understood by tracing two major lines of research. The first research line attempts to elucidate the neural correlates of conscious perception, in particular the contents of visual consciousness in humans. The second line concerns human neuroimaging using (functional) magnetic resonance imaging (f/MRI), specifically the new opportunities and challenges that come with moving to ultra-high magnetic field (UHF; 7 Tesla [T] and above). The aim of the present chapter is to provide background on each of the two research paradigms and to motivate the current work at their intersection.

As pointed out by the philosopher of science Imre Lakatos (1970) every research program can be characterized by a "hard core" and a "protective belt". The hard core is constituted by a sequence of central theses and theoretical assumptions. In order to be able to subject this theory to empirical tests, every research program also has its protective belt of auxiliary hypotheses surrounding the hard core. Researchers spend much of their time on issues related to the protective belt. Therefore, next to tracing the main tenets of the two research programs that underlie the current work, the auxiliary equipment that was necessary to perform the research presented here will also be highlighted. This includes the tools that were needed to work within the two paradigms.

A note on style in the beginning. As is common in contemporary science, a large part of the work presented in this thesis represents a group effort. For this reason, in all empirical chapters the first person plural ("we") will be used. For the [Introduction](#) and [General Discussion](#) chapters, some sections contain personal memories or opinions. In these exceptional cases, the first person singular ("I") will be used.

1.2 NEURAL CORRELATES OF CONSCIOUSNESS

The work presented here is motivated by the question of how human conscious experience relates to cortical activity in our brains. As such, the work is part of a larger research paradigm that aims to identify the neural correlates of consciousness (NCC). To better understand the aim of this paradigm, it is helpful to first clarify what is meant by the term

"neural correlates of consciousness". The philosopher of mind David Chalmers (2000) offered the following approximative definition:

"A neural system N is an NCC if the state of N correlates directly with states of consciousness." (Chalmers, 2000, p. 3)

Taking this definition as a starting point, he then elaborated on two aspects: what the relevant "states of consciousness" are and what it means for a neural state to "correlate directly" with states of consciousness. The focus here is on the states of consciousness, since the sense in which correlation will be used throughout this thesis is equivalent to the scientific notion of correlation, in other words, correlation as operationalized by Pearson's correlation coefficient.

Chalmers (2000) distinguished three broad classes of conscious states. The first class is binary and the conscious states belonging to that class are those of being conscious and of not being conscious. The second class concerns background states of consciousness, for example, being asleep, being awake, dreaming, or being under hypnosis (Chalmers, 2000). Although the neural systems correlating with states in these two classes are crucial to human everyday functioning (as we are reminded of in critical clinical conditions such as coma or death), they are not the subject of the current work.

1.2.1 Content-specific NCC

Instead, the focus of this thesis and "arguably the most interesting states of consciousness" (Chalmers, 2000, p. 4) are the states of subjective experience that humans are in at any given moment in time. This third class is called the contents of consciousness because states in this class are distinguished by their specific representational content. Example states for this class are the experience of a particular sound pattern or the experience of a visual scene. What characterizes these states is that they represent the world to be one way or another. The experience of a visual scene, for example, represents the world to contain particular arrangements of objects with particular shapes and particular spatial relationships to one another. This does not mean the experience needs to be veridical. During hallucination or when experiencing the visual illusions used in this thesis,

for example, the world is not the way that the experience represents it as being. What matters is that all conscious states in this class have distinct representational content.

The search for the neural correlates of the contents of consciousness (in brief, content-specific NCC; Koch, Massimini, Boly, and Tononi, 2016) usually proceeds in two steps (Figure 1.1). First, in a mapping step, researchers determine the receptive field (RF) of a recording site (either a neuron in single-cell electro-physiology studies or a voxel in fMRI studies). This involves determining the selectivity profile of a recording site in response to the variation of a specific parameter. This parameter will differ depending on the aspect of a conscious experience that a researcher is interested in. The parameter could, for example, be the orientation of a stimulus, its motion direction, its position, etc. For the study described in Chapter 3, the parameter of interest was visual motion axis and during the mapping step we determined whether voxels were selective to horizontal or vertical motion. Importantly, during the mapping step, stimuli are presented in an unambiguous manner, meaning that, for awake and healthy human adults, there is only a single perceptual interpretation of the stimulus.

In a second step, subjects are exposed to an ambiguous stimulus (setup) that allows for more than one perceptual interpretation. Typical methods for achieving perceptual ambiguity in visual neuroscience are multistable figures, binocular rivalry or continuous flash suppression. The underlying idea is the same: while the physical stimulus properties remain constant, the contents of consciousness alternate over time between different interpretations of the stimulus. For example, for the study described in Chapter 3 we presented participants with the bistable motion quartet (Ramachandran & Anstis, 1985) which yields alternating conscious experiences of either horizontal or vertical motion although the retinal stimulation remains constant (see Figure 1.1). Participants communicate what they are experiencing at any given moment in time, which is usually achieved via corresponding button presses, while the researcher monitors the concurrent responses in recording sites whose RFs were mapped in the first step.

This procedure allows the researcher to track the representational content in the neural system of interest and to see if it matches the content

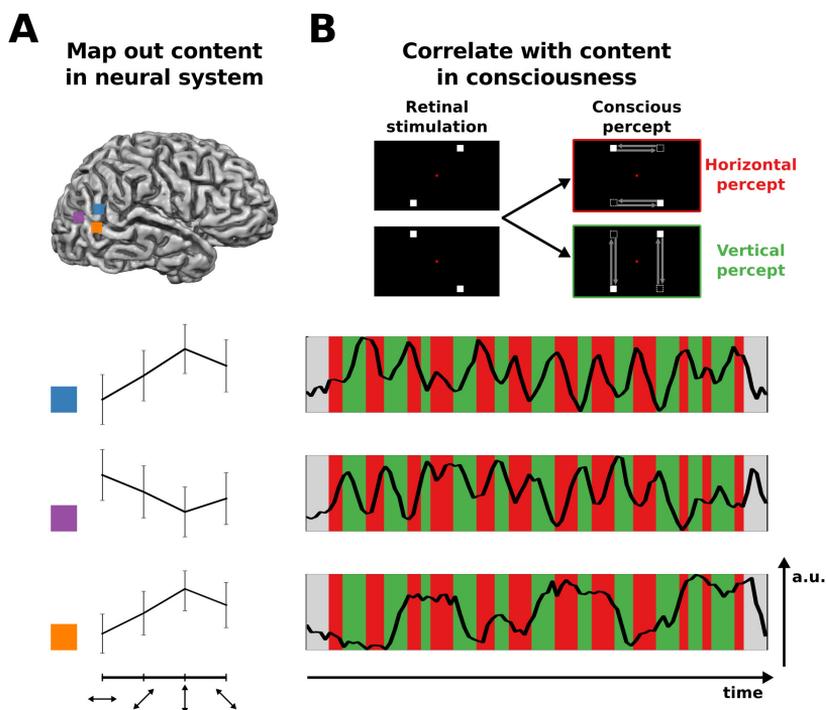


FIGURE 1.1: How to determine neural correlates of the contents of consciousness. **(A)** In a first step, researchers determine the receptive field of a recording site. In this example, the selectivity profile of three different voxels (blue, purple and orange) is determined in response to different axes of visual motion. Axis-of-motion stimuli are presented in an unambiguous manner along four motion axes (horizontal, vertical and two diagonal axes). **(B)** In a second step, subjects are exposed to an ambiguous stimulus. Here, the bistable motion quartet gives rise to the conscious percept of either horizontal (red) or vertical (green) motion although the retinal stimulation remains unchanged. While participants indicate their conscious percept (perceptual periods shown in red or green), the researcher monitors activity in the recording sites (time series). This procedure allows the researcher to track the representational content in the neural system of interest and to see if it matches the content in the subject's consciousness. The voxel in blue, e.g., appears to be tuned to vertical motion and also shows increased signal when the subject indicates perceiving vertical motion (the same applies to the voxel in purple for horizontal motion), while the voxel in orange shows a preference for vertical motion in the first step but does not modulate its signal with the conscious percept in the second step. a.u. stands for arbitrary unit.

of the subject's experience. If a recording site has been identified, for example, as preferring horizontal motion in the first step and shows responses that systematically correlate with the subject indicating the experience of horizontal motion in the second step, then it should be considered (part of) the content-specific NCC for horizontal motion experience.

Compared to the first two classes of conscious states, the search for content-specific NCC thus poses an extra requirement on the NCC. It requires that the representational content of the neural state in question matches the content of the conscious state. The additional requirement yields increased explanatory and predictive power because once a content-specific NCC has been identified it can be used to predict the presence or absence of phenomenal features in conscious experience based on the neural state of the system (Chalmers, 2000). The definition of a content-specific NCC can be updated to capture this additional requirement:

"A neural correlate of the contents of consciousness is a neural representational system N such that representation of a content in N directly correlates with representation of that content in consciousness." (Chalmers, 2000, p. 6)

1.2.2 Chalmer's caveat: content-specific NCC in humans

This definition has an important implication for the search of content-specific NCC. It is crucial that the researcher is able to map out the RF of a recording site so s/he can track the representational content in a neural system by tracking its activity. Twenty years ago, when David Chalmers proposed the NCC definitions above, he noted that, given this requirement, in practice

"...the most informative and useful results usually come from neuron-level studies on monkeys. Large claims are sometimes made for brain imaging on humans, but it is generally difficult to draw solid conclusions from such studies, especially where an NCC is concerned. We can trace the difference to the fact that neuron-level studies can monitor representational content in neural systems, whereas imaging studies cannot (or at least usually do not)." (Chalmers, 2000, p. 25)

This verdict, which I will refer to as "Chalmer's caveat" below, has been damning for fMRI research into the NCC and might explain why many researchers in the consciousness community show a fair amount of skepticism towards brain imaging methods like fMRI. I remember joining my first meeting of the Association for the Scientific Study of Consciousness (ASSC) in 2014, which took place in Brisbane, Australia. The conference committee organized a mentor lunch where young researchers had the opportunity to receive career advice from researchers that were already established in the field. In this context, I had the pleasure of meeting Sid Kouider (Kouider, de Gardelle, Sackur, & Dupoux, 2010), a respected consciousness researcher from the École Normale Supérieure in Paris. Reflecting the general turn of the consciousness community away from fMRI and towards conceptual issues and psychophysical studies, Sid Kouider did not see much potential in fMRI research for studying consciousness at that time ("with the exception of MVPA studies"). I believe he summarized the overall sentiment of the community accurately; at least, during the ASSC conference in Paris in the subsequent year, among hundreds of contributions on consciousness the keyword "fMRI" appeared only twice in the conference proceedings - one of which resulting from our own work.

However, I did not follow Sid Kouider's advice and I believe in several ways the current work is testimony to the progress in human neuroimaging that has taken place since. Although it is still not possible to monitor the activity of single neurons in humans non-invasively, I would submit that Chalmer's caveat for human brain imaging methods is no longer applicable to fMRI. Progress in MRI technology (Duyn, 2012; Ugurbil, 2014; Vaughan et al., 2001) and fMRI analysis techniques (De Martino et al., 2018; Kashyap, Ivanov, Havlicek, Poser, & Uludağ, 2017; Kemper, De Martino, Emmerling, Yacoub, & Goebel, 2018; Polimeni, Renvall, Zaretskaya, & Fischl, 2018; Wandell & Winawer, 2015) have developed fMRI into a tool that allows researchers to monitor representational content also in human brain systems. On the MRI technology side, the work presented here has in particular benefited from the development of UHF MRI scanners for neuroimaging in humans (Uğurbil et al., 2003). With regard to fMRI analysis, two techniques were particularly crucial in the context of this thesis. One is the ability to distinguish and record from distinct (columnar) clusters within human cortex (Cheng, Waggoner, & Tanaka,

2001; Yacoub & Harel, 2008; Yacoub, Shmuel, Logothetis, & Uğurbil, 2007), which was used in Chapter 3. The other one is population RF (pRF) mapping (Dumoulin & Wandell, 2008), which allows to define the RF properties of single voxels in human fMRI studies and was used in Chapter 4. These methodological advances are reviewed below after briefly highlighting why studying content-specific NCC in human, as opposed to non-human, subjects is important.

1.2.3 Why we should study NCC in humans

Studying conscious states in humans directly has obvious advantages. I remember vividly the comments that I received when writing an educational piece on the search for the NCC that was aimed at the general public. To familiarize readers with the research program, I summarized and described classic experiments using binocular rivalry in primates (Logothetis, 1989; Sheinberg & Logothetis, 1997). The comments from the general public reflected some skepticism towards the research agenda; one comment in particular ridiculed the approach: "Haha, they are trying to study consciousness ... in MONKEYS!!!" Admittedly, this comment unduly inflates the difference between human and non-human primates and, in my opinion, we should expect humans and other primates to share fundamental aspects of their conscious experience.

However, I think the comment also points to at least two important aspects. First, compared to studies on monkeys, human neuroimaging studies have the advantage that evidence for conscious content can be gathered more straightforwardly. Monkeys require additional training to learn to communicate their conscious percept. Even once monkeys have learned to press the "corresponding" buttons, drawing conclusions about human consciousness from monkey research will always require the additional assumption that monkeys perceive (visual) illusions in a way similar to us. Admittedly, the privacy of conscious experience combined with the problem of other minds require researchers to make similar assumptions for human subjects. Yet this effort is facilitated by the human ability to use verbal report to describe what they experience. Second, although we should expect many similarities in (content-specific) NCC between human and non-human primates, owing to vast commonalities in brain anatomy

and function, there might also be important differences - especially since the studied primates were macaque monkeys which are phylogenetically more distant from humans than great apes such as chimpanzees or bonobos. These differences can only be explored if appropriate human brain imaging studies are conducted.

1.3 NEUROIMAGING AT UHF: NEW OPPORTUNITIES

1.3.1 *Higher field strength, more sensitivity*

Moving to UHF MRI confers specific advantages and disadvantages to fMRI studies of the human brain. The most prominent advantage of UHF, as compared to conventional field strengths (such as 1.5, 3 and 4 T), is an increase in sensitivity. It is not straightforward to determine by how much exactly signal-to-noise ratio (SNR) increases with additional field strength since, at UHF, SNR is a complex function of resonance frequency (and thus static magnetic field strength [B₀]) but also of head size, head shape, position within the head and tissue composition (Uğurbil et al., 2003). Most derivations suggest at least linear increases in SNR above 3 T (Ladd et al., 2018) and the power factor x in B_0^x for maximum theoretically achievable SNR is estimated to be between 1.2 and 2.1, depending on whether a voxel is at the surface or deep inside the brain (Ladd et al., 2018).

The increase in sensitivity can be leveraged in several ways. It can be used to reduce scan time, to include more experimental conditions or to enable single-subject studies that were starved for SNR at conventional field strengths (Polimeni et al., 2018). The most common use, however, is to trade the gain in SNR for an increase in spatial or temporal resolution (Polimeni et al., 2018). This provides new opportunities and, as many have argued, offers new classes of experiments (for reviews, see De Martino et al., 2018; Dumoulin, Fracasso, van der Zwaag, Siero, and Petridou, 2018; Polimeni et al., 2018). Here, the focus will be on applications that have used the additional SNR from UHF to offset (some of) the SNR that is lost when moving to very high spatial resolution. This is desirable for at least two reasons. First, it allows researchers to translate image SNR gains at UHF to temporal SNR (tSNR) more efficiently. Second, it enables researchers to study how functional properties change within the cortical

ribbon, along surfaces and across cortical depths. The following sections will elaborate on these two aspects in turn.

1.3.2 Higher SNR translates to tSNR only at high-resolution

One motivation to move to higher resolutions is to use the gains in image SNR efficiently. tSNR, which is directly related to statistical power in e.g. univariate analyses, is a linear function of image SNR *only if* no physiological fluctuations exist (Murphy, Bodurka, & Bandettini, 2007). In the presence of physiological noise tSNR plateaus when a particular level of SNR is reached. However, when moving from conventional to high spatial resolutions the noise regime switches from being dominated by physiological modulations to being dominated by thermal image noise (Triantafyllou et al., 2005). This shifts the point where tSNR plateaus to higher levels. It was, for example, demonstrated empirically that, for 7 T fMRI, SNR gains translated to corresponding tSNR gains at higher resolution (e.g. $1 \times 1 \times 3 \text{ mm}^3$) while only moderate tSNR gains were observed at conventional resolution (e.g. $3 \times 3 \times 3 \text{ mm}^3$) where tSNR was limited by physiological noise (Triantafyllou et al., 2005). Thus, choosing the spatial resolution appropriately can help to efficiently translate image SNR to temporal SNR at UHF.

1.3.3 Neuroimaging at the mesoscale

The second motivation to move to high spatial resolutions (in the sub-millimeter range) is that it enables researchers to target a new organizational scale of human brain function. The organization of the human brain, and of cortex more specifically, is often divided into three different scales - the microscopic (the level of single neurons), the mesoscopic (the level of neuronal circuits) and the macroscopic scale (the level of brain areas). Until recently, non-invasive human neuroimaging with fMRI was confined to the macroscopic scale, allowing researchers to identify and study responses from different brain areas and large-scale brain networks. The increase to sub-millimeter resolution afforded by UHF allows for fMRI investigations to move to the mesoscopic scale. This is important for a number of reasons. To begin with, it offers the novel possibility to non-invasively study responses at the level of cortical columns, cortical layers, and subcortical

nuclei. fMRI research at the mesoscale thus promises a plethora of new empirical insights. Furthermore, some scientists assign particular meaning to the mesoscale in cognitive and computational neuroscience research since they believe that (some) mesoscopic structures might represent fundamental units of neural computations (Dumoulin et al., 2018; Hubel & Wiesel, 1974). Accordingly, one key to understanding brain function is to identify these fundamental units and to study their responses. More conservatively, the mesoscopic scale is a relevant organisational level since it represents an intermediate level of analysis where information from individual neurons and brain areas meet (De Martino et al., 2018; Mitra, 2014). At the very least, sub-millimeter fMRI studies might help to bridge the gap between theories of cognitive computation, which are usually formulated at the microscopic or mesoscopic level, and empirical studies in humans, which were hitherto conducted at the macroscopic scale (Marquardt, 2019).

1.3.4 Mapping cortical columnar structures

Sub-millimeter cortical fMRI studies can be divided into so-called "laminar" studies that investigate fMRI responses as a function of cortical depth and "columnar" studies that functionally map columnar structures - i.e. structures that share a singular feature in the cortical depth direction. It is important to note that the meaning of the terms "laminar" and "columnar" in the fMRI community differs from that in cellular and molecular neuroscience. Even at the very high spatial resolutions now achievable with fMRI (0.7 or 0.8 mm isotropic), it is not possible to resolve signal from distinct histological layers or microcolumns. To highlight the difference from histology, laminar fMRI studies are alternatively referred to as "cortical-depth dependent" fMRI. For columnar fMRI there is currently no established alternative although the suffix "-like" (as in "columnar-like") is often employed. Problematically, even within fundamental neuroscience there are different meanings to the term "column", resulting from differences in the defining feature (function, cell constellation, connectivity or myelin content), the extent of columns, and their spatial organization (Rakic, 2008). We have recently proposed the term "columnar clusters" to refer to functionally relevant clusters in reach of high-resolution fMRI (Schneider, Kemper, Emmerling, De Martino, & Goebel, 2019) and this

term or the term "columnar-like" will be used throughout this thesis to avoid confusion with histological definitions.

These conceptual issues should not distract from the technological strikes that now enable researchers to record responses in living human beings from different cortical depth levels (such as deep, mid, and superficial levels) or from distinct cortical clusters within the same brain area. While several studies have used UHF fMRI to investigate laminar differences with fMRI (De Martino et al., 2013; Huber et al., 2015; Huber et al., 2017; Kashyap et al., 2018; Kok, Bains, Van Mourik, Norris, & De Lange, 2016; Koopmans, Barth, & Norris, 2010; Koopmans, Barth, Orzada, & Norris, 2011; Marquardt, Schneider, Gulban, Ivanov, & Uludağ, 2018; Muckli et al., 2015; Olman et al., 2012; Polimeni, Fischl, Greve, & Wald, 2010; Ress, Glover, Liu, & Wandell, 2007), the focus here is on columnar studies.

The first fMRI studies to identify columnar-like structures in humans were confined to primary visual cortex (V1) (Cheng et al., 2001; Yacoub & Harel, 2008; Yacoub et al., 2007). This offered important simplifications. First, the ocular dominance and orientation columns in V1 were already well-established by animal electro-physiology (Hubel & Wiesel, 1974) and post-mortem histology. Second, these structures were relatively large. Third, in most human individuals V1 is relatively flat, simplifying the recording of signal at different cortical depths and enabling the acquisition even with sequences that have minimal coverage in the slice direction (even with a single slice; Yacoub and Harel, 2008). Leaving striate cortex required the development of accurate sampling techniques to deal with the folded cortex in humans (Kemper et al., 2018; Polimeni et al., 2018). Subsequently, researchers identified columnar clusters in human V2, V3, V4 (Nasr, Polimeni, & Tootell, 2016; Tootell & Nasr, 2017), V3a (Goncalves et al., 2015) and V5 (Zimmermann et al., 2011) as well as in A1 (De Martino et al., 2015).

All these studies used unambiguous sensory stimuli (first step in the NCC approach) but did *not* present ambiguous stimuli that allowed for dissociating neural signals pertaining to conscious perception from those related to sensory stimulation (second step in the NCC approach). Yet the ability to map distinct cortical clusters opens up an important avenue for the search of content-specific NCC. Since each cluster is characterized by

a shared preference for a specific feature (e.g. a stimulus orientation, a motion direction, etc.) researchers can now monitor the representational content of columnar systems. Mapping out functional cortical clusters and measuring their responses during ambiguous sensory stimulation thus fulfills the criteria for the search for content-specific NCC outlined by Chalmers (2000) about 20 years ago and is now feasible in humans with UHF fMRI.

1.3.5 More specificity at UHF

Next to the increase in sensitivity, fMRI also benefits from greater specificity of the blood oxygenation level dependent (BOLD) contrast at UHF. T2 and T2* relaxation times decrease more rapidly for blood than for neuronal tissue when moving to UHF (Uğurbil et al., 2002). As a consequence, the relative intravascular contribution to the BOLD signal decreases (Uludağ & Blinder, 2016; Uludağ, Müller-Bierl, & Uğurbil, 2009). This is desirable since intravascular signal partly originates from large ascending cortical and pial veins which are known to show less specific responses than neuronal tissue (De Martino et al., 2018; Moerel et al., 2017). The shift towards contributions from smaller diameter vessels therefore increases specificity at higher fields. As pointed out by De Martino et al., 2018, however, although the extravascular signal around smaller vessels increases more (increase with the square of B0) than signal around larger vessels (approximately linear increase with B0), the signal around larger vessels still *increases*. Especially for gradient echo (GE) sequences, which have been employed for the research presented in this thesis, signal is therefore still affected by larger, less specific vessels - an issue we will return to in the [General Discussion](#) chapter. Next to the increase in BOLD contrast, fMRI at UHF can additionally gain in neuronal specificity if images are acquired at high spatial resolution because partial volume effects are reduced (De Martino et al., 2018).

1.3.6 Population receptive field modelling

It is currently not possible to map out RF properties of single neurons in humans using non-invasive neuroimaging methods such as fMRI. However, the introduction of the population RF (pRF) mapping framework

(Dumoulin & Wandell, 2008) enables researchers to estimate the spatial RF properties of single voxels in a human fMRI study. The pRF mapping framework is based on the assumption that the neural population contained within a voxel has vastly overlapping spatial RF properties. This assumption is reasonable given that spatial RF properties are known to vary smoothly along the cortical sheet.

The pRF for a single voxel is determined in a simple procedure. The voxel time course is assumed to be generated by a 2D Gaussian model. This model can be summarized using only three parameters - two parameters to describe the position of the 2D Gaussian in visual space (conventionally called x and y parameters) and one parameter to describe its size (conventionally called σ parameter). During the pRF mapping experiment, stimuli are presented that occupy varying positions in visual space (Figure 1.2). This can, for example, be a bar stimulus that, over the course of a pRF mapping run, steps through the visual field. Conventionally, pRF mapping stimuli reveal a checkerboard pattern (Dumoulin & Wandell, 2008), although alternative carrier patterns have been suggested and used (Alvarez, de Haas, Clark, Rees, & Schwarzkopf, 2015).

Depending on the specific parameters that are assumed, the convolution of the 2D Gaussian with the spatial apertures of the stimulus presented during the pRF mapping run leads to distinct time series predictions. For example, if a 2D Gaussian is assumed to be centered on the left horizontal meridian several degrees away from fixation, then a bar stimulus stepping through the visual field from left to right should evoke activity in the beginning but not the end of the mapping run. By systematically varying the parameters of the 2D Gaussian and comparing the predicted time series against the measured time series of a voxel, the three pRF parameters can be optimized to accurately summarize the spatial pRF properties of a voxel. To account for hemodynamic delay in the fMRI response, the predicted model time course is usually convolved with a canonical double-gamma function (Friston et al., 1998).

Since its conception (Dumoulin & Wandell, 2008), pRF mapping has yielded many exciting insights into the organization of visual field maps in humans (Amano, Wandell, & Dumoulin, 2009; Fracasso, Petridou, & Dumoulin, 2016; Harvey & Dumoulin, 2011; Harvey, Fracasso, Petridou, & Dumoulin, 2015; Kay, Weiner, & Grill-Spector, 2015; Kay, Winawer, Mezer,

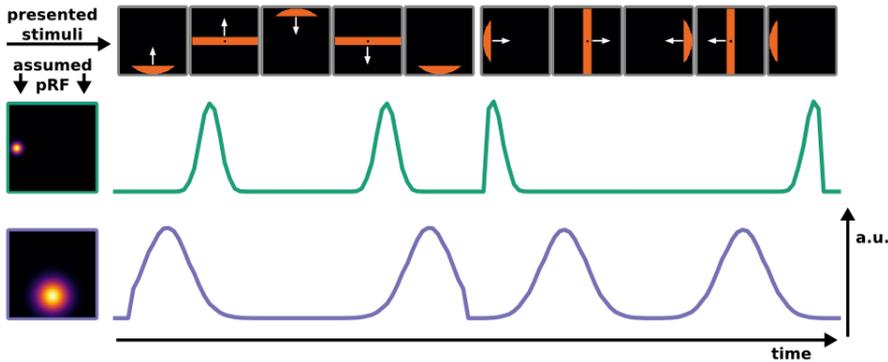


FIGURE 1.2: Overview of the pRF mapping framework. During the pRF mapping experiment, stimuli are shown at systematically varied positions. In this example, a bar stimulus first stepped from bottom to top of the visual field and back, and then from left to right and back. Exemplary bar positions are depicted in orange. Depending on the position and size of the assumed pRF, the stimulus sequence evokes a characteristic response time course. In this example, two pRFs are assumed (depicted on the left, surrounded by green or purple). The green pRF shows responses only when the bar stimulus crosses the left, central region of the visual field. The purple pRF shows broader tuning due to its larger size and shows responses when the bar crosses the central, bottom part of the visual field. a.u. stands for arbitrary unit.

& Wandell, 2013; Klein, Harvey, & Dumoulin, 2014; Winawer, Horiguchi, Sayres, Amano, & Wandell, 2010; Zuiderbaan, Harvey, & Dumoulin, 2012). Although it is possible to measure pRFs at lower resolutions and conventional field strengths, pRF mapping benefits from the increase in spatial resolution and neuronal specificity afforded by UHF since the distribution of neurons in terms of their RF becomes more homogeneous (De Martino et al., 2018). As a result, pRF estimates will be more specific. This is relevant for studies aiming to measure small changes in pRF properties in response to an experimental manipulation (see Chapter 4) since more specific pRF estimates yield increased power for picking up pRF differences.

Although pRF mapping benefits from higher spatial resolution, this does not necessarily mean that *sub-millimeter* resolution is beneficial in this context. It is sometimes forgotten in the UHF community that there is no imperative obligating researchers to acquire images at mesoscopic resolu-

tion when operating UHF scanners. Admittedly, given the considerations about tSNR and partial volume effects outlined above, it is preferable to operate high field scanners at high spatial resolutions. Furthermore, the increase in sensitivity that comes with UHF offsets the loss in SNR incurred by increasing voxel resolution. Consequently, sensitivity is roughly similar between a 2 mm isotropic voxel at 3T and a 1.6 mm isotropic voxel at 7 T (De Martino et al., 2018). However, increasing resolution even further comes at the expense of sensitivity. Depending on the research question, measuring responses at the sub-millimeter scale might therefore not be desirable. This can be the case for pRF mapping. The increase in specificity for pRF estimates motivates higher resolutions, yet if the research question does not concern laminar or columnar neuronal populations, sub-millimeter resolution unnecessarily sacrifices sensitivity and resolutions between 1.1 and 1.6 mm isotropic might represent a more appropriate choice (see Chapter 4).

1.4 NEUROIMAGING AT UHF: NEW CHALLENGES

Despite the distinct advantages that come with performing fMRI experiments at UHF, there are also several challenges that need to be addressed first to harvest the full potential of UHF fMRI, especially for fMRI studies with sub-millimeter resolution. Here, four methodological problems are highlighted since tackling them was important ground work for the empirical research presented in Chapter 3 and 4. Additional challenges for UHF analyses and their impact on statistical analyses have been discussed and reviewed elsewhere (De Martino et al., 2018; Polimeni et al., 2018).

1.4.1 Wanted: an anatomical reference

Many of the challenges at UHF are interlinked and relate to building an anatomical reference for the fMRI data (Polimeni et al., 2018). An anatomical reference is necessary, for example, in order to determine which fMRI responses originate from cortex, to sample cortical-depth dependent responses (laminar or columnar), or to reconstruct cortical surfaces for delineating specialised cortical areas such as visual field maps. With this aim in mind, next to functional images, researchers usually acquire images

that are rich in anatomical contrast. A common choice for this are images from magnetization-prepared T1-weighted pulse sequences (MPRAGE and MP2RAGE) (Marques et al., 2010; Mugler & Brookeman, 1990; Van de Moortele et al., 2009). These sequences offer good white matter (WM) / gray matter (GM) / cerebrospinal fluid (CSF) contrast and can be recorded at high spatial resolution (0.6 mm isotropic) in less than 15 minutes.

Whether to opt for the MPRAGE or MP2RAGE sequence is mostly a matter of taste; yet there are some considerations that can guide the decision. Both sequences are negatively affected by the non-uniform radiofrequency (RF) transmit field (B1+) at UHF. In the case of MPRAGE, this problem can be addressed by acquiring, next to the T1 weighted (T1W) image, a matched proton-density weighted (PDw) image (Van de Moortele et al., 2009). Division of the two images normalizes intensity bias. The MP2RAGE sequence is a modification of the MPRAGE sequence that records two different images at different inversion times and combines them, during the reconstruction process, to reduce B1+ bias (Marques et al., 2010).

These differences in acquisition and reconstruction lead to distinct advantages and disadvantages. Acquisition time for MPRAGE images can be slightly longer than for MP2RAGE images - although the difference is 3-5 minutes at most for sub-millimeter resolution. Most software packages have been bench-marked on T1w MPRAGE images, which might make processing of MPRAGE images more straightforward, although the MP2RAGE sequence also yields a "uni" image that offers similar contrast. Lastly, the MP2RAGE sequence usually yields images with superior WM contrast, while MPRAGE images allow for better delineation of the GM/CSF boundary. This is because division of T1w by PDw images highlights blood vessels (Van de Moortele et al., 2009) which can then be correctly excluded from GM (at least by manual or semi-automatic segmentation methods; see Chapter 2). As a rule of thumb, MPRAGE images can therefore be recommended when high-quality delineations of both WM/GM and GM/CSF boundaries matter (e.g. for columnar fMRI; see Chapter 3) while MP2RAGE images should be used if only a good delineation of the WM/GM border is required (e.g. for surface reconstruction and display of visual field maps; see Chapter 4).

1.4.2 Co-registration

Once functional and anatomical images have been acquired, one of the first challenges is to co-register them. The same challenge exists with images acquired at conventional field strengths but there are at least two peculiarities that make this task more challenging at UHF. The first is that the functional images will show geometric distortions ([see below](#)). The second is that functional images will be acquired even more often with partial coverage only. In such cases, boundary-based registration (BBR) (Greve & Fischl, 2009) can offer satisfactory solutions, provided that the functional images show at least some WM/GM contrast. BBR optimizes the position of the WM/GM surface reconstructed from the anatomical images with respect to the boundary between WM and GM in the fMRI data. Since BBR thus attempts to match the WM/GM boundaries in anatomical and functional data, it is ideally suited for fMRI analyses that require very accurate alignment of cortical boundaries (Polimeni et al., 2018) (which was the case for the studies in Chapter 3 and 4).

1.4.3 Geometric distortions

Two of the most commonly used sequences for the acquisition of high-resolution functional images are 2D (Feinberg et al., 2010; Moeller et al., 2010; Setsompop et al., 2012) and 3D (Poser, Koopmans, Witzel, Wald, & Barth, 2010) gradient-echo (GE) echo planar imaging (EPI) sequences (alternatives are reviewed in the [General Discussion](#) chapter). The resulting EPI data suffer from larger geometric distortions at UHF meaning that they no longer represent size and shape of the brain accurately. This is problematic since anatomical images obtained from MPAGE derivatives retain the shape of the brain relatively well and, as a result, functional and anatomical images no longer match. Different solutions exist for this problem, three of which have been reported (De Martino et al., 2018). One can record images that are matched in their distortions to the functional images but have sufficient contrast for segmentation of GM (Kashyap et al., 2017; Renvall, Witzel, Wald, & Polimeni, 2016); one can perform segmentation on functional images directly provided that they have sufficient contrast (Fracasso et al., 2016); or, one can attempt to correct the distortions in the functional images (Emmerling, Zimmermann, Sorger, Frost, & Goebel,

2016; Marquardt et al., 2018). All these solutions can be understood as a trade-off between co-registration and segmentation performance. While the first two solutions minimize co-registration error at the expense of segmentation error, the last solution minimizes segmentation error but raises the possibility of mismatches between functional and anatomical images.

1.4.4 Registration between sessions

Although the increase in SNR that comes with imaging at UHF offsets some of the SNR loss that is incurred by moving to very high resolutions, not all of the SNR is necessarily offset. This can be remedied by acquiring more data and by using averaging techniques. If one is interested in single-subject columnar or laminar effects then, depending on the effect size, even for a simple experiment with only two experimental conditions, data from more than one scanning session might be required to have sufficient statistical power (see Chapter 3). Likewise, if the resolution is not sub-millimeter but still high (between 1 to 1.2 mm isotropic) and many different experimental conditions are required, sufficient data for averaging is still crucial (see pRF mapping in Chapter 4).

Combining data from different scanning sessions introduces a new registration problem. In particular, due to small differences in head position with regard to the magnetic field and potential differences in shimming, the geometric distortions of functional EPI data acquired in different sessions can differ substantially. Some of these inter-session differences seem to persist even after functional images have been distortion corrected with session-specific field map data. There are different approaches to deal with this problem. One can, for example, use BBR to align functional images from different sessions to the same anatomical image (Nasr et al., 2016), which might yield satisfactory results especially when the registration is restricted to be driven by a smaller, ideally less distorted region of interest. Alternatively, if the functional images provide sufficient contrast one can use linear (Jenkinson, Bannister, Brady, & Smith, 2002; Jenkinson & Smith, 2001) or non-linear algorithms (Andersson, Jenkinson, & Smith, 2007) to register functional images from different sessions to one another. This registration step can take place either before or after images have been

distortion corrected. In our experience, results are more satisfactory when registration is performed first. For high-resolution data in the order of 1.1 - 1.2 mm isotropic, it seems to be possible to achieve good registration of images from different sessions using linear registration with 12 degrees of freedom (Gulban et al., 2018; Chapter 4). However, for sub-millimeter data we found the distortion differences sufficiently severe such that linear registration needed to be complemented with non-linear registration (Schneider et al., 2019; Chapter 3).

1.4.5 Tissue class segmentation

The analysis of laminar or columnar fMRI data requires an accurate definition of the cortical ribbon (i.e. of both the WM/GM and the GM/CSF boundary). Images acquired at UHF and high-resolution reveal several structures outside of the brain that were previously barely visible (Viviani et al., 2017). Such structures include the dura mater (van der Kouwe, Benner, Salat, & Fischl, 2008), medium-sized blood vessels in the sulci (Viviani, 2016) as well as draining sinuses and connective tissue adjacent to GM (Helms, Kallenberg, & Dechent, 2006). Problematically, most of the existing brain segmentation algorithms have been developed and bench-marked on images collected at 1 mm isotropic resolution or lower and at conventional field strengths (Helms, 2016). Therefore, they do not model the non-brain structures now apparent in UHF images and often falsely label them as GM. The current remedy is usually for (graduate) researchers to correct the misclassified voxels manually, which is an outspokenly laborious task. Unfortunately, a lot of time needs to be dedicated to this task and tissue class segmentations represent an important bottleneck for high-resolution UHF processing pipelines.

1.4.6 Interim summary

In summary, the introduction of UHF MRI scanners offers many opportunities for human neuroimaging and, in combination with appropriate modelling techniques, opens new avenues for human consciousness research with fMRI. It should be emphasized that none of the challenges above prevent high-field high-resolution fMRI studies. Yet the degree to which the above problems are solved varies on a continuum. Some of

these issues have found their adequate solution. For example, as long as researchers keep in mind (at the acquisition stage) that sufficient WM/GM contrast in functional images is required, BBR will yield very satisfactory results for co-registration. Other challenges, like geometric distortions, have found a workable approach for now but could benefit from further improvements in the future. However, one of the biggest unsolved issues that, in my opinion, prevents UHF fMRI to be used routinely (for example, in studies with large sample sizes or clinical populations) is tissue class segmentation. Manual correction is an unsatisfactory workaround at best since both time and potential for error scale with the number of voxels (which is high in high-resolution studies). When I started out as a young graduate student, instead of pondering the great questions of human consciousness, I found myself spending tedious hours on deciding, for many voxels, whether they belonged to GM or not. For these reasons, we found it necessary to address the issue of tissue class segmentation first in a methodological study (see Chapter 2) before turning to the empirical study of content-specific NCC in humans (see Chapter 3 and 4).

1.5 OUTLINE OF THE THESIS

This thesis comprises three studies, one methodological and two empirical. The methodological study presented in Chapter 2 addresses the important challenge of tissue class segmentation at UHF. We propose and implement a method that corrects mislabeled, non-brain voxels efficiently and semi-automatically by representing three-dimensional anatomical images in a two-dimensional histogram. We show that when using this method, improved cortical borders can be obtained in a time-efficient manner.

After laying this important groundwork, we turn to the study of visual content-specific NCC in humans, leveraging the additional sensitivity and specificity provided by UHF fMRI. In two empirical studies, we chose experimental paradigms that allowed us to dissociate cortical activity pertaining to retinal properties of the stimulus from that related to the conscious percept. Specifically, for the study presented in Chapter 3, we hypothesized that the conscious experience of a specific visual motion axis is reflected in response amplitudes of direction-selective clusters in the human motion complex. To test this hypothesis, we used neuroimaging at

sub-millimeter resolution which allowed us to map columnar clusters that were tuned to either horizontal or vertical motion presented in an unambiguous motion display. We then recorded the clusters' responses while human observers were shown the bistable motion quartet and reported the perceived axis of motion at any given point in time. We found that observer's perceptual states were dissociatively reflected in the response amplitudes of the identified horizontal and vertical columnar clusters.

For the study presented in Chapter 4, we recorded fMRI responses in early and mid-level visual cortex while human participants were presented with pRF mapping stimuli. We systematically varied two factors. We varied the motion direction of the carrier pattern (3 levels: inward, outward and flicker motion) and the contrast of the mapping stimulus (2 levels: low and high stimulus contrast). This allowed us to dissociate the retinal location of the pRF mapping stimuli from their perceived location. We observed that while physical positions were identical across all conditions, presence of low-contrast motion, but not high-contrast motion, shifted perceived stimulus position in the direction of motion. Correspondingly, we found that pRFs in early visual cortex were shifted against the direction of motion for low-contrast stimuli but not for high stimulus contrast.

1.6 REFERENCES

- Alvarez, I., de Haas, B., Clark, C. A., Rees, G., & Schwarzkopf, D. S. (2015). Comparing different stimulus configurations for population receptive field mapping in human fMRI. *Frontiers in Human Neuroscience*, 9, 1–16.
- Amano, K., Wandell, B. A., & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, 102(5), 2704–18.
- Andersson, J. L. R., Jenkinson, M., & Smith, S. (2007). *Non-linear registration, aka spatial normalisation. FMRIB Technical Report TR07JA2*. FMRIB.
- Chalmers, D. J. (2000). What is a neural correlate of consciousness? In *Neural correlates of consciousness: empirical and conceptual questions* (Chap. 1, pp. 1–27). MIT Press.
- Cheng, K., Waggoner, R. A., & Tanaka, K. (2001). Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron*, 32, 1–16.
- De Martino, F., Moerel, M., Ugurbil, K., Goebel, R., Yacoub, E., & Formisano, E. (2015). Frequency preference and attention effects across cortical depths in the human primary auditory cortex. *Proceedings of the National Academy of Sciences*, 112(52), 16036–16041.
- De Martino, F., Yacoub, E., Kemper, V., Moerel, M., Uludag, K., De Weerd, P., ... Formisano, E. (2018). The impact of ultra-high field MRI on cognitive and computational neuroimaging. *NeuroImage*, 168, 366–382.
- 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(3), e60514.
- Dumoulin, S. O., Fracasso, A., van der Zwaag, W., Siero, J. C., & Petridou, N. (2018). Ultra-high field MRI: Advancing systems neuroscience towards mesoscopic human brain function. *NeuroImage*, 168, 345–357.
- Dumoulin, S. O. & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *NeuroImage*, 39(2), 647–660.
- Duyn, J. H. (2012). The future of ultra-high field MRI and fMRI for study of the human brain. *NeuroImage*, 62(2), 1241–1248.

-
- Emmerling, T. C., Zimmermann, J., Sorger, B., Frost, M. A., & Goebel, R. (2016). Decoding the direction of imagined visual motion using 7 T ultra-high field fMRI. *NeuroImage*, *125*, 61–73.
- Feinberg, D. A., Moeller, S., Smith, S. M., Auerbach, E., Ramanna, S., Glasser, M. F., ... Yacoub, E. (2010). Multiplexed echo planar imaging for sub-second whole brain fmri and fast diffusion imaging. *Plos One*.
- Fracasso, A., Petridou, N., & Dumoulin, S. O. (2016). Systematic variation of population receptive field properties across cortical depth in human visual cortex. *NeuroImage*, *139*, 427–438.
- Friston, K. J., Fletcher, P., Josephs, O., Holmes, a., Rugg, M. D., & Turner, R. (1998). Event-related fMRI: characterizing differential responses. *NeuroImage*, *7*(1), 30–40.
- Goncalves, N. R., Ban, H., Sanchez-Panchuelo, R. M., Francis, S. T., Schluppeck, D., & Welchman, A. E. (2015). 7 Tesla fMRI Reveals Systematic Functional Organization for Binocular Disparity in Dorsal Visual Cortex. *Journal of Neuroscience*, *35*(7), 3056–3072.
- Greve, D. N. & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, *48*(1), 63–72.
- Gulban, O. F., De Martino, F., Vu, A. T., Yacoub, E., Uğurbil, K., & Lenglet, C. (2018). Cortical fibers orientation mapping using in-vivo whole brain 7T diffusion MRI. *NeuroImage*.
- Harvey, B. M. & Dumoulin, S. O. (2011). The Relationship between Cortical Magnification Factor and Population Receptive Field Size in Human Visual Cortex: Constancies in Cortical Architecture. *Journal of Neuroscience*, *31*(38), 13604–13612.
- Harvey, B. M., Fracasso, A., Petridou, N., & Dumoulin, S. O. (2015). Topographic representations of object size and relationships with numerosity reveal generalized quantity processing in human parietal cortex. *Proceedings of the National Academy of Sciences*, *112*(44), 13525–30.
- Helms, G. (2016). Segmentation of human brain using structural MRI. *Magnetic Resonance Materials in Physics, Biology and Medicine*, *29*(2), 111–124.
- Helms, G., Kallenberg, K., & Dechent, P. (2006). Contrast-driven approach to intracranial segmentation using a combination of T2- and T1-

- weighted 3D MRI data sets. *Journal of Magnetic Resonance Imaging*, 24, 790–795.
- Hubel, D. H. & Wiesel, T. N. (1974). Uniformity of monkey striate cortex: A parallel relationship between field size, scatter, and magnification factor. *Journal of Comparative Neurology*.
- Huber, L., Goense, J., Kennerley, A. J., Trampel, R., Guidi, M., Reimer, E., ... Möller, H. E. (2015). Cortical lamina-dependent blood volume changes in human brain at 7 T. *NeuroImage*, 107, 23–33.
- Huber, L., Handwerker, D. A., Jangraw, D. C., Chen, G., Hall, A., Stüber, C., ... Bandettini, P. A. (2017). High-Resolution CBV-fMRI Allows Mapping of Laminar Activity and Connectivity of Cortical Input and Output in Human M1. *Neuron*, 96, 1253–1263.
- Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2), 825–841.
- Jenkinson, M. & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2), 143–156.
- Kashyap, S., Ivanov, D., Havlicek, M., Poser, B., & Uludağ, K. (2017). Impact of acquisition and analysis strategies on cortical depth-dependent fMRI.
- Kashyap, S., Ivanov, D., Havlicek, M., Sengupta, S., Poser, B. A., & Uludağ, K. (2018). Resolving laminar activation in human V1 using ultra-high spatial resolution fMRI at 7T. *Scientific Reports*.
- Kay, K. N., Weiner, K. S., & Grill-Spector, K. (2015). Attention reduces spatial uncertainty in human ventral temporal cortex. *Current Biology*, 25(5), 595–600.
- Kay, K. N., Winawer, J., Mezer, A., & Wandell, B. A. (2013). Compressive spatial summation in human visual cortex. *Journal of Neurophysiology*, 110(2), 481–94.
- Kemper, V. G., De Martino, F., Emmerling, T. C., Yacoub, E., & Goebel, R. (2018). High resolution data analysis strategies for mesoscale human functional MRI at 7 and 9.4T. *NeuroImage*, 164, 48–58.
- Klein, B. P., Harvey, B. M., & Dumoulin, S. O. (2014). Attraction of position preference by spatial attention throughout human visual cortex. *Neuron*, 84(1), 227–237.

-
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience*, *17*(5), 307–321.
- Kok, P., Bains, L. J., Van Mourik, T., Norris, D. G., & De Lange, F. P. (2016). Selective activation of the deep layers of the human primary visual cortex by top-down feedback. *Current Biology*, *26*(3), 371–376.
- Koopmans, P. J., Barth, M., & Norris, D. G. (2010). Layer-specific BOLD activation in human V1. *Human Brain Mapping*, *31*(9), 1297–1304.
- Koopmans, P. J., Barth, M., Orzada, S., & Norris, D. G. (2011). Multi-echo fMRI of the cortical laminae in humans at 7 T. *NeuroImage*, *56*(3), 1276–1285.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*(7), 301–307.
- Ladd, M. E., Bachert, P., Meyerspeer, M., Moser, E., Nagel, A. M., Norris, D. G., ... Zaiss, M. (2018). Pros and cons of ultra-high-field MRI/MRS for human application. *Progress in Nuclear Magnetic Resonance Spectroscopy*, *109*, 1–50.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In *Criticism and the growth of knowledge: proceedings of the international colloquium in the philosophy of science, london, 1965* (Vol. 4, pp. 91–196). Cambridge University Press.
- Logothetis, N. K. (1989). Neuronal correlates of subjective visual perception. *Science*, *245*, 761–763.
- Marquardt, I. (2019). *Investigations on bottom-up and top-down processing in early visual cortex with high-resolution fmri* (Doctoral dissertation, Maastricht University).
- Marquardt, I., Schneider, M., Gulban, O. F., Ivanov, D., & Uludağ, K. (2018). Cortical depth profiles of luminance contrast responses in human V1 and V2 using 7 T fMRI. *Human Brain Mapping*, *39*(7), 2812–2827.
- Marques, J. P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P. F., & Gruetter, R. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *NeuroImage*, *49*(2), 1271–1281.
- Mitra, P. P. (2014). The circuit architecture of whole brains at the mesoscopic scale. *Neuron*, *83*, 1273–1283.

- Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., & Ugurbil, K. (2010). Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magnetic Resonance in Medicine*.
- Moerel, M., De Martino, F., Kemper, V. G., Schmitter, S., Vu, A. T., Ugurbil, K., ... Yacoub, E. (2017). Sensitivity and specificity considerations for fMRI encoding, decoding, and mapping of auditory cortex at ultra-high field. *NeuroImage*, 168.
- Muckli, L., De Martino, F., Vizioli, L., Petro, L. S., Smith, F. W., Ugurbil, K., ... Yacoub, E. (2015). Contextual Feedback to Superficial Layers of V1. *Current Biology*, 25(20), 2690–2695.
- Mugler, J. P. & Brookeman, J. R. (1990). Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE). *Magnetic Resonance in Medicine*.
- Murphy, K., Bodurka, J., & Bandettini, P. A. (2007). How long to scan? The relationship between fMRI temporal signal to noise ratio and necessary scan duration. *NeuroImage*.
- Nasr, S., Polimeni, J. R., & Tootell, R. B. H. (2016). Interdigitated color- and disparity-selective columns within human visual cortical areas V2 and V3. *Journal of Neuroscience*, 36(6), 1841–57.
- Olman, C. A., Harel, N., Feinberg, D. A., He, S., Zhang, P., Ugurbil, K., & Yacoub, E. (2012). Layer-Specific fMRI Reflects Different Neuronal Computations at Different Depths in Human V1. *Plos One*, 7(3), e32536.
- Polimeni, J. R., Fischl, B., Greve, D. N., & Wald, L. L. (2010). Laminar analysis of 7t BOLD using an imposed spatial activation pattern in human V1. *NeuroImage*, 52(4), 1334–1346.
- Polimeni, J. R., Renvall, V., Zaretskaya, N., & Fischl, B. (2018). Analysis strategies for high-resolution UHF-fMRI data. *NeuroImage*, 168, 296–320.
- Poser, B. A., Koopmans, P. J., Witzel, T., Wald, L. L., & Barth, M. (2010). Three dimensional echo-planar imaging at 7 Tesla. *NeuroImage*.
- Rakic, P. (2008). Confusing cortical columns. *Proceedings of the National Academy of Sciences*, 105(34), 12099–12100.
- Ramachandran, V. S. & Anstis, S. M. (1985). Perceptual organization in multistable apparent motion. *Perception*, 14, 135–143.

-
- Renvall, V., Witzel, T., Wald, L. L., & Polimeni, J. R. (2016). Automatic cortical surface reconstruction of high-resolution T1 echo planar imaging data. *NeuroImage*.
- Ress, D., Glover, G. H., Liu, J., & Wandell, B. (2007). Laminar profiles of functional activity in the human brain. *NeuroImage*.
- Schneider, M., Kemper, V. G., Emmerling, T. C., De Martino, F., & Goebel, R. (2019). Columnar clusters in the human motion complex reflect consciously perceived motion axis. *Proceedings of the National Academy of Sciences*, *116*, 5096–5101.
- Setsompop, K., Gagoski, B. A., Polimeni, J. R., Witzel, T., Wedeen, V. J., & Wald, L. L. (2012). Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic Resonance in Medicine*.
- Sheinberg, D. L. & Logothetis, N. K. (1997). The role of temporal cortical areas in perceptual organization. *Proceedings of the National Academy of Sciences*, *94*(7), 3408–3413.
- Tootell, R. B. H. & Nasr, S. (2017). Columnar Segregation of Magnocellular and Parvocellular Streams in Human Extrastriate Cortex. *Journal of Neuroscience*, 0690–17.
- Triantafyllou, C., Hoge, R. D., Krueger, G., Wiggins, C. J., Potthast, A., Wiggins, G. C., & Wald, L. L. (2005). Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *NeuroImage*.
- Uğurbil, K. (2014). Magnetic resonance imaging at ultrahigh fields. *IEEE Transactions on Biomedical Engineering*, *61*(5), 1364–1379.
- Uğurbil, K., Adriany, G., Andersen, P., Chen, W., Garwood, M., Gruetter, R., ... Zhu, X. H. (2003). Ultrahigh field magnetic resonance imaging and spectroscopy. *Magnetic Resonance Imaging*.
- Uğurbil, K., Adriany, G., Andersen, P., Chen, W., Gruetter, R., Hu, X., ... Ogawa, S. (2002). Magnetic Resonance Studies of Brain Function and Neurochemistry. *Annual Review of Biomedical Engineering*.
- Uludağ, K. & Blinder, P. (2016). Linking brain vascular physiology to hemodynamic response in ultra-high field MRI. *NeuroImage*, *168*.
- Uludağ, K., Müller-Bierl, B., & Uğurbil, K. (2009). An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. *NeuroImage*.

- Van de Moortele, P.-F., Auerbach, E. J., Olman, C., Yacoub, E., Uğurbil, K., & Moeller, S. (2009). T1 weighted brain images at 7 Tesla unbiased for Proton Density, T2* contrast and RF coil receive B1 sensitivity with simultaneous vessel visualization. *NeuroImage*, *46*(2), 432–446.
- van der Kouwe, A. J. W., Benner, T., Salat, D. H., & Fischl, B. (2008). Brain morphometry with multiecho MPRAGE. *NeuroImage*, *40*(2), 559–569.
- Vaughan, J. T., Garwood, M., Collins, C. M., Liu, W., Delabarre, L., Adriany, G., ... Uğurbil, K. (2001). 7T vs. 4T: RF power, homogeneity, and signal-to-noise comparison in head images. *Magnetic Resonance in Medicine*, *46*(1), 24–30.
- Viviani, R. (2016). A Digital Atlas of Middle to Large Brain Vessels and Their Relation to Cortical and Subcortical Structures. *Frontiers in Neuroanatomy*, *10*, 12.
- Viviani, R., Pracht, E. D., Brenner, D., Beschoner, P., Stingl, J. C., & Stöcker, T. (2017). Multimodal MEMPRAGE, FLAIR, and R2* segmentation to resolve dura and vessels from cortical gray matter. *Frontiers in Neuroscience*, *11*, 1–13.
- Wandell, B. A. & Winawer, J. (2015). Computational neuroimaging and population receptive fields. *Trends in Cognitive Sciences*, *19*(6), 349–357.
- Winawer, J., Horiguchi, H., Sayres, R. A., Amano, K., & Wandell, B. A. (2010). Mapping hV4 and ventral occipital cortex: The venous eclipse. *Journal of Vision*, *10*(5), 1–1.
- Yacoub, E. & Harel, N. (2008). High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences*, *105*(30), 254–264.
- Yacoub, E., Shmuel, A., Logothetis, N., & Uğurbil, K. (2007). Robust detection of ocular dominance columns in humans using Hahn Spin Echo BOLD functional MRI at 7 Tesla. *NeuroImage*.
- Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P. F., Feinberg, D., Adriany, G., ... Yacoub, E. (2011). Mapping the organization of axis of motion selective features in human area MT using high-field fmri. *Plos One*, *6*(12), 1–10.
- Zuiderbaan, W., Harvey, B. M., & Dumoulin, S. O. (2012). Modeling center – surround configurations in population receptive fields using fMRI. *Journal of Vision*, *12*(3), 1–15.

2

A scalable method to improve gray matter segmentation at ultra-high field MRI

Based on

Gulban, O. F.*, Schneider, M.*, Marquardt, I., Haast, R. A., & De Martino, F. (2018). A scalable method to improve outer gray matter boundary segmentation at ultra high field MRI. *Plos One*, 13, 1-33.

*Authors contributed equally.

2.1 ABSTRACT

High-resolution (functional) magnetic resonance imaging (MRI) at ultra-high magnetic fields (7 Tesla and above) enables researchers to study how anatomical and functional properties change within the cortical ribbon, along surfaces and across cortical depths. These studies require an accurate delineation of the gray matter ribbon, which often suffers from inclusion of blood vessels, dura mater and other non-brain tissue. Residual segmentation errors are commonly corrected by browsing the data slice-by-slice and manually changing labels. This task becomes increasingly laborious and prone to error at higher resolutions since both work and error scale with the number of voxels. Here we show that many mislabeled, non-brain voxels can be corrected more efficiently and semi-automatically by representing three-dimensional anatomical images using two-dimensional histograms. We propose a representation based on intensity values and their first spatial derivative and quantify the benefits in 7 Tesla MRI data of nine volunteers. We present an openly accessible Python implementation of this approach and demonstrate that editing cortical segmentations using two-dimensional histogram representations as an additional post-processing step aids existing algorithms and yields improved gray matter borders. By making our data and corresponding expert (ground truth) segmentations openly available, we facilitate future efforts to develop and test segmentation algorithms on this challenging type of data.

2.2 INTRODUCTION

Magnetic resonance imaging (MRI) has become one of the most important tools to study human brain function and structure *in vivo*. Moving from high (3 Tesla [T]) to ultra-high (7 and 9.4 T) magnetic fields (UHF), together with improvements in acquisition methods, leads to increases in signal and contrast-to-noise (SNR and CNR, respectively) (Duyn, 2012; Ugurbil, 2014; Vaughan et al., 2001). The increase in SNR can be leveraged to increase voxel resolution of both functional and structural images to sub-millimeter scales. Such sub-millimeter spatial resolutions allow for *in vivo* studies that probe cortical properties at the mesoscale (De Martino et al., 2018; Dumoulin, Fracasso, van der Zwaag, Siero, & Petridou, 2018;

Kemper, De Martino, Emmerling, Yacoub, & Goebel, 2018; Polimeni, Reinvall, Zaretskaya, & Fischl, 2018; Trampel, Bazin, Pine, & Weiskopf, 2017). These studies include (i) cortical-depth dependent analyses of function (De Martino et al., 2013; Huber et al., 2015; Kok, Bains, Van Mourik, Norris, & De Lange, 2016; Koopmans, Barth, Orzada, & Norris, 2011; Muckli et al., 2015; Polimeni, Fischl, Greve, & Wald, 2010) and structure (Kemper et al., 2018; Trampel et al., 2017; Waehnert et al., 2014), (ii) the mapping of cortical columnar structures (De Martino, Moerel, Ugurbil, et al., 2015; Goncalves et al., 2015; Nasr, Polimeni, & Tootell, 2016; Tootell & Nasr, 2017; Yacoub & Harel, 2008; Zimmermann et al., 2011) as well as (iii) sub-millimeter cortical topography (De Martino, Moerel, Ugurbil, et al., 2015; Fracasso, Petridou, & Dumoulin, 2016; Harvey, Klein, Petridou, & Dumoulin, 2013; Harvey, Fracasso, Petridou, & Dumoulin, 2015).

Such studies crucially depend on accurate and precise delineations of the gray matter (GM) ribbon both at the inner (white matter; WM) and outer (cerebrospinal; CSF) border. Since the aim of these studies is to investigate how the (functional) MRI (f/MRI) signal varies as a function of small position changes in GM, systematic GM segmentation errors would invalidate the conclusions drawn from these studies. Consider, as an example, an fMRI study conducted with a voxel resolution of 0.8 mm isotropic. Assume an average thickness of human cortex of 2.4 mm and a true signal change at the upper cortical depth level. In this study, under optimal conditions the functional resolution would allow a straight piece of cortical ribbon to be divided in three relative cortical depths, each of them one voxel thick. Falsely labeling an additional fourth voxel as GM would make the difference between reporting an fMRI signal change at most superficial (false voxel excluded) or mid-superficial (false voxel included) cortical depth level. This example stresses the importance of accurate and precise GM segmentations.

Obtaining accurate and precise definitions of the GM ribbon, however, is currently a difficult and time-consuming task for sub-millimeter UHF data. The increases in SNR, CNR and resolution attainable with UHF as well as analysis (Van de Moortele et al., 2009) and reconstruction (Marques et al., 2010) strategies specific to UHF reveal several structures outside of the brain that were barely visible on images obtained at conventional field strengths (1.5 and 3 T) and lower resolution (> 1 mm isotropic) (Vi-

viani et al., 2017). Such structures include the dura mater (van der Kouwe, Benner, Salat, & Fischl, 2008), medium-sized blood vessels in the sulci (Viviani, 2016) as well as draining sinuses and connective tissue adjacent to GM (Helms, Kallenberg, & Dechent, 2006). To date, many of the existing brain segmentation algorithms have been developed and benchmarked on images collected at 1 mm isotropic resolution or lower and at conventional field strengths (Helms, 2016) (but see Bazin et al., 2014). If segmentation algorithms do not model these non-brain structures they might falsely label (part of) these structures as GM. Faced with such segmentation errors, researchers commonly correct the misclassified voxels manually. However, the increase in resolution leads to an exponential increase in the number of voxels, which renders manual correction a laborious task. Furthermore, manual correction is prone to error and may introduce an observer bias, thereby reducing the reproducibility of subsequent analyses (Despotović, Goossens, & Philips, 2015). This currently leaves researchers with the dilemma of accepting the likely erroneous outcome of automatic segmentation algorithms or performing a time-consuming and error-prone manual correction.

CBS tools (Bazin et al., 2014) directly tackle many of the challenges of UHF high-resolution anatomical data by, for example, including pre-processing steps to estimate dura mater and CSF partial voluming. Consequently, these tools provide an improved initial GM segmentation compared to other solutions (Bazin et al., 2014). However, we show that in many cases the initial CBS segmentation can be further improved with the approach proposed here. Furthermore, CBS tools have been optimized for whole-brain data obtained with the MP2RAGE sequence (Marques et al., 2010). While the MP2RAGE sequence is commonly used at UHF as a basis for brain tissue class segmentations, we note that many high-resolution studies at UHF also use alternative sequences to define GM (Huber et al., 2015; Kashyap, Ivanov, Havlicek, Poser, & Uludağ, 2017; Renvall, Witzel, Wald, & Polimeni, 2016; Van de Moortele et al., 2009), some of which offering partial coverage of the brain only (Huber et al., 2015; Kashyap et al., 2017). In such cases, alternative approaches that do not depend on particular templates, atlases or other forms of prior information are useful and required.

Here, we show that non-brain voxels misclassified as GM can largely be corrected using a multi-dimensional transfer function that is specified based on a two-dimensional (2D) histogram representation (Ip, Varshney, & Jaja, 2012; Kindlmann & Durkin, 1998; Kniss, Kindlmann, & Hansen, 2001; Kniss, Kindlmann, & Hansen, 2005) of three-dimensional (3D) MRI brain data. We demonstrate that this transfer function offers an efficient way to single out non-brain tissue voxels. Removing these voxels from GM classifications found by automatic segmentation pipelines improves GM segmentations. This approach addresses the problems of an entirely manual correction, since it yields a meaningful summary representation of the data that allows to manipulate the data efficiently. As a consequence, it is both more time efficient than manual slice-by-slice correction and it reduces observer bias.

This chapter is intended as a demonstration that 2D histogram based methods are useful for improving segmentations of MRI images. In particular, we aim to show that, for images acquired at sub-millimetre resolution and very high field strength, 2D histogram based methods offer an efficient way to obtain a more refined brain mask that excludes usually undesired structures like vessels and dura mater. Several alternatives to the methods presented here exist for data visualization, dimensionality reduction, or data fusion, for example, principal component analysis (Jolliffe, 1986), multidimensional scaling (Borg & Groenen, 2005) or the t-SNE algorithm (Van Der Maaten & Hinton, 2008). A detailed quantification of the merits and disadvantages of these methods, however, is beyond the scope of this chapter which is intended to introduce a simple and fast solution. Likewise, there are alternative ways of defining clusters in an image to the normalized graph cut algorithm that we have used here (Jain et al., 2011; Liu & Jurrus, 2012; Nunez-Iglesias, Kennedy, Parag, Shi, & Chklovskii, 2013; Venkataraju, Paiva, Jurrus, & Tasdizen, 2009). While all these methods have their merit, we decided to use normalized-cut multi-level segmentation since it already has been shown to work successfully on the 2D histogram representations of volumetric data (Ip et al., 2012).

We structured the chapter as follows. In Section 2.3, we introduce the technique of specifying transfer functions based on 2D histogram representations of voxel intensity and gradient magnitude. We offer theoretical considerations for why this technique is suited to remove vessels and

dura mater voxels in high-resolution MRI data (< 1 mm isotropic voxel size). In Section 2.4 we outline required features of the input data and recommended data pre-processing steps. In Section 2.5 and 2.6 we validate the suggested method by evaluating obtained GM segmentation results against expert GM segmentations obtained for nine subjects recorded at 7 T. We demonstrate considerable improvement in segmentation performance metrics. We have implemented the method described here in a free and open Python software package (Gulban & Schneider, 2017). The package as well as validation data, corresponding expert segmentations, and processing scripts used to validate the proposed method are all openly available (see Table S2.1).

2.3 TRANSFER FUNCTIONS AND HISTOGRAMS

2.3.1 *Multi-dimensional transfer functions*

In the context of MRI data visualization, a transfer function can be understood as a mapping of voxel data to optical properties such as color and opacity. Effective transfer functions make structures of interest in the data visible. This can, for example, be achieved by assigning low opacity values to voxels that make up irrelevant structures and by highlighting desired structures with high opacity and salient color values. Multi-dimensional transfer functions assign renderable properties based on a combination of values (Kindlmann & Durkin, 1998; Kniss, Kindlmann, & Hansen, 2002; Kniss et al., 2001; Kniss et al., 2005). This is important in the context of MRI data where features of interest are often difficult to extract based on a single value alone. Considering multiple values increases the chances of uniquely isolating a feature and making it visible (Kniss et al., 2002).

In theory, multi-dimensional transfer functions could be used to perform exhaustive tissue-type segmentation of human brain MRI data. In this process, each voxel would be classified as either WM, GM, or CSF by specifying appropriate transfer functions. It has been shown, however, that this approach is less successful than other, bespoke brain segmentation algorithms (Ljung et al., 2016). Here, we propose that transfer function-based methods still have a role to play in UHF MRI brain segmentation pipelines because they are well-suited for efficient removal of mislabeled

non-brain tissue. We motivate this proposition by considering that brain and non-brain voxels become separable in 2D histogram representations.

2.3.2 2D histogram representation

2D histogram representations have been shown to greatly facilitate the process of specifying effective, 2D transfer functions (Kindlmann & Durkin, 1998; Kniss et al., 2005). 2D histograms are obtained by taking an n -dimensional data set and binning its data points along two dimensions. In principle, a 2D histogram can be obtained from any two sets of values. A 2D histogram, plotting gradient magnitude against image intensity, however, has been shown to be particularly useful for identifying tissue boundaries (Kindlmann & Durkin, 1998; Kniss et al., 2005). The term gradient magnitude here refers to the magnitude of the vector that represents the spatial intensity gradient at every MRI voxel, where the spatial intensity gradient is equal to the first spatial derivative of the image intensity values.

Figure 2.1 shows how 3D MRI data of a human brain are represented in a 2D histogram (a T1w image was divided by a PDw image [Van de Moortele et al., 2009] and brain extracted; images were acquired with 0.7 mm isotropic resolution; for more details see Section 2.5). The histogram is obtained by plotting gradient magnitude against image intensity. In this representation, different tissue types occupy different regions. CSF voxels are characterized by very low intensity and low gradient magnitude values and therefore occupy the lower-left space of the histogram. GM and WM voxels have medium to high intensities and very low gradient magnitudes. Therefore, these tissue classes form circular regions at the bottom-center of the histogram. Voxels at the GM-WM interface fall within an arc reaching from medium to high intensities, following a low to medium to low gradient magnitude trajectory. Similarly, voxels at the CSF-GM interface span an arc from low to medium intensities. Finally, blood vessels and dura mater are thin structures characterized by very high gradient magnitudes and medium to high intensities. Therefore, these structures occupy up-center and up-right parts of the 2D histogram.

Since different tissue types occupy different regions in the 2D histogram, each tissue type and boundary can, in principle, be isolated using

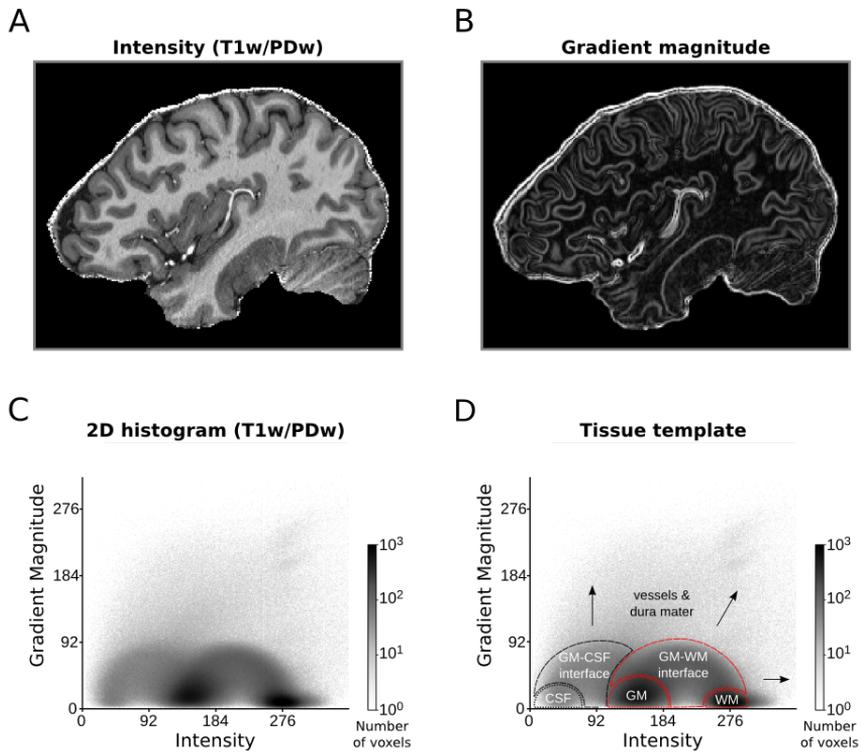


FIGURE 2.1: 2D histogram representation for MRI image of a human brain. **(A)** Intensity and **(B)** gradient magnitude values of a brain extracted T1w-divided-by-PDw MRI image are represented in a **(C)** 2D histogram. Darker regions in the histogram indicate that many voxels in the MRI image are characterized by this particular combination of image intensity and gradient magnitude. **(D)** The 2D histogram displays a characteristic pattern with tissue types occupying particular areas of the histogram. Voxels containing CSF, dura mater or blood vessels (black dashed lines and arrows) cover different regions of the histogram than voxels containing WM and GM (red dashed lines). As a result, brain tissue becomes separable from non-brain tissue.

a 2D transfer function based on image intensity and gradient magnitude. For the purposes of this chapter, we focus on the distinction between brain (WM, GM, GM-WM interface) and non-brain (CSF, CSF-GM interface, blood vessels, dura mater) voxels. The intensity-gradient magnitude his-

togram is particularly suited to distinguish non-brain tissue because voxels containing dura mater and vessels are characterized by high gradient magnitude values. Given the typical voxel sizes of current high resolution studies, gradient magnitude will be high in the entirety of these structures (see Figure 2.1B for an example) and the combination of high intensity and high gradient magnitude values renders those structures separable from WM and GM voxels.

2.3.3 Creating transfer functions

The simplest way to create a transfer function is to explore the data by moving widgets with a specified shape over the 2D histogram representation (Kniss et al., 2005). For example, Figure 2.2 shows how a circular sector could be moved on top of the 2D histogram to highlight particular regions. In this case, only MRI voxels whose intensity-gradient magnitude combination falls within the highlighted region of the 2D histogram would be selected. Position and size of the circular sector can then be refined until the desired data have been isolated.

Using such a straightforward process of exploration and refinement (Kniss et al., 2002), however, might yield slightly sub-optimal results. The shape of the widget might not capture the ideal shape given the data or the user might lack the prior knowledge that is required for this task. Alternatively, hierarchical exploration of normalized graph cut decision trees can be used (Ip et al., 2012). This graph cut method results in a set of components (i.e. clusters) of the histogram that are mutually exclusive and collectively exhaustive. This allows the user to split and merge clusters in a data-driven and intuitive way that can be aided by the immediate visualization of the resulting segmentation (Figure 2.3). The method allows for semi-automatic tissue selection, i.e. the shape of the clusters is data-driven but the decision which clusters to join and which to divide is made by the user.

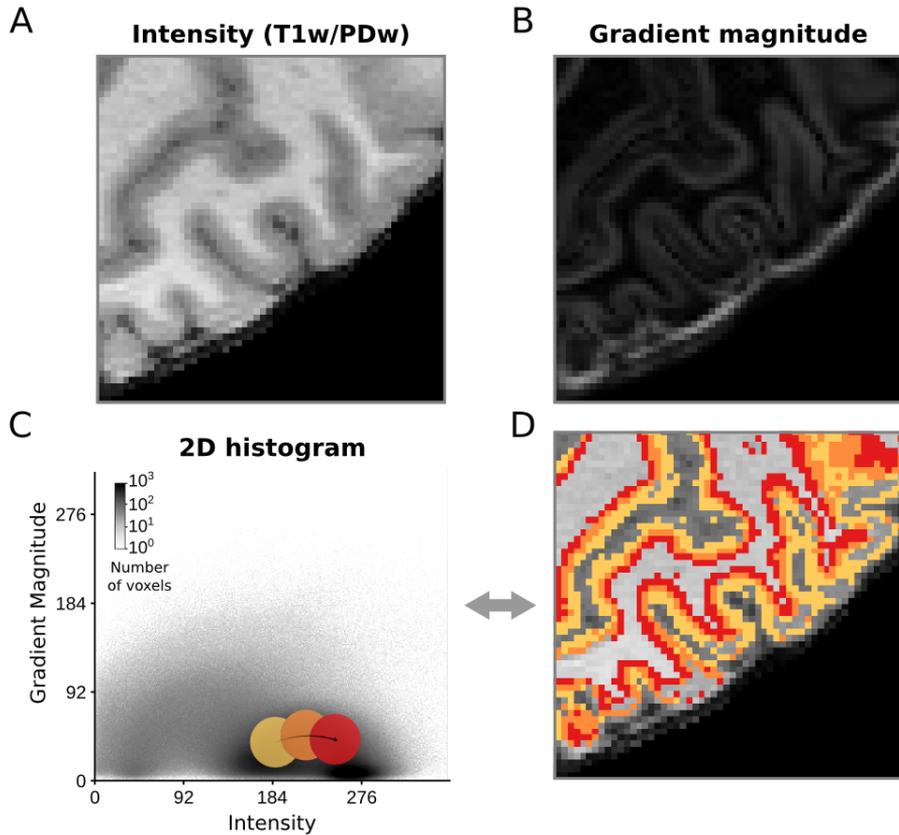


FIGURE 2.2: Creation of 2D transfer functions with pre-defined shapes. **(A)** Intensity and **(B)** gradient magnitude values of a brain extracted T1w-divided-by-PDw MRI image are represented in a 2D histogram. By moving widgets of pre-defined shape, e.g. a circle, over the **(C)** 2D histogram and **(D)** concurrent visualization of selected voxels on a 2D slice of brain, positions of different tissue types in the 2D histogram can be probed and transfer functions can be created. In this example, the different probe positions (yellow, orange and red circles) appear to contain different aspects of GM.

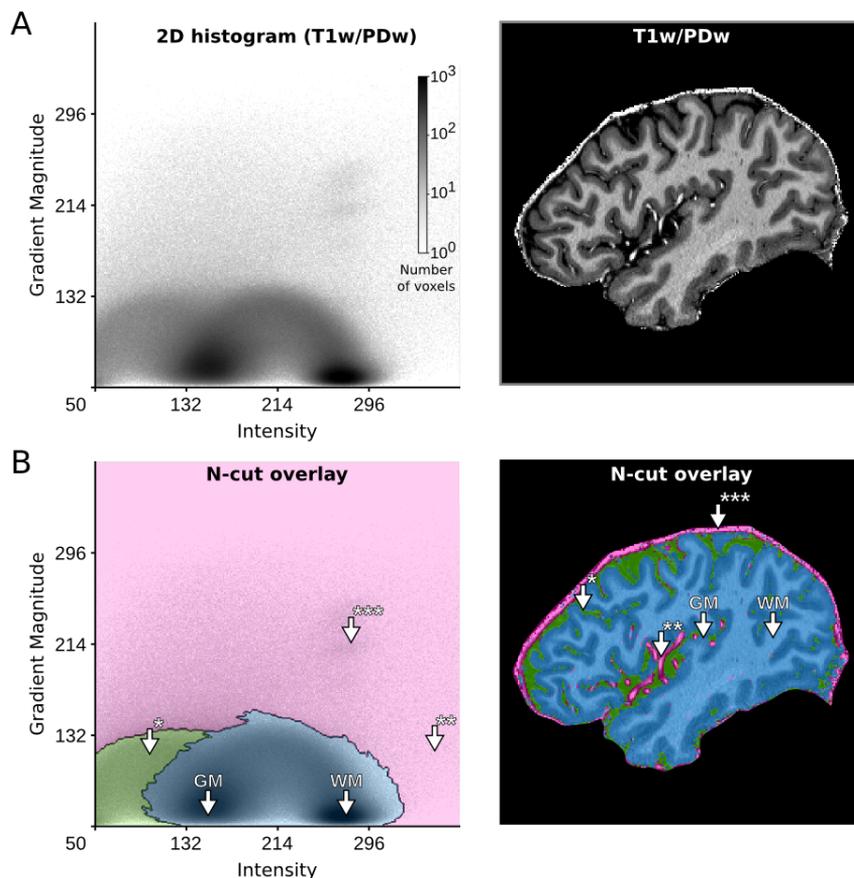


FIGURE 2.3: Creation of 2D transfer functions with data-driven shapes. **(A)** The user starts with the 2D histogram representation of image intensity and gradient magnitude (left side) and concurrent visualization of the original brain data (right side). The user can then interact with and select data in the 2D histogram to specify transfer functions. In this example, this was done with the help of a normalized graph cut decision tree. **(B)** The interaction with the 2D histogram results in data-driven shapes of selected areas, here shaded in pink, green and blue (left side). Voxels selected by those areas are highlighted in corresponding colors against the backdrop of the original brain data (right side). The visualization reveals that the area of the 2D histogram shaded in blue selects brain voxels, while the areas shaded in green and pink select CSF* and blood vessel voxels**, respectively.

2.4 INPUT DATA REQUIREMENTS AND PREPARATION

2.4.1 *Data preparation*

In order to obtain optimal results with the gradient-magnitude method, several pre-processing steps should be performed on the data. Ranging from absolutely necessary to desired (but not critical), these pre-processing steps include: (i) bias field correction, (ii) brain extraction, (iii) cerebellum removal and (iv) removal of brain stem structures. Successful bias field correction is critical to performance since otherwise intensity values for different tissue types start to mix in 2D histogram space. Brain extraction should be performed to remove irrelevant voxels from the 2D histogram representation. Removal of cerebellar and brain stem structures is recommended since it further improves conformity to ideal 2D histogram shapes (Figure 2.1D). Bias field correction and brain extraction can be performed using automatic algorithms (Ashburner & Friston, 2005; Smith, 2002). Removal of cerebellum and sub-cortical structures might require the manual creation of masks. We note, however, that generation of these masks is only desired, not strictly necessary. Furthermore, generation of these masks is often a desirable processing step for many automatic tissue class segmentation algorithms, since it improves their performance as well.

2.4.2 *Data requirements*

Suitability of the intensity-gradient magnitude histogram for separating brain from non-brain tissue will depend on the resolution and CNR of the input data. We expect a lower limit of resolution around 1 mm. At resolutions lower than that, the intensity-gradient magnitude method will yield unsatisfactory results due to partial voluming between the thin structures we are aiming to correct and surrounding CSF or tissues. We do not expect an upper resolution limit for the input data. Although, initially, values in the gradient magnitude image will no longer be high in all vessel and dura mater voxels, very high-resolution images can still be accommodated by choosing the appropriate level of smoothness on the gradient magnitude image. In Figure S2.1, we demonstrate that by setting the appropriate smoothness level of a Deriche filter (Deriche, 1987), gradient magnitude images for very high resolution data (0.25 mm isotropic; Lüsebrink, Sciarra,

Mattern, Yakupov, and Speck, 2017; Lusebrink, Sciarra, Mattern, Yakupov, and Speck, 2017) can be approximated to those observed for data at lower resolution (0.7 mm isotropic).

We furthermore expect our method to be impacted by the CNR of the input data. Figure S2.2 shows that with added Gaussian noise (i.e., decreasing CNR) the desired circular and arc-like shapes in the 2D histogram (Figure 2.1D) become less apparent. At very high noise levels separating brain from non-brain tissue in the 2D histogram space becomes challenging (see e.g. Figure S2.2). While the in-depth evaluation of additional processing tools is beyond the scope of the present chapter, we note that if the input data are very noisy, smoothing can be applied. In particular, non-linear anisotropic diffusion based smoothing (Mirebeau, Fehrenbach, Risser, & Tobji, 2015; Weickert, 1998) results in the data regaining the desired 2D histogram shapes (see Figure S2.3).

The parameter space of the input data is thus constrained by resolution and CNR. Apart from these restrictions, our methods are suitable for any 3D image and work irrespective of the field-of-view of the acquisition (partial coverage is possible) and membership to a particular species (bottle-nose dolphin brain is also possible; for examples see Figure S2.4).

2.5 VALIDATION METHODS

2.5.1 Overview

In order to validate the method proposed above, we created two validation data sets based on the acquisition of high-resolution 7 T data of nine subjects and corresponding manually-guided expert segmentations of GM. In particular, we created two validation sets based on two of the most common acquisition sequences. For five subjects, we collected MPRAGE T1w and PDw data (we refer to this data set as the MPRAGE data set below). For four different subjects, we collected MP2RAGE data, to obtain unbiased (uni) images (we refer to this data set as the MP2RAGE data set below). Both data sets can be downloaded from Gulban, Schneider, Marquardt, Haast, and De Martino, 2017.

2.5.2 Ethics statement

The experimental procedures were approved by the ethics committee of the Faculty for Psychology and Neuroscience (MPRAGE data set) and the Medical Ethical Committee at the Faculty of Health, Medicine and Life Sciences (MP2RAGE data set) at Maastricht University, and were performed in accordance with the approved guidelines and the Declaration of Helsinki. Written informed consent was obtained for every participant before conducting the experiments.

2.5.3 MRI acquisition parameters

All images were acquired on a Siemens 7 T whole body scanner (Siemens Medical Solutions, Erlangen, Germany) using a head RF coil (Nova Medical, Wilmington, MA, USA; single transmit, 32 receive channels). In all acquisitions, we used dielectric pads (Teeuwisse, Brink, & Webb, 2012).

For $n=5$ subjects (age range 24-30, 2 females, no medical condition), the MPRAGE data set consisted of: a T1w image using a 3D MPRAGE sequence (repetition time [TR] = 3100 ms; time to inversion [TI] = 1500 ms [adiabatic non-selective inversion pulse]; time echo [TE] = 2.42 ms; flip angle = 5°; generalized auto-calibrating partially parallel acquisitions [GRAPPA] = 3 (Griswold et al., 2002); field of view [FOV] = 224 × 224 mm²; matrix size = 320 × 320; 256 slices; 0.7 mm isotropic voxels; pixel bandwidth = 182 Hz/pixel; first phase encode direction anterior to posterior; second phase encode direction left to right), a PDw image (0.7 mm isotropic) with the same MPRAGE as for the T1w image but without the inversion pulse (TR = 1380 ms; TE = 2.42 ms; flip angle = 5°; GRAPPA = 3; FOV = 224 × 224 mm²; matrix size = 320 × 320; 256 slices; 0.7 mm iso. voxels; pixel bandwidth = 182 Hz/pixel; first phase encode direction anterior to posterior; second phase encode direction left to right).

For $n=4$ subjects (age range 24-58, 2 females, no medical condition) the MP2RAGE data set consisted of: 3D MP2RAGE data (TR = 5000 ms; TI1/TI2 = 900/2750 ms; TE = 2.46 ms; FA1/FA2 = 5°/3°; FOV = 224 × 224 mm²; matrix size = 320 × 320; slices = 240; 0.7 mm iso. voxels) (Marques et al., 2010). More details on the MP2RAGE data acquisition can be found in Haast, Ivanov, Formisano, and Uludağ, 2016.

2.5.4 Manually-guided expert segmentations

For every subject, we established 'ground truth' GM classifications via manually-guided expert segmentations. All segmentations were created manually by the same expert (OFG), using ITK-SNAP (Yushkevich et al., 2006) and a graphics tablet (Intuos Art; Wacom Co. Ltd; Kazo, Saitama, Japan). Segmentations were only established for cortical GM, since cerebellar and sub-cortical structures were later removed in a pre-processing step. To establish the segmentation, the expert used T1w images for the MPRAGE and uni images for the MP2RAGE data set. To avoid resulting tissue type classification to be ragged, the expert followed a particular processing sequence. The brain was first traversed in a single direction (e.g. sagittally) and the ground truth was established slice-by-slice. Subsequently, the brain was traversed in the two other directions (e.g. axially, then coronally). This sequence was repeated several times across several regions until the GM segmentation of the whole brain was considered of good quality. To further ensure the quality of the resulting segmentation, they were inspected for mistakes by two additional experts (MS and FDM).

2.5.5 Software implementation

We implemented the creation of transfer function based on 2D histograms in an open source Python package called Segmentator (2017), which is built upon several other scientific packages such as Numpy (2011), Scipy (2007), Matplotlib (2007) and Nibabel (2017). Segmentator allows for selection of data points in a 2D histogram (for example gradient magnitude over intensity) and concurrent visualization of selected brain voxels on a 2D slice. Data points can be selected using a circular sector widget with variable reflex angle and radius. Alternatively, data selection can be performed using the normalized graph cut (n-cut) method (i.e. spectral clustering) as described above. The n-cut algorithm from Scikit-image (2014) was modified to export an augmented output which provides step-wise access to independent branches of the decision tree and employed in Segmentator (the modification is available as a [GitHub repository](#)).

The package provides several options to calculate the gradient magnitude image. All the 2D histogram analyses described in this chapter were based on gradient magnitude images that were computed as the

Euclidean norm of the first spatial derivative estimated using a $3 \times 3 \times 3$ Scharr kernel (Jähne & Haußecker, 2000; Scharr, 2000). Subsequently, transfer functions were specified using the normalized graph cut algorithm and user intervention for the selection of the non-brain tissue transfer functions. Processing data for a single subject took about 10 minutes on average. The Segmentator package is openly and freely accessible as a [GitHub repository](#).

2.5.6 Segmentation procedure

For both validation data sets, we followed similar procedures, with modifications where necessary to accommodate for differences in the sequences' output. Our goal was to obtain initial GM segmentations from existing, fully-automated segmentation algorithms and to quantify the improvement in segmentation accuracy that can be obtained when using the method described here as post-processing steps. To establish the initial GM segmentations we used FSL FAST (Zhang, Brady, & Smith, 2001) and the SPM 12 "unified segmentation" algorithm (Ashburner & Friston, 2005) for the MPRAGE data set and FSL FAST and CBS tools (Bazin et al., 2014) for the MP2RAGE data set. SPM and CBS tools have been developed and benchmarked on MPRAGE and MP2RAGE images respectively. FSL FAST is suited to process either type, so we used it for both data sets. We then quantified the impact of using transfer functions based on 2D histogram representations of intensity and gradient magnitude (see Section 2.3) as an additional post-processing steps. This procedure will be referred to below as the gradient magnitude (GraMag) method. The GraMag method resulted in a mask that could be used to further refine the initial GM segmentation, e.g. by removing blood vessels and dura mater that were falsely labeled as GM initially. The analysis procedures for the two validation data sets are summarized in flow chart diagrams (Figure S2.5 and S2.6). Furthermore, in an effort to make our analyses fully reproducible, we made the Python and bash scripts used for pipeline processing openly available at Schneider and Gulban, 2018.

For the MPRAGE data set, we first computed ratio images (T1w divided by PDw) (Van de Moortele et al., 2009) to reduce inhomogeneities. Ratio images were input to either FSL FAST or SPM 12. FSL FAST was used

with default values. The FAST algorithm requires an initial brain extraction procedure that we performed using FSL BET (Smith, 2002). Additionally, we masked the images to exclude: the corpus callosum, the basal ganglia, the hippocampus, the entire brain stem and the cerebellum. Below we refer to this mask as 'NoSub mask'. The NoSub mask was created manually for every subject. In SPM 12 we used default settings with one exception. We set the number of Gaussians to be modeled to 3 for GM and 2 for WM (default values are 1 and 1), which in our experience yields improved results for 7 T data. As part of their standard segmentation routine, both FSL FAST and SPM 12 perform initial inhomogeneity correction. We inspected the bias-corrected images to ensure that the algorithms had converged on plausible solutions. We specified for the FSL FAST algorithm to output hard segmentation labels. Since SPM 12 outputs probabilities for six tissue classes, we transformed this soft output to hard segmentation labels by assigning each voxel to the tissue class with the highest posterior probability. Since the SPM segmentation algorithm works best with unmasked images, we applied the NoSub mask only to the resulting SPM GM segmentations, not to the input data. The resulting GM segmentations from FSL and SPM were saved for later evaluation.

We proceeded with bias-corrected ratio images from either SPM or FSL (Figure S2.5). Since the GraMag method works best with brain extracted images, we combined SPM's WM and GM segmentation outcomes to form a brain mask and performed brain extraction of the ratio images from SPM. After brain extraction, we also excluded cerebellum and brain stem tissue using the NoSub mask. FSL's bias-corrected ratio images did not require masking as the brain extraction (and cerebellum removal) was already performed before segmentation. We used the 2D histogram representation of intensity and gradient magnitude together with the hierarchical exploration of normalized graph cut decision trees (as described in Section 2.3) to create transfer functions. Exploration of decision trees was limited to an 8-level hierarchy. The criterion for splitting and merging clusters was subjective: a rater (MS) aimed to obtain shapes that resembled the ideal template shapes (Figure 2.1D) as closely as possible, given the 2D histogram representation and concurrent visualization of selected voxels. However, selection of voxels was well constrained by clearly-outlined shapes in the 2D histogram representation and commonly

required to move down the decision tree hierarchy by only 2 or 3 levels. Exploration of the decision tree took about 30 to 60 seconds per subject. Generation of normalized graph cut decision trees, which was done previous to exploration by a rater, took about 5 minutes on a workstation (RAM: 32 GB, 12 cores (6 virtual); CPU: 2.146 GHz; operating system: Debian 8). The transfer function resulting from this procedure was used to separate brain from non-brain tissue voxels. Non-brain tissue voxels were removed from GM if they were included in the initial FSL and SPM segmentations.

For the MP2RAGE data set, the T1 map, T1w (uni) and second inversion image from the MP2RAGE sequence were input to CBS tools (Bazin et al., 2014). Only the brain-extracted (Smith, 2002) uni image was input to FSL FAST, since this resulted in higher performance than inputting all three images. Both FSL FAST and CBS tools were run with default settings. Note that the default settings for CBS tools include removal of non brain tissue by estimating dura mater and CSF partial voluming. The resulting GM segmentations from FSL and CBS were saved for later evaluation. We proceeded with the FSL FAST bias-corrected, brain-extracted and NoSub masked uni image as for the MPRAGE data set to obtain a secondary brain mask (Figure S2.6).

2.5.7 Quantification

The segmentation procedures resulted in two different GM segmentations for each data set and initial segmentation algorithm (SPM or CBS and FSL FAST): (i) an initial segmentation without any further changes, (ii) after correction using the GraMag method. To compare segmentation quality among these two outcomes, we calculated the Dice coefficient (DICE) and the Average Hausdorff Distance (AVHD) using the openly available EvaluateSegmentation Tool (2016; VISCERAL; www.visceral.eu).

The DICE is an overlap-based metric and it is the most popular choice for validating volume segmentations (Taha & Hanbury, 2015). We included it here as a familiar reference for the reader. However, overlap-based metrics like the DICE are not recommended for validating segmentation boundaries against the ground truth, as is our aim here, since they are relatively insensitive to boundary errors. In contrast, the AVHD is a distance metric sensitive to boundary errors (Taha & Hanbury, 2015). We

therefore consider the AVHD to be a more suitable metric for our purposes and we based our conclusions on the comparisons made with the AVHD.

Given that the AVHD quantifies the similarity of two boundaries, we first extracted WM-GM and GM-CSF boundaries from the ground truth segmentations and the different GM segmentations before calculating the AVHD. Here, an AVHD of zero indicates a perfect match between the segmentation and ground truth boundaries, while values larger than zero indicate a mismatch. In this case, the value represents the average number of voxels by which the two boundaries deviate from one another.

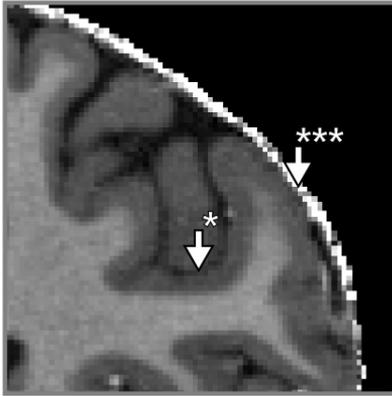
2.6 VALIDATION RESULTS

Visual inspection revealed that applying the GraMag method excluded most of the vessels and dura mater voxels and resulted in a more plausible GM matter definition for both MPRAGE and MP2RAGE data sets (see Figure 2.4).

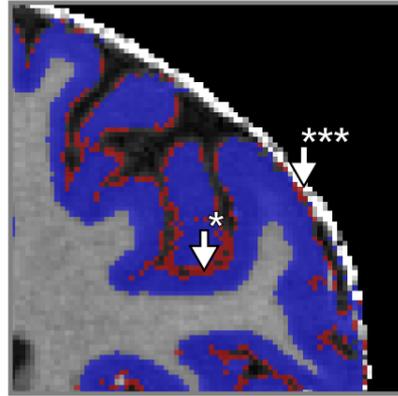
Table 2.1 compares segmentation performance before and after applying the GraMag method to the initial GM segmentations for the MPRAGE data set. The GraMag method led to an improvement of GM segmentations in all subjects for segmentations from SPM 12 and all but one subject for segmentations from FSL FAST. On average, the AVHD decreased from 0.735 ± 0.100 (mean \pm standard deviation across subjects) to 0.580 ± 0.022 for SPM 12 and from 0.601 ± 0.083 to 0.572 ± 0.073 for FSL FAST. The GraMag method did not affect the DICE coefficient. On average, it changed very little from 0.866 ± 0.014 to 0.860 ± 0.021 for SPM 12 and from 0.870 ± 0.015 to 0.873 ± 0.019 for FSL FAST.

Table 2.2 compares segmentation performances before and after applying the GraMag method to the segmentations for the MP2RAGE data set. The GraMag method decreased AVHD for all but one subject and, on average, from 0.544 ± 0.121 to 0.477 ± 0.087 for CBS tools and from 1.012 ± 0.043 to 0.830 ± 0.082 for FSL FAST. On average, it slightly increased the DICE coefficient for initial CBS tools segmentations from 0.861 ± 0.024 to 0.868 ± 0.028 and decreased the DICE coefficient for FSL FAST segmentations from 0.804 ± 0.035 to 0.785 ± 0.034 .

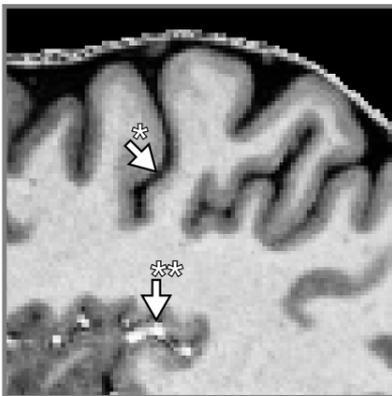
T1w/PDw



**SPM (red)
After GraMag mask (blue)**



UNI



**CBS Tools (red)
After GraMag mask (blue)**

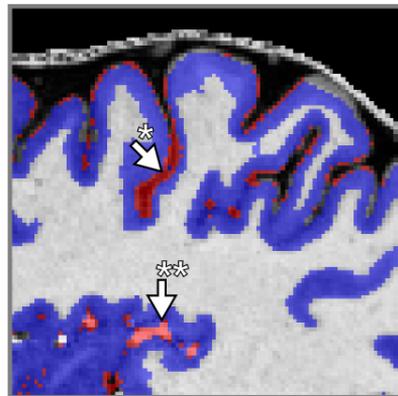


FIGURE 2.4: Comparison of GM segmentation results. GM segmentation results are shown on the transverse slices of two representative subjects for the MPRAGE (upper row) and the MP2RAGE (lower row) data set before and after applying the GraMag method. The original image that is input to the segmentation is shown on the left. The initial GM segmentation obtained with SPM 12 or CBS tools is shown in red. GM segmentations after additional polishing with brain mask obtained with the GraMag method are overlaid in blue. Additional masking removes CSF*, blood vessels**, and most of dura mater*** voxels from the initial GM definitions.

TABLE 2.1: Segmentation performance scores MPRAGE data set. The table shows the DICE (larger is better) and AVHD (less is better) for the initial SPM 12 and FSL FAST GM segmentations before and after applying the gradient magnitude method. DICE, DICE Coefficient; AVHD, Average Hausdorff Distance; Init, initial segmentation; GraMag, gradient magnitude method.

		SPM DICE	SPM AVHD	FAST DICE	FAST AVHD
S02					
	Init	0.8576	0.7413	0.8896	0.5375
	Init + GraMag	0.8594	0.5865	0.8902	0.4980
S03					
	Init	0.8644	0.8753	0.8694	0.5739
	Init + GraMag	0.8376	0.6139	0.8781	0.5612
S05					
	Init	0.8569	0.6379	0.8782	0.5230
	Init + GraMag	0.8688	0.5647	0.8909	0.5233
S06					
	Init	0.8601	0.7789	0.8635	0.6477
	Init + GraMag	0.8435	0.5797	0.8616	0.5909
S07					
	Init	0.8897	0.6403	0.8490	0.7228
	Init + GraMag	0.8893	0.5574	0.8460	0.6862

2.7 DISCUSSION

Functional and anatomical MRI studies at the mesoscale (< 1 mm isotropic) require accurate and precise definitions of the GM ribbon. Creating such definitions is currently a challenging task since sub-millimeter UHF data bring non-brain structures like blood vessels and dura mater into sharper focus. As a result, segmentation algorithms that have been benchmarked at lower resolution data might falsely label part of these structures as GM. Here we presented the GraMag method to correct many such mislabeled non-brain voxels efficiently and semi-automatically. The method was

TABLE 2.2: Segmentation performance scores MP2RAGE data set. The table shows the DICE (larger is better) and AVHD (less is better) for the initial CBS tools and FSL FAST GM segmentations before and after additional masking using the gradient magnitude method. DICE, DICE Coefficient; AVHD, Average Hausdorff Distance; Init, initial segmentation; GraMag, gradient magnitude method.

	CBS DICE	CBS AVHD	FAST DICE	FAST AVHD
S001				
Init	0.8688	0.4081	0.8157	1.0427
Init + GraMag	0.9032	0.4538	0.8095	0.8534
S013				
Init	0.8451	0.6146	0.7539	1.0551
Init + GraMag	0.8501	0.5377	0.7363	0.9299
S014				
Init	0.8389	0.6730	0.8089	0.9837
Init + GraMag	0.8410	0.5532	0.7868	0.7996
S019				
Init	0.8920	0.4798	0.8356	0.9672
Init + GraMag	0.8794	0.3639	0.8081	0.7359

based on theoretical expectations of how 3D brain data is to be represented in 2D histograms. These expectations imply that brain and non-brain tissue should become separable in 2D histogram representations that are based on gradient magnitude and intensity. We validated these expectations by implementing the suggested method in an openly available software package and by quantifying its added benefit using two new high-resolution validation data sets. We found that, in general, our suggested method offered an improvement compared to initial GM segmentations. However, we found some differences in the degree of improvement with respect to the (i) type of data and (ii) the algorithm used for initial segmentation. We will discuss these influences in turn.

First, we found that the improvements were slightly larger and more consistent across subjects for the MPRAGE than for the MP2RAGE data set.

This might be explained by the fact that the MPRAGE data conformed more to our theoretical expectations than the MP2RAGE data set. Especially, we found GM values in the MP2RAGE uni image to be less focused on one particular area of the 2D histogram (Figure S2.7) than the MPRAGE division image. This might result from differences in myelination level across cortical areas and depth (De Martino, Moerel, Xu, et al., 2015; Dick et al., 2012; Sereno, Lutti, Weiskopf, & Dick, 2013), which the MP2RAGE uni image might pick up more than MPRAGE division image (Marques & Gruetter, 2013).

Second, we observed that the performance of the initial segmentation algorithm had an influence on how much we could further improve the GM segmentation. If performance of the initial segmentation algorithm was already relatively high, the improvement obtained with our methods tended to be smaller. Differences in initial segmentation performance might be explained by whether the algorithm has been benchmarked on this particular type of data. We assume FSL FAST and CBS tools to have been benchmarked on MPRAGE and MP2RAGE data respectively, which would explain their relative high performance for these data types.

Importantly, our goal here was to aid already existing segmentation pipelines to deal with UHF sub-millimeter resolution data, not to replace those pipelines. Instead, the method presented here should be considered as an alternative to a large amount of manual slice-by-slice polishing of segmentations and thus as a time-saver. Manually correcting segmentation labels is very time-consuming and can quickly become unreliable. In contrast, our method greatly reduces the time required for manual polishing because it offers an efficient 2D summary and is more reliable because it is semi-automatic. Although the method presented here does not entirely eliminate the need for manual corrections, we estimated that for a whole brain cortical ribbon segmentation it saves on average 7.5 hours of manual work per subject (for details on this estimation see Section 2.8.1 in the Supplementary Material).

Our theoretical expectations implied that the methods presented here require high-resolution data (< 1 mm). This requirement was unfortunately not met by most available segmentation validation data sets. Simulated phantom ('BrainWeb') data (Collins et al., 1998) are available at 1 mm and thus fell short of the resolution required for our purposes.

Although an updated data set ('updated BrainWeb'; Aubert-Broche, Evans, and Collins, 2006; Aubert-Broche, Griffin, Pike, Evans, and Collins, 2006) is available at higher resolution, the simulations in this data set were based on initial 3T MRI acquisitions. As a consequence, the updated BrainWeb data revealed considerably less bright-vessel and dura-mater voxels than 7 T data usually does, and it was not suitable to validate our methods.

These considerations led us to create our own high-resolution segmentation validation data sets for which we established the 'ground truth' via manually-guided expert segmentation. While expert segmentations have well-known drawbacks (Despotović et al., 2015; Valverde, Oliver, Cabezas, Roura, & Lladó, 2015), they also have important advantages compared to alternative methods of establishing the ground truth, such as simulated phantom data. In particular, creating a validation data set based on empirical data and expert segmentations allowed us to benchmark our methods under conditions where image intensities fell into the expected range. Being aware of the problems with expert segmentations, we alleviate concerns about the quality of our expert judgment and consequently the validity of the results presented here by taking the following measures. First, the final ground truth segmentations were inspected by two additional experts. Second, we make the data sets and corresponding ground truth segmentations as well as our processing scripts available. This will allow other researchers to come up with their own judgment of the quality of the ground truth segmentations and validation data. In case changes to the ground truth are suggested and implemented, quantification could be re-run using our openly accessible work flow.

The 2D histogram method presented here is, in principle, capable of generating its own exhaustive tissue-type classifications, i.e., it does not necessarily depend on existing segmentation pipelines to derive GM and WM labels. While we expect the 2D histogram method to give no advantage over existing, fully-automated segmentation algorithms under standard conditions, the histogram method will compare well in cases where standard algorithms fail. Importantly, the 2D histogram method used here does not assume the data to conform to any atlas or template shape. Therefore, it is suitable also for acquisitions with only partial coverage and surface coils, or for specific populations like infant or even dolphin brains (see Figure S2.4).

Using histogram-based methods would be more attractive if the process of specifying transfer functions was fully automatic. We note that there is no principled obstacle to doing this. Indeed, information-theoretic measures have been suggested (Ip et al., 2012) that would make the normalized graph-cut application fully automatic, given the specification of an appropriate stopping criterion. The transfer functions (i.e. the circles and arcs applied to our 2D histograms) that we observed for the different brain tissue types were stable across subjects and conformed to expected, ideal shapes. This would allow to define probabilistic templates in the histogram space and transform the methods proposed here to a fully automatic exhaustive tissue-type classifications.

In the present work we have restricted ourselves to inputs of single 3D volumes. However, MRI can provide a multitude of informative images that weight different tissue properties to generate distinct image contrast. An obvious extension therefore is to combine information from several 3D volumes, where each volume contributes different contrast and information. Recently, a principled way to combine any number of MRI acquisitions has been suggested using the compositional data framework (Gulban, 2018). Combining the suggested method with the tools presented here could yield additional benefits. For example, in accordance with our theoretical expectation, the GraMag method identified and removed blood vessels and dura mater tissue. However, it often failed to sufficiently exclude connected tissue of the sagittal sinus because of similar intensity values and spatial proximity to GM. Given an appropriate contrast image that highlights the difference between GM and connected tissue, these structure should become separable. Naturally, performance of this method will depend on the specific choice of contrasts and whether these contrasts maximize the compositional difference between brain and non-brain tissue. Yet it can be envisioned that by selecting appropriate contrast images such a method could be used to single out targeted cortical or subcortical structures based on their compositional properties.

We understand our method as a secondary, more fine-grained brain extraction. When performing the initial brain extraction or tissue class segmentation, the user can often set parameters of the masking to be either more restrictive (at the risk of excluding brain tissue) or more liberal (at the risk of including a lot of non-brain tissue). We assume that, faced with

this trade-off, users will usually lean to the liberal choice of parameters to avoid that relevant brain tissue is excluded. In such cases, we suggest our methods will prove useful. Our methods go beyond simply choosing more restrictive parameters because they focus on information that is relevant to excluding vessels, dura mater and connective tissue (Figure 2.1D).

Our comparisons were limited to segmentations obtained from FSL, SPM and CBS tools. While several MRI studies at the mesoscale have used alternative ways of establishing tissue class segmentations (Fischl et al., 2004; Goebel, Esposito, & Formisano, 2006), we decided to limit our comparison to openly available algorithms. Furthermore, the resolution of our validation data exceeded the recommended input range for FreeSurfer (1 mm to 0.75 mm isotropic).

As is to be expected, we found our method to be impacted by the CNR of the input data (Figure S2.2). In particular, additional noise caused the 2D histogram representation to conform less to expected template shapes. However, we note that for images that were acquired with currently very common imaging parameters at UHF, we found our method to offer clear benefits in GM segmentations. Furthermore, in case acquisitions will be noisier than the ones tested here, additional processing steps like non-linear anisotropic smoothing (Mirebeau et al., 2015; Weickert, 1998) might be applied to mitigate noise issues (see Figure S2.3).

By making our validation data sets publicly available, we hope to inspire further algorithmic testing and development. There is currently a lack of validation data for the performance of tissue-type classification of MRI data acquired at UHF with sub-millimeter resolution. By publishing our data, our code, and our work flow, we invite fellow scientists to benefit from our work but also to further contribute to it. The neuro-imaging community can use our data to test the performance of entirely new methods or modifications to existing segmentation algorithms. Contributions could be made in the form of additional high-resolution data, more ground truth segmentations and algorithmic improvement. Anticipating such algorithmic improvements, we envision a future where segmentation of volumetric images will become gradually less laborious despite the increasing resolution and volume of the data.

2.8 SUPPLEMENTARY MATERIAL

2.8.1 *Time benefit estimation*

To estimate the time saved by substituting manual correction with Segmentator, we did the following calculation: Assuming that the objective of manually correcting mislabeled GM voxels is to reduce the number of false positives and to increase the number of true positives, we computed the number of false and true positives both before and after Segmentator intervention. We computed the difference in true positives and false positives before and after the application of segmentator and subtracted the true positive difference from the false positive difference. The resulting number is assumed to indicate the number of voxels which would have to be subtracted in the process of manually correcting all the voxels without using Segmentator. For the MRI data presented here, this number amounted to around 200000 voxels to be corrected (out of a total number of 1100000 voxels in the cortical ribbon). On average 35000 out of those 200000 voxels could be corrected using Segmentator. Assuming that a trained operator can manually correct one voxel per second, on average, this amounts to 7.5 hours of manual work that can be substituted with 10 minutes of Segmentator usage for the whole brain GM ribbon segmentation at 0.7 mm isotropic resolution. The time that can potentially be saved by using Segmentator will scale with the total number of GM voxels - it will be higher for high resolution acquisition (more voxels) and lower at low resolution (less voxels). The script used for our estimation is available at Schneider and Gulban, 2018.

2.8.2 Supplementary figures

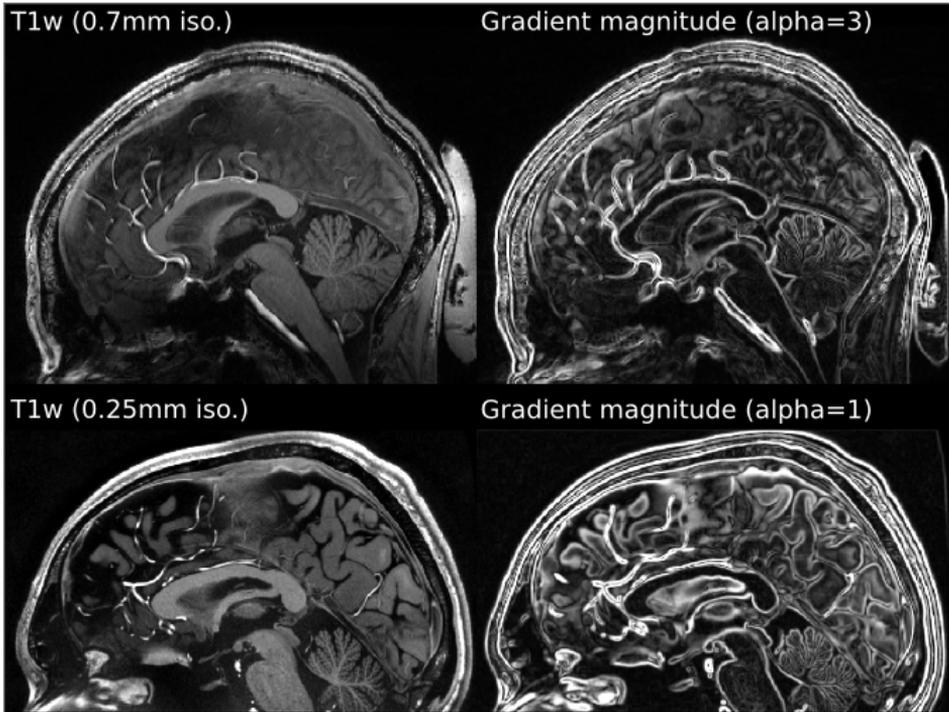


FIGURE S2.1: Appropriate kernel width approximates lower resolution. Intensity (left) and gradient magnitude (right) images are shown for T1w MRI data of a human brain (sagittal slice) that was either acquired at 0.7 mm isotropic (top) or at 0.25 mm isotropic (Lüsebrink, Sciarra, Mattern, Yakupov, & Speck, 2017; Lusebrink, Sciarra, Mattern, Yakupov, & Speck, 2017) (bottom). By choosing an appropriate kernel width for the very high resolution image (here $\alpha = 1$), the gradient magnitude image can be approximated to the lower resolution image, thus making it possible to use the gradient magnitude method also for very high resolutions.

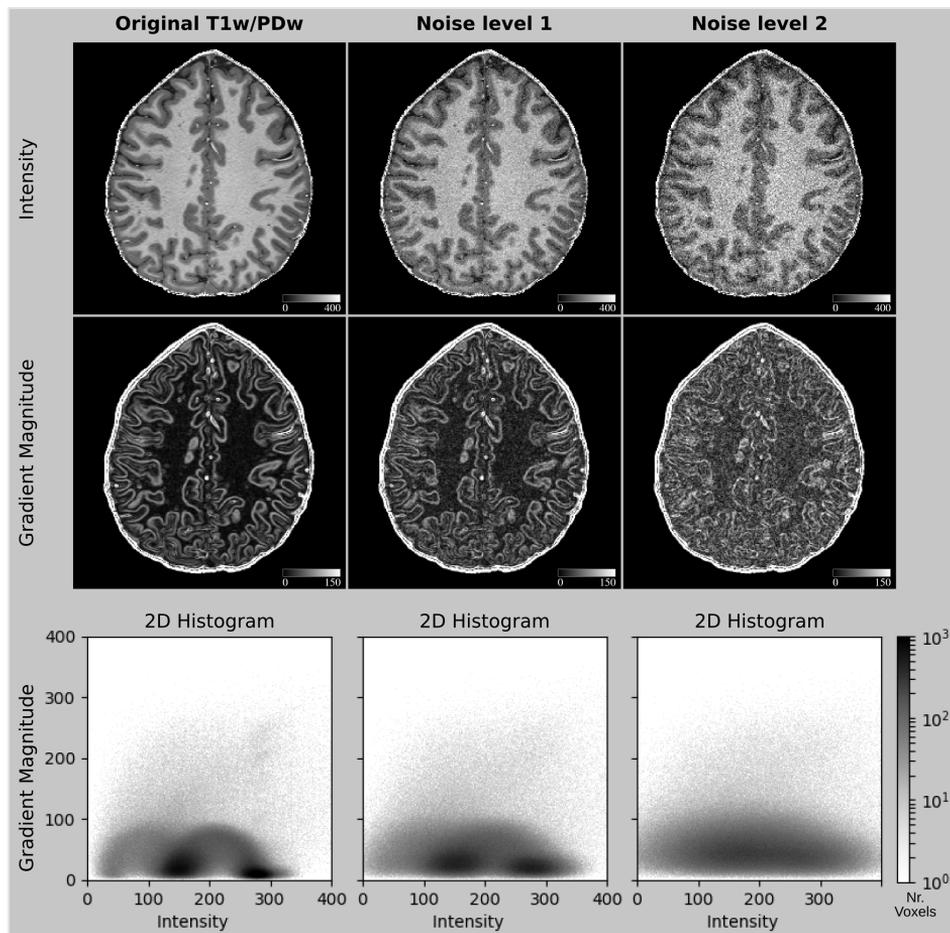


FIGURE S2.2: Impact of additional noise. Shown are intensity images (top row), gradient magnitude images (middle row) and 2D histograms (bottom row) for a T1w-divided-PDw MRI ratio image without any additional noise (left) and after applying a moderate (middle) and high (right column) amount of additive Gaussian noise with two levels of constant standard deviation of the distribution. The moderate noise ($\sigma = 25$) is 16% and high noise ($\sigma = 50$) is 32% calculated relative to the mean cortical GM intensity. Noise causes structures in the 2D histogram that are initially well-defined to spread outward and, at very high noise levels, to lose shape. Images show the transverse slice of an exemplary subject (sub-02).

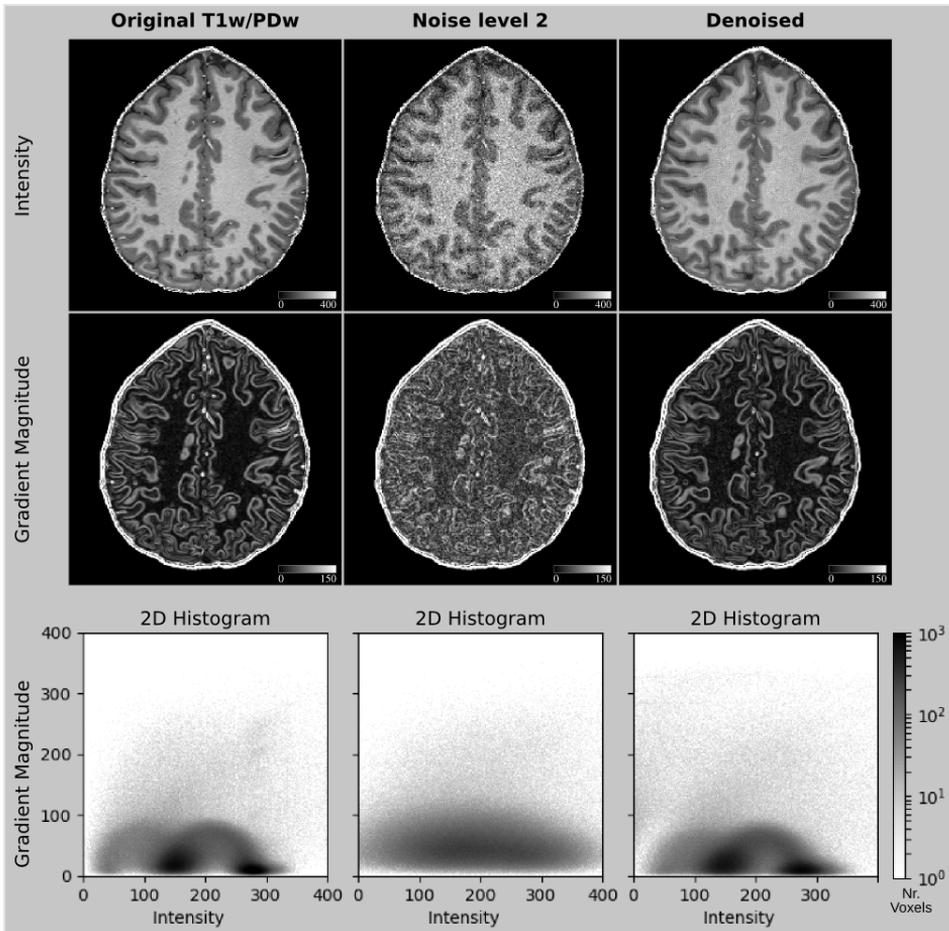


FIGURE S2.3: Noisy images can be denoised using non-linear anisotropic smoothing. Shown are intensity images (top row), gradient magnitude images (middle row) and 2D histograms (bottom row) for a T1w-divided-PDw MRI ratio image without any additional noise (left), after applying a high amount of noise (see Figure S2.2 for additional details) (middle), and after smoothing the noise-affected image (right column). As previously seen, noise causes structures in the 2D histogram to spread outward and to lose shape. This process can be reversed and noise-affected images can be recovered if a non-linear anisotropic smoothing filter (see Weickert, 1998) is applied. With smoothing, structures become more confined to the expected regions and well-defined shapes are regained. Images show the transverse slice of an exemplary subject (sub-02).

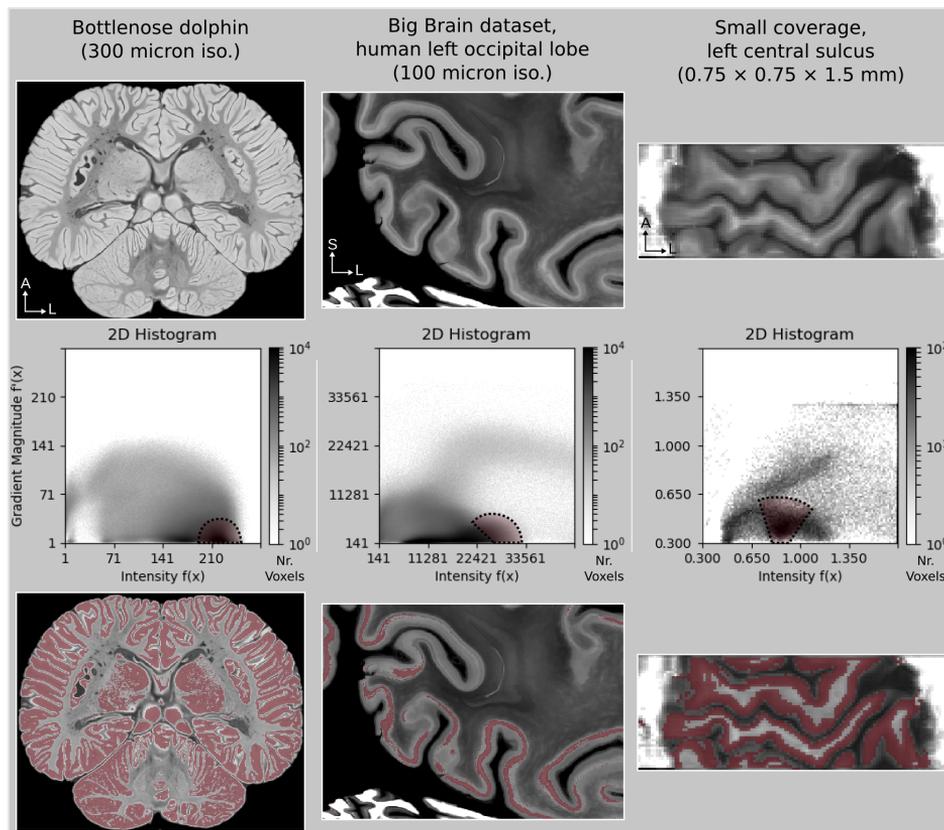


FIGURE S2.4: Application of GraMag method to extra-ordinary MR images. Shown are several examples of the variety of existing volumetric datasets for which our method appears useful. Every column represents different images: the brain of a bottle-nose dolphin (Toro, Grisanti, Herbin, & Santin, 2014) (left), the occipital lobe of a human brain with 100 micron resolution (Amunts et al., 2013) (middle) and a human motor cortex acquired with small partial coverage (T1w EPI) with anisotropic resolution (Huber et al., 2017) (right). For every image we show a slice (top row), selected voxels in the 2D histogram (middle row) and selected voxels overlaid on the slice (bottom row). These images do not contain large intensity inhomogeneities. Therefore, no bias-field correction was performed. Mild non-linear anisotropic diffusion-based smoothing was applied to enhance CNR.

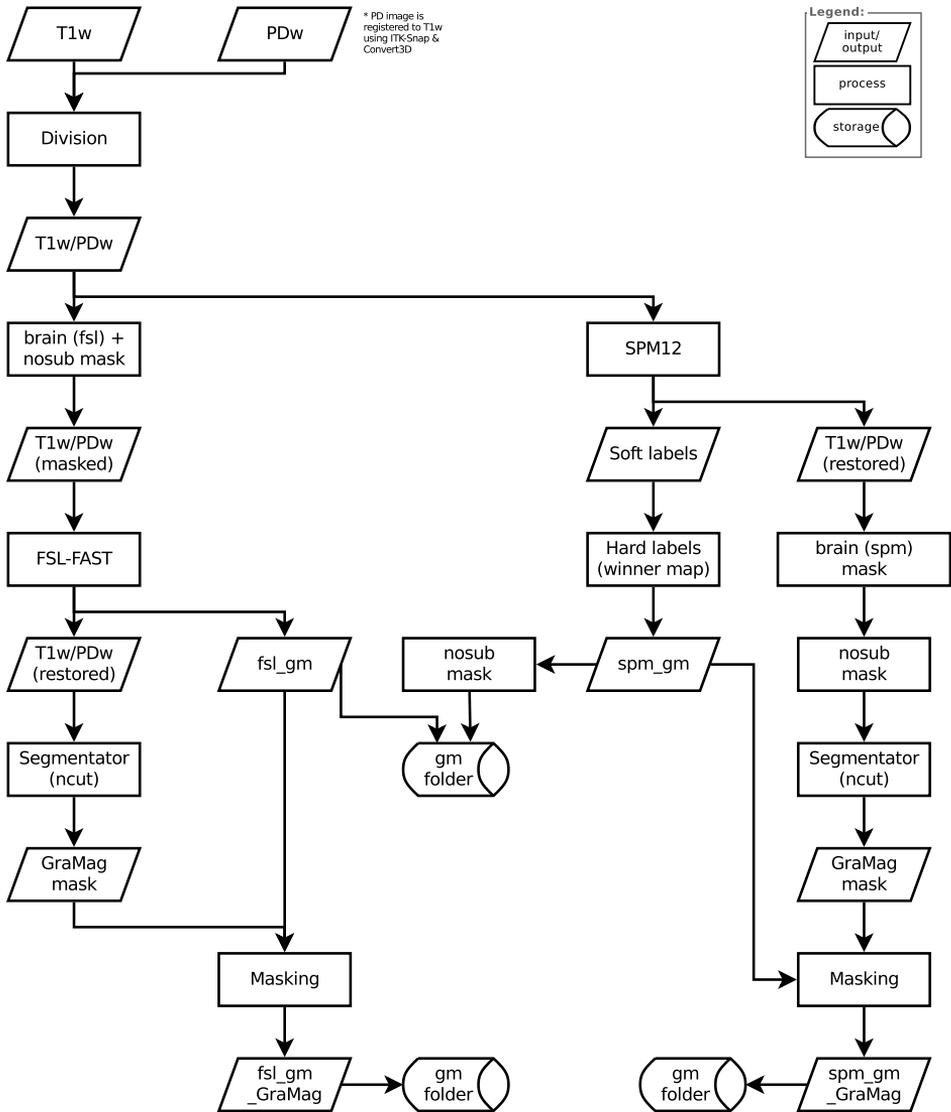


FIGURE S2.5: Flowchart diagram MPRAGE pipeline. This diagram provides a detailed overview of all the inputs, processing steps and outputs for MPRAGE pipeline. Rectangular shapes represent processing steps, rhombic shapes represent input or outputs and cylindrical shapes represent input or output locations.

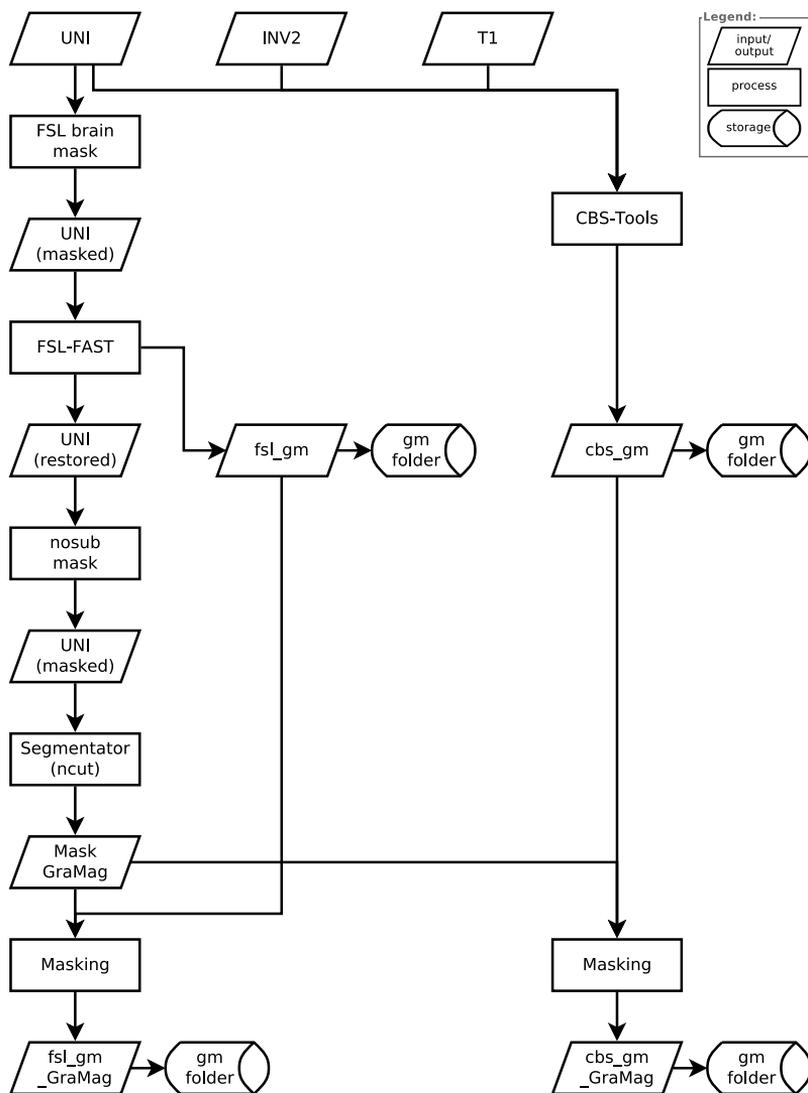


FIGURE S2.6: Flowchart diagram MP2RAGE pipeline. This diagram provides a detailed overview of all the inputs, processing steps and outputs for MP2RAGE pipeline. Rectangular shapes represent processing steps, rhombic shapes represent input or outputs and cylindrical shapes represent input or output locations.

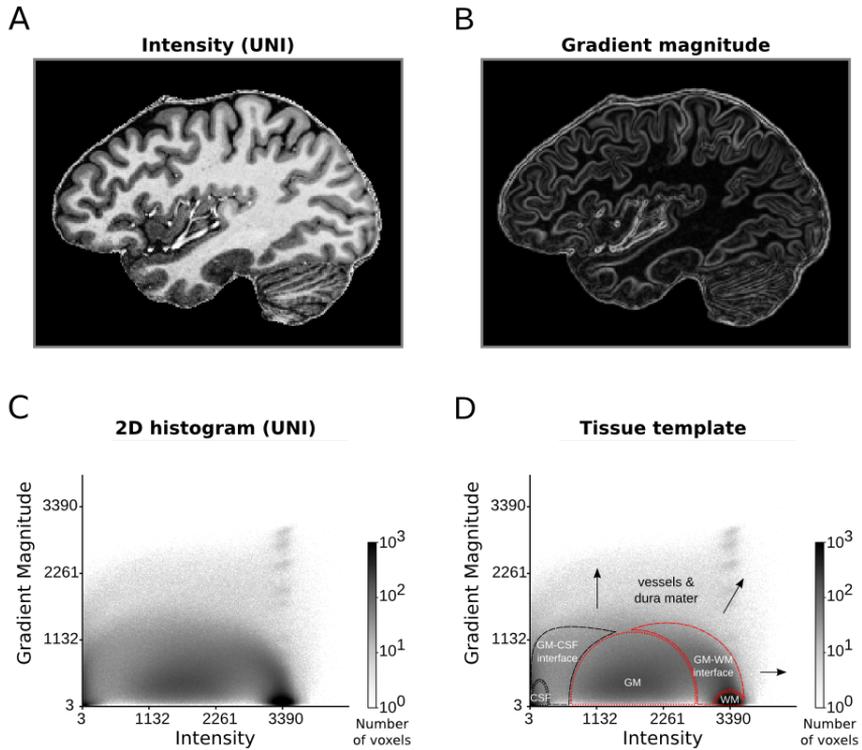


FIGURE S2.7: 2D histogram representation for MRI image of a human brain. The intensity (A) and gradient magnitude (B) values of a T1w-divided-by-PDw MRI image (MP2RAGE, 0.7 mm isotropic resolution) are represented in a 2D histogram (C). Darker regions in the histogram indicate that many voxels are characterized by this particular combination of image intensity and gradient magnitude. The 2D histogram displays a characteristic pattern with tissue types occupying particular areas of the histogram (D). Voxels containing CSF, dura mater or blood vessels (black dashed lines and arrows) cover different regions of the histogram than voxels containing WM and GM (red dashed lines). As a result, brain tissue becomes separable from non-brain tissue.

2.8.3 Supplementary tables

TABLE S2.1: Availability of validation data and code. Validation data and scripts as well as segmentation software are all openly accessible by following the corresponding links for their repositories.

What?	Where?
data	https://zenodo.org/record/1206163
scripts	https://zenodo.org/record/1219231
software	https://zenodo.org/record/1220388

2.9 ACKNOWLEDGMENTS

We would like to express our appreciation for the comments and suggestions of the reviewers which helped to improve the manuscript. We would also like to thank Thomas Emmerling for his guidance and advise on software implementation as well as Nikolaus Weiskopf, Rainer Goebel and Christophe Lenglet for valuable discussions of our early ideas.

2.10 REFERENCES

- Amunts, K., Lepage, C., Borgeat, L., Mohlberg, H., Dickscheid, T., Rousseau, M. É., ... Evans, A. C. (2013). BigBrain: An ultrahigh-resolution 3D human brain model. *Science*, *340*(6139), 1472–1475.
- Ashburner, J. & Friston, K. J. (2005). Unified segmentation. *NeuroImage*, *26*(3), 839–851.
- Aubert-Broche, B., Evans, A. C., & Collins, L. (2006). A new improved version of the realistic digital brain phantom. *NeuroImage*, *32*(1), 138–145.
- Aubert-Broche, B., Griffin, M., Pike, G. B., Evans, A. C., & Collins, D. L. (2006). Twenty new digital brain phantoms for creation of validation image data bases. *IEEE Transactions on Medical Imaging*, *25*(11), 1410–1416.
- Bazin, P. L., Weiss, M., Dinse, J., Schafer, A., Trampel, R., & Turner, R. (2014). A computational framework for ultra-high resolution cortical segmentation at 7 Tesla. *NeuroImage*, *93*, 201–209.
- Borg, I. & Groenen, P. (2005). Modern Multidimensional Scaling: Theory and Applications. *Chapter 10*, 100–131.
- Brett, M., Hanke, M., Côté, M.-A., Markiewicz, C., Ghosh, S., Wassermann, D., ... Nimmo-Smith, I. (2017). Nipy/nibabel: 2.2.0.
- Collins, D. L., Zijdenbos, a. P., Kollokian, V., Sled, J. G., Kabani, N. J., Holmes, C. J., & Evans, a. C. (1998). Design and construction of a realistic digital brain phantom. *IEEE Transactions on Medical Imaging*, *17*(3), 463–468.
- De Martino, F., Moerel, M., Ugurbil, K., Goebel, R., Yacoub, E., & Formisano, E. (2015). Frequency preference and attention effects across cortical depths in the human primary auditory cortex. *Proceedings of the National Academy of Sciences*, *112*(52), 16036–16041.
- De Martino, F., Moerel, M., Xu, J., Van De Moortele, P. F., Ugurbil, K., Goebel, R., ... Formisano, E. (2015). High-resolution mapping of myeloarchitecture in vivo: Localization of auditory areas in the human brain. *Cerebral Cortex*, *25*(10), 3394–3405.
- De Martino, F., Yacoub, E., Kemper, V., Moerel, M., Uludag, K., De Weerd, P., ... Formisano, E. (2018). The impact of ultra-high field MRI on

-
- cognitive and computational neuroimaging. *NeuroImage*, 168, 366–382.
- 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(3), e60514.
- Deriche, R. (1987). Using Canny's criteria to derive a recursively implemented optimal edge detector. *International Journal of Computer Vision*, 1(2), 167–187.
- Despotović, I., Goossens, B., & Philips, W. (2015). MRI segmentation of the human brain: challenges, methods, and applications. *Computational and Mathematical Methods in Medicine*, 2015, 450341.
- Dick, F., Taylor Tierney, A., Lutti, A., Josephs, O., Sereno, M. I., & Weiskopf, N. (2012). In Vivo Functional and Myeloarchitectonic Mapping of Human Primary Auditory Areas. *Journal of Neuroscience*, 32(46), 16095–16105.
- Dumoulin, S. O., Fracasso, A., van der Zwaag, W., Siero, J. C., & Petridou, N. (2018). Ultra-high field MRI: Advancing systems neuroscience towards mesoscopic human brain function. *NeuroImage*, 168, 345–357.
- Duyn, J. H. (2012). The future of ultra-high field MRI and fMRI for study of the human brain. *NeuroImage*, 62(2), 1241–1248.
- Fischl, B., Salat, D. H., Van Der Kouwe, A. J. W., Makris, N., Ségonne, F., Quinn, B. T., & Dale, A. M. (2004). Sequence-independent segmentation of magnetic resonance images. In *Neuroimage* (Vol. 23, SUPPL. 1).
- Fracasso, A., Petridou, N., & Dumoulin, S. O. (2016). Systematic variation of population receptive field properties across cortical depth in human visual cortex. *NeuroImage*, 139, 427–438.
- Goebel, R., Esposito, F., & Formisano, E. (2006). Analysis of Functional Image Analysis Contest (FIAC) data with BrainVoyager QX: From single-subject to cortically aligned group General Linear Model analysis and self-organizing group Independent Component Analysis. *Human Brain Mapping*, 27(5), 392–401.
- Goncalves, N. R., Ban, H., Sanchez-Panchuelo, R. M., Francis, S. T., Schluppeck, D., & Welchman, A. E. (2015). 7 Tesla fMRI Reveals Systematic

- Functional Organization for Binocular Disparity in Dorsal Visual Cortex. *Journal of Neuroscience*, 35(7), 3056–3072.
- Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., ... Haase, A. (2002). Generalized Autocalibrating Partially Parallel Acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6), 1202–1210.
- Gulban, O. (2018). The relation between color spaces and compositional data analysis demonstrated with magnetic resonance image processing applications. *Austrian Journal of Statistics*, 47(5), 34–46.
- Gulban, O. F. & Schneider, M. (2017). Segmentator v1.5.0. [Zenodo repository](#).
- Gulban, O. F., Schneider, M., Marquardt, I., Haast, R. A. M., & De Martino, F. (2017). Dataset: a scalable method to improve gray matter segmentation at ultra high field mri. [Zenodo repository](#).
- Haast, R. A. M., Ivanov, D., Formisano, E., & Uludağ, K. (2016). Reproducibility and Reliability of Quantitative and Weighted T1 and T2(*) Mapping for Myelin-Based Cortical Parcellation at 7 Tesla. *Frontiers in Neuroanatomy*, 10, 112.
- Harvey, B. M., Klein, B. P., Petridou, N., & Dumoulin, S. O. (2013). Topographic Representation of Numerosity in the Human Parietal Cortex. *Science*, 341(6150), 1123–1126.
- Harvey, B. M., Fracasso, A., Petridou, N., & Dumoulin, S. O. (2015). Topographic representations of object size and relationships with numerosity reveal generalized quantity processing in human parietal cortex. *Proceedings of the National Academy of Sciences*, 112(44), 13525–30.
- Helms, G. (2016). Segmentation of human brain using structural MRI. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 29(2), 111–124.
- Helms, G., Kallenberg, K., & Dechent, P. (2006). Contrast-driven approach to intracranial segmentation using a combination of T2- and T1-weighted 3D MRI data sets. *Journal of Magnetic Resonance Imaging*, 24, 790–795.
- Huber, L., Goense, J., Kennerley, A. J., Trampel, R., Guidi, M., Reimer, E., ... Möller, H. E. (2015). Cortical lamina-dependent blood volume changes in human brain at 7 T. *NeuroImage*, 107, 23–33.

-
- Huber, L., Handwerker, D. A., Jangraw, D. C., Chen, G., Hall, A., Stüber, C., . . . Bandettini, P. A. (2017). High-Resolution CBV-fMRI Allows Mapping of Laminar Activity and Connectivity of Cortical Input and Output in Human M1. *Neuron*, 96, 1253–1263.
- Hunter, J. D. (2007). Matplotlib: a 2d graphics environment. *Computing In Science & Engineering*, 9(3), 90–95.
- Ip, C. Y., Varshney, A., & Jaja, J. (2012). Hierarchical exploration of volumes using multilevel segmentation of the intensity-gradient histograms. *IEEE Transactions on Visualization and Computer Graphics*, 18(12), 2355–2363.
- Jähne, B. & Haußecker, H. (2000). *Handbook of computer vision and applications*. Elsevier.
- Jain, V., Turaga, S. C., Briggman, K., Helmstaedter, M. N., Denk, W., & Seung, H. S. (2011). Learning to Agglomerate Superpixel Hierarchies. *Advances in Neural Information Processing Systems*, 648–656.
- Jolliffe, I. T. (1986). Principal Component Analysis and Factor Analysis. In *Technometrics* (Vol. 30, pp. 115–128).
- Jones, E., Oliphant, T., Peterson, P., et al. (2007). SciPy: open source scientific tools for Python. *Computing in Science and Engineering*.
- Kashyap, S., Ivanov, D., Havlicek, M., Poser, B., & Uludağ, K. (2017). Impact of acquisition and analysis strategies on cortical depth-dependent fMRI.
- Kemper, V. G., De Martino, F., Emmerling, T. C., Yacoub, E., & Goebel, R. (2018). High resolution data analysis strategies for mesoscale human functional MRI at 7 and 9.4T. *NeuroImage*, 164, 48–58.
- Kindlmann, G. & Durkin, J. W. (1998). Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of the 1998 IEEE symposium on volume visualization - vvs '98* (pp. 79–86). New York, New York, USA: ACM Press.
- Kniss, J., Kindlmann, G., & Hansen, C. (2002). Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 8(3), 270–285.
- Kniss, J., Kindlmann, G., & Hansen, C. (2001). Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *Proceedings visualization 2001* (pp. 255–562). IEEE.

- Kniss, J., Kindlmann, G., & Hansen, C. D. (2005). Multidimensional transfer functions for volume rendering. *Visualization Handbook*, 189–209.
- Kok, P., Bains, L. J., Van Mourik, T., Norris, D. G., & De Lange, F. P. (2016). Selective activation of the deep layers of the human primary visual cortex by top-down feedback. *Current Biology*, 26(3), 371–376.
- Koopmans, P. J., Barth, M., Orzada, S., & Norris, D. G. (2011). Multi-echo fMRI of the cortical laminae in humans at 7 T. *NeuroImage*, 56(3), 1276–1285.
- Liu, T. & Jurrus, E. (2012). Watershed merge tree classification for electron microscopy image segmentation. *ICPR*, 3–6.
- Ljung, P., Krüger, J., Groller, E., Hadwiger, M., Hansen, C. D., & Ynnerman, A. (2016). State of the Art in Transfer Functions for Direct Volume Rendering. *Computer Graphics Forum*, 35(3), 669–691.
- Lüsebrink, F., Sciarra, A., Mattern, H., Yakupov, R., & Speck, O. (2017). Data from: t1-weighted in vivo human whole brain mri dataset with an ultrahigh isotropic resolution of 250 μm . Dryad Digital Repository.
- Lusebrink, F., Sciarra, A., Mattern, H., Yakupov, R., & Speck, O. (2017). T1-weighted in vivo human whole brain MRI dataset with an ultrahigh isotropic resolution of 250 μm . *Scientific Data*, 4, 1–11.
- Marques, J. P. & Gruetter, R. (2013). New Developments and Applications of the MP2RAGE Sequence - Focusing the Contrast and High Spatial Resolution R1 Mapping. *Plos One*, 8(7).
- Marques, J. P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P. F., & Gruetter, R. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *NeuroImage*, 49(2), 1271–1281.
- Mirebeau, J., Fehrenbach, J., Risser, L., & Tobji, S. (2015). Anisotropic diffusion in ITK. *The Insight Journal*, 1.
- Muckli, L., De Martino, F., Vizioli, L., Petro, L. S., Smith, F. W., Ugurbil, K., ... Yacoub, E. (2015). Contextual Feedback to Superficial Layers of V1. *Current Biology*, 25(20), 2690–2695.
- Nasr, S., Polimeni, J. R., & Tootell, R. B. H. (2016). Interdigitated color- and disparity-selective columns within human visual cortical areas V2 and V3. *Journal of Neuroscience*, 36(6), 1841–57.

-
- Nunez-Iglesias, J., Kennedy, R., Parag, T., Shi, J., & Chklovskii, D. B. (2013). Machine Learning of Hierarchical Clustering to Segment 2D and 3D Images. *Plos One*, 8(8).
- Polimeni, J. R., Fischl, B., Greve, D. N., & Wald, L. L. (2010). Laminar analysis of 7t BOLD using an imposed spatial activation pattern in human V1. *NeuroImage*, 52(4), 1334–1346.
- Polimeni, J. R., Renvall, V., Zaretskaya, N., & Fischl, B. (2018). Analysis strategies for high-resolution UHF-fMRI data. *NeuroImage*, 168, 296–320.
- Renvall, V., Witzel, T., Wald, L. L., & Polimeni, J. R. (2016). Automatic cortical surface reconstruction of high-resolution T1 echo planar imaging data. *NeuroImage*.
- Scharr, H. (2000). Optimale operatoren in der digitalen bildverarbeitung.
- Schneider, M. & Gulban, O. F. (2018). Processing scripts: A scalable method to improve gray matter segmentation at ultra high field MRI. *Zenodo repository*.
- Sereno, M. I., Lutti, A., Weiskopf, N., & Dick, F. (2013). Mapping the human cortical surface by combining quantitative t1 with retinotopy. *Cerebral Cortex*, 23(9), 2261–2268.
- Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3), 143–155.
- Taha, A. A. & Hanbury, A. (2015). Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging*, 15(1), 29.
- Teeuwisse, W. M., Brink, W. M., & Webb, A. G. (2012). Quantitative assessment of the effects of high-permittivity pads in 7 tesla mri of the brain. *Magnetic Resonance in Medicine*, 67(5), 1285–1293.
- Tootell, R. B. H. & Nasr, S. (2017). Columnar Segregation of Magnocellular and Parvocellular Streams in Human Extrastriate Cortex. *Journal of Neuroscience*, 0690–17.
- Toro, R., Grisanti, F., Herbin, M., & Santin, M. (2014). The brain catalogue: an open portal for comparative neuroanatomy. figshare.
- Trampel, R., Bazin, P.-L., Pine, K., & Weiskopf, N. (2017). In-vivo magnetic resonance imaging (MRI) of laminae in the human cortex. *NeuroImage*, 1–9.

- Ugurbil, K. (2014). Magnetic resonance imaging at ultrahigh fields. *IEEE Transactions on Biomedical Engineering*, 61(5), 1364–1379.
- Valverde, S., Oliver, A., Cabezas, M., Roura, E., & Lladó, X. (2015). Comparison of 10 brain tissue segmentation methods using revisited IBSR annotations. *Journal of Magnetic Resonance Imaging*, 41(1), 93–101.
- Van de Moortele, P.-F., Auerbach, E. J., Olman, C., Yacoub, E., Ugurbil, K., & Moeller, S. (2009). T1 weighted brain images at 7 Tesla unbiased for Proton Density, T2* contrast and RF coil receive B1 sensitivity with simultaneous vessel visualization. *NeuroImage*, 46(2), 432–446.
- Van Der Maaten, L. & Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9, 2579–2605.
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2), 22–30.
- van der Kouwe, A. J. W., Benner, T., Salat, D. H., & Fischl, B. (2008). Brain morphometry with multiecho MPRAGE. *NeuroImage*, 40(2), 559–569.
- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., ... Yu, T. (2014). scikit-image: image processing in Python. *PeerJ*, 2, e453.
- Vaughan, J. T., Garwood, M., Collins, C. M., Liu, W., Delabarre, L., Adriany, G., ... Ugurbil, K. (2001). 7T vs. 4T: RF power, homogeneity, and signal-to-noise comparison in head images. *Magnetic Resonance in Medicine*, 46(1), 24–30.
- Venkataraju, K. U., Paiva, A. R., Jurrus, E., & Tasdizen, T. (2009). Automatic markup of neural cell membranes using boosted decision stumps. In *Proceedings - 2009 IEEE International Symposium on Biomedical Imaging: from nano to macro, isbi 2009* (pp. 1039–1042).
- Viviani, R. (2016). A Digital Atlas of Middle to Large Brain Vessels and Their Relation to Cortical and Subcortical Structures. *Frontiers in Neuroanatomy*, 10, 12.
- Viviani, R., Pracht, E. D., Brenner, D., Beschoner, P., Stingl, J. C., & Stöcker, T. (2017). Multimodal MEMPRAGE, FLAIR, and R2* segmentation to resolve dura and vessels from cortical gray matter. *Frontiers in Neuroscience*, 11, 1–13.

-
- Waehnert, M. D., Dinse, J., Weiss, M., Streicher, M. N., Waehnert, P., Geyer, S., ... Bazin, P. L. (2014). Anatomically motivated modeling of cortical laminae. *NeuroImage*, *93*, 210–220.
- Weickert, J. (1998). Anisotropic diffusion in image processing. *Image Rochester*, *256*(3), 170.
- Yacoub, E. & Harel, N. (2008). High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences*, *105*(30), 254–264.
- Yushkevich, P. A., Piven, J., Cody Hazlett, H., Gimpel Smith, R., Ho, S., Gee, J. C., & Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *NeuroImage*, *31*(3), 1116–1128.
- Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*.
- Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P. F., Feinberg, D., Adriany, G., ... Yacoub, E. (2011). Mapping the organization of axis of motion selective features in human area MT using high-field fmri. *Plos One*, *6*(12), 1–10.

3

Columnar clusters in the human motion complex reflect consciously perceived motion axis

Based on

Schneider, M., Kemper, V. G., Emmerling, T. C., De Martino, F., & Goebel, R. (2019). Columnar clusters in the human motion complex reflect consciously perceived motion axis. *Proceedings of the National Academy of Sciences*, 116, 5096-5101.

3.1 ABSTRACT

The specific contents of human consciousness rely on the activity of specialized neurons in cerebral cortex. We hypothesized that the conscious experience of a specific visual motion axis is reflected in response amplitudes of direction-selective clusters in the human motion complex. Using sub-millimeter functional magnetic resonance imaging at ultra-high field (7 Tesla) we identified fine-grained clusters that were tuned to either horizontal or vertical motion presented in an unambiguous motion display. We then recorded their responses while human observers reported the perceived axis of motion for an ambiguous apparent motion display. Although retinal stimulation remained constant, subjects reported recurring changes between horizontal and vertical motion percepts every 7 to 13 seconds. We found that these perceptual states were dissociatively reflected in the response amplitudes of the identified horizontal and vertical clusters. We also found that responses to unambiguous motion were organized in a columnar fashion such that motion preferences were stable in the direction of cortical depth and changed when moving along the cortical surface. We suggest that activity in these specialized clusters is involved in tracking the distinct conscious experience of a particular motion axis.

3.2 INTRODUCTION

An important goal of neuroscience research is to dissociate neural signals pertaining to conscious perception from those related to sensory stimulation (Logothetis, 1989). This can be achieved by exposing observers to multistable stimuli, i.e., stimuli that lead to changes in perception despite unchanging stimulation of the senses (Brascamp, Sterzer, Blake, & Knapen, 2018; Leopold & Logothetis, 1999; Sterzer, Kleinschmidt, & Rees, 2009). For example, the bistable motion quartet (Figure 3.1A, upper panel) yields conscious perception of either horizontal or vertical motion (Ramachandran & Anstis, 1985). The perceived axis of motion switches every few seconds although the retinal stimulation remains constant. Only those neurons that modulate their activity with the perceived axis of motion should be considered content-specific neural correlates of consciousness (NCC) (Koch, Massimini, Boly, & Tononi, 2016; Metzinger, 2000).

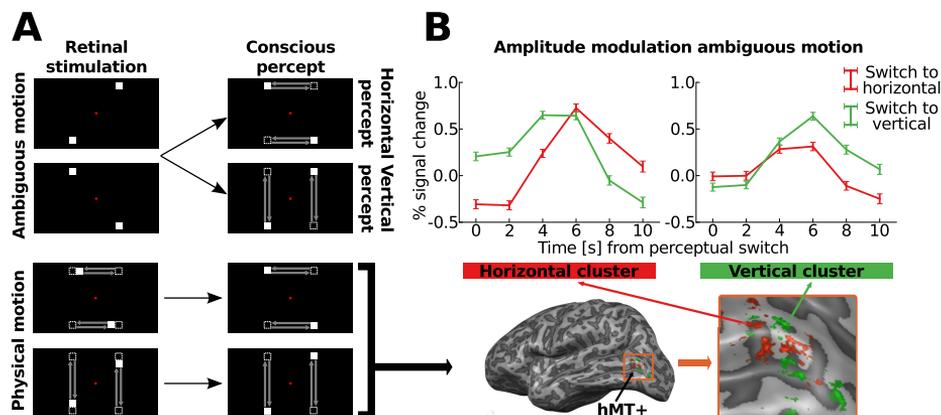


FIGURE 3.1: Experimental paradigm and event-related averages. **(A)** Overview of experimental paradigm. During the “ambiguous motion” experiment (top), the bistable motion quartet was presented. Presenting two opposed square pairs in temporal alternation gives rise to the conscious percept of either horizontal or vertical apparent motion. While the percept switches about every 10 seconds, the retinal stimulation remains constant throughout the experiment. During the “physical motion” experiment (bottom), squares moved physically along the horizontal or vertical motion paths that were previously only perceived during the ambiguous motion experiment. Physical motion led to unambiguous horizontal or vertical percepts. **(B)** The physical motion stimulus elicited responses in area hMT+, here displayed on the inflated, left hemisphere of a representative subject (bottom). These responses distinguished horizontal (red) and vertical (green) motion clusters. We studied responses in these clusters to the ambiguous motion experiment (top): Lines show event-related weighted averages across all subjects. Percent signal change is displayed during horizontal (red lines) and vertical (green lines) perceptual periods, as indicated by the subjects via button presses. At time point 0s subjects reported a perceptual switch specifying whether they now perceived horizontal or vertical motion. Error bars represent the uncertainty of the mean.

Multistable paradigms have revealed strong correlations between the experience of visual motion and responses in both the monkey and human middle temporal visual complex (MT+ and hMT+) (Leopold & Logothetis, 1999; Rees, 2007; Sterzer et al., 2009). When macaque monkeys are trained to signal their conscious percept during binocular rivalry, about 40 percent of MT+ neurons sampled during electro-physiological recordings

modulate their spiking with the perceived direction of motion (Leopold & Logothetis, 1999; Logothetis, 1989). Equally, during the presentation of bistable figures, neurons in MT (Dodd, Krug, Cumming, & Parker, 2001), medial superior temporal (MST) and parietal cortex (Williams, Elfar, Eskandar, Toth, & Assad, 2003) all show activity reflecting the consciously perceived direction of motion. Monkey areas MT and MST are also spatially organized into columns and clusters of neurons that are each tuned to a particular motion direction (Albright, Desimone, & Gross, 1984; Britten, 1998; Salzman, Britten, & Newsome, 1990). Yet it is unclear whether the NCC for motion direction is localized to these structures.

In humans, the functional magnetic resonance (fMRI) response in area hMT+ as a whole reflects conscious motion perception. Responses are increased for apparent as opposed to flicker motion (Muckli et al., 2002) and with reversals in perceived motion direction (Sterzer, Russ, Preibisch, & Kleinschmidt, 2002). When subjects indicate that they perceive horizontal component motion, the hMT+ amplitude is higher than for vertical pattern motion (Castelo-Branco et al., 2002). Furthermore, attended motion directions can be decoded from area hMT+ (Kamitani & Tong, 2006) and hMT+ signals predict perceptual states of clockwise or counterclockwise motion during ambiguous structure-from-motion (Brouwer & van Ee, 2007). Like monkey MT, hMT appears to be organized into columnar clusters (Zimmermann et al., 2011). To reveal this organization, however, advances in MRI technology (Uğurbil et al., 2003) and analysis strategies (Kemper, De Martino, Emmerling, Yacoub, & Goebel, 2018; Polimeni, Reinsel, Zaretskaya, & Fischl, 2018) were necessary that allow measuring fMRI responses at sub-millimeter spatial resolution. Until now fMRI studies on conscious motion perception have used relatively low resolutions (18 - 54 mm^3) and could not investigate whether responses in distinct clusters within hMT+ relate to the experience of a specific motion axis.

We recorded sub-millimeter fMRI responses from area hMT+ while humans viewed the bistable motion quartet stimulus and reported the perceived axis of motion via button presses. We hypothesized that (i) hMT+, like monkey MT+, has distinct functional subunits that each modulate with the perceptual state of a particular motion axis and (ii) that these units display a spatial organization into columnar clusters.

3.3 METHODS

3.3.1 *Participants*

Nine healthy participants with corrected-to-normal vision were recruited for the study. All participants were students or employees of Maastricht University and recruitment was limited to participants who were MRI compatible and had been in an MRI scanner at least once before to ensure high subject compliance. Additionally, all participants were experienced in fixation tasks. Four participants (S2, S4, S5, S6) were excluded from further data analysis (S2 did not show up for the second scanning session; S4 performed the task inappropriately during scanning, pressing buttons to indicate perceptual switches every 200 - 400 ms; S5 showed excessive head movement and S6 did not perceive stable states of horizontal and vertical apparent motion during the training session). Hence, five participants (3 females, 23-28 years old) were analyzed. All participants gave informed, written consent to participate in the experiment. The study was approved by the research ethics committee of the Faculty of Psychology and Neuroscience of Maastricht University.

3.3.2 *Experimental design and stimuli*

Visual stimuli were created and presented using the open-source application PsychoPy (version 1.82.01) (Peirce, 2007, 2008). Stimuli were projected on a frosted screen at the head end of the scanner (using Panasonic projector PT-EZ570; Newark, NJ, USA; resolution 1920 × 1200; nominal refresh rate: 60 Hz). Subjects viewed the screen via a tilted mirror attached to the head coil. Button responses were registered using an MR compatible button box (Current Designs, 8-button response device, HHSC-2 × 4-C; Philadelphia, USA). All scripts used for stimulus presentation are available as [GitHub](#) and [Zenodo repository](#).

For every subject we collected 3 runs to localize area hMT+ (240 volumes, 12 repetitions per condition per run), 5-6 runs of the ambiguous motion experiment (300 volumes, 22.30 ± 1.14 [mean \pm standard deviation] repetitions per condition per run) and 5-6 runs of the physical motion experiment (300 volumes, 24 repetitions per condition per run). These runs were acquired across two different scanning sessions on separate

days. Additionally, we performed three control experiments to exclude the possibility that observed differences between horizontal and vertical conditions were driven by retinotopic differences: For two subjects (S1 and S3), we acquired 3 runs (204 volumes, 6 repetitions per condition per run) to estimate axis-of-motion tuning curves (control experiment I) which were recorded during the first scanning session of the motion quartet experiments. Three subjects (S1, S3, S9) were recruited for an additional, third scanning session and we acquired 5-8 runs (222 volumes, 12 repetitions per condition per run) to map out responses to retinotopically identical horizontal and vertical motion conditions (control experiment II). Finally, for two subjects (S3 and S7) we organized a fourth scanning session to obtain 6 runs (172 volumes per run, in total 27 repetitions per aperture position) of data that allowed us to estimate population receptive field (pRF) parameters (control experiment III).

HMT+ LOCALIZER Area hMT+ was localized using standard moving and static dot stimuli (Amano, Wandell, & Dumoulin, 2009; Emmerling, Zimmermann, Sorger, Frost, & Goebel, 2016; Huk, Dougherty, & Heeger, 2002) presented in a circular aperture (10° of visual angle in diameter) that was either centered or displaced by 5° of visual angle to the left or right. All dots (200 dots; 0.2° in diameter; white on black background) moved at 8° per second and were either expanding or dilating, alternating direction every second. Every dot had a limited life time of 167 ms (10 frames) and was reborn at a randomized location within the aperture. 4s periods of moving dots were followed by periods of static dots (randomized duration: either 6s, 8s or 10s). The aperture position for static dots always matched the preceding aperture position for moving dots. Participants were asked to fixate on a centrally presented dot throughout the entire experimental run. The fixation dot consisted of a red circle surrounded by a yellow annulus (Marquardt, Schneider, Gulban, Ivanov, & Uludağ, 2018). The fixation dot occasionally (20 targets per run) changed color from red to yellow for 0.3 s and participants indicated a color change by pressing a button.

AMBIGUOUS AND PHYSICAL MOTION EXPERIMENT Participants were presented with two different motion quartet stimuli, an "ambiguous" and a disambiguated or "physical" version. In the ambiguous version (Experi-

ment 1) we used the apparent motion quartet (Ramachandran & Anstis, 1985) (see upper panel Figure 3.1). The quartet is composed of four blinking squares (each $1^\circ \times 1^\circ$ visual angle) in a rectangular configuration. Crucially, at any time point, only two squares at diagonally opposite corners are shown. A pair of squares was presented for 150 ms (9 frames) followed by an inter-stimulus interval of 67 ms (4 frames). Such a presentation frequency of 2.3 Hz was shown to lead to strong perception of apparent motion (Finlay & von Grünau, 1987).

In the physical version (Experiment 2) squares moved unambiguously either horizontally or vertically. This was achieved by physically moving white squares along the motion paths that were previously only perceived in Experiment 1 (see lower panel Figure 3.1A). Square positions were updated according to a harmonious oscillation, where the square location was updated as the sine of time (Muckli, Kohler, Kriegeskorte, & Singer, 2005). Stimulus parameters in Experiment 2 were matched to those in Experiment 1: Square size remained $1^\circ \times 1^\circ$ of visual angle. Outermost horizontal and vertical points of the physical motion trajectory were equivalent to square positions in Experiment 1 (individually adjusted for every subject during a training session, *see below*). Duration of a movement cycle either along the horizontal (left to right to left) or vertical trajectories (top to bottom to top) was 433 ms (26 frames), which matched the presentation frequency in Experiment 1.

The ambiguous motion quartet was shown during periods of 40s, which were interleaved with 16s baseline periods. A flicker motion quartet in which all four squares were shown synchronously served as the baseline. This local flicker stimulus did not induce apparent motion but still had the same stimulus energy as the regular motion quartet. Participants were instructed to fixate a red dot in the center of the screen throughout the experiment and to indicate the perceived axis of motion (vertical or horizontal) by button responses. The mapping of perceived motion axis to buttons (whether button "1" or button "2" was used to indicate horizontal or vertical motion) was counterbalanced across the two scanning sessions.

The physical motion stimulus was equally shown during periods of 40 s in total. However, this time the motion axis was unambiguous and alternated every 10 s between horizontal and vertical motion. During the 16 s baseline period, participants again viewed the flicker motion quartet.

They fixated the central red dot and indicated the perceived axis of motion by button responses. This was done although the motion axis was clearly determined by the stimulus, in order to engage participants in a task and to keep the task constant across Experiments 1 and 2.

CONTROL EXPERIMENT I: AXIS-OF-MOTION TUNING In order to determine the axis-of-motion tuning of voxels, we presented moving and static dot stimuli in a central aperture (aperture diameter 10° visual angle; 200 dots; 0.2° in diameter; white on black background) such that the aperture field covered the real and perceived motion trajectories in Experiments 1 and 2. There were four different conditions, each corresponding to one of four axes of motion (0° - 180° , 45° - 225° , 90° - 270° , 135° - 315° ; thus including the vertical and horizontal motion axes occupied in Experiments 1 and 2). 6s periods of moving dots were interleaved with static rest periods that lasted 8s, 10s or 12s. The order of conditions was randomized. During axis-of-motion blocks the direction of dots alternated every second, e.g., during the 0° - 180° condition, leftward motion alternated with rightward motion. Dots were moving at a speed of 8° of visual angle per second. Participants performed the same fixation task as described for the hMT+ localizer.

CONTROL EXPERIMENT II: MOVING DOT PATTERN In order to determine the tuning of voxels to motion either along the horizontal or vertical axis, we repeated control experiment I with two modifications. First, instead of four there were only two conditions, showing dots moving either along the horizontal (0° - 180°) or vertical (90° - 270°) motion axis. Second, visibility of the dots was restricted to a square annulus aperture. Shape and location of the aperture were defined as to reveal positions in visual space that had been occupied during the motion quartet experiments (simultaneously the locations of both horizontal and vertical motion trajectories as well as the four inducer squares). Importantly, this resulted in the horizontal and vertical motion conditions being retinotopically identical and only differing with regard to the motion axis of the dots revealed by the aperture. Area hMT+, unlike areas V1-V3, does not display an aperture bias resulting in higher responses at retinotopic locations where moving dots enter the stimulus aperture, making this experimental set-up suitable to study the question of interest here (Wang, Merriam, Freeman, & Heeger, 2014). The

aperture was adjusted for every control subject to match subject-specific vertical distances (Table S3.1). Other parameters were identical to those reported for control experiment I.

CONTROL EXPERIMENT III: PRF MAPPING The pRF mapping stimuli consisted of apertures that took the shape of semi rings and were presented at systematically varied positions. A single semi ring aperture had a radial extent of 3° of visual angle and subtended 180° angular degrees. Apertures were presented at four orientations (horizontal, vertical and the two diagonal orientations in-between) and at four eccentricity levels (centered on eccentricities of 1.5° , 4.5° , 7.5° , and 10.5° of visual angle). Apertures at the innermost eccentricity level (1.5°) consisted of semi circles, not semi rings. All apertures were limited to a circular region of the display, 24° of visual angle in diameter, and presented against a mean luminance gray background.

The carrier pattern of the mapping stimuli consisted of a random texture pattern as described by Kolster, Peeters, and Orban, 2010 and was presented at 98.2% Michelson contrast (which was the highest possible contrast in the scanner environment). Aperture positions were updated with every TR (3 s) and in a pseudo-random fashion such that two subsequent apertures never overlapped retinotopically. Every functional run started and ended with 15 s blank screen. To aid the estimation of large pRFs, 18×3 s periods of blank screen were inserted throughout the functional run (Amano et al., 2009).

Participants performed the same fixation task as described for the hMT+ localizer. To facilitate central fixation, an extended grid of thin lines was presented throughout the experimental run (Tyler et al., 2005). The grid consisted of two diagonal, orthogonal lines crossing behind the fixation dot as well as circles with radii of 1, 2.5, 5, 10 and 15 degrees of visual angle centered on the fixation dot.

PRE-SCAN TRAINING SESSION With central fixation, observers of the ambiguous motion quartet more frequently perceive vertical than horizontal motion (Chaudhuri & Glaser, 1991). This is thought to reflect the cost of inter-hemispheric processing, which is necessary for horizontal but not for vertical motion (Chaudhuri & Glaser, 1991; Gen, Bergmann, Singer, & Kohler, 2011). To ensure that the motion quartet stimulus was bistable,

i.e. that horizontal and vertical perceptual intervals were approximately equal in length, we scheduled a 20-30 minutes training session for every participant before the first scanning session. During this training session, we kept the horizontal distance between squares of the motion quartet constant (3° visual angle from square center to central fixation point), while we adjusted the vertical distance of the squares for every subject to obtain a bistable stimulus display. Table S3.1 shows the resulting subject-specific vertical distances. The ratio of vertical-to-horizontal distances was between 1.26 and 1.30 for all subjects, which is in agreement with previously reported ratios for motion quartet stimuli to achieve equal frequencies (Chaudhuri & Glaser, 1991). Once the vertical distance had been calibrated for a subject, it was kept constant across Experiment 1 and 2 and the two scanning sessions.

3.3.3 MRI acquisition

Data acquisition was performed on a whole-body Magnetom scanner (nominal field strength 7 Tesla [T], Siemens Medical Systems, Erlangen, Germany) at the Maastricht Brain Imaging center, The Netherlands. All images were acquired using a 32-channel head-coil (NovaMedical Inc.; Wilmington, MA, USA).

HMT+ LOCALIZER To aid the positioning of the sub-millimeter slab for Experiments 1-2 and to determine our region of interest, we acquired one hMT+ localizer scan in the first session and two hMT+ localizer scans in the second session. We used a 2D gradient echo (GE) echo planar imaging (EPI) sequence (1.6 mm isotropic nominal resolution; TE/TR = 18/2000 ms; in-plane field of view [FoV] 150×150 mm; matrix size 94 × 94; 28 slices; nominal flip angle [FA] = 69° ; echo spacing = 0.71 ms; GRAPPA factor = 2, partial Fourier = 7/8; phase encoding direction head - foot). We ensured that the area of acquisition had bilateral coverage of the posterior inferior temporal sulci, where we expected the hMT+ areas. Before acquisition of the first functional run, we collected 10 volumes for distortion correction - 5 volumes with the settings specified above and 5 more volumes with identical settings but opposite phase encoding (foot - head).

AMBIGUOUS AND PHYSICAL MOTION EXPERIMENT For the sub-millimeter measurements, we used a 2D GE EPI sequence (TE/TR = 25.6/2000 ms; in-plane FoV 148×148 mm; matrix size 186 × 186; slices = 28; nominal FA = 69°; echo spacing = 1.05 ms; GRAPPA factor = 3, partial Fourier = 6/8; phase encoding direction head - foot), yielding a nominal resolution of 0.8 mm isotropic (Moeller et al., 2010; Setsompop et al., 2012). Placement of the small functional slab was guided by online analysis of the hMT+ localizer data recorded immediately at the beginning of the first session. This allowed us to ensure bilateral coverage of area hMT+ for every subject. In the second scanning session, the slab was placed using Siemens auto-align functionality and manual corrections. Before acquisition of the first functional run, we collected 10 volumes for distortion correction (5 volumes with opposite phase encoding: foot - head). During acquisition, runs for the ambiguous and physical motion experiments were interleaved.

CONTROL EXPERIMENTS The acquisition parameters for control experiments I and II were identical to those reported for the motion quartet experiments. For the pRF mapping, we used the same 2D GE EPI sequence as in the motion quartet experiments with slightly adjusted parameters to obtain a larger field of view (in particular, multi band factor = 2; TE/TR = 23.2/3000 ms; in-plane FoV 140×140 mm; matrix size 176 × 176; slices = 82; nominal resolution = 0.8; nominal FA = 82°; echo spacing = 1.04 ms; GRAPPA factor = 3; partial Fourier = 6/8; phase encoding direction head - foot) (Feinberg et al., 2010; Moeller et al., 2010; Setsompop et al., 2012).

ANATOMICAL SCANS For visualization of the functional results, we acquired scans with structural information in the first scanning session. At high magnetic fields, MR images exhibit high signal intensity variations that result from heterogeneous RF coil profiles (Van de Moortele et al., 2009). We therefore acquired both T1w images and PDw images using a magnetization-prepared 3D rapid gradient-echo (3D MPRAGE) sequence (TR: 3100 ms [T1w] or 1440 ms [PDw], voxel size = 0.6 mm isotropic, FOV = 230 × 230 mm², matrix = 384 × 384, slices = 256, TE = 2.52 ms, FA = 5°). Acquisition time was reduced by using 3× GRAPPA parallel imaging and 6/8 Partial Fourier in phase encoding direction (acquisition time [TA]: 8 min 49 s [T1w] and 4 min 6 s [PDw]).

3.3.4 Behavioral data analysis

For every subject we calculated the mean length of horizontal and vertical perceptual periods during the ambiguous motion experiment. To test for differences between horizontal and vertical periods, we conducted an independent-samples t-test for every subject ($p < .05$, two-sided). We also calculated the mean, standard deviation and range across subjects for the two types of perceptual periods.

3.3.5 Structural data analysis and segmentation

Structural images were processed using advanced segmentation tools in BrainVoyager 20.0 (Brain Innovation, Maastricht, The Netherlands), SPM's bias correction (Ashburner & Friston, 2005), ITK-SNAP (Yushkevich et al., 2006), FSL BET (Smith, 2002), morphological operations (Jones, Oliphant, Peterson, et al., 2007) and Segmentator (Gulban, Schneider, Marquardt, Haast, & De Martino, 2018). Where not specified otherwise, default settings were used. Since the processing included many different steps and programs, a diagram of the input/output relationships for all processing steps is provided in Figure S3.1. We first registered the PDw image for every subject to their T1w image, using ITK-SNAP's automatic co-registration tools. In order to reduce the B1 bias field, we divided the T1w image by the PDw image and obtained a ratio image (Van de Moortele et al., 2009). The ratio image was brain-masked with a mask obtained by inputting the co-registered PDw image to FSL BET. We then used SPM's anatomical bias correction to further reduce inhomogeneities in the ratio image.

The resulting image was input to BrainVoyager's advanced segmentation routine to obtain a white matter (WM) definition. This initial WM image was inspected and manually polished in ITK-SNAP using the adaptive paint brush mode in combination with a graphics tablet (Intuos Art; Wacom Co. Ltd; Kazo, Saitama, Japan). Special emphasis was placed on corrections in the region of interest (bilateral hMT+). In two additional steps, using the round paint brush mode, the parts of the cerebellum and sagittal sinus that were falsely included in the WM definition were removed and brain stem structures and ventricles were masked.

The thus polished WM definition and ratio image were input to BrainVoyager's advanced segmentation routine to obtain a gray matter (GM) definition. This GM definition tended to be too inclusive, containing dura mater and blood vessels. For this reason, we masked the ratio image with the WM and GM definitions and input the resulting image to Segmentator. This allowed us to further exclude non-brain voxels from the GM definition based on their 2D histogram profile along the dimensions of image intensity and gradient magnitude. The thus improved GM definition was inspected and further polished manually in ITK-SNAP using the round paint brush mode. For later mesh visualization, the WM-GM segmentation images were separated in two hemispheres using ITK-SNAP's scalpel tool. Remaining topological errors were corrected using the bridge removal option in BrainVoyager (Kriegeskorte & Goebel, 2001) and manual correction. Note that all segmentation steps were performed at the original resolution of the anatomical images (0.6 mm isotropic).

3.3.6 Functional data - preparation

Functional data were processed using BrainVoyager 20.0, SPM 12 (Friston, 2006), FSL 5.0 (Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012) as well as custom code in Python 2.7 (Hunter, 2007; Jones et al., 2007; Van Der Walt, Colbert, & Varoquaux, 2011) and in Matlab R2014a (The Mathworks Inc.; Natick, MA, USA). Where not specified otherwise, default settings were used. All functional pre-processing steps were scripted and scripts are available as a [GitHub](#) and [Zenodo repository](#). Figure S3.2 provides an overview of all functional pre-processing steps.

PRE-PROCESSING STEPS Pre-processing for all sub-millimeter images was performed in the following order: slice scan time correction (BrainVoyager), motion-correction (SPM 12), linear trend removal and high-pass filtering (5 cycles) using a general linear model (GLM) Fourier basis set (BrainVoyager) and distortion correction (FSL topup). For motion and distortion correction, we deviated from default settings. Functional images from the first session were motion-corrected using SPM 12 in three steps. First, the first image of each run was realigned to the first image of the first run. Second, the images within each run were aligned to the first image of the run. Third, to avoid local minima, after these first two steps, a mean of all images was calculated

and images were realigned to this mean. Motion correction was limited to voxels inside the brain based on an intensity-thresholded and manually corrected brain mask of the functional images. Note that the results of the three steps were combined into a single transformation that was applied to functional images to minimize interpolation artefacts. Functional images from the second scanning session were motion-corrected in a similar vein, with the difference that the first functional image of the second session served as reference image in the first correction step. Steps 2 and 3 were identical. For details on the alignment across scanning sessions, please see below. EPI distortions of the functional images were corrected using FSL topup (Andersson, Skare, & Ashburner, 2003; Smith et al., 2004). The pairs of opposite phase encoding images acquired at the beginning of the first session were input to topup to estimate the susceptibility-induced off-resonance field. The estimated field was then used to correct the distortions for all functional images of the first and second scanning session. For the MT+ localizer images, pre-processing steps included slice scan time correction, motion-correction, linear trend removal combined with high-pass filtering (all performed in BrainVoyager) and distortion correction (FSL topup). In addition to the pre-processing steps described above, we applied Gaussian spatial smoothing (smoothing kernel = 0.8 mm) to the pRF mapping data, reflecting the prior assumption that pRF parameters vary smoothly over space.

REGISTRATION TO FUNCTIONAL IMAGES All statistical analyses were conducted in the space of the sub-millimeter functional images in the first scanning session (hereafter called "voxel space"). To this end, the pre-processed anatomical T1w-divided-PDw ratio image was registered to voxel space by exploiting the scanner's positional information and fine-tuning co-registration with boundary-based registration (Greve & Fischl, 2009) as implemented in BrainVoyager. The result was visually inspected for each subject by overlaying functional and anatomical images (anatomical images were displayed with inverted image intensities for better visualization). Co-registration and down-sampling of the anatomical image were performed in one step.

In order to register the high-resolution segmentation images (0.6 mm isotropic) to lower-resolution voxel space (0.8 mm isotropic) we proceeded

as follows. Problematically, simply transforming an image with only two segmentation values (for WM and GM) results in interpolation artefacts and how to bin the resulting distribution of values to regain only two segmentation values is not obvious. For that reason, we performed a surface reconstruction of both the inner and outer GM border in BrainVoyager. This resulted in two different surface meshes to which we applied the established transformation from anatomical images to voxel space. Transformed meshes were smoothed, using advanced mesh smoothing tools (restricting vertex displacement to 0.5). Importantly, the advanced mesh smoothing is restricted to high-frequencies and leaves low-frequency changes such as cortical folds intact. Smoothed meshes were projected back into volume space and area filling tools were used to regain a segmentation image with only two values. We inspected the resulting WM-GM segmentation images in voxel space separately for every subject and found this procedure to preserve important features of the segmentation and to minimize the re-introduction of topological errors (see Figure S3.3A). Re-introduced errors were manually corrected using *itksnap*. This was limited to 20-30 mislabeled voxels in our region of interest.

hMT+ localizer functional images acquired in the second session were first registered to the anatomical images in the first session based on the scanner's positional information, manual corrections, and boundary-based registration (Greve & Fischl, 2009) for fine-tuning. Subsequently, the already established transformation from anatomical images to voxel space was used. MT+ localizer functional images were up-sampled from 1.6 mm isotropic to 0.8 mm isotropic using nearest-neighbor interpolation. All transformations were combined and applied in a single step to avoid unnecessary interpolation.

To align the sub-millimeter functional images acquired in the second scanning session (and in the control experiments) to voxel space, we chose a more complex procedure than for the hMT+ localizer functional images. This was because we found EPI distortions to differ across scanning sessions for sub-millimeter functional images and simply co-registering both (distortion-corrected) functional sessions to the anatomy did not meet our quality demands for sub-millimeter analyses. For that reason, we calculated a mean image across time, separately for every scanning session (before EPI-distortion correction). We then ran an initial affine registration

between the mean image of the first session and each of the mean images of the other sessions using FSL FLIRT (Jenkinson, Bannister, Brady, & Smith, 2002; Jenkinson & Smith, 2001) with 12 degrees of freedom (dof) and, since we noticed residual distortion differences, a subsequent non-linear registration using FSL FNIRT (Andersson, Jenkinson, & Smith, 2007). The resulting linear and non-linear transformations were combined into a single transform, which was applied to all sub-millimeter images recorded outside the first scanning session. We visually inspected the quality of the resulting alignment between the first run in the first session and all runs in the other sessions. Only then EPI distortion correction was applied to the transformed images, using the off-resonance field estimated in session 1. We applied a mask to all images, excluding voxels that had a mean intensity value below 12 in at least one of the sessions. This excluded voxels at the fringes of the slabs (due to minimal differences in slab placement across the two sessions) and we verified that voxels in our region of interest (ROI) were not affected.

ROI DEFINITIONS In order to define our ROI (bilateral hMT+), we calculated a voxel-wise GLM on a single-subject level for the hMT+ localizer data. The GLM was corrected for temporal auto-correlation (AR2). The model contained a separate predictor for the three stimulus conditions (moving dots in the left, central and right aperture). We selected voxels that showed a significant response to the central condition (using a threshold (q) corrected for multiple comparisons using false discovery rate; $q(\text{FDR}) < .05$; S1: $t(474) > 3.00$; S3: $t(474) > 2.80$; S7: $t(474) > 2.68$; S8: $t(474) > 2.81$; S9: $t(474) > 2.75$). For each hemisphere, we projected this selection of voxels onto the inflated surface reconstructed along the middle of GM (for every vertex, the maximum statistical value in a range from -2 mm to +2 mm from the middle of GM was displayed). This allowed us to delineate area hMT+ on each hemisphere by manual drawing. Although drawing introduced a component of subjective judgment, the degree of subjectivity was minimal given that the areas were clearly outlined by significant responses (for an example, see Figure S3.4). Significant responses needed to be located at the posterior/dorsal limb of the inferior temporal sulcus to be included in the ROI, in keeping with previous empirical work reporting the location of area hMT+ (Dumoulin, 2000; Huk et al., 2002; Kolster et al.,

2010). Figure S3.5 shows the resulting definitions on the surface for left and right hemispheres. The surface patches were projected back into volume space (from -2 mm to +2 mm from the middle of GM) and restricted to GM. Delineations on the left and right hemisphere were grouped together to constitute a single ROI per subject. Table S3.2 lists the number of voxels included in the ROIs as well as the estimated areas of left and right hMT+ on the surface. In summary, three criteria thus determined the selection of voxels for our ROI: (i) significant responses to central condition of independent hMT+ localizer data, (ii) anatomical position at the posterior/dorsal limb of ITS, (iii) voxels needed to be in GM.

3.3.7 Functional data - statistical analyses

We used the physical motion quartet data to identify clusters within our ROI that preferred either horizontal or vertical motion. We computed a GLM with predictors for horizontal and vertical motion and calculated t-values for the contrast horizontal >vertical. We observed voxels with a significant ($q(\text{FDR}) < .05$) preference for either horizontal or vertical motion (S1: $t(1492) > 2.80$; S3: $t(1492) > 3.01$; S7: $t(1492) > 2.69$; S8: $t(1492) > 2.94$; S9: $t(1792) > 2.71$). Based on these responses, we formed horizontal and vertical motion clusters for every subject (Figure S3.6).

To form single-subject event-related averages for the ambiguous motion experiment (Experiment 1), we converted every voxel's response to percent signal change where the mean of the respective run time series served as baseline. Separately for each cluster, signals were first averaged across voxels and then grouped by time point during horizontal and vertical perceptual periods. This resulted in event-related averages from time point 0 s (subjects indicated a perceptual switch) until 10 s later (average length of a perceptual period), in steps of 2 s (our TR). In total, there were four event-related averages per subject (2 clusters \times 2 perceptual periods). Event-related averages for the physical motion experiment (Experiment 2) were created in a similar fashion, with the following modification. To avoid circularity, we used a leave-one-run-out cross-validation scheme where all runs but one were used to assign voxels to either the horizontal or vertical cluster and responses from the left-out run were averaged across the cross-validation folds.

To obtain event-related averages across all subjects and to take differences in number of perceptual periods per subject into account, we calculated a weighted mean per time point (Cohen, 1998), according to

$$\bar{x}_{tp} = \frac{\sum_{i=1}^n \frac{1}{\sigma_i^2} x_i}{\sum_{i=1}^n \frac{1}{\sigma_i^2}} \quad (3.1)$$

where n is the number of subjects, x_i is the mean per subject per time point, and σ_i is the standard deviation across perceptual periods per subject per time point. Uncertainty in \bar{x}_{tp} for display of the error bars was calculated using error propagation (Cohen, 1998) as

$$\sigma_{tp} = \frac{1}{\sqrt{\sum_{i=1}^n \frac{1}{\sigma_i^2}}} \quad (3.2)$$

To test for statistical significance of the amplitude modulations observed during the ambiguous motion experiment level, we computed a GLM on a single-subject level containing predictors for horizontal and vertical motion and calculated t-values for the contrast horizontal >vertical in horizontal clusters and for the contrast vertical >horizontal in the vertical cluster (empirical t-values). We then obtained a null distribution of t-values by randomly permuting condition labels and rerunning the GLM analysis (1000-fold permutation testing). If the empirical t-value was above the 97.5th percentile of the null distribution, the modulation for a cluster was declared significant. This amounts to a two-sided hypothesis test at .05. To test how robust the observed effect was to varying the number of voxels in the clusters, we systematically varied number of included voxels (100, 200, 300, 400, 500, 1000) and redid the analyses described above.

In order to quantify the event-related signal increases and decreases observed during ambiguous and physical motion (Figure 3.2), we extracted the 12s time periods (6 TRs) after subjects indicated a perceptual switch from one motion axis to the other. For ambiguous motion, the average signal across all voxels in either the horizontal or vertical cluster was

extracted. Assignment of voxels to clusters was based on responses to the physical motion stimulus. In order to extract signal for physical motion without introducing circularity, we employed a leave-one-run-out cross-validation scheme. We used all but one run to assign voxels to either the horizontal or vertical cluster and extracted perceptual periods from the left-out run.

We expected average event-related signal to increase for horizontal/vertical motion percepts in the horizontal/vertical cluster. Reversely, we expected signal to decrease for vertical/horizontal motion percepts in the horizontal/vertical clusters. When we expected an increase, we calculated the difference between the maximum of the last two time points of the extracted time period (volume 5 or 6) and the minimum of the first two time points (volume 1 or 2). When we expected a decrease, we calculated the difference between the minimum of the last two time points of the extracted time period and the maximum of the first two time points. This amounts to $n = 4$ multiple comparisons. We chose the first two time points based on our expectation that signal in these volumes reflected the percept before the perceptual switch (due to the hemodynamic delay of the fMRI signal). We chose the last time points because we expected the signal to reflect the new percept after the switch.

We expected (a subset of) voxels to show consistent preferences across Experiments 1 and 2 for either horizontal or vertical motion. To test this, we selected voxels whose time courses were significantly modulated by moving dots presented in the central aperture ($q(\text{FDR}) < .05$) during the hMT+ localizer experiment. For every selected voxel we ran a GLM containing predictors for horizontal and vertical motion and calculated the t-values for the contrast horizontal > vertical, separately for Experiment 1 and 2. Voxels showing a preference for horizontal motion thus had positive t-values, while voxels preferring vertical motion had negative values. We calculated the correlation (Pearson's r) between t-values for Experiment 1 and 2, treating every voxel as a data point, and tested for statistical significance at an alpha level of .05. We estimated variability of correlation coefficients by bootstrapping the population of voxels included in the calculation (20,000 re-samples). From the resulting distribution of 20,000 coefficients, we calculated the 2.5th and 97.5th percentile for the display of error bars. To test how robust correlation was to changing the

number of voxels included in the analysis, we systematically varied the number of voxels and redid the analyses. The number of voxels we varied over was: 100, 200, 300, 400, 500, 1000.

To visualize and quantify the degree of columnarity, we used Brain-Voyager's grid sampling approach. This approach takes as inputs a cortical thickness map and a ROI definition in order to create separate regular 2D grids at specified relative cortical depth levels (Kemper et al., 2018; Zimmermann et al., 2011). Resulting points in a grid column fall on corresponding points in the cortical depth direction, taking the varying curvature and volume of cortex into account (equi-volume approach; Waehnert et al., 2014). We used the grid points to sample t-values calculated for the contrast horizontal >vertical for the physical motion experiment in our hMT+ ROI. We then calculated the correlation (Pearson's r) between t-values for corresponding deep (0.9) and superficial (0.1) relative depth levels, treating every pair of corresponding points as a data point (Nasr, Polimeni, & Tootell, 2016). We restricted the analysis to profiles where cortical thickness exceeded 1.6 mm to prevent that the same voxel contributed to both deep and superficial grid points. Mean cortical thickness values were between 2.56 and 2.97 mm (Table S3.3). As a control, we calculated the correlation between t-values for each grid point at the deep depth level and a randomly chosen grid point at the same (deep) depth level that was located, on average, 2.0 mm apart from the target grid point (Nasr et al., 2016). We compared the two resulting correlation coefficients for each subject separately using Fisher's method for comparing correlation coefficients, adjusted for correlated coefficients (Meng, Rosenthal, & Rubin, 1992), and evaluated statistical significance at a level of .05.

As an additional test for columnarity, we compared spatial autocorrelation of physical motion preferences in the cortical depth and in the cortical plane direction. If motion preferences were organized in a columnar fashion, we would expect that motion preferences display a higher autocorrelation coefficient in the cortical depth than in the cortical plane

direction. To quantify the degree of spatial autocorrelation, we calculated Moran's I (Moran, 1950; Rosenberg, 2010) as:

$$I = \frac{n \sum_{i \neq j} w_{ij} z_i z_j}{\sum_{i \neq j} w_{ij} \sum_i z_i^2} \quad (3.3)$$

where n is the total number of grid points; $\sum_{i \neq j}$ is the double summation of all grid points i from 1 to n and j from 1 to n where $i \neq j$; $z_i = Y_i - \bar{Y}$, where Y_i is the value of motion preference at grid point i and \bar{Y} is the mean of Y for all grid points; finally, w_{ij} is the weight describing the relationship between grid point i and grid point j and was based on the distance between the two grid points. In particular, we constructed several distance bins. Bins ranged from 1 mm to 4 mm cortical distance with a bin width of 0.6 mm. Below 1 mm and above 4 mm we had insufficient data points given that cortical depth levels were separated by at least 0.4 x relative cortical thickness (lower bound) and at most 0.8 x cortical thickness (upper bound). We then iterated over the distance bins and assigned a weight of 1 if the distance of a pair of grid points fell within the bin and a zero otherwise and calculated the autocorrelation coefficients for each bin separately.

To determine whether there was a difference in autocorrelation between the cortical depth and the cortical plane direction, we calculated directional autocorrelation (Oden & Sokal, 1986; Rosenberg, 2010). This was achieved by updating the weights w_{ij} as follows: To determine autocorrelation in the cortical plane direction, we only considered relationships between grid points that were located at the same depth level (i.e. the same grid surface) by setting all weights to zero for relationships between grid points located on different depth levels. Conversely, in order to calculate autocorrelation in the cortical depth direction, we only considered relationships between grid points that were in the same grid column and set all weights to zero for relationships between grid points that were not located in the same grid column.

3.3.8 *Functional data - statistical analysis of control experiments*

CONTROL EXPERIMENT I In order to obtain axis-of-motion tuning curves, we followed the same procedure as described in Emmerling, 2016. Briefly, we fit a GLM with predictors for each of the four motion axes and calculated four contrasts testing for modulation in response to each of the motion axes, which resulted in four t-values per voxel. We demeaned the t-values and averaged responses per motion axis and cluster. To test for statistical significance of the obtained tuning curve, we proceeded in two steps. First, we converted the four responses in every cluster to ranks (from high to low) and compared these ranks to idealized ranks for either a horizontally or vertically tuned cluster. The idealized rank pattern for a horizontally tuned cluster was coded as [1, 2.5, 3, 2.5], expressing the expectation that the highest rank should be obtained for the horizontal motion axis, the lowest rank for the vertical motion axis, and tied ranks for diagonal motion axes. Reversely, the idealized rank pattern for a vertically tuned cluster was coded as [3, 2.5, 1, 2.5]. We calculated the Kendall rank correlation coefficient (Tau-b statistic) between observed and idealized ranks (empirical correlation). Second, we obtained a null distribution of tuning curves by randomly permuting condition labels and rerunning the analysis described above (1000-fold permutation testing). This resulted in 1000 axis-of-motion tuning curves as well as 1000 correlation coefficients. We compared the empirically observed correlation coefficient to the null distribution of coefficients. If the empirical correlation coefficient was above the 97.5th percentile of the null distribution, the motion tuning was declared significant. Ties in correlation coefficients were handled using the 'weak' method as implemented in the 'percentileofscore' function in scipy (Jones et al., 2007).

CONTROL EXPERIMENT II The data obtained in control experiment II allowed us to conduct two types of analysis. First, we again investigated tuning of responses in the horizontal and vertical clusters. For this analysis, we proceeded in a similar fashion as for control experiment I, with the following modifications. We fit a GLM with only two predictors for horizontal and vertical motion conditions and calculated two contrasts, resulting in two t-values per voxel. To test for statistical significance of

the tuning, in each cluster we first formed differences between empirically observed average t-values for horizontal and vertical motion (empirical difference). We then obtained a null distribution of differences by randomly permuting condition labels and re-running the analysis described above (1000-fold permutation testing). If the empirical difference was above the 97.5th percentile of the null distribution, the motion tuning was declared significant.

Second, we investigated whether amplitude modulations during ambiguous motion would still reflect the perceived motion axis even if voxel selection was based on retinotopically identical conditions. To this end, we assigned voxels within the ROI used in the main experiment to either horizontal or vertical clusters, depending on their responses during control experiment II. We formed clusters based on the 100 voxels with the largest absolute t-values for the contrast horizontal >vertical. For each cluster, we obtained average time courses for the ambiguous motion experiment and calculated t-values for the contrast horizontal >vertical in the horizontal clusters and for the contrast vertical >horizontal in the vertical clusters. These empirical values were compared to a null distribution from a 1000-fold permutation analysis and modulation of a cluster's time course was considered significant if its empirical t-value was above the 97.5th percentile of the null distribution. To see if our results were robust to a varying number of voxels included in the cluster, we systematically varied the number of included voxels (from 25 to 50, 75, 100, 250 or 500).

CONTROL EXPERIMENT III The pRF mapping data was analyzed as described in Dumoulin and Wandell, 2008. Briefly, every voxel time course was assumed to be generated by an isotropic 2D Gaussian in visual space. Model responses were calculated by transforming the visual stimuli presented during the pRF mapping experiment to binary apertures (indicating presence or absence of a stimulus in space) and spatially convolving the apertures with an instance of the 2D Gaussian model. The Gaussian model includes parameters for its position (x- and y-position) and its size (sigma; standard deviation of the 2D Gaussian) in visual space and the predicted model time course will vary with assumed parameters. To account for hemodynamic delay in the fMRI response, we convolved every model

time course with a canonical double-gamma function (Friston et al., 1998). We determined x , y and σ parameters for every voxel by minimizing the sum of squared errors (SSE) between the predicted model and the observed fMRI time course. In particular, we systematically varied parameters and created candidate model time courses on a polar grid: 48 radial position (0° - 12° eccentricity) \times 24 polar positions (0 - 2π) \times 40 pRF sizes (0° - 22°). We then scaled the model time courses to every voxel time course using a GLM and minimized SSE in a brute force search across candidate models. In order to avoid over-fitting, a leave-one-run-out cross-validation procedure was employed: the GLM was fit on $n - 1$ runs and the SSE were calculated between the predicted and the fMRI time series of the left-out run. The winner model for every voxel was determined by minimizing the mean SSE across all cross-validation folds. Exhaustively searching a very dense grid was achieved in a timely manner using multi threading and optimized Cython code, which is publicly available as part of the PyPrf feature package (Schneider, Marquardt, Gulban, & Isik, 2018).

Based on the resulting pRF parameters, we created a representation of the visual field coverage, separately for voxels in the vertical and horizontal cluster. Clusters included voxels from both left and right hemispheres. For every voxel we recreated its estimated 2D Gaussian pRF in visual space (modeled as a 512×512 pixels square). For this analysis, the 2D Gaussians are normalized such that the maximum value of the Gaussian was equal to 1. For every cluster, we summed up the resulting values across all voxels at every pixel and divided by the total number of voxels included in the cluster. This results in coverage values between 0 and 1, where 1 would indicate that all voxels in a cluster had their receptive field center on this particular pixel of the visual field representation. These steps were implemented in the openly available ViFiCov package (Schneider, 2018).

We also calculated for every voxel in the clusters by how much their estimated 2D Gaussian pRF overlapped with either the horizontal or vertical motion trajectories that were shown or perceived in the motion quartet experiments. Motion trajectories were represented as binary images. This resulted in two overlap values per voxel, for overlap with either horizontal or vertical trajectories. Using a paired samples t-test, we tested for every cluster separately for a difference in overlap with the horizontal and the

vertical motion trajectory. We also formed the difference of overlap values for every voxel and compared the mean difference values between horizontal and vertical clusters using an independent samples t-test. For visualization of the visual field coverage as well as the statistical tests, we employed an R2 threshold such that voxels with an R2 value lower than .05 were excluded. We also report the statistical results with all voxels included in the analysis.

3.4 RESULTS

Figure 3.1B (top) shows amplitude modulations in horizontal and vertical clusters during ambiguous motion averaged across participants. Amplitude modulations reflected the consciously perceived motion axis. When subjects indicated a perceptual switch from vertical to horizontal motion, the response increased in horizontal clusters and decreased in the vertical ones. Reversely, when subjects reported a switch from horizontal to vertical motion, signal increased in the vertical clusters and decreased in the horizontal ones.

Figure 3.2A demonstrates that we consistently found the same pattern of fMRI signal for each subject. Switches to the preferred percept of a given cluster led to increases in fMRI signal. Our permutation testing revealed that in 4 out of 5 horizontal clusters (S1: $t=3.86$, $p=.036$; S3: $t=9.90$, $p<.001$; S7: $t=4.21$, $p=.022$; S8: $t=1.63$, $p=.444$; S9: $t=2.73$, $p=.042$; 1000-fold permutation) and in 4 out of 5 vertical clusters (S1: $t=3.59$, $p=.054$; S3: $t=3.15$, $p=.046$; S7: $t=4.20$, $p=.018$; S8: $t=7.09$, $p<.001$; S9: $t=6.09$, $p<.001$; 1000-fold permutation) the amplitude modulations reflected perceived motion axis (also see Figure S3.7). This finding was robust in 7 out of 10 clusters also when we varied the number of voxels included in the horizontal and vertical clusters (Figure S3.8).

To determine to which extent the cluster signals reflected the perceived as opposed to the physical stimulus, we compared signal changes during physical and ambiguous motion. For the physical motion quartet both retinal and perceived motion axis changed, while for the ambiguous motion display only the perceived motion axis changed. Figure 3.2B shows the average fMRI time course for every subject during the physical motion experiment. Comparison of responses during the two experiments

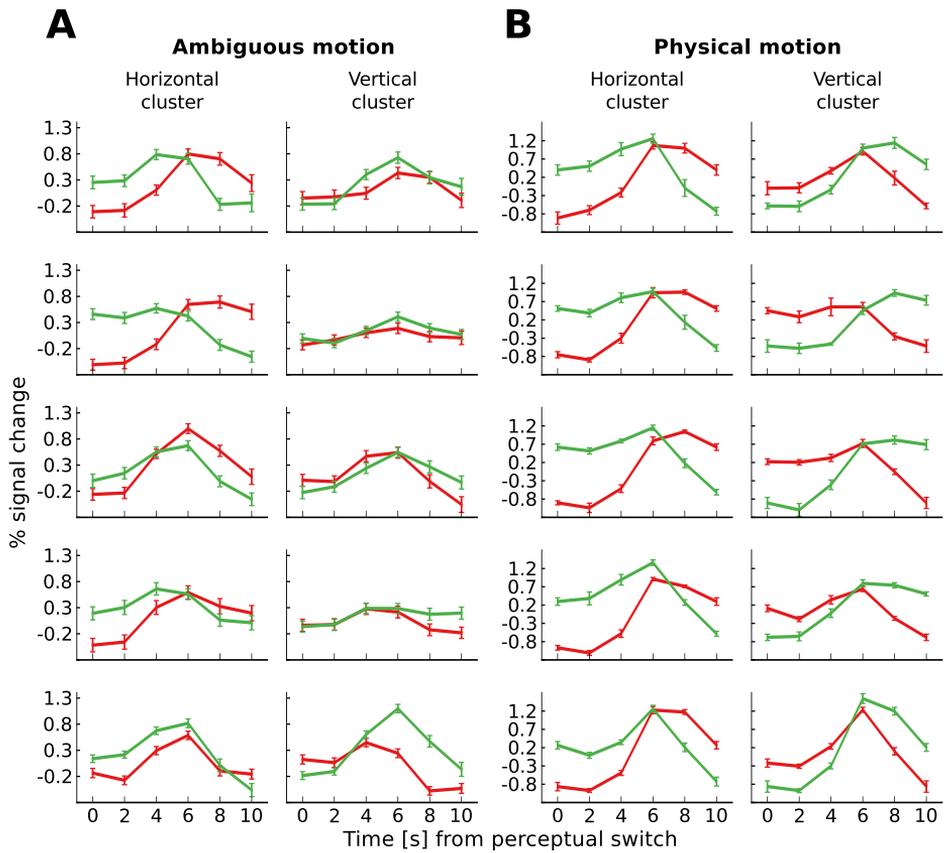


FIGURE 3.2: Single-subject event-related averages for ambiguous and physical motion. **(A)** Lines show event-related sample means across perceptual periods during ambiguous motion. Error bars represent standard deviation of the mean across perceptual periods. **(B)** Lines show event-related sample means across cross-validation folds during physical motion. Error bars represent standard deviation of the mean across cross-validation folds. Both panels show responses for either horizontal clusters (on the respective left) or vertical clusters (on the respective right of the panel), separately for every subject (rows). Time courses (normalized to run average) are displayed during horizontal (red lines) and vertical (green lines) perceptual periods. At time point 0s subjects reported a perceptual switch specifying whether they now perceived horizontal or vertical motion.

revealed a similar qualitative pattern. Quantitatively, the amplitude modulations were larger for physical than for ambiguous motion (Table S3.4).

Upon visual inspection we found voxel preferences between physical and ambiguous motion to be spatially overlapping (Figure 3.3A). Thus, when voxels showed a horizontal preference during physical motion, they often showed the same preference during ambiguous motion. To assess the consistency in preference we computed correlations between preferences during the two experiments and found significant, positive correlations (Figure 3.3B) in all subjects (S1: $r=0.25$, $p<.001$; S3: $r=0.27$, $p<.001$; S7: $r=0.27$, $p<.001$; S8: $r=0.21$, $p<.001$; S9: $r=0.22$, $p<.001$; two-sided test). This consistency often stayed stable even when we varied the number of voxels included in our region of interest over a larger range (Figure S3.9).

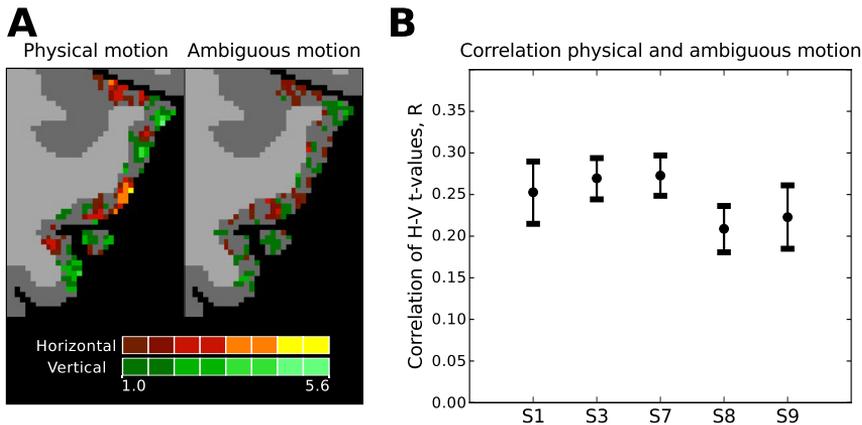


FIGURE 3.3: Consistency of preferences between physical and ambiguous motion. **(A)** Qualitative consistency. Preferences for either horizontal (red) or vertical (green) motion are shown on the transverse slice of a selected subject (S1, left area hMT+) during either physical (left) or ambiguous (right) motion. Perceptually brighter colors indicate a higher degree of preference. Preferences are overlaid on segmentation labels of white (white-gray voxels) and gray matter (dark-gray voxels). **(B)** Quantitative consistency. Displayed are the median bootstrapped correlation coefficients for t-values of the contrast horizontal >vertical found for physical and ambiguous motion. Voxels were selected based on independent data from the hMT+ localizer. Every point represents the result of a single subject (S1, S3, S7, S8, S9). Error bars represent the 2.5th and 97.5th percentile of the bootstrapped correlation coefficients.

To exclude the possibility that differences in retinotopic preferences are responsible for our observed effects, we conducted several control experiments (Section 3.6.2). First, we showed that the identified clusters display expected axis-of-motion tuning on two independently acquired data sets (Figure S3.10A,B). Second, we estimated the visual field coverage for both horizontal and vertical clusters and did not find a difference in their coverage of the horizontal and vertical motion trajectory (S3: $t(205)=0.066$, $p=.948$; S7: $t(729)=1.26$, $p=.209$) (Figure S3.11). Third, we selected voxels for horizontal and vertical clusters based on horizontal and vertical motion conditions that were retinotopically matched (Figure S3.10C) and obtained responses in these clusters during ambiguous motion. In 3 out of 3 horizontal clusters (S1: $t=4.22$, $p=.002$; S3: $t=9.35$, $p>.001$; S9: $t=5.26$, $p>.001$; 1000-fold permutation) and in 2 out of 3 vertical clusters (S1: $t=3.86$, $p=.048$; S3: $t=-4.44$, $p=.018$; S9: $t=3.44$, $p=.014$; 1000-fold permutation), the amplitude reflected perceived motion axis. This finding was robust to a varying number of voxels included in the clusters (Figure S3.10D).

Finally, we wondered whether preferences for horizontal and vertical motion showed a form of cortical organization akin to columns selective for a particular motion axis. To address this question, we zoomed in to the observed responses (Figure 3.4A) by constructing regular, high-resolution grids that covered the identified hMT+ areas (Figure 3.4B) (Kemper et al., 2018). The grid allowed us to visualize and quantify how motion preferences change in the direction of cortical depth and along the cortical surface. If motion preferences are organized in a columnar fashion, they should stay stable in the direction of cortical depth and change when moving along the cortical surface. Using Meng's (1992) z -test, we found that physical motion preferences at corresponding points along the cortical depth were more correlated to each other than preferences at grid points separated by a similar distance along the cortical plane (Figure 3.4C, S3.12) (Nasr et al., 2016) in all subjects (S1: $z=12.76$, $p<.001$; S3: $z=11.54$, $p<.001$; S7: $z=18.11$, $p<.001$; S8: $z=3.80$, $p<.001$; S9: $z=8.33$, $p<.001$; two-sided test) (Figure 3.4D). A comparison of directional spatial autocorrelation corroborated this finding (see Section 3.6.3 and Figure S3.13). Preferences during ambiguous motion showed smaller and less consistent correlation differences (S1: $z=-3.26$, $p=.001$; S3: $z=6.09$, $p<.001$; S7: $z=3.08$, $p=.002$; S8: $z=7.73$, $p<.001$; S9: $z=-0.21$, $p=.837$; two-sided test).

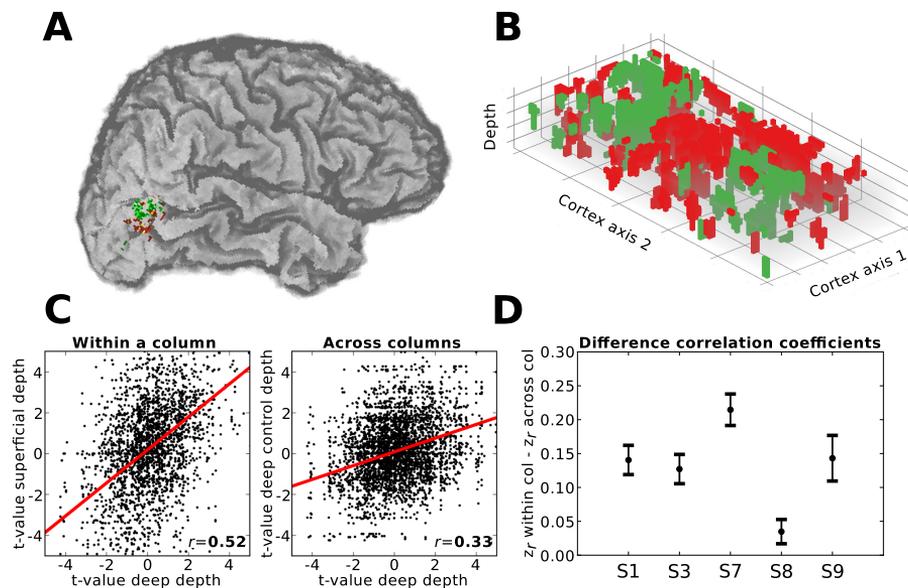


FIGURE 3.4: Horizontal and vertical columnar-like clusters for physical motion. Based on data from the physical motion experiment, voxels were colored in either red (horizontal motion preference) or green (vertical motion preference). **(A)** Functional hMT+ clusters shown on an individual brain. **(B)** Functional responses to physical motion were sampled using a regular 3D grid. Cortex axis 1 and 2 represent the cortical plane; the z-axis represents cortical depth. **(C)** Correlation of axis preference sampled for deep and corresponding superficial depth level (left) or at different nearby locations in deep depth level only. Results are shown for a selected subject and hemisphere (S7, left hemisphere). **(D)** Subject-wise 95 percent confidence intervals for the difference in z-transformed correlation coefficients within and across grid columns. Every interval represents the result of a single subject (S1, S3, S7, S8, S9). "Col" stands for column.

3.5 DISCUSSION

Our results support findings of previous human fMRI studies (Brouwer & van Ee, 2007; Castelo-Branco et al., 2002; Kamitani & Tong, 2006; Muckli et al., 2002; Sterzer et al., 2002) that area hMT+ makes up part of the content-specific NCC for visual motion by demonstrating a link between the experience of motion and response amplitudes in area hMT+. Importantly, our findings extend this idea by demonstrating involvement of specific horizontal and vertical hMT+ clusters in tracking conscious experience of a particular motion axis. Earlier human fMRI studies demonstrated a coupling between categorical contents of consciousness and activity in macroscopic human areas (Andrews, Schluppeck, Homfray, Matthews, & Blakemore, 2002; Hasson, Hendler, Bashat, & Malach, 2001; Tong, Nakayama, Vaughan, & Kanwisher, 1998). In those studies the two competing perceptual states pertained to two different categories (e.g. face vs. building or face vs. object) and experience of these categories was shown to be reflected in the amplitude of different category-specialized areas like FFA or PPA. By contrast, in our study the two perceptual states fell within the same category (motion) and were distinguishable only by subordinate categorical differences (horizontal vs. vertical).

fMRI studies that have compared sub-category perceptual states (Castelo-Branco et al., 2002; Muckli et al., 2002; Sterzer et al., 2002) were limited by spatial resolution and could only distinguish the states by global amplitude differences of the same area. Although studies using multivariate pattern analysis (Brouwer & van Ee, 2007; Kamitani & Tong, 2006) detected small activation differences between two perceptual states, these studies did not unequivocally reveal the underlying spatial organization of such activation patterns (Logothetis, 2008) since the existence of two types of patterns does not imply the existence of two different neural populations (Bartels, Logothetis, & Moutoussis, 2008). Problematically, this complicates interpretation since coarse-scale biases might account for found differences (Wang et al., 2014). In comparison, using sub-millimeter resolution, our study could link sub-category contents of consciousness to dissociative amplitude modulations in distinct populations within the same brain area.

Our study follows recent advances in ultra-high field MRI which now allow for probing functional responses at the level of mesoscopic structures such as columns and layers (De Martino et al., 2018; Kemper et al., 2018; Polimeni et al., 2018). Previous studies have identified columnar-like structures in humans in V1 (Cheng, Waggoner, & Tanaka, 2001; Yacoub & Harel, 2008), V2, V3 and V4 (Nasr et al., 2016; Tootell & Nasr, 2017), V3a (Goncalves et al., 2015) and hMT (Zimmermann et al., 2011). Yet these studies did not use stimuli that allowed for dissociating neural signals pertaining to conscious perception from those related to sensory stimulation. Using a multistable stimulus, our study indicates that response modulations in columnar structures relate to sub-category contents of consciousness.

An alternative explanation of our findings would be that the identified clusters do not reflect a preference for motion axis but for the retinotopic location of the two illusory motion paths. However, our control experiments showed (i) that the selected clusters display characteristic tuning to either the horizontal or vertical motion axis, (ii) that there is no evidence for a difference in population receptive field coverage of horizontal and vertical motion paths and (iii) that even if voxels are selected based on a stimulus with two retinotopically identical conditions only distinguished by motion axis, we still observe modulations with the conscious percept. We believe that, taken together, the results rule out this alternative explanation.

Several studies indicate that not all signals that modulate with a reported conscious state reflect the conscious experience itself. Instead, at least in part, these signals reflect processes that precede or follow the experience, among them monitoring the experimental task, planning and reporting (De Graaf, Hsieh, & Sack, 2012; Frassle, Sommer, Jansen, Naber, & Einhauser, 2014; Koch et al., 2016). Although we cannot exclude the possibility, that some of the response modulations we observed resulted from button presses to report the perceived axis of motion, this interpretation is not in line with findings that activity for reporting and task monitoring is localized to frontal, executive areas, and not to occipital visual cortex (Frassle et al., 2014). Furthermore, button presses were counterbalanced across the two scanning sessions, making it unlikely that they caused the

observed response modulations, which by design needed to be consistent across sessions.

The signal modulations, that we found to reflect the conscious percept, were based on average cluster responses. It is therefore possible that, although as a whole the cluster showed modulation in the expected direction, a subgroup of voxels did not (or only to a lesser degree) modulate its signal with the percept. Indeed, when we compared amplitude modulations for ambiguous and physical motion (Figure 3.2, Table S3.4), we found larger modulations in response to the latter. One explanation for this difference is that the physical motion stimulus gives rise to both physical bottom-up and perceptual lateral and top-down processes, while the ambiguous motion display invokes primarily perceptual lateral and top-down processes. Earlier electro-physiology findings that only about 40% of identified MT neurons modulated their spiking with the dominant motion direction in binocular rivalry (Logothetis, 1989) are in line with this interpretation. One idea inspired by animal physiology (Felleman & Van Essen, 1991; Markov et al., 2014) is that bottom-up and lateral / top-down processes recruit different laminar parts of a column. This idea could account for our observation that preferences during ambiguous motion displayed less consistent columnarity than during physical motion. To fully illuminate this possibility, one would need to distinguish the specific laminar contributions to activity in a column, which is beyond the scope of this chapter and subject to future studies.

The concept of a cortical column has been subject to substantial debate, starting with the first discovery of columns in primary somatosensory cortex of cats (Mountcastle, 1956), and its functional significance has been doubted (see General Discussion). Problematically, the term is used in different ways and as a result the various types of columns that have been identified differ in their defining feature (function, cell constellation, connectivity or myelin content), their extent, and spatial organization (Rakic, 2008). Here, we have used an operational definition of columnarity, where functional preferences were required to be more stable along the vertical than the horizontal extent of cortex. To highlight the difference with microcolumns, which would be beyond the available resolution of current fMRI, or idealized hypercolumns, we have used the term columnar clusters instead of columns throughout.

Based on our findings, we suggest that the hMT+ clusters identified here constitute part of the content-specific neural correlate for experiencing motion axis. Future studies could examine whether similar, clustered organizations exist for sub-category contents of consciousness other than motion axis. The activity in many cortical areas has been shown to be linked to general stimulus categories, including orientations, body parts, houses, faces, words, bigrams and letters. Some of these cortical areas are also known to display columnar organization (Tanaka, 2003). Such research would clarify how the activity of these known functional subunits relates to phenomenal distinctions in conscious content. Furthermore, we pave the way for studies to investigate if columnar clusters represent subordinate dimensions in other high-level phenomena like attention and memory.

3.6 SUPPLEMENTARY MATERIAL

3.6.1 Behavioral results

All participants reported regular perceptual switches between horizontal and vertical percepts during the ambiguous motion experiment. Across all participants, mean perceptual periods ranged from 7.2s to 13.3s. The average length of horizontal and vertical perceptual periods was $10.8\text{s} \pm 1.9\text{s}$ and $10.8\text{s} \pm 2.4\text{s}$ (mean \pm standard deviation across subjects), respectively. Goal of the calibration of the motion quartet stimulus during the training session had been 10s long percepts. Figure S3.14 shows the average length of horizontal and vertical perceptual periods individually for every subject. For subjects S3 and S9 we found a statistical difference between the length of horizontal and vertical periods ($t(249)=-3.78$, $p<.001$ and $t(336)=4.80$, $p<.001$, respectively). For subjects S1, S7 and S8 we found no statistical difference (S1: $t(252)=-0.41$, $p=.685$; S7: $t(223)=-0.07$, $p=.947$; S8: $t(179)=1.14$, $p=.255$). We thus did not observe bias towards one of the two motion axes that was consistent across participants.

3.6.2 Results control experiments

CONTROL EXPERIMENT I Figure S3.10A shows the axis-of-motion tuning curves derived from control experiment I in clusters used in the main experiment for two control participants. Clusters showed expected axis-of-motion tuning. Horizontal clusters showed the largest average t-value for the horizontal motion condition, lower t-values for the diagonal motion conditions and the lowest t-value for the vertical motion condition. For vertical clusters we observed a reversed pattern. The permutation testing showed significant correlation for ranks in the horizontal cluster and an idealized horizontally tuned cluster (S1: $\tau=0.91$, $p<.001$; S3: $\tau=0.91$, $p<.001$). Ranks of the vertical cluster correlated significantly with ranks of an idealized vertically tuned cluster (S1: $\tau=0.91$, $p<.001$; S3: $\tau=0.91$, $p<.001$).

CONTROL EXPERIMENT II Figure S3.10B shows motion tuning derived from control experiment II for three control participants. All clusters showed tuning in the expected direction, i.e., horizontal clusters show

larger average t-values for the horizontal than for the vertical condition and vertical clusters show a reversed pattern. The permutation testing showed significant differences between average t-values for the horizontal and vertical motion condition in both the horizontal (S1: $\Delta t=0.43$, $p<.001$; S3: $\Delta t=0.25$, $p<.001$; S9: $\Delta t=0.27$, $p<.001$) and vertical cluster (S1: $\Delta t=-0.65$, $p<.001$; S3: $\Delta t=-1.21$, $p<.001$; S9: $\Delta t=-0.68$, $p<.001$).

CONTROL EXPERIMENT III Upon visual inspection, we did not find a difference in visual field coverage between horizontal and vertical clusters (Figure S3.11A). All clusters showed coverage of the area foveal to the inducer squares. We found that clusters had more coverage of the vertical than the horizontal trajectory (Figure S3.11B), which was significant in all clusters (S3 horizontal cluster: $t(165)=-4.21$, $p<.001$; S3 vertical cluster: $t(40)=-6.35$, $p<.001$; S7 horizontal cluster: $t(522)=-22.15$, $p<.001$; S7 vertical cluster: $t(207)=-9.40$, $p<.001$). This is partially to be expected since the retinotopic trajectory is larger for the vertical than for the horizontal condition (by a factor of 1.26 for both control participants). We did not find a difference between horizontal and vertical clusters in the amount they overlapped with either the horizontal or vertical motion trajectory (S3: $t(205)=0.066$, $p=.948$; S7: $t(729)=1.26$, $p=.209$). We obtained similar findings when no R2 thresholding was performed and all voxels were included in the analysis (S3: $t(549)=-1.56$, $p=.119$; S7: $t(1213)=1.40$, $p=.163$).

3.6.3 Directional spatial autocorrelation

Figure S3.13 plots spatial autocorrelation of preferences for physical motion both in the cortical depth and the cortical plane direction over geometric distance. Spatial autocorrelation is consistently higher in the cortical depth than in the cortical plane direction, across participants and across a wide range of cortical distances (with the exception of small distances in S9). At low cortical distances, the autocorrelation is high both in the cortical depth and plane direction. Autocorrelation decreases with cortical distance for both directions; however, in several participants (S1, S3, S9) remains high in the cortical depth direction even at large cortical distances. Taken together, these observations indicate that motion preferences were more similar in the cortical depth than the cortical plane direction and often remained similar over large cortical distances in the cortical depth direction.

3.6.4 Supplementary figures

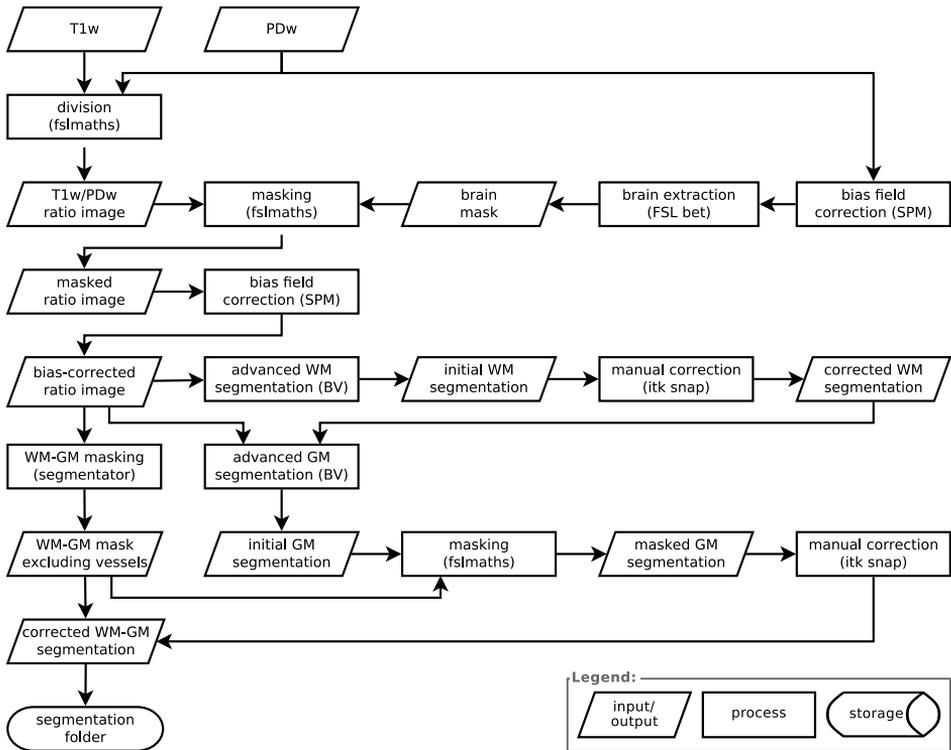


FIGURE S3.1: Overview of pre-processing pipeline for structural images. Rectangular shapes represent processing steps, rhombic shapes represent input or outputs and cylindrical shapes represent input or output locations.

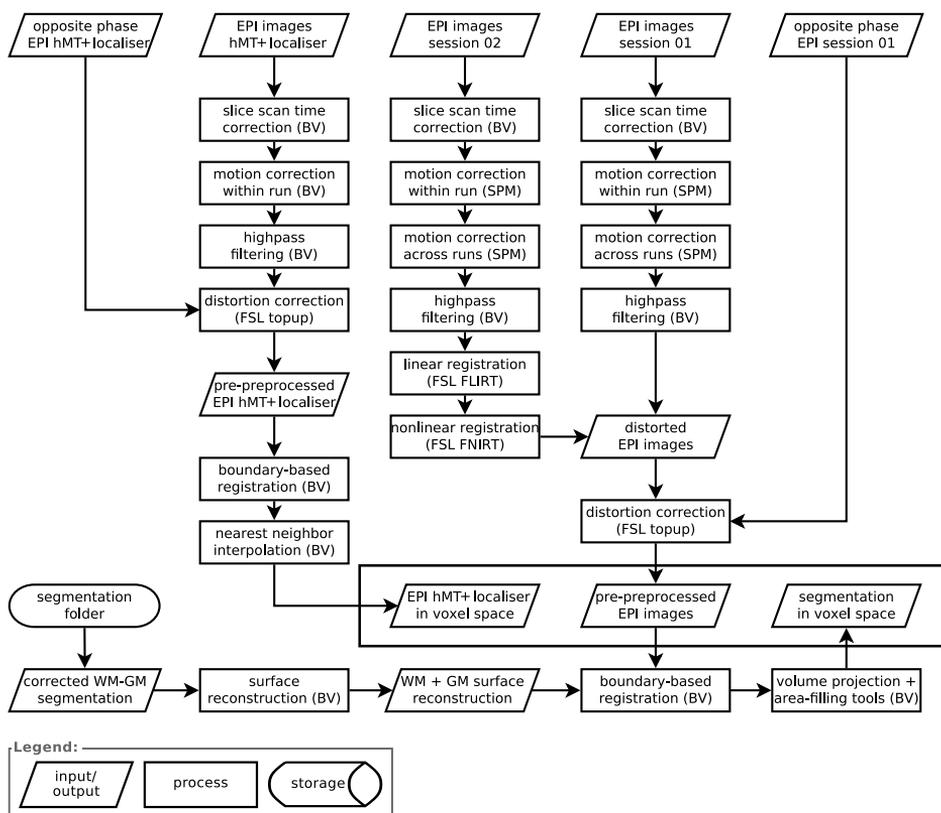


FIGURE S3.2: Overview of pre-processing pipeline for functional images. Rectangular shapes represent processing steps, rhombic shapes represent input or outputs and cylindrical shapes represent input or output locations.

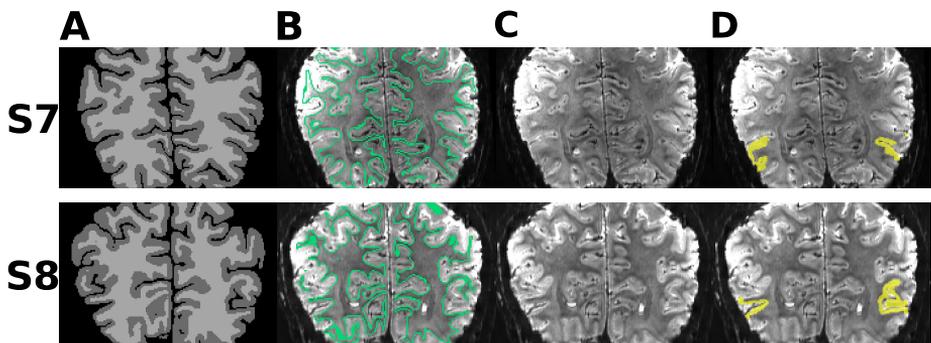


FIGURE S3.3: Segmentation and co-registration quality. All images show coronal slices in radiological convention. **(A)** Segmentation image showing WM (light gray) and GM (dark gray). Although the segmentation has been down-sampled to the slightly lower resolution of the functional space (0.6 mm isotropic to 0.8 mm isotropic), desirable segmentation features are preserved due to our surface reconstruction pipeline. **(B)** Overlay of the WM-GM boundary (mint) on the mean functional image in voxel space. High congruence between projected WM-GM boundary (mint) and inherent contrast of functional image (lower intensity values in WM, higher intensities in GM) indicate little segmentation and co-registration error. **(C)** Mean functional image without segmentation overlay. **(D)** Mean functional image overlaid with ROI (yellow).

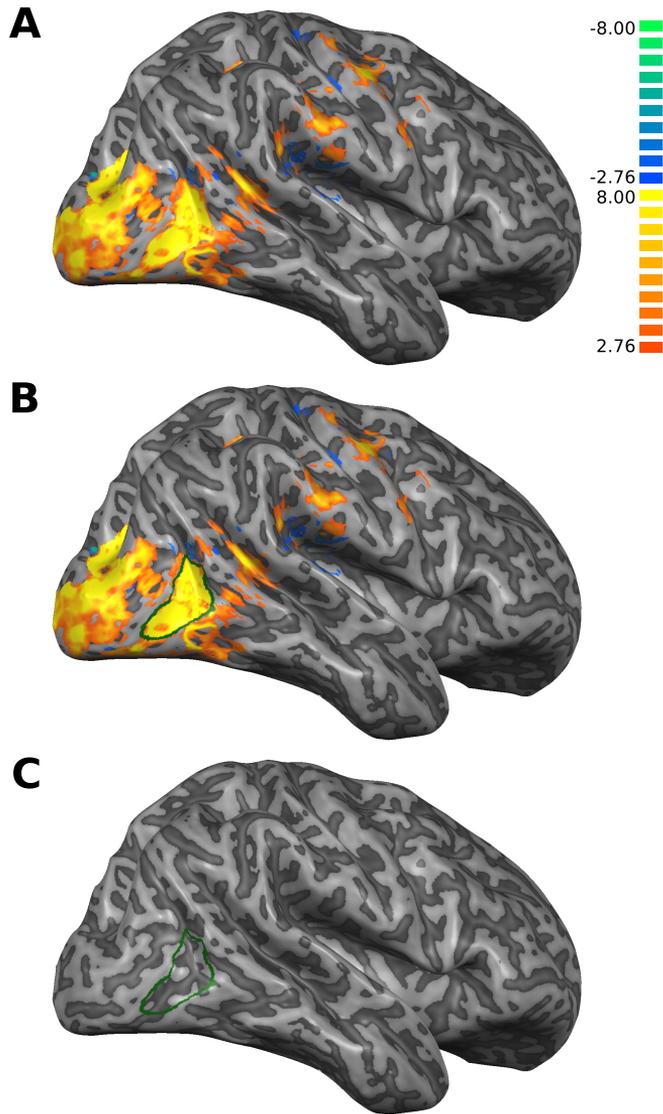


FIGURE S3.4 (previous page): Example of delineation process for area hMT+. Displayed is the smoothed and inflated right hemisphere of one representative subject (S3). Dark-gray color indicates sulci and light-gray color indicates gyri. **(A)** Statistical map of responses (t-values) to moving stimuli presented in central apertures (thresholded at $q(\text{FDR}) < .05$). **(B)** Delineation of area hMT+ (green) in the right hemisphere overlaid on top of the statistical map. Selection of area by manual drawing was clearly constrained by the activation map, leaving little room for subjective judgment. **(C)** Delineation of area hMT+ (green) overlaid without statistical map.

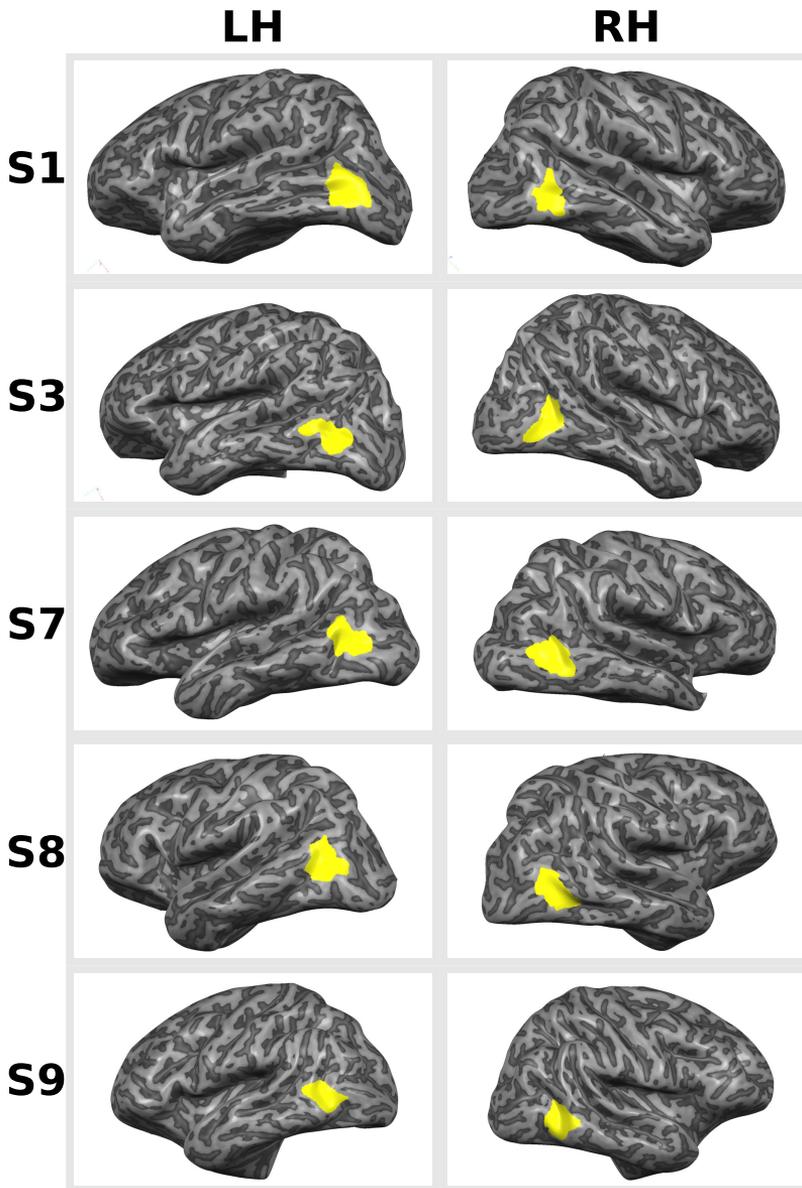


FIGURE S3.5 (previous page): Delineations of area hMT+ for all subjects. Displayed are the smoothed and inflated left (left side) and right (right side) hemispheres for all subjects (different rows: S1, S3, S7, S8, S9). Dark-gray color indicates sulci and light-gray color indicates gyri. The delineations of areas hMT+ are overlaid in yellow.

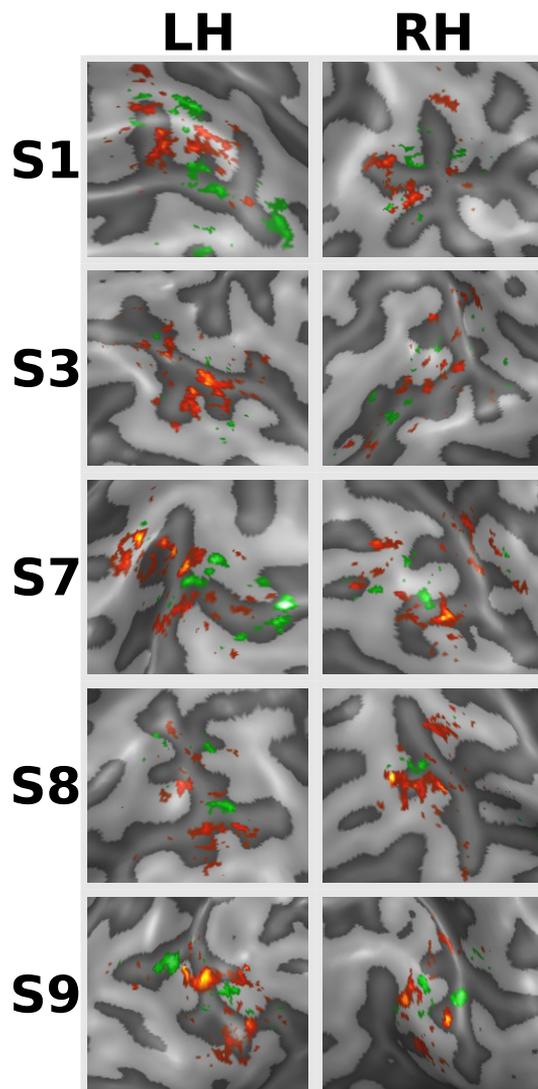


FIGURE S3.6 (previous page): Horizontal and vertical clusters for physical motion. Horizontal (red) and vertical (green) clusters determined by responses during the physical motion experiment are displayed on the smoothed and inflated left (left side) and right (right side) mid-GM reconstructions of area hMT+ for all subjects (different rows: S1, S3, S7, S8, S9). Dark-gray indicates sulci and light-gray indicates gyri.

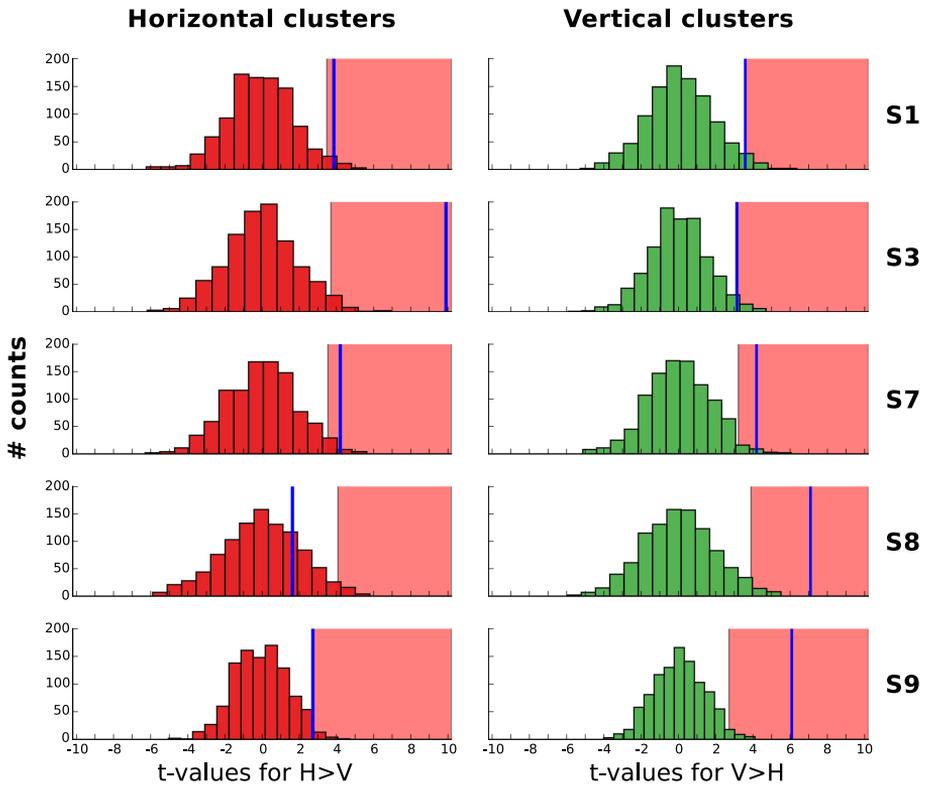


FIGURE S3.7: Results for statistical significance testing of the amplitude modulations during ambiguous motion. The histograms represent a null distribution of t-values obtained by randomly permuting condition labels (1000-fold permutation testing) for either the horizontal (left side, red) or vertical (right side, green) clusters. Blue lines indicate the empirical t-value. Red shading indicates areas above the 97.5th percentile of the null population obtained with permutation testing (rejection area). "H" stands for horizontal, "V" for vertical. Each row represents the results from a different subject.

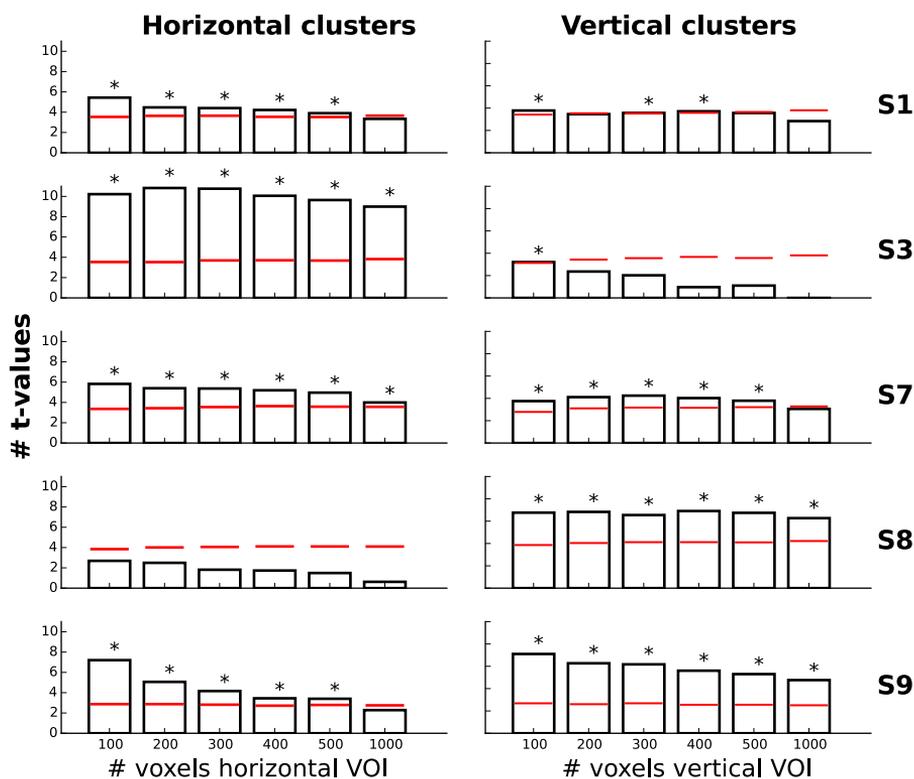


FIGURE S3.8: Observed effects of amplitude modulation during ambiguous motion are robust to changing numbers of voxels included in the ROI. Black bars indicate empirical t-values either for the horizontal (left side) or vertical (right side) cluster. Different bars indicate the results for a particular number of voxels included in the ROI (100, 200, 300, 400, 500, 1000). Red values indicate the 97.5th percentile of a null population obtained with 1000-fold permutation testing. Stars indicate that empirical values fell above the 97.5th percentile of the null distribution. Each row represents the results from a different subject.

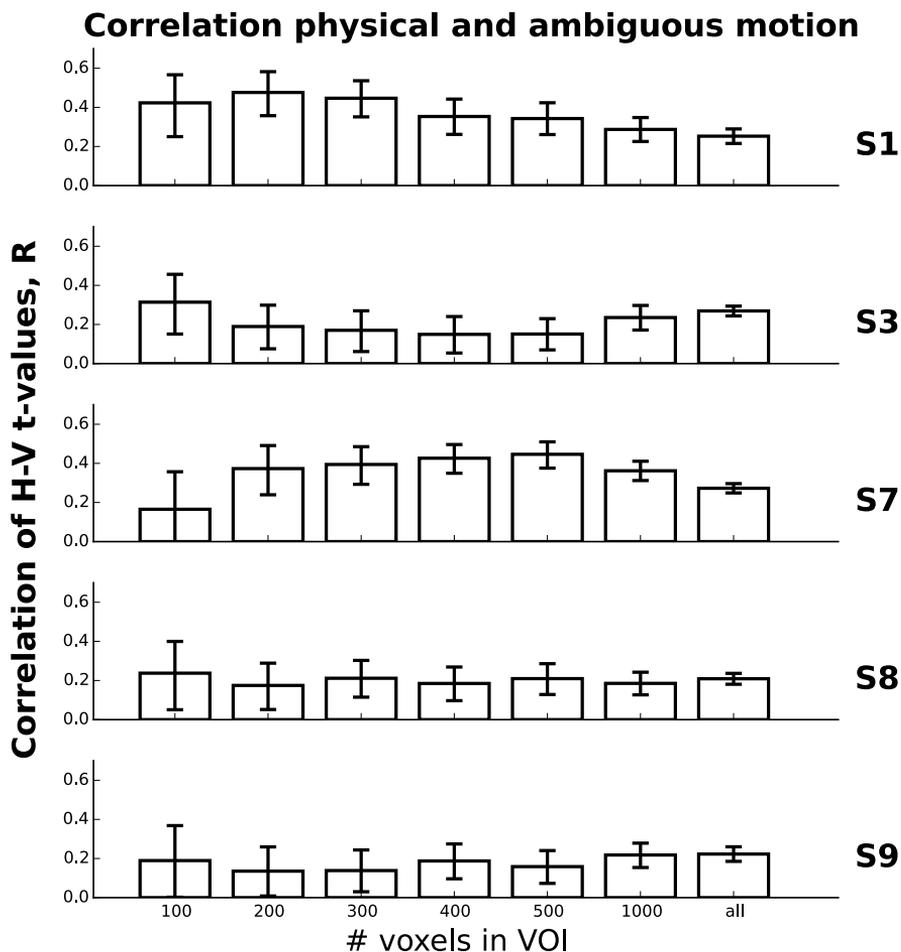


FIGURE S3.9: Consistency in voxel preference was robust to a varying number of voxels included in our ROI. Black bars indicate the median bootstrapped correlation coefficient between t-values for physical and ambiguous motion for a bootstrapped population of voxel values (20,000 re-samples). Error bars represent the 2.5th and 97.5th percentile of the bootstrapped correlation coefficients. Different bars indicate the results for a particular number of voxels included in the ROI (100, 200, 300, 400, 500, 1000). Each row represents the results from a different subject.

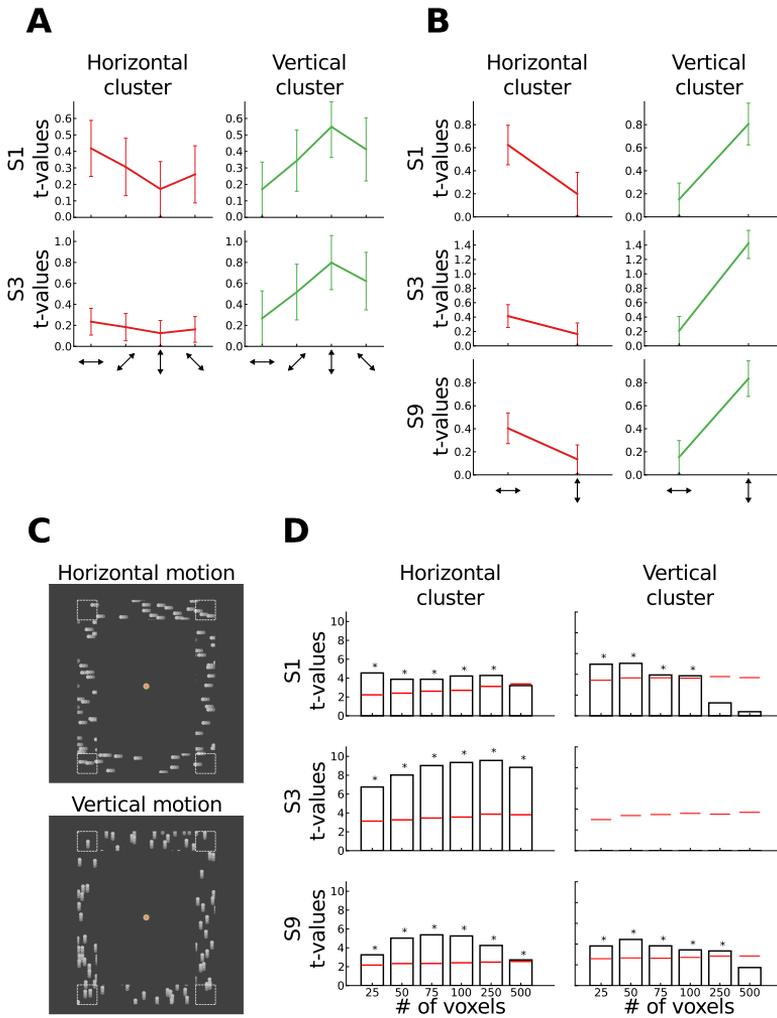


FIGURE S3.10 (previous page): Clusters show expected motion tuning. **(A)** Experimental results for control experiment I. Plots show axis-of-motion tuning curves for horizontal (left) and vertical (right) clusters, as used in the main experiment. Lines depict average t-values in response to four presented axes of motion. Error bars represent standard error of the mean. To facilitate visual comparison of tuning curves between clusters, all t-values were normalized such that the lowest response (including its standard error) in a given cluster was equal to zero. This did not change differences between t-values and this step was not performed for data entering statistical analyses. The clusters for the two control participants (S1 and S3, different rows) show expected motion tuning with highest responses to the preferred axis of motion and gradually lower responses to non-preferred axes. **(B)** Experimental results for control experiment II. Same conventions as for panel **(A)**, just that in this experiment only horizontal and vertical motion axes were presented and three subjects were recorded (S1, S3 and S9, different rows). Although the underlying data were obtained in independent experiments, the same tuning as in **(A)** was observed, with highest t-values to the preferred axis of motion. **(C)** Time-lapsed image showing sum of three frames for stimuli used in control experiment II. Retinotopically identical dot fields were moving either horizontally (upper row) or vertically (lower row). Visibility of dots was restricted to the aperture defined by the inducer squares and motion trajectories in the motion quartet experiments. White dashes indicate positions of inducer squares in the motion quartet experiments and were not shown during the actual experiment. **(D)** Amplitude modulations during ambiguous motion when data from control experiment II were used to assign voxels to either the horizontal (left) or vertical (right) cluster. Black bars indicate empirical t-values. Different bars indicate results for different numbers of voxels included in the ROI (25, 50, 75, 100, 250, 500). Red bars indicate the 97.5th percentile of a null population obtained with 1000-fold permutation testing. Stars indicate empirical values above the 97.5th percentile of the null distribution. With the exception of one cluster, we can replicate the modulations observed in the main experiment, even though assignment of voxels was based on retinotopically identical conditions.

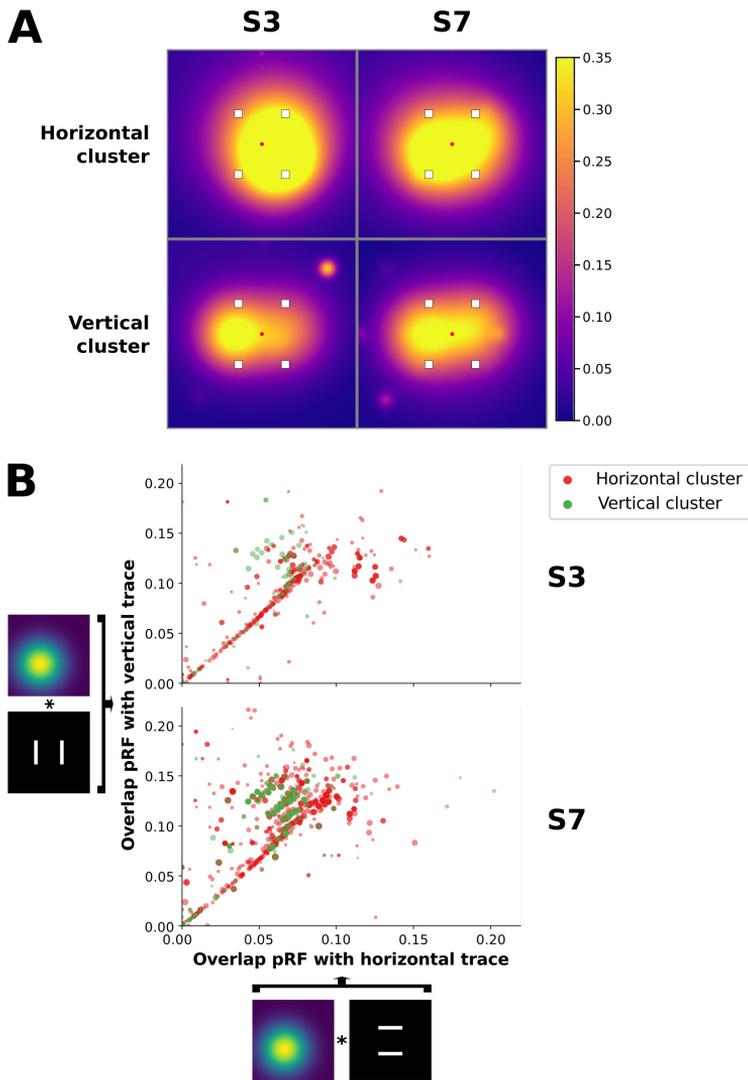


FIGURE S3.11 (previous page): Clusters show no difference in retinotopic preference. **(A)** Visual field coverage for either the horizontal (upper row) or vertical (lower row) cluster, as used in the main experiment, for subject S3 (left) and S7 (right). Coverage is computed as the average value across all voxels' population receptive fields (pRF) in a given cluster at every pixel of the visual field (theoretical maximum equal to 1). The subject-specific configuration of the motion quartet is overlaid on top to indicate where inducer squares were presented. Generally, clusters show coverage of the area between inducer squares and foveal to inducer squares. Both horizontal and vertical clusters show a retinotopic bias to vertical motion trajectories but retinotopic bias did not differ between horizontal and vertical clusters. **(B)** Calculated overlap between the pRF of every voxel in the horizontal (red) and vertical (green) cluster for the horizontal (x-axis) and vertical (y-axis) motion trajectories. Every dot represents a single voxel and dot size is scaled with the fitting accuracy of the pRF model. Results are shown for subject S3 (upper panel) and S7 (lower panel). Many voxels are close to or even on the diagonal, indicating similar or equal overlap with both motion trajectories. The bias towards vertical trajectories observed in **(A)** is apparent as a cluster of voxels above the diagonal but horizontal and vertical clusters show no difference in this bias.

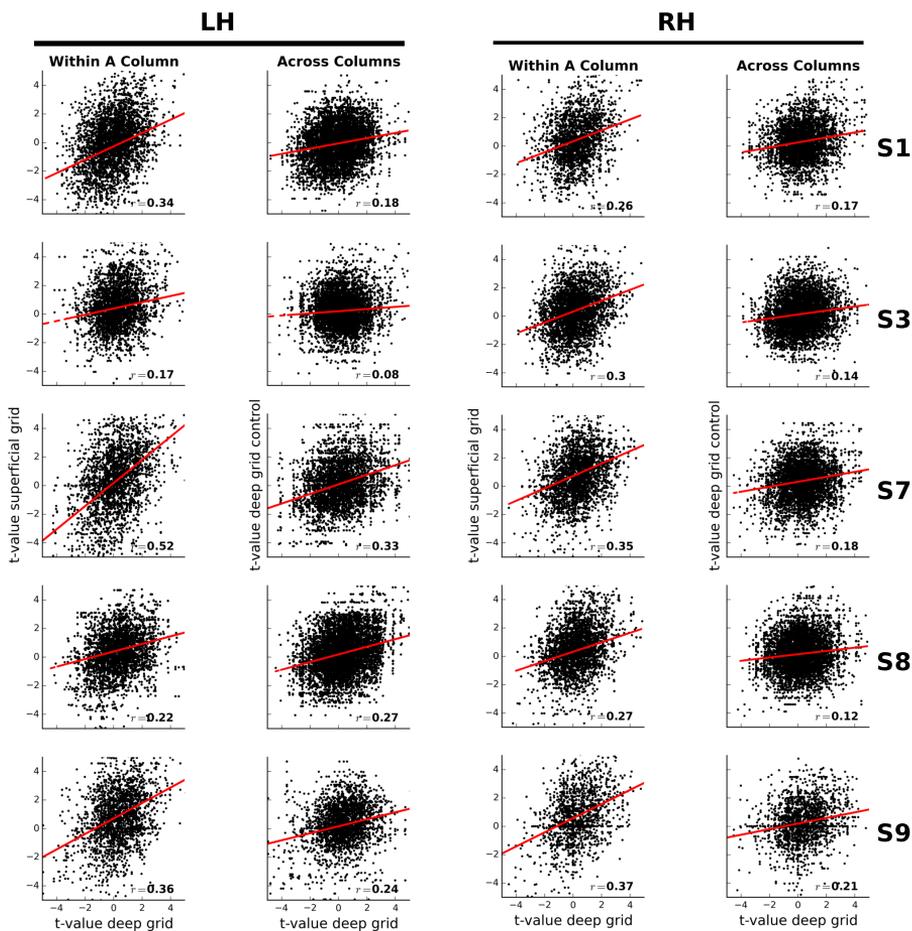


FIGURE S3.12: Motion preferences were more stable in the direction of cortical depth than along the cortical surface. Results are shown for all subjects (rows) and two hemispheres (left hemisphere on the left). Scatter plots show the correlation of axis preference sampled for deep and corresponding superficial depth level ("within a column", respective left side) compared to the correlation for different nearby locations in deep depth level only ("across columns", respective right side).

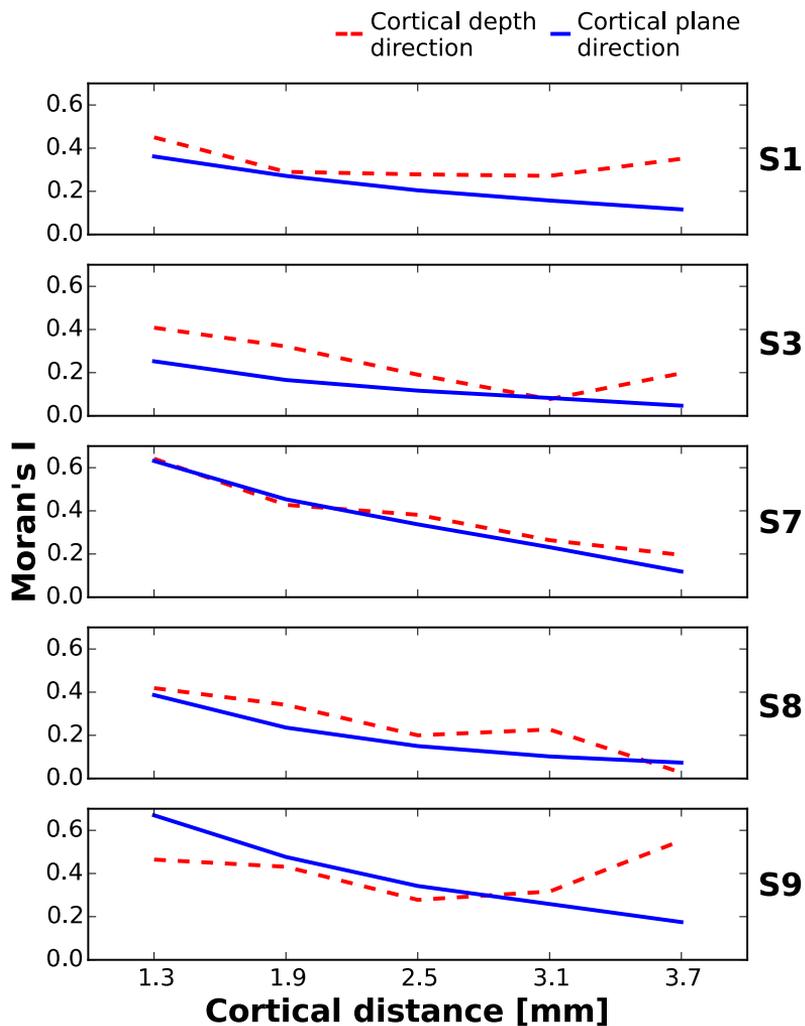


FIGURE S3.13: Greater autocorrelation in cortical depth than in cortical plane direction. Moran's I, a measure of spatial autocorrelation, is plotted over cortical distance (in mm). The red, dashed line shows autocorrelation coefficients in cortical depth (columnar) direction; the blue, solid line shows coefficients in cortical plane direction. Each row shows the results from a different subject.

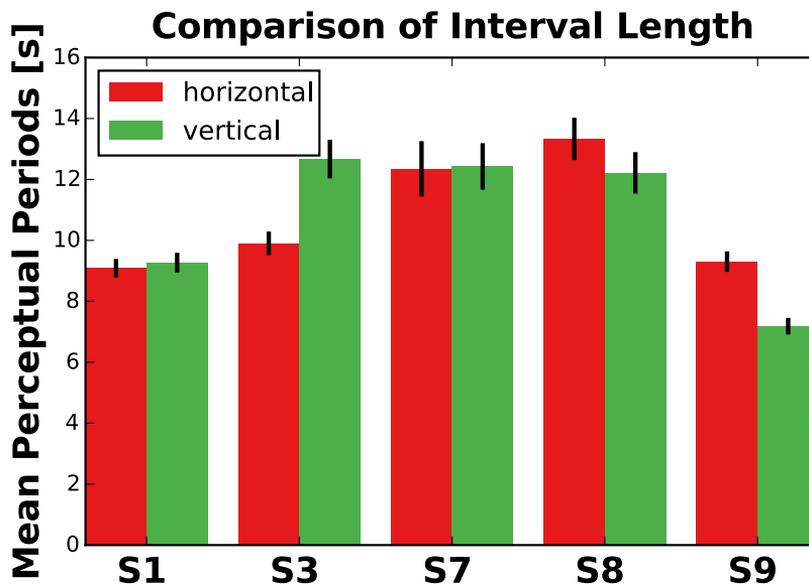


FIGURE S3.14: Average length of perceptual periods. Mean length of perceptual periods in seconds (s) during the ambiguous motion experiment are shown for every analyzed subject (S1, S3, S7, S8, S9), separately for horizontal (red) and vertical (green) perceptual periods. Error bars represent standard error across perceptual periods.

3.6.5 Supplementary tables

TABLE S3.1: Subject-specific vertical distances for the motion quartet display. Distances are in degrees of visual angle from fixation point to center of the square in the motion quartet. Horizontal distance was kept constant at 3 degrees of visual angle.

Subject	Vertical distance
S1	3.8
S3	3.9
S7	3.9
S8	3.9
S9	3.8

TABLE S3.2: Size of the hMT+ region of interest for every subject. LH and RH indicate left and right hemisphere, respectively. Surface area is in mm^2 .

Subject	Number of voxels	Surface area LH	Surface area RH
S1	8341	855	517
S3	10267	781	783
S7	8394	710	857
S8	7923	720	632
S9	7361	626	720

TABLE S3.3: Average cortical thickness of hMT+ region of interest for every subject. LH and RH indicate left and right hemisphere, respectively. 2H means that both hemispheres were included. Std indicates standard deviation.

Subject	<i>Mean ± std</i> LH	<i>Mean ± std</i> RH	<i>Mean ± std</i> 2H
S1	2.74 ± 0.55	2.73 ± 0.66	2.74 ± 0.59
S3	2.72 ± 0.56	2.84 ± 0.65	2.77 ± 0.61
S7	2.78 ± 0.56	2.82 ± 0.59	2.80 ± 0.57
S8	2.58 ± 0.53	2.53 ± 0.52	2.56 ± 0.53
S9	3.35 ± 1.18	2.47 ± 0.46	2.97 ± 1.04

TABLE S3.4: Event-related signal modulations during ambiguous and physical motion. Mean trough-to-peak differences \pm standard error of the mean. Positive values indicate that signal increased from the minimum of the first two time points to the maximum of the last two time points. Negative values indicate a signal decrease from the maximum of the first two time points to the minimum of the last two time points. H and V cluster represent the horizontal and vertical clusters defined in the main experiment. H and V percept indicate time periods during which subjects indicated a horizontal or vertical percept. S1, S3, S7, S8 and S9 indicate the different subjects.

	Ambiguous Motion		Physical Motion		
	H percept	V Percept	H percept	V percept	
S1					
	H cluster	1.07 \pm 0.17	-0.42 \pm 0.26	1.92 \pm 0.20	-1.23 \pm 0.21
	V cluster	-0.15 \pm 0.17	0.68 \pm 0.22	-0.50 \pm 0.25	1.73 \pm 0.16
S3					
	H cluster	1.27 \pm 0.16	-0.89 \pm 0.15	1.85 \pm 0.18	-1.08 \pm 0.16
	V cluster	0.11 \pm 0.15	0.30 \pm 0.13	-0.96 \pm 0.26	1.51 \pm 0.21
S7					
	H cluster	0.89 \pm 0.16	-0.65 \pm 0.19	2.09 \pm 0.12	-1.23 \pm 0.15
	V cluster	-0.37 \pm 0.15	0.46 \pm 0.18	-1.13 \pm 0.19	1.91 \pm 0.20
S8					
	H cluster	0.77 \pm 0.17	-0.43 \pm 0.18	1.82 \pm 0.12	-0.97 \pm 0.12
	V cluster	-0.12 \pm 0.14	0.29 \pm 0.20	-0.79 \pm 0.15	1.42 \pm 0.13
S9					
	H cluster	0.24 \pm 0.13	-0.75 \pm 0.16	2.11 \pm 0.13	-0.98 \pm 0.13
	V cluster	-0.53 \pm 0.14	0.79 \pm 0.16	-0.66 \pm 0.15	2.14 \pm 0.12

3.7 ACKNOWLEDGMENTS

We would like to thank Ingo Marquardt, Faruk Gulban, Sriranga Kashyap, Agustin Lage and Nikolaus Weiskopf for valuable advice.

3.8 REFERENCES

- Albright, T. D., Desimone, R., & Gross, C. G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *Journal of Neurophysiology*, *51*(1), 16–31.
- Amano, K., Wandell, B. A., & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, *102*(5), 2704–18.
- Andersson, J. L. R., Jenkinson, M., & Smith, S. (2007). *Non-linear registration, aka spatial normalisation. FMRIB Technical Report TR07JA2*. FMRIB.
- Andersson, J. L., Skare, S., & Ashburner, J. (2003). How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. *NeuroImage*, *20*(2), 870–888.
- Andrews, T. J., Schluppeck, D., Homfray, D., Matthews, P., & Blakemore, C. (2002). Activity in the fusiform gyrus predicts conscious perception of Rubin’s vase-face illusion. *NeuroImage*, *17*(2), 890–901.
- Ashburner, J. & Friston, K. J. (2005). Unified segmentation. *NeuroImage*, *26*(3), 839–851.
- Bartels, A., Logothetis, N. K., & Moutoussis, K. (2008). fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. *Trends in Neurosciences*, *31*(9), 444–453.
- Brascamp, J., Sterzer, P., Blake, R., & Knapen, T. (2018). Multistable Perception and the Role of Frontoparietal Cortex in Perceptual Inference. *Annual Review of Psychology*, *69*(1), 77–103.
- Britten, K. H. (1998). Clustering of response selectivity in the medial superior temporal area of extrastriate cortex in the macaque monkey. *Visual Neuroscience*, *15*(3), 553–558.
- Brouwer, G. J. & van Ee, R. (2007). Visual cortex allows prediction of perceptual states during ambiguous structure-from-motion. *Journal of Neuroscience*, *27*(5), 1015–1023.
- Castelo-Branco, M., Formisano, E., Backes, W., Zanella, F., Neuenschwander, S., Singer, W., & Goebel, R. (2002). Activity patterns in human motion-sensitive areas depend on the interpretation of global motion. *Proceedings of the National Academy of Sciences*, *99*(21), 13914–13919.
- Chaudhuri, A. & Glaser, D. A. (1991). Metastable Motion Anisotropy. *Visual Neuroscience*, *7*(5), 397–407.

- Cheng, K., Waggoner, R. A., & Tanaka, K. (2001). Human ocular dominance columns as revealed by high-field functional magnetic resonance imaging. *Neuron*, 32, 1–16.
- Cohen, E. R. (1998). An Introduction to Error Analysis: The Study of Uncertainties in Physical Measurements. *Measurement Science and Technology*, 9(6).
- De Graaf, T. A., Hsieh, P. J., & Sack, A. T. (2012). The 'correlates' in neural correlates of consciousness. *Neuroscience and Biobehavioral Reviews*, 36(1), 191–197.
- De Martino, F., Yacoub, E., Kemper, V., Moerel, M., Uludag, K., De Weerd, P., ... Formisano, E. (2018). The impact of ultra-high field MRI on cognitive and computational neuroimaging. *NeuroImage*, 168, 366–382.
- Dodd, J. V., Krug, K., Cumming, B. G., & Parker, A. J. (2001). Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *Journal of Neuroscience*, 21(13), 4809–21.
- Dumoulin, S. O. (2000). A New Anatomical Landmark for Reliable Identification of Human Area V5/MT: a Quantitative Analysis of Sulcal Patterning. *Cerebral Cortex*, 10(5), 454–463.
- Dumoulin, S. O. & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *NeuroImage*, 39(2), 647–660.
- Emmerling, T. C., Zimmermann, J., Sorger, B., Frost, M. A., & Goebel, R. (2016). Decoding the direction of imagined visual motion using 7 T ultra-high field fMRI. *NeuroImage*, 125, 61–73.
- Emmerling, T. (2016). *Imaging imagery: an investigation of visual cognition using high-resolution fmri* (Doctoral dissertation, Maastricht University).
- Feinberg, D. A., Moeller, S., Smith, S. M., Auerbach, E., Ramanna, S., Glasser, M. F., ... Yacoub, E. (2010). Multiplexed echo planar imaging for sub-second whole brain fmri and fast diffusion imaging. *Plos One*.
- Felleman, D. J. & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*.
- Finlay, D. & von Grünau, M. (1987). Some experiments on the breakdown effect in apparent motion. *Perception and Psychophysics*, 42(6), 526–534.
- Frassle, S., Sommer, J., Jansen, A., Naber, M., & Einhauser, W. (2014). Binocular Rivalry: Frontal Activity Relates to Introspection and Action But Not to Perception. *Journal of Neuroscience*, 34(5), 1738–1747.

-
- Friston, K. J., Fletcher, P., Josephs, O., Holmes, a., Rugg, M. D., & Turner, R. (1998). Event-related fMRI: characterizing differential responses. *NeuroImage*, 7(1), 30–40.
- Friston, K. J. (2006). *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Academic Press.
- Gen, E., Bergmann, J., Singer, W., & Kohler, A. (2011). Interhemispheric connections shape subjective experience of bistable motion. *Current Biology*, 21(17), 1494–1499.
- Goncalves, N. R., Ban, H., Sanchez-Panchuelo, R. M., Francis, S. T., Schluppeck, D., & Welchman, A. E. (2015). 7 Tesla fMRI Reveals Systematic Functional Organization for Binocular Disparity in Dorsal Visual Cortex. *Journal of Neuroscience*, 35(7), 3056–3072.
- Greve, D. N. & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48(1), 63–72.
- Gulban, O., Schneider, M., Marquardt, I., Haast, R., & De Martino, F. (2018). A scalable method to improve outer gray matter boundary segmentation at ultra high field MRI. *Plos One*, 13, 1–33.
- Hasson, U., Hendler, T., Bashat, D. B., & Malach, R. (2001). Vase or face? A neural correlate of shape-selective grouping processes in the human brain. *Journal of Cognitive Neuroscience*, 13(6), 744–753.
- Huk, A. C., Dougherty, R. F., & Heeger, D. J. (2002). Retinotopy and functional subdivision of human areas MT and MST. *Journal of Neuroscience*, 22(16), 7195–7205.
- Hunter, J. D. (2007). Matplotlib: a 2d graphics environment. *Computing In Science & Engineering*, 9(3), 90–95.
- Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2), 825–841.
- Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, 62(2), 782–790.
- Jenkinson, M. & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2), 143–156.
- Jones, E., Oliphant, T., Peterson, P., et al. (2007). SciPy: open source scientific tools for Python. *Computing in Science and Engineering*.

- Kamitani, Y. & Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Current Biology*, 16(11), 1096–1102.
- Kemper, V. G., De Martino, F., Emmerling, T. C., Yacoub, E., & Goebel, R. (2018). High resolution data analysis strategies for mesoscale human functional MRI at 7 and 9.4T. *NeuroImage*, 164, 48–58.
- Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience*, 17(5), 307–321.
- Kolster, H., Peeters, R., & Orban, G. a. (2010). The retinotopic organization of the human middle temporal area MT/V5 and its cortical neighbors. *Journal of Neuroscience*, 30(29), 9801–20.
- Kriegeskorte, N. & Goebel, R. (2001). An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical MR volumes. *NeuroImage*, 14(2), 329–346.
- Leopold, D. A. & Logothetis, N. K. (1999). Multistable phenomena: Changing views in perception. *Trends in Cognitive Sciences*, 3(7), 254–264.
- Logothetis, N. K. (1989). Neuronal correlates of subjective visual perception. *Science*, 245, 761–763.
- Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197), 869–878.
- Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., . . . Kennedy, H. (2014). Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology*.
- Marquardt, I., Schneider, M., Gulban, O. F., Ivanov, D., & Uludağ, K. (2018). Cortical depth profiles of luminance contrast responses in human V1 and V2 using 7 T fMRI. *Human Brain Mapping*, 39(7), 2812–2827.
- Meng, X.-l., Rosenthal, R., & Rubin, D. B. (1992). Comparing correlated correlation coefficients. *Psychological Bulletin*, 111(1), 172–175.
- Metzinger, T. (2000). Neural correlates of consciousness: Empirical and conceptual questions. *MIT Press*.
- Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., & Uğurbil, K. (2010). Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application

-
- to high spatial and temporal whole-brain fMRI. *Magnetic Resonance in Medicine*.
- Moran, P. A. (1950). Notes on continuous stochastic phenomena. *Biometrika*.
- Mountcastle, V. (1956). Modality and Topographic Properties of Single Neurons of Cats's Somatic Sensory Cortex. *Journal of Neurophysiology*, 20(4), 408–434.
- Muckli, L., Kohler, A., Kriegeskorte, N., & Singer, W. (2005). Primary visual cortex activity along the apparent-motion trace reflects illusory perception. *Plos Biology*, 3(8).
- Muckli, L., Kriegeskorte, N., Lanfermann, H., Zanella, F. E., Singer, W., & Goebel, R. (2002). Apparent motion: event-related functional magnetic resonance imaging of perceptual switches and states. *Journal of Neuroscience*, 22(9), 219–219.
- Nasr, S., Polimeni, J. R., & Tootell, R. B. H. (2016). Interdigitated color- and disparity-selective columns within human visual cortical areas V2 and V3. *Journal of Neuroscience*, 36(6), 1841–57.
- Oden, N. L. & Sokal, R. R. (1986). Directional Autocorrelation - an Extension of Spatial Correlograms to 2 Dimensions. *Systematic Zoology*.
- Peirce, J. W. (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.
- Polimeni, J. R., Renvall, V., Zaretskaya, N., & Fischl, B. (2018). Analysis strategies for high-resolution UHF-fMRI data. *NeuroImage*, 168, 296–320.
- Rakic, P. (2008). Confusing cortical columns. *Proceedings of the National Academy of Sciences*, 105(34), 12099–12100.
- Ramachandran, V. S. & Anstis, S. M. (1985). Perceptual organization in multistable apparent motion. *Perception*, 14, 135–143.
- Rees, G. (2007). Neural correlates of the contents of visual awareness in humans. *Philosophical Transactions of the Royal Society*, 362(1481), 877–886.
- Rosenberg, M. S. (2010). The Bearing Correlogram: A New Method of Analyzing Directional Spatial Autocorrelation. *Geographical Analysis*, 32(3), 267–278.

- Salzman, C. D., Britten, K. H., & Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346(6280), 174–177.
- Schneider, M. (2018). Vificov - visual field coverage visualization in python. *Zenodo repository*.
- Schneider, M., Marquardt, I., Gulban, O. F., & Isik, I. (2018). Pyprf feature - population receptive field estimation with feature weights. *Zenodo repository*.
- Setsompop, K., Gagoski, B. A., Polimeni, J. R., Witzel, T., Wedeen, V. J., & Wald, L. L. (2012). Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic Resonance in Medicine*.
- Smith, S., Jenkinson, M., Woolrich, M., Beckmann, C., Behrens, T., Johansen-Berg, H., ... Matthews, P. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*.
- Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3), 143–155.
- Sterzer, P., Kleinschmidt, A., & Rees, G. (2009). The neural bases of multi-stable perception. *Trends in Cognitive Sciences*, 13(7), 310–318.
- Sterzer, P., Russ, M. O., Preibisch, C., & Kleinschmidt, A. (2002). Neural correlates of spontaneous direction reversals in ambiguous apparent visual motion. *NeuroImage*, 15(4), 908–916.
- Tanaka, K. (2003). Columns for complex visual object features in the inferotemporal cortex: Clustering of cells with similar but slightly different stimulus selectivities. *Cerebral Cortex*, 13(1), 90–99.
- Tong, F., Nakayama, K., Vaughan, J. T., & Kanwisher, N. (1998). Binocular rivalry and visual awareness in human extrastriate cortex. *Neuron*, 21(4), 753–759.
- Tootell, R. B. H. & Nasr, S. (2017). Columnar Segregation of Magnocellular and Parvocellular Streams in Human Extrastriate Cortex. *Journal of Neuroscience*, 0690–17.
- Tyler, C. W., Likova, L. T., Chen, C.-C., Kontsevich, L. L., Schira, M. M., & Wade, A. R. (2005). Extended Concepts of Occipital Retinotopy. *Current Medical Imaging Reviews*, 1, 319–329.

-
- Uğurbil, K., Adriany, G., Andersen, P., Chen, W., Garwood, M., Gruetter, R., . . . Zhu, X. H. (2003). Ultrahigh field magnetic resonance imaging and spectroscopy. *Magnetic Resonance Imaging*.
- Van de Moortele, P.-F., Auerbach, E. J., Olman, C., Yacoub, E., Uğurbil, K., & Moeller, S. (2009). T1 weighted brain images at 7 Tesla unbiased for Proton Density, T2* contrast and RF coil receive B1 sensitivity with simultaneous vessel visualization. *NeuroImage*, *46*(2), 432–446.
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, *13*(2), 22–30.
- Waehnert, M. D., Dinse, J., Weiss, M., Streicher, M. N., Waehnert, P., Geyer, S., . . . Bazin, P. L. (2014). Anatomically motivated modeling of cortical laminae. *NeuroImage*, *93*, 210–220.
- Wang, H. X., Merriam, E. P., Freeman, J., & Heeger, D. J. (2014). Motion Direction Biases and Decoding in Human Visual Cortex. *Journal of Neuroscience*, *34*(37), 12601–12615.
- Williams, Z. M., Elfar, J. C., Eskandar, E. N., Toth, L. J., & Assad, J. A. (2003). Parietal activity and the perceived direction of ambiguous apparent motion. *Nature Neuroscience*, *6*(6), 616–623.
- Yacoub, E. & Harel, N. (2008). High-field fMRI unveils orientation columns in humans. *Proceedings of the National Academy of Sciences*, *105*(30), 254–264.
- Yushkevich, P. A., Piven, J., Cody Hazlett, H., Gimpel Smith, R., Ho, S., Gee, J. C., & Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *NeuroImage*, *31*(3), 1116–1128.
- Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P. F., Feinberg, D., Adriany, G., . . . Yacoub, E. (2011). Mapping the organization of axis of motion selective features in human area MT using high-field fmri. *Plos One*, *6*(12), 1–10.

4

Motion displaces population receptive fields in the direction opposite to motion

Based on

Schneider, M., Marquardt, I., Sengupta, S., De Martino, F., & Goebel, R. (2019). Linking apparent position to population receptive field estimates using a visual field projection model. *Proceedings of the Cognitive Computational Neuroscience Conference*, 3, 1006-1009.

Schneider, M., Marquardt, I., Sengupta, S., De Martino, F., & Goebel, R. (in preparation). Motion displaces population receptive fields in the direction opposite to motion.

4.1 ABSTRACT

Motion signals can bias the perceived position of visual stimuli. While the apparent position of a stimulus is biased in the direction of motion, electro-physiological studies have shown that the receptive field (RF) of neurons is shifted in the direction opposite to motion, at least in cats and macaque monkeys. In humans, it remains unclear how motion signals affect population RF (pRF) estimates. We addressed this question using psychophysical measurements and functional magnetic resonance imaging (fMRI) at 7 Tesla. We systematically varied two factors: the motion direction of the carrier pattern (inward, outward and flicker motion) and the contrast of the mapping stimulus (low and high stimulus contrast). We observed that while physical positions were identical across all conditions, presence of low-contrast motion, but not high-contrast motion, shifted perceived stimulus position in the direction of motion. Correspondingly, we found that pRF estimates in early visual cortex were shifted against the direction of motion for low-contrast stimuli but not for high stimulus contrast. We offer an explanation in form of a model for why apertures are perceptually shifted in the direction of motion even though pRFs shift in the opposite direction.

4.2 INTRODUCTION

An important task of the visual system is to infer the location of objects in our environment. A wide range of psychophysical studies shows that motion signals lead to systematic localisation biases (Arnold, Thompson, & Johnston, 2007; Chung, Patel, Bedell, & Yilmaz, 2007; De Valois & De Valois, 1991; Eagleman, Eagleman, Sejnowski, & Sejnowski, 2007; Fu, Shen, & Dan, 2001; Kwon, Tadin, & Knill, 2015; McGraw, Whitaker, Skillen, & Chung, 2002; Nishida & Johnston, 1999; Ramachandran & Anstis, 1990). In illusions called motion-induced position shifts (MIPS), a coherent motion signal shifts the apparent location of a stimulus (Ramachandran & Anstis, 1990). For example, when drifting Gabor patches are presented within a stationary aperture, the stimulus appears shifted in the direction of motion (Arnold et al., 2007; Chung et al., 2007; De Valois & De Valois, 1991). Such illusions raise the question how our visual system encodes

location and how, in the case of MIPS, the apparent position shift can be explained. Furthermore, they offer a dissociation between the physical and the perceived position of a stimulus that can clarify which neuronal processes correspond to the apparent position of the stimulus.

The magnitude of MIPS is known to depend on spatial and temporal properties of the stimulus. MIPS are larger when the stimulus is shown for longer duration (tested up to 453 ms; Chung et al., 2007), presented at higher speed (Chung et al., 2007; Kwon et al., 2015) or at higher eccentricities (Chung et al., 2007; Fu, 2004; Kwon et al., 2015). The magnitude of MIPS furthermore depends on spatial blurring of the presented stimulus. Blurred stimulus edges lead to larger perceptual displacements than sharp edges (Fu et al., 2001; Kwon et al., 2015) and increasing the size of the Gaussian envelope of a Gabor stimulus yields larger MIPS (Fu et al., 2001). Arnold et al. (2007) have suggested that MIPS are driven by modulation of apparent contrast of the stimulus. Supporting this suggestion, they reported perceived position shifts when observers were asked to match the extremities of two contrast envelopes (low-contrast region), but no corresponding shifts when observers were instructed to match the phase within two gratings (high-contrast region).

While the stimulus-related factors that influence MIPS are well studied, the neuronal mechanisms underlying MIPS in humans have remained controversial. Based on the results of a functional magnetic resonance imaging (fMRI) study, Whitney and colleagues (2003) initially suggested that the perceptual shift was accompanied by a shift in retinotopic representation of the stimulus. They presented Gabor stimuli that contained motion drifting either towards or away from fixation. Although Gabor stimuli were presented at the same physical location, when contrasting inward and outward motion conditions, the authors observed positive signal differences in an anterior V1 region (representing higher eccentricities) and negative signal differences in a more posterior V1 region (representing lower eccentricities). Based on this observation, as well as several control experiments, they concluded that retinotopic representation is flexible and shifted by the motion direction signals of a stimulus.

This account has since been challenged by Liu and colleagues (2006). In a series of experiments, they convincingly demonstrated that retinotopic representations in V1-V3 of stimuli presented at the same location but with

opposing motion are, to a large degree, overlapping. Furthermore, in one of their control experiments, they determined the preferred eccentricity for V1 voxels presenting a traveling wave stimulus that contained either inward or outward motion. Eccentricity estimates did not vary significantly between the two motion directions. Based on these results, they argued that the activation differences reported by Whitney et al., 2003 were caused by small signal differences at retinotopic locations that represent the trailing edge of motion and were amplified by contrasting motion conditions directly (since the hemodynamic response in fMRI acts like a compressive nonlinearity). Thus, although the center of mass of the retinotopic representations was slightly shifted against the direction of motion, the two representations showed substantial overlap.

In light of recent studies, the results presented by Whitney et al., 2003 and Liu et al., 2006 can both be interpreted in terms of a phenomenon that has been termed "aperture-inward" bias (Wang, Merriam, Freeman, & Heeger, 2014). Wang et al., 2014 demonstrated that the trailing edge of a motion stimulus evokes larger responses than the leading edge, at least in early visual cortex (V1-V3) but not in V5/hMT+. Furthermore, several studies showed that responses decrease gradually with distance from the trailing edge (Schellekens, Ramsey, & Raemaekers, 2015; Schellekens, van Wezel, Petridou, Ramsey, & Raemaekers, 2016; Wang et al., 2014).

Although the effects initially reported by Whitney et al., 2003 are now better understood, it still remains elusive which neuronal mechanisms underlie perceptual position shifts in humans. When monkeys and cats were presented with motion stimuli that lead to perceptual shifts for human observers, electro-physiological recordings revealed that the receptive field (RF) of V1 (Fu, 2004) and V4 (Sundberg, Fallah, & Reynolds, 2006) neurons were shifted against the direction of motion. Importantly, these RF shifts cannot be understood in terms of the aperture-inward bias. If response amplitudes are larger at the trailing edge of the stimulus and an RF is assumed to be composed of inputs from pre-synaptic neurons with smaller RF fields, then RFs should shift in the direction of motion, which is the opposite of what was observed empirically (Fu, 2004).

Using fMRI in humans it is not possible to measure RF properties of individual neurons. However, a quantitative framework for estimating the population RF (pRF) of voxels has been developed (Dumoulin & Wandell,

2008). Using this framework, Harvey et al. (2016) reported that pRF size and eccentricity depended on the speed and direction of a motion stimulus. In particular, higher speed led to higher pRF size and eccentricity estimates in early- and mid-level visual areas if motion was in the direction of the pRF measurement but not if motion was in an orthogonal direction. However, since opposing motion directions were presented during the same stimulus block, it was not possible to disentangle the effects of opposing motion directions on pRF eccentricity and size. Consequently, it is still unclear whether pRFs are displaced with motion in humans (and, if so, in which direction).

As a result, there is an unsatisfactory gap between human psychophysics, animal electro-physiology, and human neuroimaging. In an attempt to bridge this gap, we identified several factors that could explain the mismatch between results in human and non-human subjects. These factors motivated the current study. First, to determine preferred eccentricity for V1 voxels, Liu et al., 2006 presented ring apertures, while we decided to employ annulus sector apertures instead. This difference might be important since the two aperture types lead to perceptually different phenomena: While the former has been reported to affect perceived depth (Edwards & Badcock, 2003), the latter affects perceived position of the stimulus. Second, in addition to presenting our stimuli at high-contrast, we also presented them at lower contrast. This was motivated by the findings that MIPS magnitude scales with blurring of the stimulus and reduction in apparent contrast. We therefore expected larger perceptual displacements for low-contrast stimuli. Third, Liu et al., 2006 reported results for early visual areas V1 to V3. In the current study we included several additional visual areas, including V3ab, hV4, LO, VO and TO (which comprises V5/hMT+; see Amano, Wandell, and Dumoulin, 2009) to be able to study the effect on retinotopic representation across a wider range of areas. Fourth, potentially the difference in signal that comes with the shift in retinotopic representation is too small to be picked up by human neuroimaging methods. Yet advances in ultra-high field MRI (Uğurbil et al., 2003) now allow for reliable recordings of human subjects at higher field strength which yield an improved signal-to-noise ratio. For the present study, we combined 7 Tesla [T] MRI with a custom-made surface

coil optimized for high-resolution imaging of the human visual cortex (Sengupta et al., 2016) to obtain state-of-the-art recordings.

We recorded fMRI responses from early- and mid-level visual areas in humans while they were presented with pRF mapping stimuli that varied along two factors: the motion direction of the carrier pattern (3 levels: inward, outward and flicker motion) and the contrast of the mapping stimulus (2 levels: low and high stimulus contrast)(Figure 4.1). We hypothesized that carrier motion should displace population receptive field estimates against the direction of carrier motion. Furthermore, we expected greater and more reliable displacements for lower as compared to higher contrast.

4.3 METHODS

4.3.1 *Participants*

Six healthy participants with corrected-to-normal vision were recruited for the study. All participants were students of Maastricht University. Recruitment was limited to participants who were MRI compatible and had been in an MRI scanner before to ensure high subject compliance. All participants were experienced in fixation tasks. The data for one participant (S1) were excluded from further analysis because mirror and projector settings in the scanner were updated between scanning sessions, rendering the data for different scanning sessions of this subjects incompatible in terms of eccentricity of the presented stimuli. Hence, five participants (4 females, 24-30 years old) were analyzed. All participants gave informed, written consent to participate in the experiment. The study was approved by the research ethics committee of the Faculty of Psychology and Neuroscience of Maastricht University.

4.3.2 *Experimental design and stimuli*

Visual stimuli were created and presented using the open-source software package PsychoPy (version 1.83.01) (Peirce, 2007, 2008). Stimuli were projected on a frosted screen (using Panasonic projector PT-EZ570; Newark, NJ, USA; resolution 1920 x 1200; nominal refresh rate: 60 Hz). Subjects viewed the screen via a mirror attached to the surface coil. Button re-

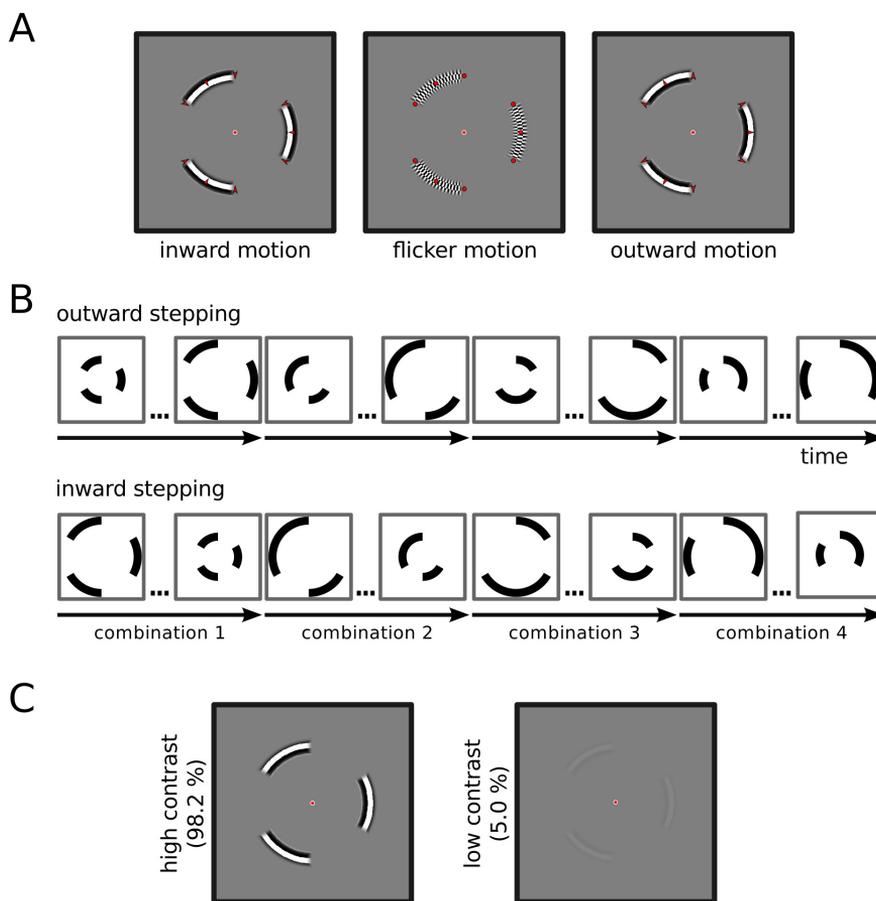


FIGURE 4.1: Overview of the experimental design.

sponses were registered using an MR compatible button box (Current Designs, 8-button response device, HHSC-2 x 4-C; Philadelphia, USA). Subjects participated either in two (S6) or three (S2, S3, S4, S5) scanning sessions of approximately 2 hours duration. Scanning sessions were on separate days. In session 01, we collected 4 runs of population receptive field mapping which was used to define regions of interest (ROI) as well as to derive area- and subject-specific hemodynamic response functions (HRF) (308 volumes each; independent ROI and HRF data). In sessions 02

FIGURE 4.1 (previous page): Overview of the experimental design. We systematically varied carrier pattern, aperture position and stimulus contrast. **(A)** Three types of carrier pattern. During inward and outward motion conditions, apertures revealed a radial square wave grating which drifted either inward or outward, while the aperture remained stationary. During flicker condition, apertures revealed a dartboard pattern, as classically used for pRF mapping. Red arrows and dots illustrate motion direction and were not shown during the experiment. **(B)** Aperture positions. Stimuli were presented behind stationary annulus sector apertures, indicated schematically as black sectors. Aperture position was systematically varied. There were four different annulus sector combinations (combination 1-4), covering different polar angles. These sector combination either stepped from low to high eccentricities (outward stepping) or from high to low eccentricities (inward stepping). Consecutive dots indicate in-between aperture position. **(C)** Stimulus contrast. All stimuli were presented at both high (98.2%) and low (5.0%) Michelson contrast.

and 03, we acquired 9 to 12 runs of pRF data for the main experiment at low and high luminance contrast (222 volumes each; low- and high-contrast pRF data). The different contrast levels were presented in separate sessions.

LOW- AND HIGH-CONTRAST PRF DATA The pRF mapping stimuli consisted of annulus sector apertures presented at systematically varied positions. The choice of annulus sectors, as opposed to entire annuli (Liu et al., 2006), was motivated by our aim to manipulate perceived position, not perceived depth (Edwards & Badcock, 2003). A single sector had a radial extent of 1.7° of visual angle and subtended 60 angular degrees such that the visual field was divided into six sectors in the angular direction, starting at the upper vertical meridian. At any moment, a combination of three sectors with equal eccentricity was shown. There were four possible sector combinations and each of the six sectors was shown during two of these combinations (see Figure 4.1B). Combinations were chosen to make the presentation of one sector maximally independent from presentation of other sectors. This was done by following a binary decision tree with four levels, picking the branches that led to exactly two successes and two failures. A binary decision tree with four levels has exactly six such branches

and a single branch determined for one of the six sectors whether it was shown in a given combination (shown if success; not shown if failure).

For every combination, the eccentricity was systematically varied such that a combination started at a minimum radius of 3.4° and stepped up to a maximum of 8.5° , updating every 2 s (our TR) in 15 continuous steps. For the presentation of the first two and last two apertures the inner and outer edge, respectively, remained fixed. This resulted in an aperture step size of 0.34° . Combinations could either step from low to high eccentricities or, reversely, from high to low eccentricities and stepping direction was counter-balanced across functional runs. After completion of either 15 inward or outward steps, a different sector combination completed the same sequence of steps.

During inward and outward motion conditions, apertures revealed a radial square wave grating (spatial frequency: 1 cycle per degree) which drifted 5° of visual angle per s inward or outward, respectively, while the aperture remained stationary (Figure 4.1A). During the flicker condition, apertures revealed a dartboard pattern, as classically used for retinotopy and pRF mapping (Dumoulin & Wandell, 2008). In radial direction, a dartboard element had an extent of 0.5° of visual angle, thus matching the spatial frequency of the square wave gratings during the inward and outward motion condition (1 cycle per degree). In the angular direction, dartboard elements subtended 3.36 angular degrees ($360 / 107 = 3.36$). This extent was chosen such that even at maximum eccentricity (8.5° of visual angle), the extent of the dartboard elements in angular direction was smaller than the extent in radial direction. This ensured that when the dartboard was flickered, elements were perceived to move in angular rather than radial direction. Elements were perceived to move either clockwise or counter-clockwise and perceived direction changed randomly over time (Amano et al., 2009). The dartboard was flickered at a frequency of 5 on-off cycles per s such that the frequency was matched to the inward/outward motion conditions.

Gratings were presented either at 5% (low-contrast condition) or at 98.2% (high-contrast condition) Michelson contrast, which was the highest possible contrast in the scanner environment. In order to reduce increased responses to the aperture edge (Schellekens et al., 2016; Wang et al., 2014), contrast was ramped following a raised cosine over a distance of 0.6 deg

from the aperture edge. Luminous intensity was determined by varying the Psychopy settings for color from the most negative setting ([-1, -1, -1], i.e. "black") to the most positive setting ([1, 1, 1], i.e. "white") in 17 linear steps and concurrently measuring luminance inside the scanner environment using a photometer. Every of the 17 measurements was repeated twice and a cubic polynomial function was fit to the mean across measurements. The resulting polynomial function allowed us to estimate the intensities, given the settings in Psychopy, and to calculate the Michelson contrast according to $(L1 - L2)/(L1 + L2)$, where $L1$ and $L2$ refer to the luminance of "white" and "black" segments, respectively.

To avoid adaptation to motion direction (Fu, 2004; Schellart, Trindade, Reits, Verbunt, & Spekreijse, 2004), all stimulus presentations in the main experiment were restricted to brief periods of 400 ms. During the first and last 50 ms of the 400 ms presentation period contrast was ramped on and off following a raised cosine in order to reduce on- and offset effects. This was followed by 600 ms blank screen. During a TR (2 s), this sequence was shown twice. Every functional run consisted of 222 volumes and was divided into 3 parts (74 volumes). Each part showed one of the 3 conditions (outward drifting, inward drifting, flicker motion) and consisted of 7 volumes blank screen in the beginning and end as well as a sequence of 15 steps (either inward or outward stepping) for each of the four sector combinations (4 x 15 volumes = 60 volumes).

Participants were asked to fixate on a centrally presented dot throughout the entire experimental run. The fixation dot consisted of a red circle of 0.125° radius. To increase visibility of the fixation dot, it was surrounded by a white annulus of 0.19° radius (Harvey & Dumoulin, 2016). To ensure that participants were complying with the fixation task, occasionally the central fixation dot changed color for 0.3 s to perceptually either darker or brighter shades of red. Participants were asked to indicate the direction of the change by pressing one of two buttons (button 1 for "darker", button 2 for "brighter").

INDEPENDENT ROI AND HRF DATA In an additional session, we obtained data to estimate population receptive field parameters independent of the main experiment. These data were used to define visual field maps, to select voxels for the main experiment and to determine area- and subject-

specific HRF. We again systematically varied position and shape of the stimulus aperture, this time using three types of apertures: vertical bars, horizontal bars and circular sectors ("wedges"). All apertures were limited to a circular region of the display, 24° of visual angle in diameter, and presented against a mean luminance gray background. Vertical or horizontal bars traversed the visual field in 31 discrete steps. Vertical/horizontal bars stepped either from the left/top to right/bottom and back, or in reverse order. Both horizontal and vertical bars subtended 3° of visual angle in the stepping directions. In the direction orthogonal to stepping they were limited by the global circular aperture. The wedge aperture (angle width 45° , radius 12° of visual angle) completed a full circle in 32 steps, moving either clock-wise then counter-clockwise, or in reverse order.

The carrier pattern for all aperture types consisted of a dynamic, high-contrast pseudo-checkerboard as described by Schwarzkopf, Anderson, de Haas, White, and Rees, 2014. In comparison to traditional flickering checkerboard patterns, such a carrier pattern of concentric shapes varies in spatial frequency and phase over time and ensures stimulation of a broad range of visually responsive neurons (Alvarez, de Haas, Clark, Rees, & Schwarzkopf, 2015). The carrier pattern was always shown at 98.2% Michelson contrast (i.e., the highest possible contrast in the scanner environment). For every aperture type, contrast was ramped following a raised cosine over a range of 0.3° from the aperture border.

Aperture positions were updated with every TR (2 s). To avoid the appearance of apertures "jumping" from one position to the next, the stimulus was shown for 1.5 s only, followed by 500 ms blank screen. During the first and last 100 ms of the 1.5 s stimulation, contrast faded in or out following a raised cosine. Every functional run started and ended with 14 s blank screen. During every run, all three aperture types were shown. Vertical and horizontal bars each traversed the visual field twice; the wedge aperture traversed the visual field four times. The order of aperture types and traversing directions was counter-balanced across runs according to a balanced Latin square. Every two traversals of the visual field (back and forth) by one of the aperture types was followed by a period of 28 s blank screen.

As in the main experiment, participants were asked to maintain central fixation. Make-up of the fixation dot was identical to the one in the

main experiment. The fixation dot occasionally changed color from red to yellow for 0.3 s and participants indicated a color change by pressing a button. When the aperture stimuli passed behind the fixation point, they spared out a circular region of 0.4° radius that revealed the fixation dot and part of the gray mean-luminance background. This was to avoid eye motion tracking the moving stimulus aperture (Harvey & Dumoulin, 2016). To furthermore facilitate central fixation, an extended grid of thin lines was presented throughout the experimental run. The grid consisted of two diagonal, orthogonal lines crossing behind the fixation dot as well as circles with radii of 1° , 2.5° , 5° , 10° and 15° of visual angle centered on the fixation dot. Such a stable grid stabilizes fixation (Tyler et al., 2005). This grid was not shown in the main experiment to prevent participants from using it as a reference for location of the apertures.

4.3.3 Psychophysics experiment

To determine the perceptual displacement of the stimuli used in the fMRI experiment, all participants took part in a separate psychophysics session. This session followed a similar procedure as described by Harvey and Dumoulin, 2016, with modifications where necessary due to differing stimuli. Like the fMRI experiment, all psychophysics stimuli were generated in PsychoPy (version 1.83.01) (Peirce, 2007, 2008). Luminance measurements were repeated in the environment of the psychophysics lab, the polynomial function was refit and stimulus settings were adjusted to match the contrast levels employed in the fMRI experiment (5% and 98.2%).

Stimuli in the psychophysics session consisted of a radial sector aperture which subtended 120 angular degrees and had an extent of 1.7° . The aperture was centered on the upper vertical meridian at a distance of 5.95° from fixation, which represented the center eccentricity during the fMRI experiment. The carrier patterns were the same as in the fMRI experiment (outward, inward or flicker motion) and stimuli were again presented at two contrast levels (5% and 98.2% Michelson contrast).

In order to derive psychometric functions, we used the method of constant stimuli and a two-interval forced choice paradigm. On a given trial we presented two stimuli in temporal succession and then systematically varied the vertical distance between them. There was either

no offset or the second aperture was 0.05° , 0.1° , 0.2° , 0.4° , or 0.8° above or below the first one. Stimuli were shown for the same duration as in the fMRI experiment (400 ms, of which contrast was ramped during the first and last 50 ms) and an interval of 500 ms blank screen was inserted between two subsequently shown stimuli. Subjects were asked to judge whether they perceived the aperture that was presented first or second as higher and communicated their judgment via button presses.

There were five possible combinations of subsequently presented stimulus pairs: outward then flicker motion, flicker then outward motion, inward then flicker motion, flicker then inward motion, flicker then flicker motion. The flicker then flicker combination was shown twice. An experiment block contained all vertical offsets crossed with all carrier pattern pairs, resulting in 11 (offset conditions) \times 6 (motion combinations) = 66 trials per block. The order of trials was randomized. The two contrast-levels were presented in separate blocks. Each subject completed 8-10 blocks per contrast level.

4.3.4 *Psychophysics analysis*

Combinations that contained the same relative carrier patterns during the presentation of the two stimuli were combined. This reduced the number of combinations to three: flicker-flicker, outward-flicker and inward-flicker. For every vertical offset, we divided the number of times the subject judged the second aperture as higher by the total number of presentations at this offset. This yielded the proportion of perceiving the second aperture as higher which constituted our dependent measure. We then fitted a cumulative Gaussian psychometric function to these proportions using Bayesian inference based on samples from the posterior distribution as implemented in the Python version of the *psignifit* toolbox (Frund, Haenel, & Wichmann, 2011). We used *psignifit*'s default priors for the 'YesNo' option, which reflect the belief that the point of subjective equality (PSE) falls within the range of measured data and, with decreasing probability, up to half the range above or below the measured data (Frund et al., 2011). The PSE in our setting corresponds to the point where subjects perceived the two consecutively presented stimuli to be at equal height.

Based on individual psychometric functions, we determined the PSE for every carrier motion combination and subject. These individually fitted PSEs were used for inferential statistical analyses. We used the scaled PSE that is corrected for "guessing" and "lapse" rates estimated for every individual psychometric function. We averaged the scaled PSEs across subjects and calculated the standard error of the mean and 95% confidence interval.

4.3.5 MRI acquisition

Data acquisition was performed on a whole-body Magnetom scanner (nominal field strength 7 T, Siemens Medical Systems, Erlangen, Germany) at the Maastricht Brain Imaging Centre, The Netherlands. All images were acquired using a surface coil (4 channel transmit, 16 channel receive) custom-made for high resolution fMRI of the human visual cortex (Sengupta et al., 2016). Images with anatomical contrast were acquired using the 3D MP2RAGE sequence (TR = 5000 ms; T11/T12 = 900/2750 ms; TE = 2.55 ms; FA1/FA2 = 5°/3°; FOV = 230 x 230 mm²; matrix size = 384 x 384; slices = 240; voxel resolution = 0.6 mm isotropic; GRAPPA factor = 2, partial Fourier = 6/8; phase encoding direction anterior - posterior) (Marques et al., 2010). T2*-weighted functional data were acquired using a 2D gradient echo (GE) echo planar image (EPI) sequence (TE/TR = 20/2000 ms; in-plane FoV 144 x 144 mm²; matrix size 120 x 120; slices = 35; voxel resolution = 1.2 mm isotropic; nominal FA = 82°; echo spacing = 0.78 ms; GRAPPA factor = 2, partial Fourier = 6/8; phase encoding direction head - foot) using a protocol modified from Feinberg et al., 2010; Moeller et al., 2010; Setsompop et al., 2012.

4.3.6 Structural data preparation

Structural images were processed using FSL FAST (Zhang, Brady, & Smith, 2001), SPM's bias correction (Ashburner & Friston, 2005), ITK-SNAP (Yushkevich et al., 2006), FSL BET (Smith, 2002) and morphological operations (Jones, Oliphant, Peterson, et al., 2007). Where not specified otherwise, default settings were used.

Since structural images were recorded using the visual surface coil, they had high signal close to the coil elements in the posterior part of

the brain and gradually lower signal at anterior positions. We observed that this inhomogeneity affected the quality of the T1 map obtained from the MP2RAGE sequence: We found good quality in posterior parts and increasingly noisy responses and signal drop-out in anterior parts. Since we were only interested in early- and mid-level visual areas up to visual field map TO, and since noisy responses can negatively affect automatic segmentation algorithms, we first generated a mask to exclude voxels outside the occipital lobe.

We created the occipital lobe mask by performing anatomical bias correction on the second inversion image and inputting the resulting image to FSL BET to obtain a first tentative brain mask. This mask often included the most posterior two thirds of the brain, where no signal drop-out was observed. We then used ITK-SNAP's adaptive paint brush mode in combination with a graphics tablet (Intuos Art; Wacom Co. Ltd; Kazo, Saitama, Japan) to restrict this mask to the occipital lobe. An opening morphological operation (erosion followed by dilation) was applied to the mask after manual correction in order to remove high-frequency variations between slices that result from manual, slice-by-slice corrections.

The occipital lobe mask was applied to the T1 image and the masked T1 image was input to FSL FAST to obtain a white matter (WM) definition. This WM definition was used for boundary-based registration between functional and structural images (see below). Using the scalpel tool in ITK-SNAP, we divided the WM definition into two separate definitions for the left and right hemisphere. We dilated these WM definitions by one voxel and used the dilated definitions for surface reconstruction in BrainVoyager. Since voxel size was 1.2 mm isotropic and occipital gray matter (GM) is 2-3 mm thick, this resulted in surface reconstructions along the middle of GM. The surface meshes were smoothed using BrainVoyager's advanced mesh smoothing tools (restricting vertex displacement to 0.5) and inflated (500 fold). The resulting surfaces were used to display polar angle estimates obtained from the independent ROI and HRF data in order to delineate visual field maps.

4.3.7 Functional data preparation

Functional data were processed using BrainVoyager 20.0, SPM 12 (Friston, 2006), FSL 5.0 (Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012) as well as custom code in Python 2.7 (Hunter, 2007; Jones et al., 2007; Van Der Walt, Colbert, & Varoquaux, 2011) and in Matlab R2014a (The Mathworks Inc.; Natick, MA, USA). Where not specified otherwise, default settings were used. Functional images were pre-processed using the following operations: slice scan time and motion correction (SPM), high-pass filtering (FSL) and distortion correction (FSL topup). An overview of the pre-processing steps is provided by Figure S4.1.

For motion and distortion correction, we deviated from default settings. Functional images from each scanning session were motion-corrected using SPM 12 in three steps. First, the first image of runs in a given session was realigned to the first image of the first run of that session. Second, the images within each run were aligned to the first image of the run. Third, to avoid local minima, after these first two steps, a mean of all images in a session was calculated and images of that session were realigned to this mean. Motion correction was limited to voxels inside the brain based on an intensity-thresholded and manually corrected brain mask of the functional images. Note that the results of the three steps were combined into a single transformation that was applied to functional images to minimize interpolation artefacts.

EPI distortions of the functional images were corrected using FSL topup (Andersson, Skare, & Ashburner, 2003; Smith et al., 2004). The pairs of opposite phase encoding images acquired at the beginning of the first session of the main experiment (second scanning session) were input to topup to estimate the susceptibility-induced off-resonance field. The estimated field was then used to correct the distortions for all functional images of all scanning sessions.

To align the functional images acquired in different scanning sessions to each other, we calculated a mean image across time, separately for every scanning session (before EPI-distortion correction). We then ran an affine registration between the mean image of the second scanning session and each of the mean images of the other sessions using FSL FLIRT (Jenkinson, Bannister, Brady, & Smith, 2002; Jenkinson & Smith, 2001) with

12 degrees of freedom. The resulting linear transformation was applied to all images recorded outside the second scanning session. We visually inspected the quality of the resulting alignment between the first run in the second session and all runs in the other sessions. Only then EPI distortion correction was applied to the transformed images, using the off-resonance field estimated in session 2. We applied a mask to all images, excluding voxels that had a mean intensity value below 30 in at least one of the sessions. This excluded voxels at the fringes of the slabs (due to minimal differences in slab placement across the three sessions).

Each run for the low- and high-contrast pRF data was split in 3 parts corresponding to the presented motion conditions. Since during each run (222 volumes) each of the 3 motion conditions were presented, this resulted in segments of $222 / 3 = 74$ volumes. The division point between two motion conditions fell in the middle of a 14 volumes rest block such that carry-over effects from one motion condition to the next were minimal. All time courses were demeaned and averaged across runs. The low- and high-contrast pRF data were averaged separately for every motion condition.

To co-register functional data with the T1 image, we obtained an initial transformation estimate using the scanner's positional information and fine-tuned this estimate using boundary-based registration (Greve & Fischl, 2009) as implemented in BrainVoyager. The result was visually inspected for each subject by overlaying functional and anatomical images. Co-registration was implemented using nearest neighbor interpolation.

4.3.8 pRF modelling

OVERVIEW The pRF mapping data were analyzed as described in Dumoulin and Wandell, 2008; Harvey and Dumoulin, 2011; Kay, Weiner, and Grill-Spector, 2015; Kay, Winawer, Mezer, and Wandell, 2013; Zuiderbaan, Harvey, and Dumoulin, 2012. In brief, voxel time courses were assumed to be generated by an isotropic 2D Gaussian in visual space. Model responses were calculated by transforming the presented pRF stimuli to binary apertures (indicating presence or absence of a stimulus in space) and spatially convolving the apertures with an instance of the 2D Gaussian model. The Gaussian model included parameters for its position (μ_x, μ_y)

and its size (σ ; standard deviation of the 2D Gaussian) in visual space and the predicted model time course varied with assumed parameters (Dumoulin & Wandell, 2008). In what follows, an exact description of the employed models is provided.

EMPLOYED MODELS A voxel time course $y(t)$ was modelled using the predicted fMRI time course $p(t)$ of a given pRF model as predictor in a linear regression model according to

$$y(t) = \beta p(t) + e \quad (4.1)$$

where $p(t)$ is the predicted fMRI response for $t = 1, \dots, n$ with n being the number of acquired volumes, β is a scaling factor that accounts for the unknown units of the fMRI signal, and e is measurement noise (Dumoulin & Wandell, 2008). We solved for β using the ordinary least squares method.

The predicted fMRI time course $p(t)$ for a pRF model was obtained according to

$$p(t) = r(t) * h(t) \quad (4.2)$$

where $r(t)$ is the time course response for a given pRF model. This time course was convolved with an HRF $h(t)$ to account for the hemodynamic delay in the fMRI response. The employed HRF was either the canonical double-gamma function (Friston et al., 1998) or estimated subject- and area-specific functions (see below). The time course response of a pRF model was calculated as linear multiplication between the stimulus and the model pRF

$$r(t) = \sum_{x,y} s(x, y, t) g(x, y | \mu_x, \mu_y, \sigma) \quad (4.3)$$

where $s(x,y,t)$ describes the effective stimuli presented over time t and x and y represent positions in the visual field.

The pRF was modelled as an isotropic 2D Gaussian following

$$g(x, y | \mu_x, \mu_y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-[(x-\mu_x)^2 + (y-\mu_y)^2] / (2\pi\sigma^2)} \quad (4.4)$$

where μ_x is the preferred x-position, μ_y is the preferred y-position, σ is the size of the pRF. The parameters μ_x , μ_y , σ were determined for every voxel by minimizing the sum of squared errors (SSE) between the predicted and the observed fMRI time course (Dumoulin & Wandell, 2008). Note that SSE was calculated on the same data as used for optimizing the parameters.

$$SSE = \sum_t [y(t) - \beta p(t)]^2 \quad (4.5)$$

We obtained an approximation of percent signal change (% signal change) for every voxel by determining the maximum of the predicted time course and dividing by the mean signal time course over time.

$$\% \text{ signal change} = \max_t [\beta p(t)] / \frac{\sum_{t=1}^T y(t)}{T} \quad (4.6)$$

We also employed established extensions of the standard pRF model (Kay et al., 2013; Zuiderbaan et al., 2012). These extensions included a compressive non-linearity as described by Kay et al., 2015; Kay et al., 2013. For this extension the predicted fMRI time course was calculated as

$$p(t) = r(t)^n * h(t) \quad (4.7)$$

Raising the response element-wise to the power of n will act as a compressive nonlinearity if $n < 1$. Another extension was to add a suppressive surround using a difference-of-Gaussians (DoG) model as described by Zuiderbaan et al., 2012. For this extension a voxel time course was modelled as

$$y(t) = \beta_1 p_+(t) + \beta_2 p_-(t) + e \quad (4.8)$$

where $p_+(t)$ describes the response of the excitatory center and $p_-(t)$ the suppressive surround. We required that (i) $\beta_1 > \beta_2$, (ii) $\beta_1 > 0$ and (ii) $\beta_2 < 0$. Since we did not expect that all voxels show a center-surround organization, we used the classical pRF model for all voxels in which the surround model did not yield a lower SSE (equation 4.5). Subtraction of the surround Gaussian changes the effective size of the excitatory center.

Therefore, instead of reporting σ for the pRF size we reported the half width at half maximum (HWHM). For voxels with the classical model, the HWHM maximum could be easily calculated as $HWHM = \sqrt{2 \ln 2} \sigma \approx 1.18\sigma$. For voxels with a suppressive surround, HWHM was approximated empirically. The size of the suppressive surround was calculated as the distance between the point where the pRF reaches its minimum amplitude and the pRF center (Zuiderbaan et al., 2012).

OPTIMISATION Parameters of the pRF model were optimized in a two-step procedure. In an initial step, we created candidate model time courses on a grid and minimized SSE in a brute force search across candidate models (see below for specific parameter choices of the grid). Exhaustively searching a dense grid was achieved in a timely manner using multi threading and optimized Cython code. In the second step, we used non-linear optimization (Nelder-Mead algorithm) with constraints as implemented in the Python package `lmfit` (Jones et al., 2007; Newville et al., 2018). Using this strategy of combining an initial grid search with non-linear optimization mitigates the problem of finding *local* optima.

4.3.9 Analysis of independent ROI and HRF data

The independent ROI and HRF data served three purposes. First, we used them to obtain independent pRF estimates which were used to delineate early- and mid-level visual field maps. Second, we used them to derive subject- and field-map specific HRF. Finally, the data were used for voxel selection in the main experiment.

VISUAL FIELD MAP DELINEATION For this analysis only, time courses were not averaged across runs and we applied Gaussian temporal (kernel width = 2 s) and spatial (kernel width = 1.2 mm) smoothing to the pre-processed EPI images. The spatial smoothing was motivated by the prior assumption that pRF parameters vary smoothly over space, which aids the process of manually delineating visual field maps.

To delineate visual field maps, we estimated a voxel-wise pRF model with a compressive non-linearity. The code implementing this estimation is publicly available (Schneider & Marquardt, 2019). Candidate pRF model time courses were created on the following grid: 64 radial position (0°-12°

eccentricity) \times 64 polar positions ($0-2\pi$) \times 40 pRF sizes ($0.2^\circ-16^\circ$) \times 5 exponents (0.1, 0.2, 0.4, 0.7, 1.0). The resulting pRF parameters were projected onto the inflated, reconstructed hemispheres. This allowed us to delineate visual field maps on every hemisphere by manual drawing following established strategies. In particular, visual field maps V1, V2v, V2d, V3v, V3d were clearly described by reversals in polar angle. A combined map of V3a/b was delineated using criteria described in Larsson and Heeger, 2006; Wandell, Brewer, and Dougherty, 2005 and no effort was made to obtain separate maps for V3a and V3b. To delineate hV4 as well as VO1 and VO2, we followed landmarks and strategies described in Winawer and Witthoft, 2017. Field maps LO1, LO2, TO1 and TO2 were delineated following procedures described in Amano et al., 2009. Additionally, we observed that visual field maps TO1 and TO2 were often marked by a substantial drop in the exponent for the compressive nonlinearity. To increase the number of participating data points, we collapsed the identified visual field maps into V1, V2, V3, hV4 (Winawer, Horiguchi, Sayres, Amano, & Wandell, 2010), V3ab, LO, VO and TO. Resulting delineations, along with the pRF properties that guided the delineations, are shown for one representative hemisphere in Figure S4.2. The combined field maps were projected back into volume space (-1mm to $+1\text{mm}$ from the reconstructed surface) and voxels for a common field map were grouped across the left and right hemisphere.

HRF ESTIMATION To obtain subject- and field-map specific HR functions, we proceeded in two steps (Harvey & Dumoulin, 2011; Harvey & Dumoulin, 2016; Harvey, Fracasso, Petridou, & Dumoulin, 2015). First, we estimated the pRF parameters using the canonical double-gamma function (Friston et al., 1998) as the HRF. For this estimation, we employed the difference-of-Gaussians model with suppressive surround and a brute force search for optimization (64 radial position [$0^\circ-12^\circ$ eccentricity] \times 24 polar positions [$0-2\pi$] \times 40 pRF sizes [$0.2^\circ-16^\circ$] \times 15 surround ratios [1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.8, 2.0, 2.2, 2.4, 2.6, 2.9, 3.2, 3.5, 4.0], where a ratio indicates the size of the surround compared to the size of the center). For every visual field map we then selected voxels that had a coefficient of determination (R^2) of at least 0.4 on the training data. This was done to exclude visually unresponsive voxels from the hrf estimation. In case a

visual field map had more than 500 voxels with R2 greater than 0.4, the 500 voxels with the highest R2 value were selected to keep the computing time practicable.

Second, for selected voxels we kept the estimated pRF parameters fixed and varied the parameters of a double-gamma function with the objective of minimizing the SSE between the predicted and observed voxel responses. Hrf parameters were optimized using a brute force search on a grid of hrf parameters (5 grid positions for the delay of the peak [5.0 - 7.0 s], 10 positions for the delay of the undershoot [16.0 - 20.0 s], 5 positions for dispersion of the peak [0.1 - 1.6], 5 positions for dispersion of the undershoot [0.1 - 1.6] and 16 positions for the peak to undershoot ratio [5.0 to 1.0]). The code implementing this optimization is publicly available as part of the HrfOpt package (Schneider, 2019). To obtain a single field map-specific HRF, we averaged the optimized hrf parameters across all voxels in the field map that were included in the analysis (maximally 500 voxels). The resulting HR parameters that were employed for the analysis of the low- and high-contrast pRF data are shown in Figure S4.3.

VOXEL SELECTION We were interested in selecting only voxels that were visually responsive and, in particular, responsive to the eccentricities employed in the main experiment. We implemented these desiderata by performing a pRF analysis with compressive non-linearity on the independent ROI and HRF data to determine μ_x, μ_y, σ and n for every voxel. We then only selected voxels for the main experiment that had an R2 value of at least 0.4 which served as an indicator of "visual responsiveness". Based on the pRF estimates we also calculated the expected maximal response \bar{r} to the stimuli $s(x,y,t)$ presented in the main experiment according to:

$$\bar{r} = \max_t \left[\sum_{x,y} s(x, y, t) g(x, y | \mu_x, \mu_y, \sigma) \right]^n \quad (4.9)$$

Expected values could range from 0 (unresponsive) to 1 (fully responsive). We only considered voxels with a predicted response of at least 0.5. Finally, we required all voxels to be located in one of the delineated visual field maps.

Three criteria thus determined the selection of voxels for the main experiment: (i) voxels needed to have an R2 value of at least 0.4 on the

independently acquired pRF data, (ii) voxels needed to have an expected response to the stimuli in the main experiment of at least 0.5, and (iii) voxels needed to be located in one of the field maps.

4.3.10 *Analysis of low- and high-contrast pRF data*

We estimated a voxel-wise pRF model with suppressive surround separately for every carrier motion (inward, outward, flicker) and for the low- and high-contrast data. The code implementing this estimation is publicly available as part of the PyPrf feature package (Schneider, Marquardt, Gulban, & Isik, 2018). We constructed candidate pRF model time courses on a grid (64 radial position [0° - 12° eccentricity] \times 24 polar positions [0 - 2π] \times 40 pRF sizes [0.2° - 16°] \times 15 surround ratios [1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.8, 2.0, 2.2, 2.4, 2.6, 2.9, 3.2, 3.5, 4.0]) and minimized SSE in a brute force search. Subsequently, we used non-linear optimization to further adjust the μ_x , μ_y and σ parameters. The ratio of size-of-surround to size-of-center was carried over from the brute force search and kept fixed during non-linear optimization. During both brute-force and non-linear optimization the center of the estimated pRF was constrained to be in the visually stimulated part of the visual field. During all analyses for the main experiment we employed the subject- and area-specific HR functions derived above. In a second, independent analysis we repeated the analysis outlined above but this time without the suppressive surround option.

To compare pRF parameters across motion conditions, voxels in each field map were pooled across subjects. Results per field map were summarized using the median as a measure of central tendency and 68% confidence intervals were obtained by 50,000 fold bootstrapping of medians to summarize variability across voxels.

4.3.11 *Simulation of shifted population position codes*

To investigate the potential effect of changes in pRF position on perceptual position shifts, we implemented a model-based decoding analysis. For this analysis we grouped the pRF parameters found in the independent pRF mapping experiment across all subjects and then applied systematic shifts to the pRF centers in the order that was observed empirically: [0.1, 0.1, 0.1, 0.3, 0.2, 0.6, 0.17, 0.0] across V1, V2, V3, V3a/b, hV4, LO, VO,

TO to simulate inward motion and [-0.1, -0.1, -0.1, -0.3, -0.2, -0.6, -0.17, -0.0] to simulate outward motion. To simulate flicker motion, no shifts were applied. We used the resulting pRF to simulate the responses \bar{r}_v of each voxel that would be produced by the different stimuli shown during the psychophysics experiment. The simulated response patterns were projected into the visual field by calculating the 2D projection M as

$$M = \frac{\sum_{v=1}^V g_v \bar{r}_v}{\sum_{v=1}^V \bar{r}_v} \quad (4.10)$$

where $g_v = g(x, y | \mu_x, \mu_y, \sigma)$ is the 2D Gaussian for a pRF with μ_x , μ_y , and σ , \bar{r}_v is the simulated response and V is the total number of voxels in a given field map. g_v was assumed to be fixed across different motion conditions, reflecting the assumption that retinotopy is not flexible (Liu et al., 2006). Instead \bar{r}_v varied across flicker, inward and outward motion as described above.

This resulted in one visual field projection M per psychophysics stimulus, motion condition and field map. For every projection we calculated the expected eccentricity value $E(Ecc)$ (also referred to as the center of mass; Liu et al., 2006) according to

$$E(Ecc) = \sum_{p=1}^P ecc_p p(ecc)_p \quad (4.11)$$

where P is the total number of pixels in a projection image, ecc_p is the eccentricity value at a given pixel p , and $p(ecc)_p = M_p$ is the normalized value of a projection at a pixel p (normalized such that $\sum_{p=1}^P M_p = 1$).

For a given projection, we also calculated the variance of the expected eccentricity value according to

$$Var(Ecc) = \sum_{p=1}^P (ecc_p - E(Ecc))^2 p(ecc)_p \quad (4.12)$$

Intuitively, when we compare two projections M_1 and M_2 for two stimuli s_1 and s_2 then the stimulus for which the projection has a higher expected eccentricity value should be perceived as higher on a proportion of trials $p > 0.5$. This intuition was captured via a cumulative Gaussian

function (Van Humberg, Putzeys, & Wagemans, 2016) and we calculated the proportion p of trials where the second stimulus is perceived as higher than the first according to

$$p = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx \quad (4.13)$$

with $\mu = 0$, $x = E(Ecc_2) - E(Ecc_1)$

and $\sigma = \frac{1}{\epsilon} \sqrt{Var(Ecc_1) + Var(Ecc_2)}$

where ϵ is an efficiency parameter that was set to 0.1 (Van Humberg et al., 2016). This allowed us to calculate a proportion p for all the stimulus pairs shown in the psychophysics experiment. The same fitting as described in the psychophysics methods was used on these proportions to estimate the PSEs for every motion condition combination and visual field map.

In a second pass, we repeated the analysis above but this time, in addition to assuming shifts in pRF center, we also assumed differences in pRF size between motion conditions. In particular, we assumed an increase in pRF size by [0.08, 0.11, 0.26, 0.51, 0.66, 0.51, 0.61, 0.54] across V1, V2, V3, V3a/b, hV4, LO, VO, TO to simulate inward motion and a decrease of the same order to simulate outward motion. To simulate flicker motion, no changes were applied.

4.4 RESULTS

4.4.1 How do motion and contrast influence perceived position?

Figures 4.2A and B show group-averaged psychometric functions for high- and low-contrast stimuli, respectively. For high-contrast stimuli, the fitted psychometric functions are very similar for the different carrier motion combinations and we do not observe a systematic shift in perceived position. By contrast, for low-contrast stimuli the fitted psychometric functions are clearly separated and the arrangement of the curves is consistent with previous studies that reported a shift of perceived position in the direction of motion (Arnold et al., 2007; Chung et al., 2007; De Valois & De Valois, 1991; Fu et al., 2001; Kwon et al., 2015; McGraw et al., 2002; Nishida & John-

ston, 1999; Ramachandran & Anstis, 1990). This pattern is also reflected in the individual psychometric functions for single-subject data (see Figure S4.4 and S4.5): While there is a systematic perceptual displacement in the direction of motion for low-contrast stimuli, no such pattern is evident for high-contrast stimuli.

Figures 4.2C and D plot the group mean PSEs (and 95% confidence intervals) obtained from psychometric functions that were fitted to each participant and pattern combination individually for high- and low-contrast stimuli. The group mean PSEs confirm the trend that was already apparent from the group-averaged psychometric functions. In the low-contrast condition, mean PSE for the outward-flicker combination was $0.20^\circ \pm 0.14^\circ$ and was significantly different from the veridical position difference (since the 95% CI does not include 0). Reversely, mean PSE for the inward-flicker combination was $-0.21^\circ \pm 0.07^\circ$ and also significantly different from 0. In the high-contrast condition, no mean PSE for any combination was significantly different from 0. Furthermore, the mean PSEs for the outward-flicker as well as the inward-flicker combination were both positive with $0.10^\circ \pm 0.12^\circ$ and $0.07^\circ \pm 0.11^\circ$, respectively. We thus did not find the perceived position to be shifted in any systematic direction for the high-contrast data.

4.4.2 *How do motion and contrast influence pRF properties?*

We found that the effect of carrier motion on pRF properties differed depending on the stimulus contrast and visual field map (Figure 4.3). For high stimulus contrast, we found no substantial changes in eccentricity for areas V1 - V3, replicating the finding by Liu et al., 2006. For higher visual field maps, starting from V3ab, however, we observed higher preferred eccentricity values for outward compared to inward motion - with the highest difference, 0.7° , in field map TO. Flicker motion led to values intermediate to those obtained for inward and outward motion. For pRF sizes we observed a systematic decrease for outward as compared to inward motion across all visual field maps, with increasingly higher differences for higher visual field maps. Beyond V3, these effects were substantial with differences from 0.76° in V3a/b to 2.83° in TO. pRF size values for flicker motion were in-between values for inward and outward motion and often closer to outward motion. Percent signal change was

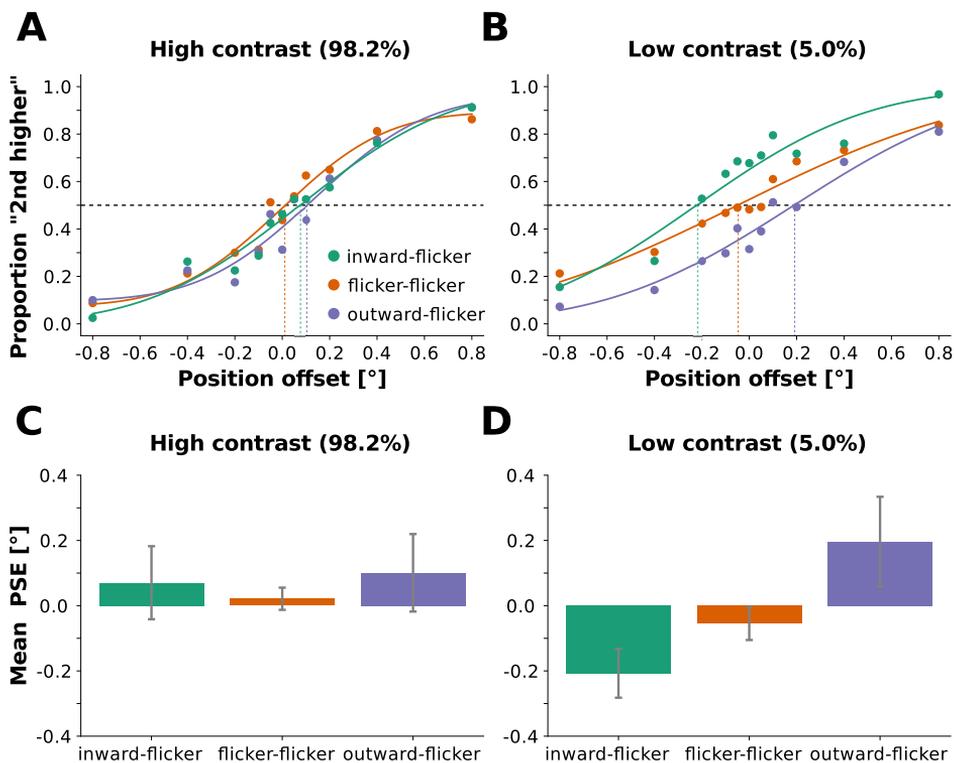


FIGURE 4.2: Overview of psychophysics results.

decreased for outward compared to inward motion, which is consistent with a bias for motion towards fixation that has previously been reported for eccentricities below 14° in humans (Wang et al., 2014).

For low stimulus contrast, we found lower preferred eccentricity values for outward as compared to inward motion across all visual field maps, with the notable exception of field map TO. This corresponds to a shift of the pRF center in the direction opposite to motion, which is consistent with results from the animal literature (Fu, 2004; Sundberg et al., 2006). Eccentricity decreases were in the range from 0.20° to 1.23° and the decrease was very consistent across voxels as indicated by small error bars. As for high-contrast stimuli, we observed a systematic decrease of

FIGURE 4.2 (previous page): Displacement of perceived stimulus position in the direction of motion only observed for low-contrast stimuli. We varied the physical position offset between two subsequently presented stimuli and the carrier pattern of the stimulus (flicker, inward and outward motion). Participants indicated whether they perceived the first or second stimulus as higher. **(A)** Data and fitted psychometric functions for high-contrast stimuli averaged across participants. Dots represent data and curves represent fitted functions. The fitted curves represent responses for the sequence of inward then flicker motion (green), flicker then flicker motion (orange) and outward then flicker motion (violet). The psychometric functions largely overlap. **(B)** Data and fitted psychometric functions for low-contrast stimuli. Same conventions as in **(A)**. Responses for sequences of inward and flicker stimuli (green) were shifted upward, indicating that subjects more frequently perceived the second (flicker) stimulus as higher than the first (inward). By contrast, responses for sequences of outward and flicker stimuli (violet) were shifted downward, indicating that subjects more frequently perceived the second (flicker) stimulus as lower than the first (outward). Flicker-flicker stimulus sequences perceptually assumed an in-between position (orange). **(C)** Point of subjective equality (PSE) for high stimulus contrast. Bars represent average PSE across subjects. Error bars represent 95% confidence interval. We did not find systematic perceptual displacements. **(D)** PSE for low stimulus contrast. Same conventions as in **(C)**. We found no perceptual displacements for flicker-flicker combinations (orange). Outward motion displaced the perceived stimulus position in the direction of motion (violet), increasing the PSE. Reversely, inward motion led to a decrease in perceived position (green).

pRF sizes across all visual field maps for outward as compared to inward motion. These decreases were higher for higher visual areas and in the range from 0.15° to 1.53° . As before, the flicker motion condition usually assumed values in-between the inward and outward motion conditions. As for high-contrast stimuli, percent signal change was decreased for outward compared to inward motion. Percent signal change for flicker motion was much higher than for either inward or outward motion, especially in visual field maps V1, V2, and V3.

The results described above were obtained for fitting the pRF model with the suppressive surround option. This allowed us to compare the effect of carrier motion also on suppressive surround size and the suppression index. Since the number of voxels that had a suppressive surround

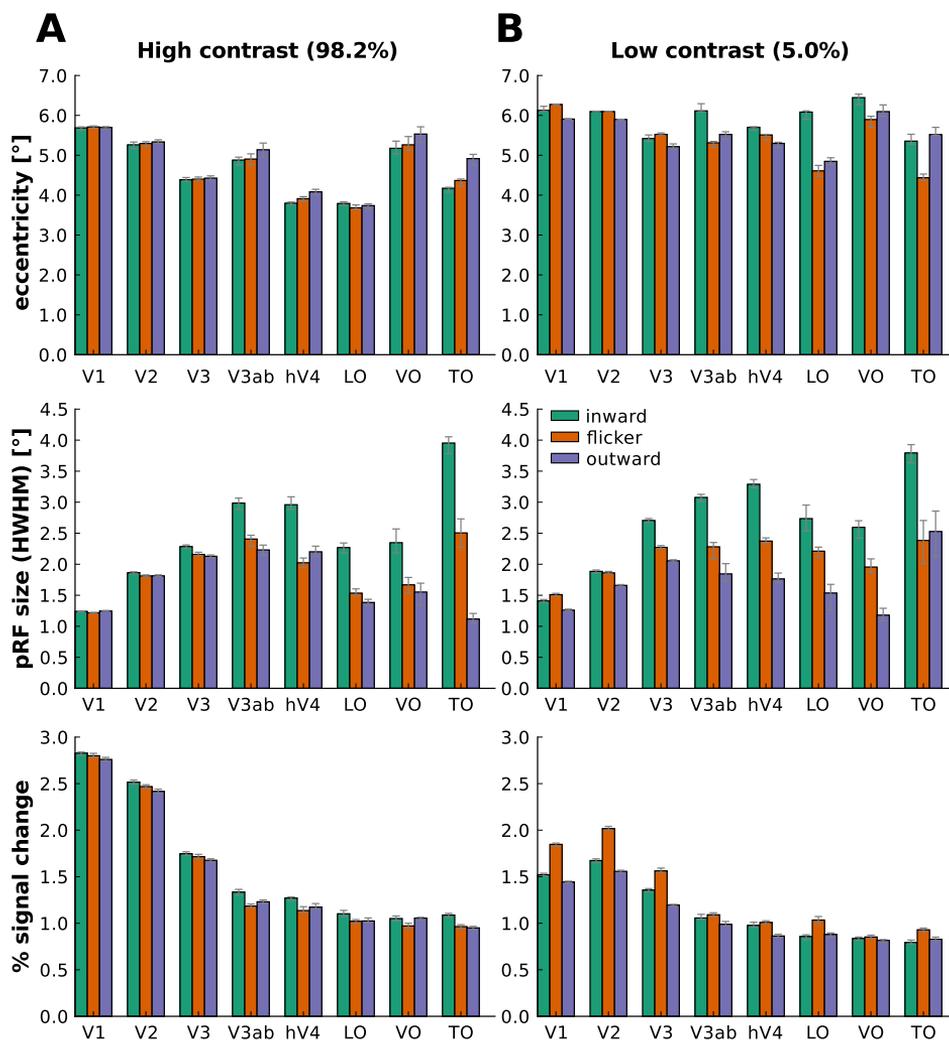


FIGURE 4.3: Effect of carrier motion on pRF properties depends on stimulus contrast and visual field map. Each bar represents a pRF property (median across voxels \pm 68% confidence interval) in different visual field maps (V1 to TO) during one of three motion conditions: inward (green), flicker (orange) and outward (violet) motion. Panels (A) and (B) show results for high- and low-stimulus contrast. Different rows show different pRF properties: eccentricity of the pRF center (in ° of visual angle), pRF size (HWHM, ° of visual angle) and approximated % signal change.

was low for visual field maps outside early visual cortex (less than 25 voxels for all subjects together), we only show the results for V1-V3. As shown in Figure S4.6, we observed larger suppression indices for outward as compared to inward motion. Furthermore, mirroring the observations for pRF size, we found lower surround sizes for outward motion than for inward motion. Finally, when fitting the pRF model *without* suppressive surround, we obtained very similar results as before (Figure S4.7).

4.4.3 Can pRF center shifts account for perceptual shifts?

If the pRF center shifts in the direction opposite to motion, as they do for low-contrast stimuli, how can we account for the observation that, perceptually, the apertures appear displaced in the direction of motion? Figure 4.4 offers an explanation in form of a model for these seemingly contradictory findings. The model makes two essential assumptions: (i) during the MIPS, pRF centers in lower-level visual areas are shifted against the direction of motion; (ii) at the same time, retinotopy is *not* flexible in the sense that connections from lower- to higher-level visual areas are not dynamically updated during the 400 ms of stimulus presentation.

Figure 4.4A and B demonstrate how, given pRF shifts opposite to motion direction, lower-level visual areas send position codes that indicate a shift in the direction of motion. During natural viewing conditions, pRF centers are not displaced and lower-level visual voxels send position codes congruent with the physical position of a stimulus (middle column). However, during inward motion, pRFs will be shifted outward (left column) meaning that units that usually respond to stimuli at an eccentricity of, for example, 5° of visual angle (indicated by "orange" stimulation), will now respond to stimuli at an eccentricity of 5.2° visual angle ("green" stimulation). At the population level, the code that lower-level voxels of different eccentricities will send will signal a position of the stimulus that is slightly more foveal than the actual physical stimulus (see Figure 4.4B). During stimulation with outward motion, reverse conditions and predictions apply (see Figure 4.4A and B, right column).

Figure 4.4C simulates the perceptual consequences of shifted population position codes. We assume the activation of pRF units is connected to (perceived) stimulus position via a visual field projection. The projection

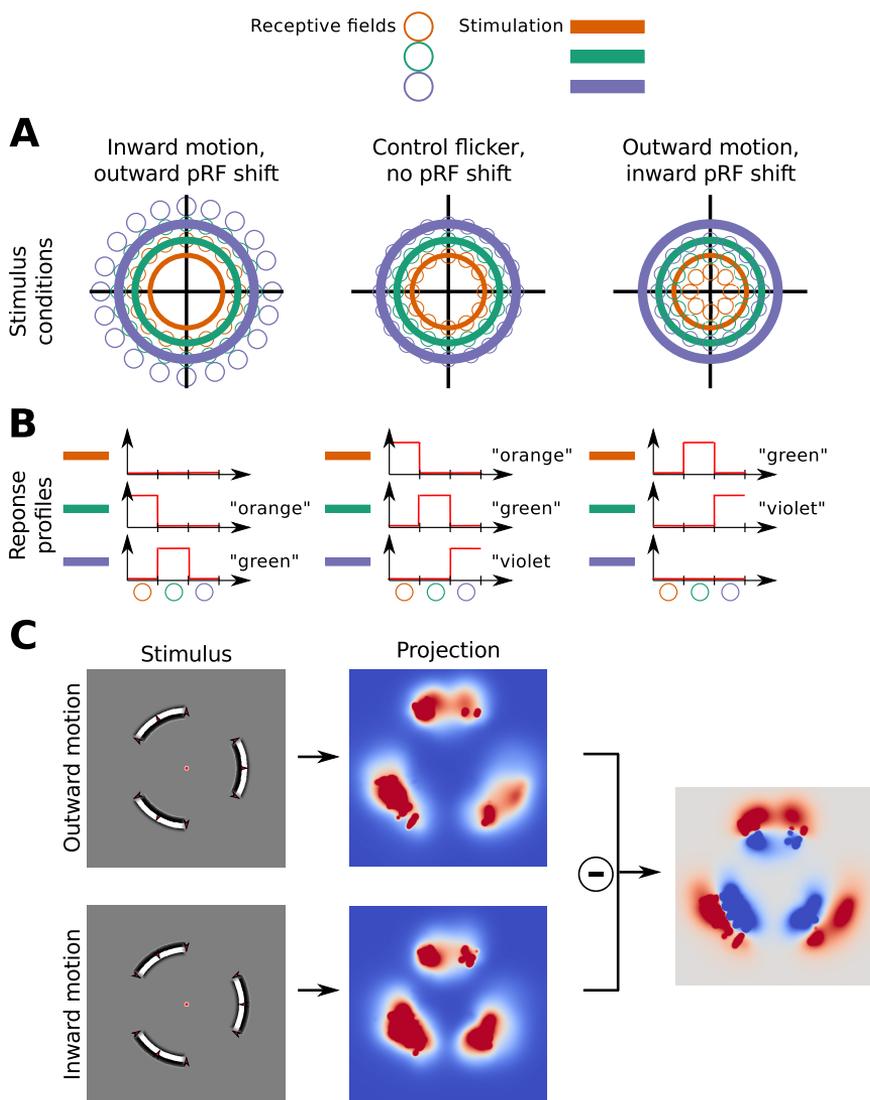


FIGURE 4.4: Simulated perceptual consequences of pRF shifts.

FIGURE 4.4 (previous page): Simulated perceptual consequences of pRF shifts in the direction opposite to motion. **(A)** During natural viewing conditions, pRF centers are not displaced and lower-level visual neurons send position codes congruent with the actual position of a stimulus (middle column). However, if pRF centers are shifted outward during inward motion (left column), neurons that usually respond to stimuli at an eccentricity of, for example, 5° of visual angle (indicated by "orange" stimulation), will now respond to stimuli at an eccentricity of 5.2° visual angle ("green" stimulation). Reversely, if pRF centers are shifted inward during outward motion (right column), they will respond to more foveal stimulation than usual. **(B)** During inward (outward) motion neurons send a population level code that signals a position of the stimulus at slightly lower (higher) eccentricities than its actual position. For example, during peripheral ("violet") stimulation, the neuron population will send a more foveal ("green") code. **(C)** The shifted population level codes observed in **(B)** can be turned into quantifiable, perceptual predictions via visual field projections.

describes the mapping between the activation of a unit and the (perceived) presence of a stimulus at a certain position of the visual field. The mapping is learned during natural viewing conditions and is assumed not to change during inward (outward) motion, because connections between low-level and high-level neurons do not reorganise at the time scale of stimulus presentation. As a consequence, if during inward (outward) motion a population of units in a lower-level visual area signals a more foveal (peripheral) code than usual and connections between low- and high-level areas do not reorganise, then this will be interpreted (perceived) as the presence of a more foveal (peripheral) stimulus. This corresponds to a shift of perceived position in the direction of motion of the stimulus.

We performed a model-based decoding analysis to test whether the mechanism of shifted population position codes can reproduce perceptual shifts in the order that we observed empirically. In the low-contrast condition, we observed empirical pRF center shifts between inward and outward motion in the order of 0.20° to 1.23° across different field maps, while the perceptual shifts between inward and outward motion were in the order of 0.4° . Figure S4.8A shows the results of our simulation. We obtained perceptual shifts in the order of 0.1° to 0.5° depending on the visual field map. For early visual field maps V1-V3 the simulated

perceptual shifts were too small to account for our empirical findings. However, the simulated perceptual shifts obtained in higher visual field maps approximate the perceptual shifts found empirically. We also investigated the effect of pRF size changes on perceptual shifts since we observed larger pRF sizes for inward compared to outward motion. Figure S4.8B shows that if we assume pRF size differences in addition to pRF center shifts, the simulated perceptual shift is enhanced. We discuss the potential contribution of pRF size differences to perceptual shifts below.

4.5 DISCUSSION

We found that coherent motion signals impact both the perceived position of a stimulus as well as the estimated pRF position and identified stimulus contrast as an important moderating factor. Our study replicates findings from a previous fMRI study showing that preferred eccentricity does *not* change with carrier motion presented at high stimulus contrast in early visual areas V1 - V3. However, our study extends earlier findings in four important ways. First, we performed psychophysical measurements and only found MIPS when stimuli were presented at low-contrast but not at high-contrast. Second, we show that at low contrast, where perceptual displacements occur, pRF are shifted in the direction opposite to motion, which is consistent with findings from the animal literature (Fu, 2004). Third, we study the effects of carrier motion on field maps beyond V3. Finally, we find effects of carrier motion on other pRF properties such as pRF size and suppressive surround size. We discuss these extensions in turn.

We observed substantial perceptual displacements only for low-contrast stimuli, not for high-contrast stimuli. Based on this finding, we propose contrast as an important factor that influences the magnitude of perceptual position shifts. This proposal can account for previous findings of larger MIPS with blurred stimulus edges (Fu et al., 2001; Kwon et al., 2015) or increased spatial blurring of the stimulus (Fu et al., 2001) since blurring decreases the effective stimulus contrast. Furthermore, this proposal is consistent with the object-tracking model that was suggested as a unifying framework for position and motion perception (Kwon et al., 2015). Accordingly, the human visual system implements an instance of an

object-tracking system that attempts to infer the generative causes of retinal motion signals (Kwon et al., 2015). Lower stimulus contrast decreases the available sensory evidence and increases signal uncertainty. With increased uncertainty, the object-tracking system is more likely to falsely attribute retinal motion signals to object motion (resulting in the perceived position shift) instead of attributing it to the actual cause, pattern motion (i.e. the drifting square wave grating). Reversely, increased stimulus contrast reduces signal uncertainty, allowing the object-tracking system to correctly attribute retinal motion to pattern motion.

We found that when substantial perceptual displacements occurred, this was accompanied by pRF center shifts in the direction opposite to motion. Based on this finding and a corresponding simulation, we suggest that the perceptual shifts can be accounted for, neuronally, by shifted population position codes in early visual cortex. This result fills an important empirical gap between psychophysical, electro-physiology and fMRI research. Previously, it was unclear whether differences in findings between humans and cats/monkeys were due to differences in recording methods or inter-species differences in the underlying neuronal mechanism. We propose that the neuronal mechanism underlying MIPS is shared across human and non-human animals. The reason that this mechanism had previously not been discovered in humans was that perceptual shifts in response to high-contrast motion and their underlying neuronal mechanisms were probably too small to be measured using neuroimaging methods.

Our work suggests a synthesis of two different accounts that have been used to explain MIPS. Arnold et al. (2007) have termed these two explanations the "RF account" and the "Contrast Modulation account". According to the RF account, motion changes the region of retina to which a population of neurons responds. The Contrast Modulation account suggests that motion influences spatial coding via interactions that modulate apparent contrast (Arnold et al., 2007). These accounts were previously viewed as opposing each other but we propose that they should be considered connected. We demonstrate that perceptual displacements are linked to contrast (Contrast Modulation account). At the same time, we show that these perceptual displacements coincide with shifts of the pRF in the direction opposite to motion (RF account). This implies that any

comprehensive account of motion and position coding in the human visual system should unite both explanations.

Two previous fMRI studies (Fischer, Spotswood, & Whitney, 2011; Maus, Fischer, & Whitney, 2013) found that patterns of fMRI responses in early visual areas reflected physical stimulus position while higher areas like V3a, LO or hMT+ (included in TO) represented the perceived position. At face value, these findings are consistent with the pRF eccentricity changes that we observed for high-contrast stimuli. However, the previous studies used multivariate pattern analysis (MVPA). Therefore, it is difficult to clarify how exactly the results obtained in these studies and our study are related since several kinds of ambiguity are inherent to MVPA (Naselaris & Kay, 2015). In particular, due to geometric ambiguity it is unclear whether dis/similarity of response patterns was based on overall activation alone or also orientation of the multivariate vectors. In addition, there is representational ambiguity which makes it impossible to determine whether position was the source of variation that was driving the observed degree of correlation between response patterns in those studies. By contrast, using a voxel-wise analysis, the pRF model allowed us to explicitly model position in visual space as the source of variation that was driving fMRI responses and avoided the ambiguities surrounding MVPA (Naselaris & Kay, 2015).

Several studies have found that the trailing edge of a motion stimulus evokes larger responses than its leading edge, at least for early visual cortex (V1-V3) (Liu et al., 2006; Schellekens et al., 2016; Wang et al., 2014; Whitney et al., 2003) but not for V5/hMT+ (Wang et al., 2014). Could this trailing edge effect or "aperture-inward" bias account for some of the pRF property changes that we have observed here? We think that several considerations speak against a strong influence of the aperture-inward bias on our results.

First, it is unclear how the aperture-inward bias scales with the size of the aperture. We note that in our experiments, we employed apertures with an extent of 1.7° in the direction of motion, while studies that reported the aperture-inward bias used apertures that were considerably larger (4° - 14° in Schellekens et al., 2016 and 11° - 22° in Wang et al., 2014). Several studies (Schellekens et al., 2015; Schellekens et al., 2016; Wang et al., 2014) have reported that responses decrease gradually with distance from the aperture edge. This decrease appeared to be approximately linear (Schellekens et al.,

2015) at the employed aperture sizes. When employing smaller apertures, as we did here, this effect might scale to the new size of the aperture in which case a linear decrease from trailing to leading edge would be expected. However, it is equally possible that no decrease in responses is observed at all for small apertures since all receptive fields are in proximity to the aperture edge and the resulting difference between trailing and leading edge is minimal. This could be tested in future research by keeping one side of the aperture fixed and systematically increasing the extent of the aperture for different conditions.

Second, our stimulus settings and experimental design should have minimized aperture-inward bias. Moving dot stimuli seem to lead to larger aperture bias (in the order of 1.0 - 1.5 % signal change) (Schellekens et al., 2016; Wang et al., 2014) than drifting grating stimuli (in the order 0.3 - 0.5 % signal change) (Liu et al., 2006; Whitney et al., 2003). This can be attributed to the fact that moving dot stimuli unavoidably show a small asymmetry in dot onsets and offsets at the trailing and leading edge. Drifting grating stimuli, like we used in the current study, are unaffected by this problem and should show lower bias. Likewise, our experimental design and analysis should have reduced the bias. Directly contrasting the effects of opposing motion directions leads to larger effects than contrasting the effects of each motion direction with a mean-luminance background first and only then calculating the difference between the two effects (Liu et al., 2006). This is because the hemodynamic nature of the fMRI signal acts like a compressive nonlinearity (Liu et al., 2006).

Third, all previous studies reporting aperture-inward bias presented only a single stimulus. It is therefore an open question how aperture bias affects pRF mapping experiments with varying stimulus apertures. To understand how the bias might affect the estimation of pRF properties, we ran a simulation that is summarized in Figure 4.5. The results of the simulation are shown in Figure S4.9. Overall, the effects on pRF properties were much smaller than the changes we observed empirically although we simulated bias-induced response changes as much as 100% from leading to trailing edge. We found the bias to shift the pRF center in the direction of motion which is *incompatible* with our empirical results for low-contrast stimuli (but compatible with results for high-contrast stimuli). Furthermore, pRF size decreased for outward as compared to

inward motion (compatible with empirical results), and gain, which is comparable to the signal change we reported, increased (incompatible with empirical results).

For all these reasons, we do not believe that the aperture-inward bias had a large effect on the pRF properties reported here - especially not on the eccentricity values for the low-contrast condition. This does not exclude the possibility, however, that it might have affected other pRF properties such as the pRF size estimation.

For the high-contrast data we observed pRF center shifts in the direction of motion, at least for mid-level visual areas. If we assume that these effects were mostly unaffected by aperture-inward bias, then how can we explain these shifts? One possible explanation is offered by the motion-direction asymmetry (MDA) model (Liu et al., 2006). According to this model, neurons selective to inward motion should always show a spatial response profile that is higher in the inward part than the outward part of the receptive field, independently of which motion direction is presented. Conversely, outward-selective neurons should show higher responses in the outward part of the receptive field. A voxel that is assumed to show an averaged response profile of both inward- and outward-preferring neurons will then show an inward bias in case of inward motion and an outward bias for outward motion. In line with this model, we observed pRF shifts in the direction of motion for V3a/b and TO - areas that are known to contain a high proportion of motion-direction selective cells (Born & Bradley, 2005; Dubner & Zeki, 1971).

If this explanation is correct, why did we not observe similar effects for the low-contrast data? Since we measured data for the same voxels, the same directional biases should have led to similar pRF shifts in the direction of motion for the low-contrast stimuli. Possibly, the effects did occur but were over-shadowed by the pRF shifts in the direction opposite to motion that underlie the perceptual shifts. This would explain why in the low-contrast condition for TO, where the effect of the hypothesized MDA effect was strongest, we no longer observed a shift in the direction opposite to motion. Reversely, in LO where we did not find such an MDA effect, the shift in the direction opposite to motion was very large.

We observed pRF shifts for the low-contrast stimulus in all visual field maps with the exception of TO. This may be surprising as TO, which

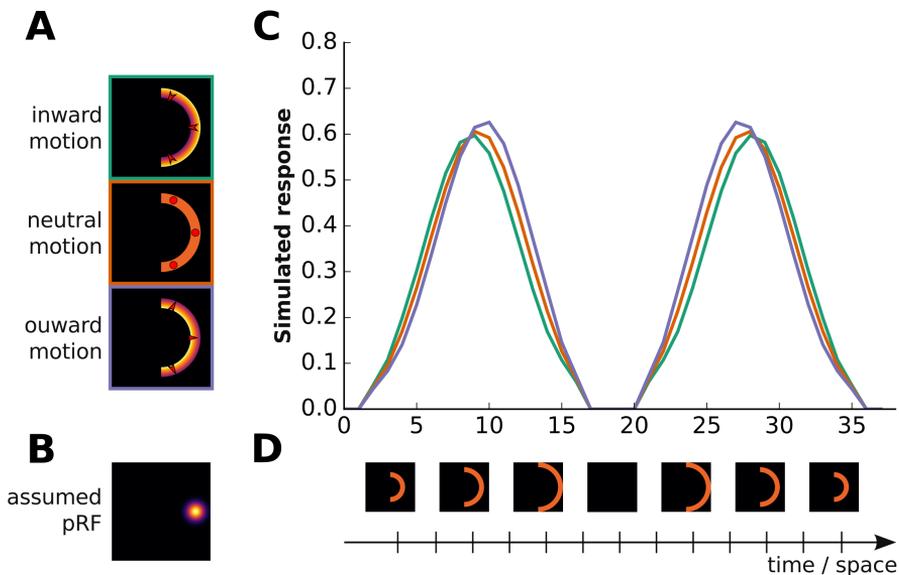


FIGURE 4.5: Simulation of the effects of aperture-inward bias on pRF response. When assuming higher responses to the trailing edge of motion, the pRF center is shifted in the direction of motion. **(A)** We simulated an aperture-inward bias by assuming higher responses to the trailing edge of motion. This implied higher responses to the outer aperture edge for inward motion (green), equal responses across the entire aperture for flicker motion (orange), and higher responses to the inner aperture edge for outward motion (violet). Red arrows indicate motion direction, red dots indicate absence of a particular motion direction. **(B)** We assumed a pRF that was centered on the right horizontal meridian with an eccentricity of 5.95° of visual angle and a pRF size of 1° (shown), 2° or 3° (not shown). **(D)** We assumed that a half-ring aperture first stepped from low (3.4°) to high (8.5°) eccentricity (15 steps, in 30s) and then, after some rest period (8s), stepped back from high to low eccentricities. **(C)** Responses were obtained by spatially convolving the assumed pRF with the different apertures presented in the simulation. Inward (green), flicker (orange) and outward (violet) motion yielded different responses, as shown by the three correspondingly coloured curves.

comprises human V5/hMT+, is most prominently known for its responses to motion and one might thus expect TO to show motion-dependent effects. There are (at least) two possible explanations for why we did not

find pRF shifts in TO. First, it is possible that the MDA effect hypothesized by Liu (2006) also affected the low-contrast data. For high-contrast stimuli, we observed the highest displacement in the direction of motion for area TO. It is thus possible that this effect counteracted the displacement of the pRF against the direction of motion in our measurements, and as a result, we found no eccentricity difference between inward and outward moving stimuli. Second, to link displacements in pRF to position shifts in perception, we assumed above that lower-level visual areas signal displaced position codes that are read-out by hierarchically higher populations. Given that TO is hierarchically the highest visual area included in our study, it is possible that TO represents a higher-order read-out area.

For both high-contrast and low-contrast pRF mapping data we found that inward motion evoked higher signal change than outward motion in early visual areas. This indicates that, at least up to the eccentricities probed in this study (maximum 8.5°), there might be a bias for motion towards the fovea (centripetal motion bias). Several fMRI studies (Maloney, Watson, & Clifford, 2014; Raemaekers, Lankheet, Moorman, Kourtzi, & Van Wezel, 2009; Schellekens, Van Wezel, Petridou, Ramsey, & Raemaekers, 2013) previously attempted to investigate radial motion biases but findings were confounded by aperture-inward bias. Wang et al. (2014) found evidence for a weak centripetal motion bias in V1–V3 although they noted that their measured responses were dominated by the aperture-inward bias. Our results are fully consistent with the findings from Wang et al. (2014) and offer additional empirical evidence for a centripetal motion bias in areas V1 to V3 without the complications of dominant aperture-inward bias.

Some fMRI studies also reported larger responses for motion towards the fovea in area hMT+ (Giaschi, Zwicker, Young, & Bjornson, 2007; Wang et al., 2014), yet other studies did not find a difference in response amplitude between motion towards and away from the fovea (Beckett, Peirce, Sanchez-Panchuelo, Francis, & Schluppeck, 2012; Raemaekers et al., 2009). We observed a centripetal bias in TO for high-contrast stimuli but found no difference between inward and outward motion for low-contrast stimuli. Future studies are therefore needed to further clarify radial biases in TO.

We consistently observed larger pRF size for inward compared to outward motion for both high- and low-contrast stimuli. We believe that several factors contributed to these differences. First, our simulations on the link between perceptual displacements and pRF shifts showed that the observed pRF size difference might contribute to the perceptual shift in the direction of motion. Second, our simulations on the influence of aperture-inward bias on pRF properties revealed that if aperture-inward bias were to affect our findings they would have resulted in pRF size difference in the direction observed here. Finally, Zeidman et al. (2017) noted that for the pRF estimation procedure described in their paper, there was a high positive correlation between the model parameter for pRF size and the parameter for scaling the response. Given that we observed higher responses for inward as compared to outward motion, at least in early visual areas, this might have affected the estimation of pRF sizes. The expected effect is consistent with our results, with higher pRF sizes for inward motion than outward motion.

Although, we found no perceptual displacements at high stimulus contrast, we cannot conclude with certainty that position shifts are absent during high-contrast motion. It is possible that perceptual displacements were simply too small to be picked up by our current study design. For example, Fu et al. (2004) found small displacements in the order of 0.03° to 0.10° for stimuli presented at 100% contrast. This suggests that perceptual displacements probably vary continuously with stimulus contrast and the exact relationship between stimulus contrast and perceptual displacements should be addressed in future studies by studying a wide range of contrast levels.

The choice to use checkerboard flicker in the current study, as opposed to counter-phased gratings, was motivated by two considerations. First, by setting the angular width of the checkerboard cubes appropriately, we could avoid that black and white elements are perceived to be displaced in radial direction for flicker reversals. By contrast, counter-phased gratings can give the impression of rings spontaneously "jumping" in radial direction when contrast reverses. Second, checkerboard flicker corresponds to the carrier pattern that is traditionally chosen for retinotopy (Serenio et al., 1995) and pRF mapping (Dumoulin & Wandell, 2008). It

thus offered a direct reference point to other retinotopy and pRF mapping studies.

The choice of checkerboard flicker came with the downside that spatial frequency in the angular direction could not be matched to the inward and outward conditions. This usually did not have a visible impact on our results since the flicker condition tended both perceptually and in terms of the pRF properties to assume a position in-between the inward and outward conditions. However, there are two observations that might be explained by this design choice. First, at low contrast we observed that flicker elicited much higher responses than either of the other two conditions, at least in early visual areas. Second, at least for low-contrast stimuli, we observed that the slope for psychometric functions in the flicker condition was often flatter than for the other two conditions.

4.6 SUPPLEMENTARY MATERIAL

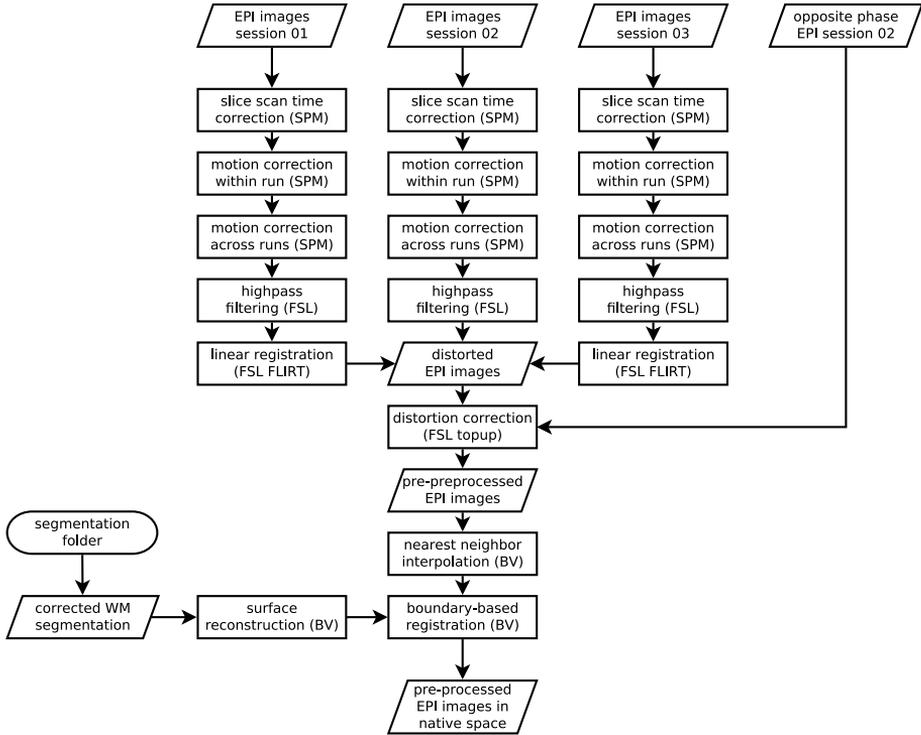


FIGURE S4.1: Overview of pre-processing pipeline for functional images. Rectangular shapes represent processing steps, rhombic shapes represent input or outputs and cylindrical shapes represent input or output locations.

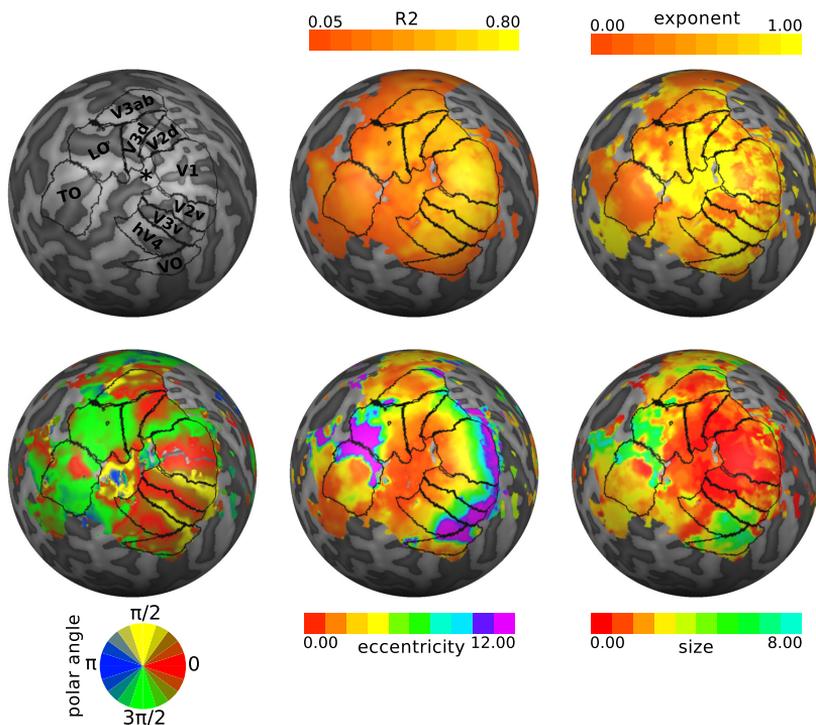


FIGURE S4.2: Overview of pRF mapping results used for delineation of visual field maps. The pRF properties are shown on the left hemisphere of an exemplary subject (S02). For better visualization the hemisphere reconstruction was morphed into a sphere. Dark-gray color indicates sulci and light-gray color indicates gyri.

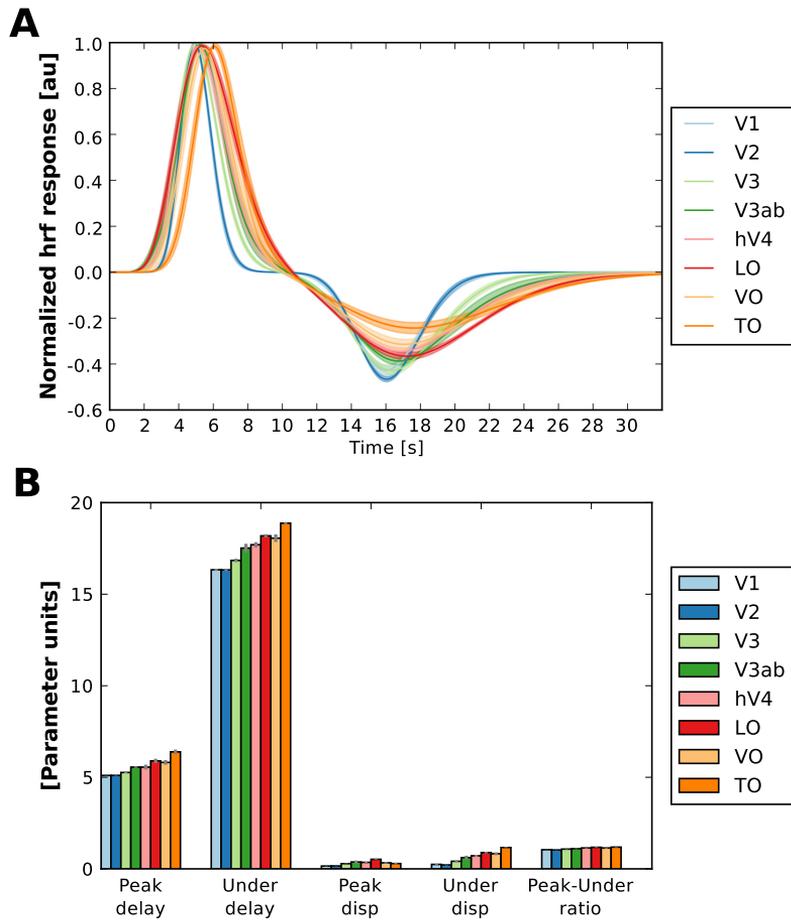


FIGURE S4.3: Fitted hrf parameters and hr functions. **(A)** Fitted hr functions. For visual display and improved comparison only, hr functions were normalized to set the peak for every hrf function to 1. **(B)** Fitted hrf parameters. Shown are the averages across all subjects. Error shading and error bars indicate the standard error of the mean.

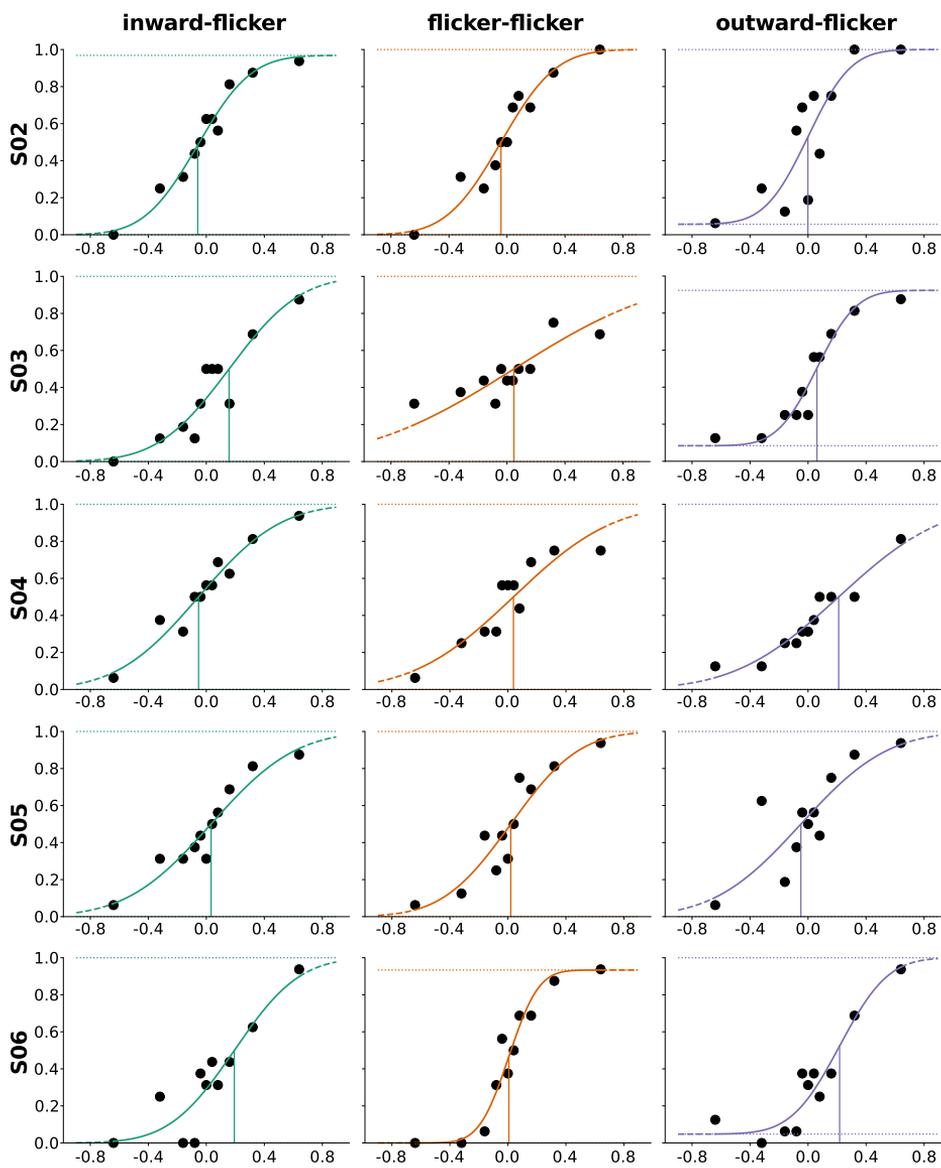


FIGURE S4.4: Single-subject psychometric data and functions for high contrast stimuli.

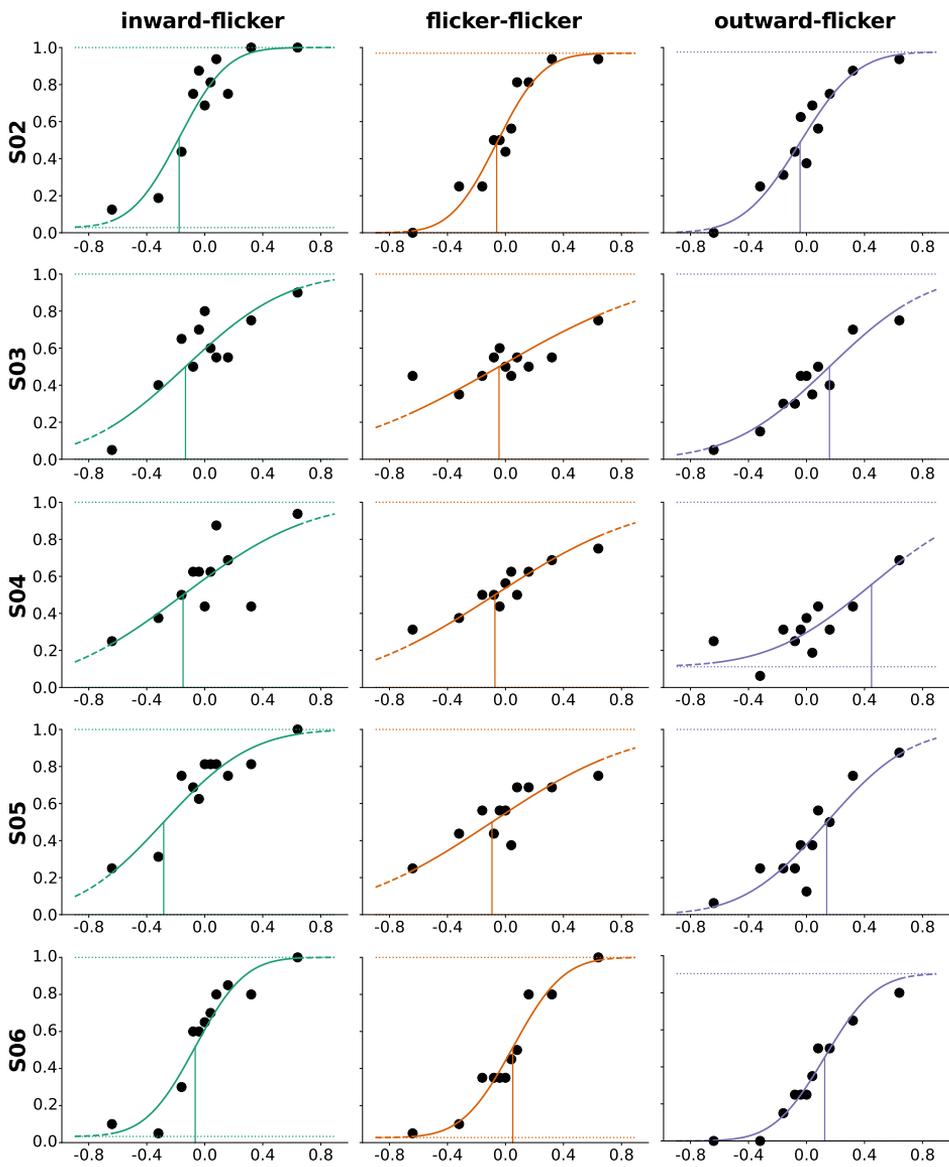


FIGURE S4.5: Single-subject psychometric data and functions for low contrast stimuli.

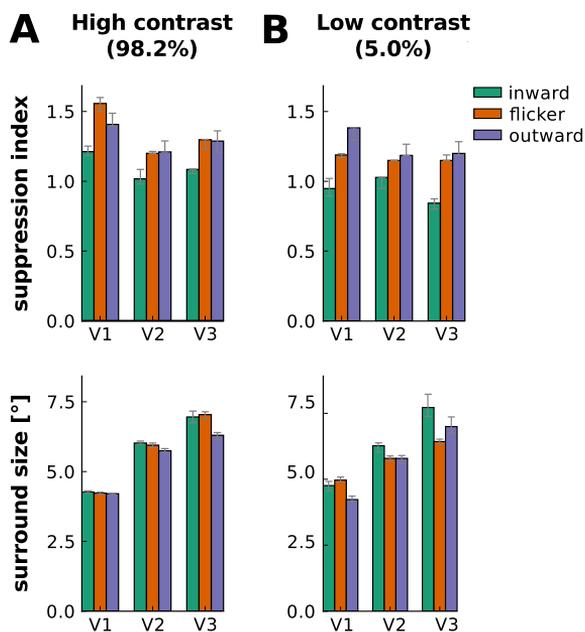


FIGURE S4.6: Effect of carrier motion on suppression index and suppressive surround size in early visual cortex. Same conventions as in Figure 4.3.

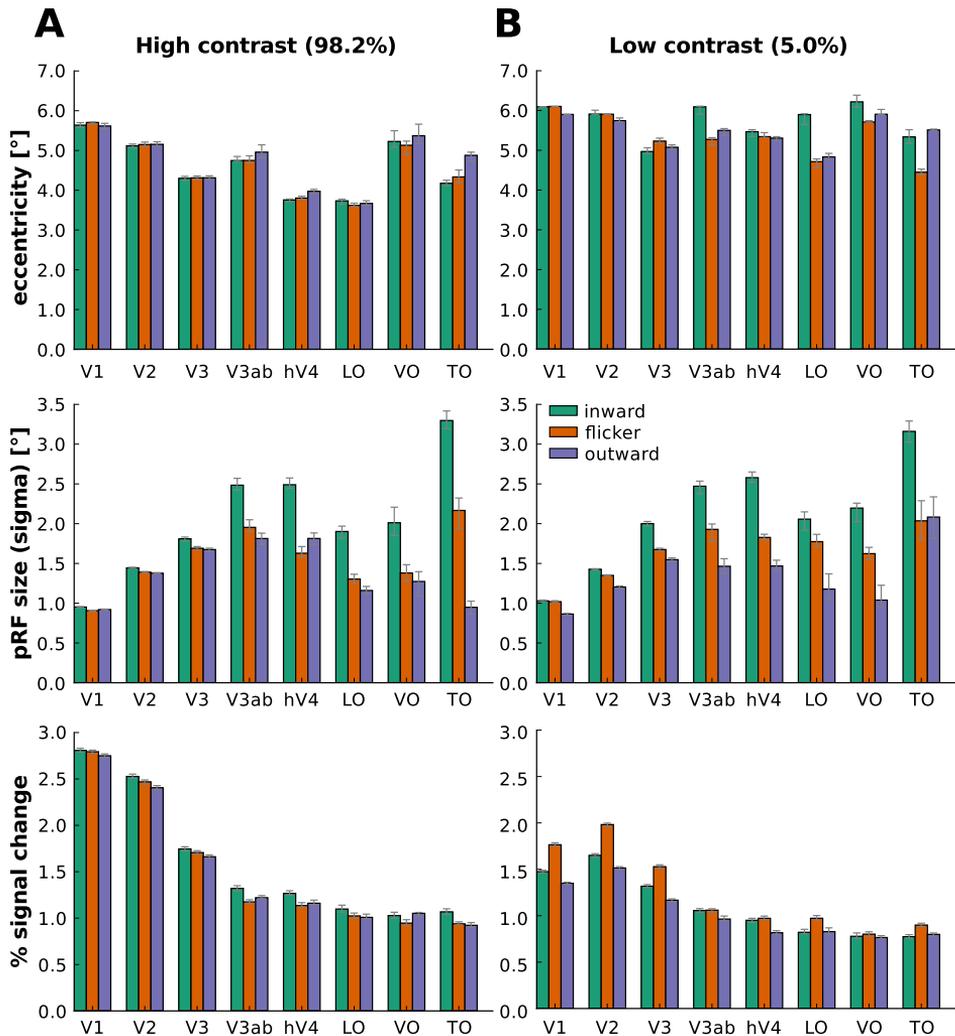


FIGURE S4.7: Effect of carrier motion on pRF properties when classical pRF model without suppressive surround is used. Same conventions as in Figure 4.3 with the difference that for pRF size the σ parameter of the 2D Gaussian is reported (in $^{\circ}$ of visual angle).

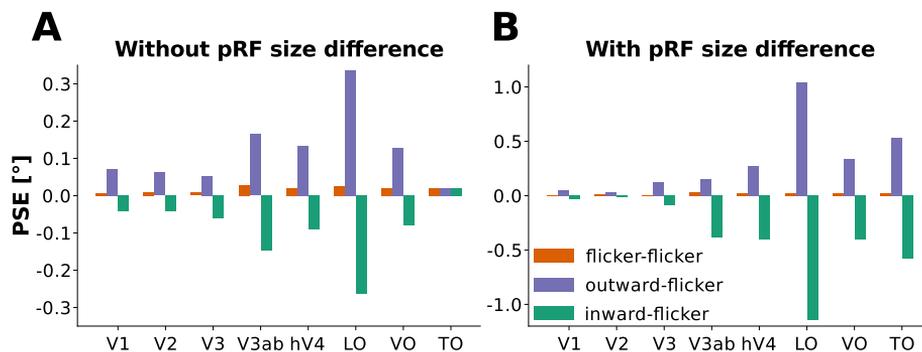


FIGURE S4.8: Simulated perceptual consequences of pRF shifts opposite to motion direction. A pRF center shift in the direction opposite to motion leads perceptual shifts in the direction of motion. **(A)** Point of subjective equality (PSE) for simulated psychophysics data. The three bar conditions represent simulated responses for the sequence of inward then flicker motion (green), flicker then flicker motion (orange) and outward then flicker motion (violet). We assumed that the eccentricity of the pRF shifted with the same magnitude that was observed empirically (see Figure 4.3, upper row). The pRF size was assumed to remain fixed. **(B)** PSE for simulated psychophysics data. Same conventions as in **(A)** but this time, in addition to shifting the pRF center, we also changed the pRF size in accordance with the pRF size difference that we observed empirically (see Figure 4.3, middle row). Note the different scale between the y-axes in **(A)** and **(B)**.

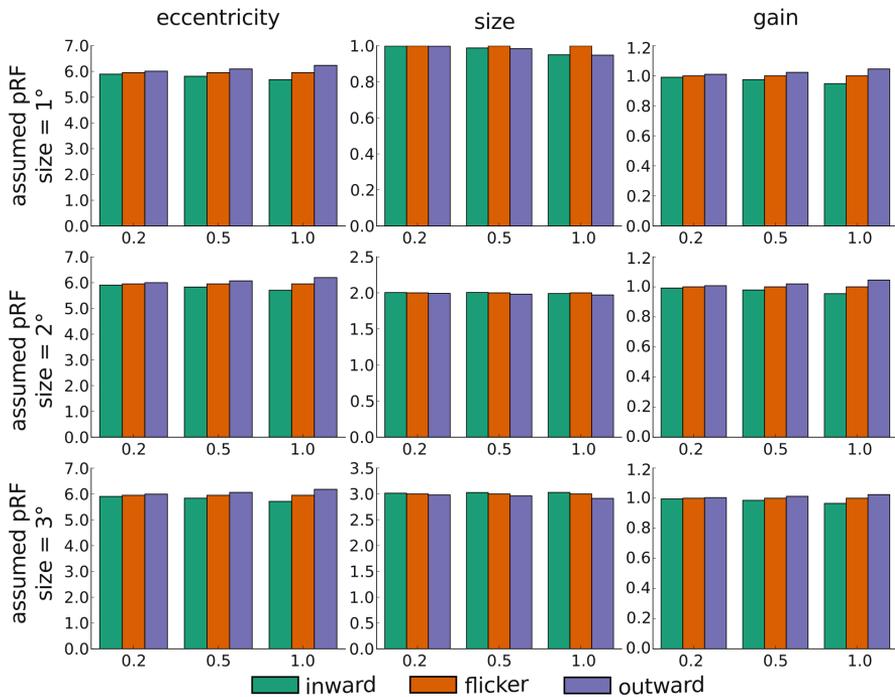


FIGURE S4.9: Simulated effects of aperture-inward bias on pRF response. We assumed a pRF that was centered on the right horizontal meridian with an eccentricity of 5.95° of visual angle. We varied the pRF size assuming either 1° , 2° or 3° , as shown in the different rows. The different columns show the effects of aperture-inward bias on different pRF properties: eccentricity, size and gain. The different colors represent the effect of inward motion (green), flicker motion (orange) or outward motion (violet). Different sections on the x-axis of each panel reflect the steepness of the modelled linear increase/decrease in response resulting from the aperture. The response was simulated to be 1 at the center of the aperture and we simulated an increase/decrease in response of either 0.2, 0.4, or 1.0 from aperture center to aperture edges.

4.7 ACKNOWLEDGMENTS

We would like to thank Dimo Ivanov for help with adjusting the MP2RAGE sequence and Mario Senden and Brian Wandell for valuable advice.

4.8 REFERENCES

- Alvarez, I., de Haas, B., Clark, C. A., Rees, G., & Schwarzkopf, D. S. (2015). Comparing different stimulus configurations for population receptive field mapping in human fMRI. *Frontiers in Human Neuroscience*, *9*, 1–16.
- Amano, K., Wandell, B. A., & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, *102*(5), 2704–18.
- Andersson, J. L., Skare, S., & Ashburner, J. (2003). How to correct susceptibility distortions in spin-echo echo-planar images: Application to diffusion tensor imaging. *NeuroImage*, *20*(2), 870–888.
- Arnold, D. H., Thompson, M., & Johnston, A. (2007). Motion and position coding. *Vision Research*, *47*(18), 2403–2410.
- Ashburner, J. & Friston, K. J. (2005). Unified segmentation. *NeuroImage*, *26*(3), 839–851.
- Beckett, A., Peirce, J. W., Sanchez-Panchuelo, R. M., Francis, S., & Schluppeck, D. (2012). Contribution of large scale biases in decoding of direction-of-motion from high-resolution fMRI data in human early visual cortex. *NeuroImage*, *63*(3), 1623–1632.
- Born, R. T. & Bradley, D. C. (2005). Structure and Function of Visual Area Mt. *Annual Review of Neuroscience*, *28*(1), 157–189.
- Chung, S. T. L., Patel, S. S., Bedell, H. E., & Yilmaz, O. (2007). Spatial and temporal properties of the illusory motion-induced position shift for drifting stimuli. *Vision Research*, *47*(2), 231–243.
- De Valois, R. L. & De Valois, K. K. (1991). Vernier acuity with stationary moving Gabors. *Vision Research*, *31*(9), 1619–1626.
- Dubner, R. & Zeki, S. M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*.
- Dumoulin, S. O. & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *NeuroImage*, *39*(2), 647–660.
- Eagleman, D. M., Eagleman, D. M., Sejnowski, T. J., & Sejnowski, T. J. (2007). Motion signals bias localization judgments: A unified explanation for the ash-lag, ash-drag, ash-jump, and Frohlich illusions. *Neurobiology*, *7*, 1–12.

- Edwards, M. & Badcock, D. R. (2003). Motion distorts perceived depth. *Vision Research*, 43(17), 1799–1804.
- Feinberg, D. A., Moeller, S., Smith, S. M., Auerbach, E., Ramanna, S., Glasser, M. F., . . . Yacoub, E. (2010). Multiplexed echo planar imaging for sub-second whole brain fmri and fast diffusion imaging. *Plos One*.
- Fischer, J., Spotswood, N., & Whitney, D. (2011). The emergence of perceived position in the visual system. *Journal of Cognitive Neuroscience*, 23(1), 119–136.
- Friston, K. J., Fletcher, P., Josephs, O., Holmes, a., Rugg, M. D., & Turner, R. (1998). Event-related fMRI: characterizing differential responses. *NeuroImage*, 7(1), 30–40.
- Friston, K. J. (2006). *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Academic Press.
- Frund, I., Haenel, N. V., & Wichmann, F. A. (2011). Inference for psychometric functions in the presence of nonstationary behavior. *Journal of Vision*, 11(6), 16–16.
- Fu, Y. X., Shen, Y., & Dan, Y. (2001). Motion-induced perceptual extrapolation of blurred visual targets. *Journal of Neuroscience*, 21(20), RC172.
- Fu, Y.-X. (2004). Asymmetry in Visual Cortical Circuits Underlying Motion-Induced Perceptual Mislocalization. *Journal of Neuroscience*, 24(9), 2165–2171.
- Giaschi, D., Zwicker, A., Young, S. A., & Bjornson, B. (2007). The role of cortical area V5/MT+ in speed-tuned directional anisotropies in global motion perception. *Vision Research*, 47(7), 887–898.
- Greve, D. N. & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration. *NeuroImage*, 48(1), 63–72.
- Harvey, B. M. & Dumoulin, S. O. (2011). The Relationship between Cortical Magnification Factor and Population Receptive Field Size in Human Visual Cortex: Constancies in Cortical Architecture. *Journal of Neuroscience*, 31(38), 13604–13612.
- Harvey, B. M. & Dumoulin, S. O. (2016). Visual motion transforms visual space representations similarly throughout the human visual hierarchy. *NeuroImage*, 127, 173–185.
- Harvey, B. M., Fracasso, A., Petridou, N., & Dumoulin, S. O. (2015). Topographic representations of object size and relationships with numerosity reveal generalized quantity processing in human parietal

-
- cortex. *Proceedings of the National Academy of Sciences*, 112(44), 13525–30.
- Hunter, J. D. (2007). Matplotlib: a 2d graphics environment. *Computing In Science & Engineering*, 9(3), 90–95.
- Jenkinson, M., Bannister, P., Brady, M., & Smith, S. (2002). Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2), 825–841.
- Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, 62(2), 782–790.
- Jenkinson, M. & Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5(2), 143–156.
- Jones, E., Oliphant, T., Peterson, P., et al. (2007). SciPy: open source scientific tools for Python. *Computing in Science and Engineering*.
- Kay, K. N., Weiner, K. S., & Grill-Spector, K. (2015). Attention reduces spatial uncertainty in human ventral temporal cortex. *Current Biology*, 25(5), 595–600.
- Kay, K. N., Winawer, J., Mezer, A., & Wandell, B. A. (2013). Compressive spatial summation in human visual cortex. *Journal of Neurophysiology*, 110(2), 481–94.
- Kwon, O.-S., Tadin, D., & Knill, D. C. (2015). Unifying account of visual motion and position perception. *Proceedings of the National Academy of Sciences*, 112(26), 8142–8147.
- Larsson, J. & Heeger, D. J. (2006). Two Retinotopic Visual Areas in Human Lateral Occipital Cortex. *Journal of Neuroscience*.
- Liu, J. V., Ashida, H., Smith, A. T., & Wandell, B. A. (2006). Assessment of stimulus-induced changes in human V1 visual field maps. *Journal of Neurophysiology*, 96(6), 3398–3408.
- Maloney, R. T., Watson, T. L., & Clifford, C. W. (2014). Determinants of motion response anisotropies in human early visual cortex: The role of configuration and eccentricity. *NeuroImage*, 100, 564–579.
- Marques, J. P., Kober, T., Krueger, G., van der Zwaag, W., Van de Moortele, P. F., & Gruetter, R. (2010). MP2RAGE, a self bias-field corrected sequence for improved segmentation and T1-mapping at high field. *NeuroImage*, 49(2), 1271–1281.

- Maus, G. W., Fischer, J., & Whitney, D. (2013). Motion-dependent representation of space in area MT+. *Neuron*, 78(3), 554–562.
- McGraw, P. V., Whitaker, D., Skillen, J., & Chung, S. T. L. (2002). Motion adaptation distorts perceived visual position. *Current Biology*, 12(23), 2042–2047.
- Moeller, S., Yacoub, E., Olman, C. A., Auerbach, E., Strupp, J., Harel, N., & Uğurbil, K. (2010). Multiband multislice GE-EPI at 7 tesla, with 16-fold acceleration using partial parallel imaging with application to high spatial and temporal whole-brain fMRI. *Magnetic Resonance in Medicine*.
- Naselaris, T. & Kay, K. N. (2015). Resolving Ambiguities of MVPA Using Explicit Models of Representation. *Trends in Cognitive Sciences*, 19(10), 551–554.
- Newville, M., Otten, R., Nelson, A., Ingargiola, A., Stensitzki, T., Allan, D., ... Almarza, A. (2018). Lmfit. non-linear least-square minimization and curve-fitting for python.
- Nishida, S. & Johnston, A. (1999). Influence of motion signals on the perceived position of spatial pattern. *Nature*, 397(6720), 610–612.
- Peirce, J. W. (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.
- Raemaekers, M., Lankheet, M. J., Moorman, S., Kourtzi, Z., & Van Wezel, R. J. (2009). Directional anisotropy of motion responses in retinotopic cortex. *Human Brain Mapping*, 30(12), 3970–3980.
- Ramachandran, V. S. & Anstis, S. M. (1990). Illusory displacement of equi-luminous kinetic edges. *Perception*, 1(2).
- Schellart, N. A., Trindade, M. J., Reits, D., Verbunt, J. P., & Spekreijse, H. (2004). Temporal and spatial congruence of components of motion-onset evoked responses investigated by whole-head magneto-electroencephalography. *Vision Research*, 44(2), 119–134.
- Schellekens, W., Ramsey, N. F., & Raemaekers, M. (2015). Predictions to motion stimuli in human early visual cortex: Effects of motion displacement on motion predictability. *NeuroImage*, 118, 118–125.
- Schellekens, W., Van Wezel, R. J., Petridou, N., Ramsey, N. F., & Raemaekers, M. (2013). Integration of Motion Responses Underlying Direc-

-
- tional Motion Anisotropy in Human Early Visual Cortical Areas. *Plos One*, 8(6).
- Schellekens, W., van Wezel, R. J., Petridou, N., Ramsey, N. F., & Raemaekers, M. (2016). Predictive coding for motion stimuli in human early visual cortex. *Brain Structure and Function*, 221(2), 879–890.
- Schneider, M. (2019). Hrfopt - hrf optimization with cython. *Zenodo repository*.
- Schneider, M. & Marquardt, I. (2019). Pyprf motion - population receptive field estimation with static nonlinearity. *Zenodo repository*.
- Schneider, M., Marquardt, I., Gulban, O. F., & Isik, I. (2018). Pyprf feature - population receptive field estimation with feature weights. *Zenodo repository*.
- Schwarzkopf, D. S., Anderson, E. J., de Haas, B., White, S. J., & Rees, G. (2014). Larger Extrastriate Population Receptive Fields in Autism Spectrum Disorders. *Journal of Neuroscience*, 34(7), 2713–2724.
- Sengupta, S., Roebroek, A., Kemper, V. G., Poser, B. A., Zimmermann, J., Goebel, R., & Adriany, G. (2016). A specialized multi-transmit head coil for high resolution fMRI of the human visual cortex at 7T. *Plos One*, 11(12), 1–22.
- Sereno, M. I., Dale, A. M., Reppas, J. B., Kwong, K. K., Belliveau, J. W., Brady, T. J., . . . Tootell, R. B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*.
- Setsompop, K., Gagoski, B. A., Polimeni, J. R., Witzel, T., Wedeen, V. J., & Wald, L. L. (2012). Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magnetic Resonance in Medicine*.
- Smith, S., Jenkinson, M., Woolrich, M., Beckmann, C., Behrens, T., Johansen-Berg, H., . . . Matthews, P. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*.
- Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3), 143–155.
- Sundberg, K. A., Fallah, M., & Reynolds, J. H. (2006). A motion-dependent distortion of retinotopy in area V4. *Neuron*, 49(3), 447–457.
- Tyler, C. W., Likova, L. T., Chen, C.-C., Kontsevich, L. L., Schira, M. M., & Wade, A. R. (2005). Extended Concepts of Occipital Retinotopy. *Current Medical Imaging Reviews*, 1, 319–329.

- Uğurbil, K., Adriany, G., Andersen, P., Chen, W., Garwood, M., Gruetter, R., ... Zhu, X. H. (2003). Ultrahigh field magnetic resonance imaging and spectroscopy. *Magnetic Resonance Imaging*.
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. (2011). The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2), 22–30.
- Van Humbeeck, N., Putzeys, T., & Wagemans, J. (2016). Apparent Motion Suppresses Responses in Early Visual Cortex: A Population Code Model. *Plos Computational Biology*, 12(10), 1–25.
- Wandell, B. A., Brewer, A. A., & Dougherty, R. F. (2005). Visual field map clusters in human cortex. *Philosophical Transactions of the Royal Society*.
- Wang, H. X., Merriam, E. P., Freeman, J., & Heeger, D. J. (2014). Motion Direction Biases and Decoding in Human Visual Cortex. *Journal of Neuroscience*, 34(37), 12601–12615.
- Whitney, D., Goltz, H., Thomas, C., & Goodale, M. (2003). Flexible retinotopy: Motion dependent position coding in visual cortex. *Journal of Vision*, 3(9), 111a.
- Winawer, J., Horiguchi, H., Sayres, R. A., Amano, K., & Wandell, B. A. (2010). Mapping hV4 and ventral occipital cortex: The venous eclipse. *Journal of Vision*, 10(5), 1–1.
- Winawer, J. & Witthoft, N. (2017). Identification of the ventral occipital visual field maps in the human brain. *F1000Research*, 6, 1526.
- Yushkevich, P. A., Piven, J., Cody Hazlett, H., Gimpel Smith, R., Ho, S., Gee, J. C., & Gerig, G. (2006). User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *NeuroImage*, 31(3), 1116–1128.
- Zeidman, P., Silson, E. H., Schwarzkopf, D. S., Baker, C. I., & Penny, W. (2017). Bayesian population receptive field modelling. *NeuroImage*, 180, 1–15.
- Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*.
- Zuiderbaan, W., Harvey, B. M., & Dumoulin, S. O. (2012). Modeling center – surround configurations in population receptive fields using fMRI. *Journal of Vision*, 12(3), 1–15.

5

General Discussion

5.1 OVERVIEW

The goal of this thesis was to leverage the new opportunities in functional magnetic resonance imaging (fMRI) research afforded by higher field strength to study the neural correlates of visual conscious content in humans (content-specific NCC). In order to achieve this goal, we first saw the need to develop a tool that aids tissue class segmentation of ultra-high field (UHF) images (Chapter 2). We then performed two high-field high-resolution fMRI experiments to clarify how representations in topographically organized cortical structures are related to the perception of apparent motion (Chapter 3) and apparent position (Chapter 4). After summarizing the results of each study, the presented work will be contextualized by relating it to the conceptual framework proposed by De Graaf, Hsieh, and Sack, 2012 and by highlighting common considerations that motivated and guided the experimental design of all studies. Subsequently, limitations of the presented studies will be discussed and taken as motivation to outline a new vision for UHF fMRI and consciousness research.

5.2 SUMMARY

Neuroimaging at UHF offers increased sensitivity and specificity. Yet these advantages can only take full effect if certain challenges are appropriately addressed - including co-registration of anatomical and functional images with partial coverage, geometric distortions, registration of functional images across scanning sessions and tissue class segmentation. While there are solutions or adequate workarounds for most of these issues, we found tissue class segmentation at UHF to constitute such a severe bottleneck that we decided to tackle this issue first in the form of a new tool. In Chapter 2, we proposed that 2D histograms offer a suitable representation of 3D anatomical images that allows for efficient editing of segmentation images. In particular, motivated by theoretical expectations, we expected non-brain voxels to become separable from brain voxels in the 2D space represented by image intensity and gradient magnitude. Based on previous results, we suggested that separation could either be obtained using appropriate widgets with expert knowledge or using more automated techniques such

as normalized graph cuts. We then created our own, openly accessible implementation of the proposed methods and two annotated 7 Tesla MRI data sets for validation. We demonstrated that our method improves cortical grey matter definitions when used as an additional post-processing step to existing segmentation algorithms and frees a considerable amount of researcher time that would otherwise need to be spent on manual corrections.

Equipped with this new tool, we went on, in Chapter 3, to clarify how responses in columnar clusters of the human motion complex (hMT+) relate to the conscious experience of a specific visual motion axis (hereafter called "motion quartet study"). Leveraging the increased sensitivity of UHF, we first mapped out distinct, fine-grained clusters that either preferred horizontal or vertical motion. We then measured their response modulations during ambiguous stimulation with the bistable motion quartet. We found that when participants indicated the conscious percept of horizontal motion the response amplitude increased in the "horizontal" cluster and decreased in the "vertical" cluster, while we found the reverse pattern during vertical motion. Finally, we also demonstrated that these clusters were organized in a columnar fashion such that preferences for vertical or horizontal motion were stable in the direction of cortical depth and changed when moving along the cortical surface. These results support the idea that hMT+ makes up part of the content-specific NCC for visual motion and, more importantly, extend this idea by demonstrating involvement of specific horizontal and vertical columnar clusters in tracking conscious experience of a particular motion axis.

While this study imposed important constraints on the location of content-specific NCC, it did not offer a strong constraint on mechanisms that link cortical activity to conscious visual perception. In Chapter 4, we therefore explored the ability of the population receptive field (pRF) mapping paradigm to connect responses in early and mid-level visual cortex to the perceived position of visual stimuli (hereafter called "motion-dependent pRF study"). We first identified stimulus contrast as an important factor that influences the strength of motion-induced position shifts. We observed that at high stimulus contrast, when sufficient sensory evidence is available, little to no perceptual displacements occur, while at low stimulus contrast, when signal uncertainty is increased, substantial per-

ceptual displacements occur. Correspondingly, we found that pRF centers were substantially shifted against the direction of motion for low-contrast, but not high-contrast, stimuli. Based on these observations, we proposed a model that links pRF estimates to perceived stimulus position. This model is based on the idea that early visual areas send shifted population codes in response to motion stimuli that bias the conscious percept of stimulus position in the direction of motion.

5.3 PAST AND CURRENT CONTEXT

In this section, we will first discuss how the obtained results can be interpreted and what they imply for the search of content-specific NCC. We will then highlight two common threads that run through the presented research. Specifically, we will argue that the studies can be understood as a reaction to both the earlier introduction of multivariate pattern analysis (MVPA) techniques and the more recent debate about reproducibility.

5.3.1 *Prerequisites, substrates and consequences*

De Graaf and colleagues (2012) proposed a conceptual distinction for NCC that offers a framework for interpreting the results obtained in this thesis. The authors pointed out that an NCC may fall into one of three distinct categories: It may be a prerequisite, a consequence, or the actual substrate of a conscious experience. Neural prerequisites are neural events or processes that are necessary for a particular conscious experience; yet they are not themselves the "neural instantiation" of that conscious experience (De Graaf et al., 2012). Examples of neural prerequisites are neural events associated with wakefulness, arousal and attention. If these events did not occur, a conscious experience would not occur. By contrast, neural consequences are the result of neural events underlying the conscious experience. There may be content-specific consequences (when the conscious experience of a rose gives rise to associative responses such as the memory of our grandmother who grew roses in her garden) and content-invariant consequences (neural events that follow any kind of conscious experience) (De Graaf et al., 2012). Thus, although prerequisites and consequences by definition always co-occur with a conscious experience, they do not

directly underlie its subjective phenomenal content - as the neural substrate does. Only the neural substrate is the neural state that is directly underlying the experience and is neither a result nor a precondition for a specific conscious percept.

The conceptual distinctions proposed by de Graaf et al. are useful for interpreting the empirical results presented in this thesis. Conversely, our findings have implications for the strategies suggested to dissociate prerequisites, substrates and consequences of conscious experience from each other. With regard to the findings presented in Chapter 3, it is difficult to determine with certainty whether the observed response modulations in the vertical or horizontal clusters are prerequisites, substrates or consequences. Given that the identified NCC are content-specific and applying Occam's razor, however, we propose that the observed response modulations are neural *substrates* for the experience of a specific visual axis of motion. In this interpretation, we benefit from the increased explanatory power that comes with the research approach of content-specific NCC, as opposed to general NCC (see [General Introduction](#)). After all, in an independent analysis step, we first mapped out the representational content for hMT+ and showed that its states represented horizontal and vertical visual motion. One would need to introduce additional assumptions in order to explain that these states also correlated with conscious content of exactly horizontal and vertical visual motion if they are just general prerequisites or consequences. This would not be parsimonious and we thus favour the hypothesis that the columnar clusters identified in Chapter 3 are neural substrates. However, as pointed out by De Graaf et al. (2012) and as shown below, content-specific NCC are not always neural substrates since sometimes neural prerequisites and consequences can also be content-specific.

In the context of the study presented in Chapter 3, it also becomes apparent that one of the strategies suggested by De Graaf et al., 2012 for dissociating neural prerequisites/consequences from substrates might be problematic, at least in the context of traditional macroscopic fMRI studies. The authors proposed that brain areas which activate during unconscious perception should be discarded as content-specific candidates for substrates of conscious experience. However, given that we demonstrated that correlation with conscious motion perception can be restricted to

fine-grained clusters within an area, it is conceivable that only part of the area (such as hMT+) is a substrate of conscious experience. It is possible that unconscious processes correlate with neural events in some sub-parts of hMT+ that are unrelated to conscious processing. Using macroscopic fMRI studies it is difficult to distinguish different sub-parts of an area since signals are often averaged across all voxels in an area or spatial smoothing is applied to increase statistical power (or both). We therefore caution against discarding an area as a neural substrate candidate all together simply based on finding activation during unconscious processes in a macroscopic fMRI study.

Regarding the findings presented in Chapter 4, we suggest to interpret the observed pRF displacements for low-contrast stimuli as neural prerequisites for conscious position perception. Furthermore, if the hypothesized mechanism for linking pRF shifts to perceptual displacements is correct, then the pRF shifts should be considered *content-specific* neural prerequisites. To elaborate, the observed shifts in early visual areas cannot be the neural substrate of perceived position because the neural and perceptual displacements are in opposite directions. Therefore, the content represented in the neural system does not match that represented in consciousness. The identified NCCs, although only prerequisites, can still be content-specific if, as we have proposed, the shifted population code they produce is read-out by higher-order areas. Based on the population-level code higher-order areas misattribute the stimulus position (by a particular) amount such that neural representation of position and the perceived position of the stimulus now match. If this (speculative) scenario is correct, then these higher-order areas should be considered the neural substrate of conscious position perception and the early visual areas their content-specific prerequisite.

The analysis conducted in Chapter 4 also exemplifies a new strategy for dissociating neural prerequisites from substrates. De Graaf and colleagues (2012) suggested that if neural events temporally precede a given conscious experience, they cannot be substrates but must be prerequisites. By analogy, we would propose that, in addition to temporal precedence, there is also mechanistic precedence that allows us to tell prerequisites from substrates. A mechanistic explanation always specifies a sequence of steps in order to account for a given phenomenon. In Chapter 4, for

example, the proposed mechanism links the pRF shifts in early visual cortex to perceptual displacements in position via a number of steps. Since the pRF shifts in early visual cortex mechanistically precede the read-out by higher-order areas and the corresponding conscious experience, they cannot be neural substrates but must be prerequisites. Note that mechanistic precedence here implies temporal precedence - with the difference that temporal precedence can only be established empirically, while mechanistic precedence can be derived from mechanistic model assumptions.

5.3.2 *Historical context: Reaction to MVPA studies*

In several ways, the studies presented in this thesis are a reaction to the introduction and wide-spread use of MVPA techniques for the analysis of fMRI data since the early years of the millennium (Haxby, Connolly, & Guntupalli, 2014). MVPA is a powerful analysis method and has several advantages compared to classical activation-based analysis methods (Mur, Bandettini, & Kriegeskorte, 2009). MVPA is more sensitive than univariate GLM analyses because the response of many voxels is jointly analyzed. This means that MVPA is able to classify patterns of brain activity into distinct experimental stimuli, tasks or states even if average signal change in a voxel or ROI does not differ across conditions. Furthermore, MVPA can offer an intuitive decoding of brain states that can be beneficially combined with real-time fMRI methods (LaConte, Peltier, & Hu, 2007).

However, the interpretation of MVPA results is fraught with ambiguities (Bartels, Logothetis, & Moutoussis, 2008; Logothetis, 2008; Naselaris & Kay, 2015; Wang, Merriam, Freeman, & Heeger, 2014). Naselaris and Kay (2015) identified three types of ambiguity for MVPA - geometric, spatial and representational ambiguity. Geometric ambiguity refers to the fact that, once MVPA has identified a difference between two experimental conditions, it is unclear whether the corresponding brain activity patterns, represented in the analysis as multivariate vectors, differed in their orientation or in their length. It is considered the most benign type of ambiguity since it can be resolved via additional activation analyses. Spatial ambiguity refers to the limitation that successful classification of experimental conditions does not offer any information about the underlying spatial organization and structure that may have supported classification

- such as, for example, topographic organization, columnar, or laminar structures (but see Kriegeskorte, Goebel, and Bandettini, 2006). Lastly, representational ambiguity refers to the problem that MVPA cannot identify the different sources of variation that might have driven classification performance.

The problems of spatial and representational ambiguity are illustrated by the debate surrounding hyperacuity or subvoxel resolution that MVPA was thought to confer to fMRI. Initially, the successful categorization with MVPA of stimulus orientation (Haynes & Rees, 2005; Kamitani & Tong, 2005) and motion direction (Kamitani & Tong, 2006) in early visual areas was assumed to be driven by small differences in responses at the level of different orientation or direction-of-motion columns even though the spatial resolution employed in these studies was too low to resolve columnar structures directly. Yet later studies convincingly demonstrated that the reported response pattern differences were driven by large-scale spatial biases such as activity elicited by the edge of an oriented grating (Carlson, 2014; Freeman, Brouwer, Heeger, & Merriam, 2011; Roth, Heeger, & Merriam, 2018) or the trailing edge of a motion stimulus (Wang et al., 2014). These examples caution against inferring any particular spatial organization from response differences identified in MVPA and highlight that successful decoding of a stimulus feature does not imply that this feature is actually represented or encoded in the measured population.

In the present work, the combination of high-resolution imaging with explicit models of representation (Naselaris & Kay, 2015) allowed us to resolve (some of) this spatial and representational ambiguity. In the case of the motion quartet study we combined sub-millimeter resolution imaging with the simplest possible encoding model (Naselaris, Kay, Nishimoto, & Gallant, 2011), a GLM with only two predictors, to encode axis of motion. This approach allowed us to demonstrate the existence of two distinct neural populations representing different motion axes and to show their spatial columnar organization. Although previous MVPA studies had found small activation differences between perceptual states related to different motion directions (Brouwer & van Ee, 2007; Kamitani & Tong, 2006), they could not make claims about the underlying spatial organization of such activation patterns (Bartels et al., 2008; Logothetis, 2008).

In case of the motion-dependent pRF study we used pRF mapping, which directly encodes position in visual space as the relevant stimulus feature driving an individual voxel's response. The pRF mapping benefited from high-resolution imaging since, with decreased voxel size, the distribution of neurons contained in a voxel becomes more homogeneous in terms of their RF (De Martino et al., 2018). Previous studies (Fischer, Spotswood, & Whitney, 2011; Maus, Fischer, & Whitney, 2013) found that patterns of fMRI responses in early visual areas reflected physical stimulus position while higher areas like V3a, LO or hMT+ represented the perceived position. However, due to representational ambiguity it is unclear which feature was driving the observed correlations in response patterns. Using the voxel-wise pRF model and explicitly encoding position enabled us to show that pRFs are displaced against the direction of motion under conditions of perceived position shift.

5.3.3 *Current context: Reaction to reproducibility debate*

The studies presented in this thesis share several experimental design strategies. Many of these common strategies can be understood by reference to the recent debate surrounding reproducibility in human neuroimaging research. Neuroimaging researchers are becoming increasingly aware of the importance of reproducibility and some have openly expressed concern that conclusions drawn from many human neuroimaging studies might not be reproducible (Poldrack et al., 2017). Poldrack and colleagues (2017) identified several factors that can contribute to false-positive findings. These factors include low statistical power, flexibility in data analysis, insufficient study reporting, software errors and a lack of direct replication. Over the next few paragraphs, the current work will be discussed in the light of these factors in order to highlight how (some of) these problems have been addressed.

One major factor increasing the risk of false-positive findings is low statistical power. As pointed out by Poldrack et al. (2017), even detecting the large effect that results from the comparison of motor activity to rest requires a study with more than 20 subjects. In this light, the prospects for conducting a reproducible fMRI study and identifying much smaller cognitive effects appear bleak. To make things worse, UHF fMRI studies

pose additional practical restrictions that prevent large sample sizes. Since the price of a single scanning session at a 7 T scanner is currently in the order of several hundreds of Euros, many researchers need to content themselves with as little as 20 scanning hours per study. In addition to monetary constraints, other limitations apply. As highlighted in this thesis, automatic segmentation algorithms yield results that often do not meet the quality demands of high-resolution analyses. Consequently, many hours of manual corrections are needed per subject and tissue class segmentations become a bottleneck for studies with large sample size.

As a result of these practical limitations, the results reported in this thesis are based on a handful of individuals only. In order to alleviate concerns about power, we used several strategies. First, instead of collecting little data for many participants, we studied a few individuals in great detail by acquiring many hours of scanning for every individual (five to eight hours of MRI data per individual) (Poldrack, 2017). This approach is common in visual neuroscience and, in particular, electro-physiology studies of primate brains because the need for extensive training and eventual sacrifice of the monkeys makes large sample sizes impractical and unethical. The approach is made possible by relatively large responses in visual cortex that usually yield sufficient signal-to-noise (SNR) to study effects in individual brains. Proponents of such small-N designs point out that neural and behavioural studies usually face a trade-off between sampling and measurement error (Kolossa & Kopp, 2018; Smith & Little, 2018). Sampling error results from variance across participants, while measurement error results from variance across different measurements in the same individual. Scarce research time can either be used to sample more subjects or to sample more data per subject. The decision to prioritize one option over the other should be made by considering which variance is expected to be higher. For the UHF high-resolution fMRI studies presented in Chapter 3 and 4, it can be argued that measurement error is more severe. On the one hand, high-spatial resolution is afforded by trading in SNR which makes high-resolution studies SNR-limited (higher measurement error). On the other hand, the visual areas studied here represent basic sensory systems which are expected to be similar across individuals (lower sampling error). For these reasons, it can be justified to prioritize data quality over data quantity (Kolossa & Kopp, 2018). It is important to keep in mind, however,

that this approach limits inference to the studied individuals and does not allow for generalization of results to the population level. This important limitation applies to all studies presented in this thesis and, consistent with it, we did not employ and report standard group statistics and showed single-subject results instead.

A second strategy that we used to improve statistical power was to limit our analysis to ROIs that were defined using independent functional localizer data. This approach increases power by limiting the search space and reducing multiple comparisons. Fortunately, this strategy dovetails well with the standard approach for the search for NCC, outlined in the [General Introduction](#). The NCC approach requires that data analysis is divided into two steps. Before correlation between the representations in a neural system and in consciousness can be investigated, the representational states of a neural system first need to be mapped out. Similarly, to avoid circularity (Kriegeskorte, Simmons, Bellgowan, & Baker, 2009), ROIs first need to be explicitly delineated based on independent data before any other statistical analysis steps can be conducted (Poldrack et al., 2017).

Next to low statistical power, flexibility in data analysis can also hamper reproducibility. All the findings reported in this thesis are the result of many different pre-processing and analysis steps. For each step, one particular method and set of parameters needed to be chosen out of a multitude of potential alternatives. This results in variability of observed effects across all theoretically possible analysis settings. Such variability can be substantial (Carp, 2012) and is problematic if the reported effect is highly dependent on a specific configuration of settings or if researchers intentionally modify the analysis settings until significant results are obtained. This analytical flexibility has been termed "researcher degrees of freedom" (Simmons, Nelson, & Simonsohn, 2011).

Some authors have suggested pre-registration of methods and analysis plans as the main solution to this problem (Poldrack et al., 2017). However, it can be argued that pre-registration unduly restricts explorative efforts, which is essential to scientific discovery, and hampers the development of novel analysis techniques. This is particularly relevant in the context of UHF studies, since harvesting the opportunities afforded by UHF is only possible after the development of appropriate analysis methods (see Chapter 1 and 2).

To reduce researcher degrees of freedoms, we instead combined several strategies throughout this thesis. First, especially for data preparation and pre-processing, we aimed to restrict ourselves to default parameters and deviation from default settings was clearly motivated and reported in the methods description. Second, we usually explored the impact of analysis settings in a single subject and, after having established an analysis pipeline, we validated analysis and results in the other subjects. Third, even after an analysis pipeline was established and the result figures for each project were created, we still systematically varied some analysis settings to check that the observed effects are robust over a wider range of settings. For example, we commonly varied the number of voxels included in a ROI to see how variable the obtained effects were (see Chapter 3) and to determine the parameter space for which the obtained results hold. Although all these strategies limit analytical flexibility, the research presented in this thesis should still be considered exploratory and further validation of the findings presented here is important.

Other obstacles to reproducibility are insufficient study reporting and software errors (Eklund, Nichols, & Knutsson, 2016). Insufficient study reporting means that the methods described in a scientific paper are not enough for other researchers to redo the analyses that gave rise to the reported results and conclusions. One remedy to this problem is adherence to established reporting guidelines, such as the COBIDAS report (Nichols et al., 2017) of the Organization for Human Brain Mapping (OHBM). We tried to adhere to these guidelines throughout this thesis, where possible. However, there is an inherent conflict between the readability and reproducibility of a methods section. For example, listing all the parameters that are used for distortion correction (around 30 parameters) makes research more reproducible but would also make a manuscript entirely unreadable.

There are at least two alternative and potentially more efficient ways for communicating analysis steps, which have been used in all studies presented in this thesis. First, flow chart diagrams offer a comprehensive overview of all the relevant (pre-)processing steps of a study. It turned out that this was not only helpful to communicate the processing pipeline to other researchers but was a welcome memory aid to ourselves as well. Second, making both data and analysis code available appears to be the best solution to increase reproducibility without sacrificing readability of

the manuscript. For this, the use of services such as GitHub and Zenodo have proven very helpful. Since GitHub (and other services) combines distributed version control with a web-based hosting service, it both reduces the chance of software bugs and improves study reporting. Version control with Git enabled us to organize analysis code and to keep track of code changes. It also enabled collaborative coding efforts among co-authors. All this facilitated tracing and detecting software errors at an early stage. Since repositories are web-based, analysis code can be easily released upon publication or can be shared already during software and analysis development.

5.4 LIMITATIONS AND FUTURE DIRECTIONS

Several limitations apply to the findings reported in this thesis that relate to experimental design, MRI methodology, analysis methodology or interpretation of our results. These limitations reflect the two big research lines that were introduced in the [General Introduction](#) - human neuroimaging at UHF and the search for content-specific NCC. We will start by discussing limitations that are specific to a particular study but then go on to formulate, more generally, desirable directions for future UHF fMRI and consciousness research.

5.4.1 Tissue class segmentation at UHF

We estimated that the 2D histogram method presented in Chapter 2 saves on average 7.5 hours of manual work for a whole brain cortical ribbon segmentation of a single subject (Gulban, Schneider, Marquardt, Haast, & De Martino, 2018). While this is a valuable time-saver, the presented method does not entirely eliminate the need for manual corrections and even more automation of the segmentation task would be desirable. The main problem with tissue class segmentation of UHF MRI images is that most of the segmentation algorithms have been benchmarked using low-field, lower resolution data (Helms, 2016). Therefore, these algorithms often do not deal well with the additional anatomical detail apparent in UHF images. A straightforward approach would be to train the same algorithms on UHF data. However, many segmentation algorithms (Ash-

burner & Friston, 2005; Bazin et al., 2014) rely on particular templates, atlases or other forms of prior information. This constitutes an important hindrance whenever new forms of data become available (for example partial acquisitions) because they might not conform to the pre-supposed shape. By contrast, a human expert can perform segmentation very flexibly and adjust to new forms of data quickly.

For these reasons, it is desirable to develop new algorithms that deploy skills akin to human perception and action. One very promising direction is the use of deep convolutional neural networks (CNN) for image segmentation. It was demonstrated that CNNs trained for image classification constitute a good model of the ventral visual pathway system for object recognition in humans (Kriegeskorte, 2014; Yamins & DiCarlo, 2016). Thus, there appear to be important parallels between the neural implementation for image classification in humans and CNNs. Image classification can be extended to image segmentation by not only identifying the objects in an image but also determining for every pixel in the image to which object it belongs. A convolutional network called U-net (because it takes the shape of a U) currently offers close-to-human performance in image segmentation tasks, at least for segmenting microscopic images of neuronal structures in fruit flies (Ronneberger, Fischer, & Brox, 2015) and kidneys of clawed frogs (Çiçek, Abdulkadir, Lienkamp, Brox, & Ronneberger, 2016). An obvious direction for future research is the implementation of a 3D U-net for tissue class segmentation of UHF MRI images of the human brain and comparison of its performance to conventional segmentation algorithms used for this task.

5.4.2 *Columnar fMRI and large draining veins*

One important limitation of the columnar fMRI study presented in Chapter 3 is the potential contribution of large draining veins to the reported fMRI signal. Two different types of veins should be distinguished. In the columnar and laminar fMRI literature, when authors speak of "draining veins" they frequently refer to pial veins that are located on the cortical surface. These pial veins render the signal close to the grey matter (GM) / cerebrospinal fluid (CSF) boundary less specific (Moerel et al., 2017; Polimeni, Fischl, Greve, & Wald, 2010) and are a major contributor to an

increase in fMRI signal towards the superficial cortical depth levels, a finding that is very robust across many laminar fMRI studies (De Martino et al., 2013; Koopmans, Barth, & Norris, 2010; Koopmans, Barth, Orzada, & Norris, 2011; Marquardt, Schneider, Gulban, Ivanov, & Uludağ, 2018; Polimeni et al., 2010; Ress, Glover, Liu, & Wandell, 2007). Yet there is a second meaning to the term "draining vein" that describes ascending cortical veins which originate in deep cortical depth and, as they ascend towards the pial surface, increase in diameter due to the collection of venous blood (Duvernoy, Delon, & Vannson, 1981). These ascending cortical veins are problematic for laminar fMRI studies since they also contribute to a cortical depth profile that is biased towards superficial cortical depths. However, these veins pose a particular problem to the interpretation of columnar fMRI studies because they usually run radially to the cortical surface, just like columns, and may mimic columnar-like neuronal activity (De Martino et al., 2018).

An obvious remedy to this limitation is to use sequences other than gradient echo echo planar imaging (GE EPI) that have less sensitivity to large veins. Alternative sequences that have been proposed and employed in the literature for laminar and columnar fMRI are 3D GRASE (De Martino et al., 2013; Feinberg, Harel, Ramanna, Ugurbil, & Yacoub, 2008; Kemper et al., 2015; Moerel et al., 2017; Zimmermann et al., 2011) and VASO (Huber et al., 2015; Huber et al., 2017). Uludag and colleagues (2009) showed that spin echo (SE) sequences in particular become less macrovasculature weighted when moving to 7 T. Since the 3D GRASE sequence is a combination of SE and GE elements, it is expected to be less macrovasculature weighted and thus more specific - a prediction that was supported experimentally by De Martino et al., 2013. The VASO sequence offers a measure of cerebral blood volume and therefore can provide higher spatial specificity (Huber et al., 2015; Huber et al., 2017). Next to an increase in specificity, these sequences have other advantages for mesoscopic fMRI. As a result of their small spatial coverage, 3D GRASE and VASO images tend to show less geometric distortions than images acquired with GE EPI sequences, simplifying one of the major challenges for current UHF fMRI studies (see Chapter 1). Furthermore, for the VASO sequence, it is relatively straightforward, via a modification of the sequence, to generate images with sufficient white matter (WM)/GM/CSF contrast that allow

for tissue class segmentation. Thus, some of the problems with building an anatomical reference (see Chapter 1) can be circumvented.

However, 3D GRASE and VASO also each have their distinct disadvantages. In particular, both sequences have very limited coverage compared to 2D or 3D GE EPI sequences (Huber et al., 2015; Huber et al., 2017; Zimmermann et al., 2011), which poses severe restrictions on the region of interest to be studied as well as on subject motion during scanning. Furthermore, as a consequence of the reduced macrovasculature weighting, both 3D GRASE and VASO are less sensitive than 2D GE EPI (demonstrated empirically for 3D GRASE in De Martino et al., 2013). Unfortunately, I can fully attest to the correctness of this finding based on my own experience. To further validate the findings presented in Chapter 3, we initially also recorded data with the 3D GRASE sequence (five participants with one scanning session each). For the GE EPI data reported in Chapter 3 at least two scanning sessions were necessary to robustly observe the reported effects. However, a trend in the expected direction was already apparent after a single scanning session. This was not the case for the 3D GRASE data, and as a young researcher new to fMRI data analysis it took me the first half of my PhD to realize and comprehend that the temporal SNR of the 3D GRASE images was too low to draw valid and reliable conclusions from the data. Now, with the experience and insights gained over the years, I would record data from only one or two subjects but with four to five scanning sessions each. Such a setup would offer important validation for the claims about columnarity made in Chapter 3.

Another remedy to the limitation of ascending cortical veins could become available in the future by applying an accurate spatial deconvolution model that takes signal contribution of these veins into account (Markuerkiaga, Barth, & Norris, 2016). We have used such a model in a laminar fMRI study investigating contrast-dependence in V1 and V2 (Marquardt et al., 2018). However, it is not straightforward to extend this approach to columnar fMRI studies since the deconvolution parameters will need to vary when moving along the cortical plane direction in order to account for the varying distribution of veins. Furthermore, currently the deconvolution model relies on vasculature estimates obtained from animal studies which are only available for primary visual cortex. Therefore, a challenge for the future will be to extend this model to mid-level

areas (such as hMT+) and even higher-order areas for which no homologous area in non-human animal might exist. Both these problems could be addressed by developing accurate in-vivo methods for mapping out ascending cortical veins in humans (De Martino et al., 2018).

5.4.3 *Laminar contributions within a column*

For the study presented in Chapter 3, our approach was to identify distinct columnar clusters and then to test whether the average signal across all voxels in a given cluster reflected the conscious percept. We believe that this approach was an important first step and demonstrates the novel empirical opportunities offered by columnar fMRI to probe content-specific NCC in humans. However, several additional insights for perceptual processing could have been gained if we had been able to study more precisely where in a given columnar cluster the conscious percept is reflected or formed. This could be achieved by measuring distinct laminar contributions to the identified columnar clusters. Several theoretical and empirical observations suggest that this will represent a fruitful avenue for future research. We will illustrate our considerations by referring to the specific example of the motion quartet study in Chapter 3 but we believe that these points hold more generally for columnar fMRI studies probing perception.

Results from animal physiology indicate that different computational processes occur across different cortical layers (Felleman & Van Essen, 1991; Markov et al., 2014; Rockland & Pandya, 1979). In particular, feedback connections appear to originate from neurons mainly in deep layers (layer VI), and terminate mainly in layer I or in a combination of layers, but not in layer IV (granular layer) (Rockland, Saleem, & Tanaka, 1994; Salin & Bullier, 2017). By contrast, forward connections originate in superficial cortical layers (layer II and III) and terminate in layer IV. Inspired by these insights from animal physiology, we predict that bottom-up and lateral / top-down processes will recruit different laminar parts of a column.

In the specific example of the motion quartet study, during the physical motion experiment, physical bottom-up and perceptual top-down signals are in agreement. We would therefore predict that these processes should activate most (or even the entirety) of a column. During ambiguous apparent motion, however, bottom-up and top-down signals are dissoci-

ated which means that we expect different parts of a column to be activated to different degrees. For example, if ambiguous motion was based on feedback that enters primarily in superficial layers, we would expect that part of a column to be more activated than other parts.

This example illustrates the potential benefit for understanding the cortical implementation of perceptual processes that could come from exploiting distinct termination patterns and studying the neuronal circuits implemented at different depth levels along a column. This will only be possible once we can reliably measure the distinct laminar contributions to particular columns. Although this has been achieved in a pioneering study (De Martino et al., 2015), current resolution limits hinder routine and reliable application, which will likely only become available once spatial resolution approaches 300 to 400 microns.

5.4.4 Functional significance of columns

Further validation of columnar fMRI studies and the ability to determine laminar contributions within a column will thus represent important steps for future UHF research. These developments will hopefully also shed light on the controversy surrounding the functional significance of columnar organization. Some authors have argued that cortical columns are structures devoid of function (Horton & Adams, 2005) based on the observation that no difference in visual performance is apparent between species that exhibit columnar organization in visual areas and those that do not. By contrast, other authors have successfully used columns as a fundamental functional unit to account for information processing in cortical computations (Bastos et al., 2012).

We propose that currently a good working assumption for columnar fMRI studies is that the column offers a way to organize feedforward and feedback information around a common computational feature. For example, feedforward and feedback processing can be integrated for a particular orientation in V1 or a particular frequency in A1, which offers an efficient manner to process attentional or contextual effects (De Martino et al., 2015). We believe that, ultimately, empirical studies are the best way to advance the debate about the functional significance of columns. Studies should map out columnar structures in species that have them

and test their relation to perception and cognition. In the motion quartet study, when mapping horizontal and vertical preferences in response to physical motion, response preferences were clustered in a columnar fashion. Furthermore, the representational content in these columnar structures appeared to be functionally relevant for conscious perception as each structure was involved in tracking a conscious percept of a particular motion axis.

5.4.5 Richer representational content

In this thesis, the representational content was either restricted to two sub-categorical features (horizontal or vertical motion; Chapter 3) or to one feature with several levels (different positions; Chapter 4). An obvious extension for future work is to reveal the neural substrate for conscious content with multiple features and levels. The development of complex representational models such as deep neural networks (Bashivan, Kar, & DiCarlo, 2019; Kriegeskorte, 2014, 2015; Yamins & DiCarlo, 2016) together with methods that connect these models with brain activity (Diedrichsen & Kriegeskorte, 2017; Diedrichsen, Yokoi, & Arbuckle, 2018; Kriegeskorte, 2008; Naselaris et al., 2011) now allow for mapping representational content with fMRI in cortical systems of humans along many stimulus features and feature dimensions. These developments lift most of the restrictions on the first step for studying content-specific NCC in humans and enable the mapping of rich representational content.

Instead, we anticipate future challenges in the second step - finding appropriate methods for ambiguous stimulation in order to dissociate neural processes that pertain to conscious processing from those that pertain to sensory stimulation. For visual perception, several psychophysical methods are available to achieve this dissociation, including masking (Breitmeyer & Ogmen, 2010; Kouider & Dehaene, 2007), bistable figures and binocular rivalry (Kim & Blake, 2005) as well as continuous flash suppression (Tsuchiya & Koch, 2005). Some psychophysical methods might be more suitable for revealing content-specific NCC than others. For example, visual masking is not well suited for processes requiring sustained stimulation (such as motion processing, perceptual learning, temporal integration, etc) since stimuli become visible when presented for more pro-

longed periods (Faivre, Berthet, & Kouider, 2014). By contrast, continuous flash suppression can render stimuli invisible for long periods and offers reliable control of timing (Faivre et al., 2014; Tsuchiya & Koch, 2005). As an alternative to psychophysical approaches, imagery or occlusion (Muckli et al., 2015) studies can be used to dissociate sensory and conscious processing because these studies lack the relevant sensory input (see below). Regardless of the specific method being used, it will be challenging to disentangle the exact contributions of single features and their distinct levels and many pairwise comparisons might be required to achieve this goal.

Another impediment to mapping richer representational content will be statistical power. Even mapping out just one feature with two states (Chapter 3), already required a single scanning sessions per participant (not including ambiguous stimulation). For the motion-dependent pRF study (Chapter 4), we benefited from the large effect of visual field position on responses in early visual cortex. When moving to more complex content, limitations in statistical power can be anticipated since many different conditions will be required. Adding considerations about reproducibility, future consciousness research will benefit from small-N studies with many scanning sessions per participant.

5.4.6 Reconstructing conscious experience

For the study presented in Chapter 4, we grouped pRF and psychophysical estimates across all participants to increase statistical power. Although the group average was reflective of individual trends, there were also inter-individual differences that we did not account for; for example, some subjects showed larger perceptual displacements or larger pRF shifts than others. For future studies, it would be desirable to link the change in pRF estimates and perceptual shifts on a single-subject level. Such a model should be used to predict how much perceptual displacement a subject will experience based on the measured shift in pRF. Assuming that the hypothesized link between pRF estimates and position perception is correct, this should become feasible by collecting more data per participant and thus increasing statistical power.

More generally, the long-term goal for consciousness science should be to make accurate prediction from the brain activity of an individual to her or his conscious experience. The goal should not only be to decode different stimulus categories seen during training from brain activity (as has been done with MVPA). Instead, consciousness science should aim to make quantifiable predictions about the ongoing subjective experience. Such predictions should be made in a specific format by means of an appropriate reference space. For example, if visual consciousness is concerned, the reference space could be image space - with fluctuations in visual experience being reflected in time-varying images.

Such a goal might appear far-fetched or outright unreachable. Obvious obstacles in the way are that neural representations are noisy and high-dimensional (St-Yves & Naselaris, 2019). Moreover, this goal re-emphasizes the need to distinguish neural processes pertaining to sensory information from those related to conscious experience. For example, it has been shown that the orientation of a stimulus can be decoded from response patterns in primary visual cortex even when there is no conscious awareness of the stimulus due to masking (Haynes & Rees, 2005). Clearly, it is undesirable for information pertaining to unconscious processes to leak into the reconstruction of conscious experience. Problematically, it is conceivable that the involvement of cortical structures in conscious processing varies over time. Thus, the biggest obstacle will be the demarcation problem - clarifying which brain processes are conscious and which ones are not.

However, several considerations are also encouraging the pursuit of this goal. First of all, several successful attempts have shown that it is possible to reconstruct presented images from brain activity (Miyawaki et al., 2008; Naselaris, Prenger, Kay, Oliver, & Gallant, 2009; Nishimoto et al., 2011; Schoenmakers, Barth, Heskes, & van Gerven, 2013; Senden, Emmerling, Hoof, Frost, & Senden, 2018; Thirion et al., 2006; van Gerven, Güçlütürk, Seeliger, Ambrogioni, & Güçlü, 2018; St-Yves & Naselaris, 2019). In an important way, the task of reconstructing conscious experience from brain activity is more straightforward than that of reconstructing shown images: By most scientific and by many philosophical accounts of consciousness we should expect all information relevant for conscious experience to be available in the brain (but see O'Regan and Noë, 2001).

By contrast, information about image details are incomplete given that not all aspects of an image are necessarily encoded by our senses (St-Yves & Naselaris, 2019); these aspects of an image can never be recovered from brain activity. Furthermore, a few studies have shown that it is possible to reconstruct imagined images (Naselaris, Olman, Stansbury, Ugurbil, & Gallant, 2015) and letters (Senden et al., 2018) from brain activity. This is particularly encouraging for reconstruction of conscious experience since during imagery no sensory information from the stimulus was available (thus bypassing the demarcation problem).

Second, the necessary ingredients for reconstruction of conscious experience can already be specified. All of the above studies of (mental) image reconstruction shared one aspect: they all relied upon an encoding model that specified how stimuli in the external world are encoded in brain responses. Examples of used encoding models are the Gabor wavelet model (Kay, Naselaris, Prenger, & Gallant, 2008), the pRF model (Dumoulin & Wandell, 2008) or the feature-weighted RF model (St-Yves & Naselaris, 2017). All these models specify a mapping from image space to patterns of brain activity. The feature-weighted RF model is neutral with regard to the features that are encoded and can be employed with any model that specifies such a mapping, including deep CNNs. All of the above studies collected independent data that could be used to fit the parameters of the encoding model to every individual participant. For reconstruction, the mapping then needs to be inverted such that it points from brain response patterns to image space, for which different solutions exist (Naselaris et al., 2009; Rezende, Mohamed, & Wierstra, 2014; Senden et al., 2018; St-Yves & Naselaris, 2019).

Reconstruction of conscious experience can proceed in a similar manner with the additional requirement that appropriate methods (such as continuous flash suppression) are employed that allow for dissociating brain responses for conscious and unconscious processing. The idea is to obtain brain responses that result from the contrast between conscious and unconscious processing of the same stimulus. An interesting question is whether these brain responses should be used for fitting parameters of the encoding model (step A) or during reconstruction (step B) (or both). For letter and image imagery, parameters were fitted using data from conventional presentation of letters or natural images (Naselaris et al.,

2015; Senden et al., 2018) and only reconstruction was based on imagery data. This approach relies on the assumption that neural processes are shared between perception and imagery, which is supported by evidence (Pearson, Naselaris, Holmes, & Kosslyn, 2015). However, we know that unconscious and conscious processes do not share the same neuronal processes. Alternatively, brain responses resulting from the contrast between conscious and unconscious processing can be used in both step A and step B, as long as the two steps rely on data sets that are independent from one another. Finally, a third possibility is to use these brain responses only to fit the parameters of the encoding model and then to use an inversion of the established mapping to reconstruct the stream of consciousness experience without employing any dissociative techniques in step B. Future research should explore these two possibilities.

5.5 CONCLUSION: CHALMER'S FALLEN CAVEAT

Neuroimaging in humans was previously thought to be ill-suited to investigate content-specific NCC given difficulties to reliably map and track representational content for neural systems in humans. We have named this provision for human neuroimaging methods "Chalmer's caveat" after the author that originally formulated it (Chalmers, 2000). This thesis can be understood as a demonstration that, with the introduction of UHF scanners and the development of appropriate fMRI analysis techniques, Chalmer's caveat no longer holds for fMRI studies. It is now possible to map sub-categorical representational content in human cortical systems and to formulate and test mechanisms that link these representations to conscious experience. The future will see even richer representational content and potentially even the ability to reconstruct conscious experience from our brain responses.

5.6 REFERENCES

- Ashburner, J. & Friston, K. J. (2005). Unified segmentation. *NeuroImage*, 26(3), 839–851.
- Bartels, A., Logothetis, N. K., & Moutoussis, K. (2008). fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. *Trends in Neurosciences*, 31(9), 444–453.
- Bashivan, P., Kar, K., & DiCarlo, J. (2019). Neural Population Control via Deep ANN Image Synthesis. *Science*, 9436.
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical Microcircuits for Predictive Coding. *Neuron*, 76(4), 695–711.
- Bazin, P. L., Weiss, M., Dinse, J., Schafer, A., Trampel, R., & Turner, R. (2014). A computational framework for ultra-high resolution cortical segmentation at 7 Tesla. *NeuroImage*, 93, 201–209.
- Breitmeyer, B. & Ogmen, H. (2010). *Visual Masking: Time Slices Through Conscious and Unconscious Vision*. Oxford University Press.
- Brouwer, G. J. & van Ee, R. (2007). Visual cortex allows prediction of perceptual states during ambiguous structure-from-motion. *Journal of Neuroscience*, 27(5), 1015–1023.
- Carlson, T. A. (2014). Orientation Decoding in Human Visual Cortex: New Insights from an Unbiased Perspective. *Journal of Neuroscience*, 34(24), 8373–8383.
- Carp, J. (2012). On the plurality of (methodological) worlds: Estimating the analytic flexibility of fmri experiments. *Frontiers in Neuroscience*.
- Chalmers, D. J. (2000). What is a neural correlate of consciousness? In *Neural correlates of consciousness: empirical and conceptual questions* (Chap. 1, pp. 1–27). MIT Press.
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., & Ronneberger, O. (2016). 3D U-net: Learning dense volumetric segmentation from sparse annotation. *BioArxiv*, 9901 LNCS, 424–432.
- De Graaf, T. A., Hsieh, P. J., & Sack, A. T. (2012). The ‘correlates’ in neural correlates of consciousness. *Neuroscience and Biobehavioral Reviews*, 36(1), 191–197.
- De Martino, F., Moerel, M., Ugurbil, K., Goebel, R., Yacoub, E., & Formisano, E. (2015). Frequency preference and attention effects across cortical

- depths in the human primary auditory cortex. *Proceedings of the National Academy of Sciences*, 112(52), 16036–16041.
- De Martino, F., Yacoub, E., Kemper, V., Moerel, M., Uludag, K., De Weerd, P., . . . Formisano, E. (2018). The impact of ultra-high field MRI on cognitive and computational neuroimaging. *NeuroImage*, 168, 366–382.
- 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(3), e60514.
- Diedrichsen, J. & Kriegeskorte, N. (2017). Representational models: A common framework for understanding encoding, pattern-component, and representational-similarity analysis. *Plos Computational Biology*, 13(4), 1–33.
- Diedrichsen, J., Yokoi, A., & Arbutckle, S. A. (2018). Pattern component modeling: A flexible approach for understanding the representational structure of brain activity patterns.
- Dumoulin, S. O. & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *NeuroImage*, 39(2), 647–660.
- Duvernoy, H. M., Delon, S., & Vannson, J. L. (1981). Cortical blood vessels of the human brain. *Brain Research Bulletin*.
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*.
- Faivre, N., Berthet, V., & Kouider, S. (2014). Sustained invisibility through crowding and continuous flash suppression: A comparative review.
- Feinberg, D., Harel, N., Ramanna, S., Ugurbil, K., & Yacoub, E. (2008). Sub-millimeter Single-shot 3D GRASE with Inner Volume Selection for T2 weighted fMRI applications at 7 Tesla 1. *Magnetic Resonance Medicine*.
- Felleman, D. J. & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*.
- Fischer, J., Spotswood, N., & Whitney, D. (2011). The emergence of perceived position in the visual system. *Journal of Cognitive Neuroscience*, 23(1), 119–136.

-
- Freeman, J., Brouwer, G. J., Heeger, D. J., & Merriam, E. P. (2011). Orientation Decoding Depends on Maps, Not Columns. *Journal of Neuroscience*, 31(13), 4792–4804.
- Gulban, O., Schneider, M., Marquardt, I., Haast, R., & De Martino, F. (2018). A scalable method to improve outer gray matter boundary segmentation at ultra high field MRI. *Plos One*, 13, 1–33.
- Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding Neural Representational Spaces Using Multivariate Pattern Analysis. *Annual Review of Neuroscience*.
- Haynes, J. D. & Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nature Neuroscience*.
- Helms, G. (2016). Segmentation of human brain using structural MRI. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 29(2), 111–124.
- Horton, J. C. & Adams, D. L. (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society*, 360(1456), 837–862.
- Huber, L., Goense, J., Kennerley, A. J., Trampel, R., Guidi, M., Reimer, E., ... Möller, H. E. (2015). Cortical lamina-dependent blood volume changes in human brain at 7 T. *NeuroImage*, 107, 23–33.
- Huber, L., Handwerker, D. A., Jangraw, D. C., Chen, G., Hall, A., Stüber, C., ... Bandettini, P. A. (2017). High-Resolution CBV-fMRI Allows Mapping of Laminar Activity and Connectivity of Cortical Input and Output in Human M1. *Neuron*, 96, 1253–1263.
- Kamitani, Y. & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*.
- Kamitani, Y. & Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Current Biology*, 16(11), 1096–1102.
- Kay, {KN}., Naselaris, T., Prenger, {RJ}., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452(7185), 352–355.
- Kemper, V. G., De Martino, F., Vu, A. T., Poser, B. A., Feinberg, D. A., Goebel, R., & Yacoub, E. (2015). Sub-millimeter T₂ weighted fMRI at 7 T: Comparison of 3D-GRASE and 2D SE-EPI. *Frontiers in Neuroscience*, 9.

- Kim, C. Y. & Blake, R. (2005). Psychophysical magic: Rendering the visible 'invisible'.
- Kolossa, A. & Kopp, B. (2018). Data quality over data quantity in computational cognitive neuroscience. *NeuroImage*.
- Koopmans, P. J., Barth, M., & Norris, D. G. (2010). Layer-specific BOLD activation in human V1. *Human Brain Mapping*, 31(9), 1297–1304.
- Koopmans, P. J., Barth, M., Orzada, S., & Norris, D. G. (2011). Multi-echo fMRI of the cortical laminae in humans at 7 T. *NeuroImage*, 56(3), 1276–1285.
- Kouider, S. & Dehaene, S. (2007). Levels of processing during non-conscious perception: A critical review of visual masking. In *Philosophical transactions of the royal society b: biological sciences*.
- Kriegeskorte, N., Goebel, R., & Bandettini, P. (2006). Information-based functional brain mapping. *Proceedings of the National Academy of Sciences*.
- Kriegeskorte, N. (2008). Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2, 1–28.
- Kriegeskorte, N. (2014). Deep Supervised , but Not Unsupervised , Models May Explain IT Cortical Representation. *Plos One*, 10(11).
- Kriegeskorte, N. (2015). Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing. *Annual Review of Vision Science*, 1(1), 417–446.
- Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S., & Baker, C. I. (2009). Circular analysis in systems neuroscience: The dangers of double dipping. *Nature Neuroscience*.
- LaConte, S. M., Peltier, S. J., & Hu, X. P. (2007). Real-time fMRI using brain-state classification. *Human Brain Mapping*, 28(10), 1033–1044.
- Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197), 869–878.
- Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., . . . Kennedy, H. (2014). Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology*.

-
- Markuerkiaga, I., Barth, M., & Norris, D. G. (2016). A cortical vascular model for examining the specificity of the laminar BOLD signal. *NeuroImage*.
- Marquardt, I., Schneider, M., Gulban, O. F., Ivanov, D., & Uludağ, K. (2018). Cortical depth profiles of luminance contrast responses in human V1 and V2 using 7 T fMRI. *Human Brain Mapping*, 39(7), 2812–2827.
- Maus, G. W., Fischer, J., & Whitney, D. (2013). Motion-dependent representation of space in area MT+. *Neuron*, 78(3), 554–562.
- Miyawaki, Y., Uchida, H., Yamashita, O., Sato, M. a., Morito, Y., Tanabe, H. C., ... Kamitani, Y. (2008). Visual image reconstruction from human brain activity using a combination of multiscale local image decoders. *Neuron*, 60(5), 915–929.
- Moerel, M., De Martino, F., Kemper, V. G., Schmitter, S., Vu, A. T., Ugurbil, K., ... Yacoub, E. (2017). Sensitivity and specificity considerations for fMRI encoding, decoding, and mapping of auditory cortex at ultra-high field. *NeuroImage*, 168.
- Muckli, L., De Martino, F., Vizioli, L., Petro, L. S., Smith, F. W., Ugurbil, K., ... Yacoub, E. (2015). Contextual Feedback to Superficial Layers of V1. *Current Biology*, 25(20), 2690–2695.
- Mur, M., Bandettini, P. A., & Kriegeskorte, N. (2009). Revealing representational content with pattern-information fMRI - An introductory guide. *Social Cognitive and Affective Neuroscience*, 4(1), 101–109.
- Naselaris, T. & Kay, K. N. (2015). Resolving Ambiguities of MVPA Using Explicit Models of Representation. *Trends in Cognitive Sciences*, 19(10), 551–554.
- Naselaris, T., Kay, K. N., Nishimoto, S., & Gallant, J. L. (2011). Encoding and decoding in fMRI.
- Naselaris, T., Olman, C. A., Stansbury, D. E., Ugurbil, K., & Gallant, J. L. (2015). A voxel-wise encoding model for early visual areas decodes mental images of remembered scenes. *NeuroImage*, 105, 215–228.
- Naselaris, T., Prenger, R. J., Kay, K. N., Oliver, M., & Gallant, J. L. (2009). Bayesian Reconstruction of Natural Images from Human Brain Activity. *Neuron*, 63(6), 902–915.
- Nichols, T. E., Das, S., Eickhoff, S. B., Evans, A. C., Glatard, T., Hanke, M., ... Yeo, B. T. (2017). Best practices in data analysis and sharing in neuroimaging using MRI.

- Nishimoto, S., Vu, A. T., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. L. (2011). Report Reconstructing Visual Experiences from Brain Activity Evoked by Natural Movies. *Current Biology*, 1641–1646.
- O'Regan, J. K. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*.
- Pearson, J., Naselaris, T., Holmes, E. A., & Kosslyn, S. M. (2015). Mental Imagery: Functional Mechanisms and Clinical Applications. *Trends in Cognitive Sciences*.
- Poldrack, R. A. (2017). Precision Neuroscience: Dense Sampling of Individual Brains. *Neuron*, 95(4), 727–729.
- Poldrack, R. A., Baker, C. I., Durnez, J., Gorgolewski, K. J., Matthews, P. M., Munafò, M. R., ... Yarkoni, T. (2017). Scanning the horizon: Towards transparent and reproducible neuroimaging research. *Nature Reviews Neuroscience*, 18(2), 115–126.
- Polimeni, J. R., Fischl, B., Greve, D. N., & Wald, L. L. (2010). Laminar analysis of 7t BOLD using an imposed spatial activation pattern in human V1. *NeuroImage*, 52(4), 1334–1346.
- Ress, D., Glover, G. H., Liu, J., & Wandell, B. (2007). Laminar profiles of functional activity in the human brain. *NeuroImage*.
- Rezende, D. J., Mohamed, S., & Wierstra, D. (2014). Stochastic Back-propagation and Variational Inference in Deep Latent Gaussian Models. *BioArxiv*.
- Rockland, K. S. & Pandya, D. N. (1979). Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey. *Brain Research*.
- Rockland, K. S., Saleem, K. S., & Tanaka, K. (1994). Divergent feedback connections from areas V4 and TEO in the macaque. *Visual Neuroscience*.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. *BioArxiv*, 9351, 234–241.
- Roth, Z. N., Heeger, D. J., & Merriam, E. P. (2018). Stimulus vignetting and orientation selectivity in human visual cortex. *eLife*, 7, 1–19.
- Salin, P. A. & Bullier, J. (2017). Corticocortical connections in the visual system: structure and function. *Physiological Reviews*.
- Schoenmakers, S., Barth, M., Heskes, T., & van Gerven, M. (2013). Linear reconstruction of perceived images from human brain activity. *NeuroImage*, 83, 951–961.

-
- Senden, M., Emmerling, T., Hoof, R. V., Frost, M., & Senden, M. (2018). Reconstructing imagined letters from early visual cortex reveals tight topographic correspondence between visual mental imagery and perception. *Brain Structure and Function*, 1–37.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*.
- Smith, P. L. & Little, D. R. (2018). Psychonomic Bulletin and Review Small is beautiful: In defense of the small-N design. *Psychonomic Bulletin and Review*, 25, 2083–2101.
- Thirion, B., Duchesnay, E., Hubbard, E., Dubois, J., Poline, J. B., Lebihan, D., & Dehaene, S. (2006). Inverse retinotopy: Inferring the visual content of images from brain activation patterns. *NeuroImage*, 33(4), 1104–1116.
- Tsuchiya, N. & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8(8), 1096–1101.
- Uludağ, K., Müller-Bierl, B., & Uğurbil, K. (2009). An integrative model for neuronal activity-induced signal changes for gradient and spin echo functional imaging. *NeuroImage*.
- van Gerven, M., Güçlütürk, Y., Seeliger, K., Ambrogioni, L., & Güçlü, U. (2018). Generative adversarial networks for reconstructing natural images from brain activity. *NeuroImage*, 181, 775–785.
- Wang, H. X., Merriam, E. P., Freeman, J., & Heeger, D. J. (2014). Motion Direction Biases and Decoding in Human Visual Cortex. *Journal of Neuroscience*, 34(37), 12601–12615.
- Yamins, D. L. & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3), 356–365.
- St-Yves, G. & Naselaris, T. (2017). The feature-weighted receptive field: An interpretable encoding model for complex feature spaces. *NeuroImage*, 1, 1–15.
- St-Yves, G. & Naselaris, T. (2019). Generative Adversarial Networks Conditioned on Brain Activity Reconstruct Seen Images. In *Proceedings - 2018 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2018*.

Zimmermann, J., Goebel, R., De Martino, F., van de Moortele, P. F., Feinberg, D., Adriany, G., ... Yacoub, E. (2011). Mapping the organization of axis of motion selective features in human area MT using high-field fmri. *Plos One*, 6(12), 1–10.

Summary

Functional magnetic resonance imaging (fMRI) research at ultra-high field (UHF) strength offers new opportunities and challenges for the study of the neural correlates of visual conscious content in humans (content-specific NCC). The goal of this thesis was to leverage these opportunities and to overcome some of the challenges. One such challenge is tissue class segmentation of UHF images. In Chapter 2, we describe and implement a method that removes non-brain voxels efficiently and semi-automatically by representing three-dimensional anatomical images in a two-dimensional histogram. We demonstrate that this method improves cortical grey matter definitions when used as a post-processing step to existing segmentation algorithms and frees a considerable amount of researcher time that would otherwise need to be spent on manual corrections.

Equipped with this new tool, we went on to perform two high-field high-resolution fMRI experiments to clarify how representations in cortical structures are related to the perception of apparent motion and apparent position. In Chapter 3, we study how responses in columnar clusters of the human motion complex (hMT+) relate to the conscious experience of a specific visual motion axis. Leveraging the increased sensitivity of UHF, we first map out distinct, fine-grained clusters that either preferred horizontal or vertical motion. We then measure their response modulations during ambiguous stimulation with the bistable motion quartet. We find that when participants indicate the conscious percept of horizontal motion the response amplitude increases in the "horizontal" cluster and decreases in the "vertical" cluster, while we find the reverse pattern during vertical motion. In addition, we demonstrate that these clusters are organized in a columnar fashion such that preferences for vertical or horizontal motion are stable in the direction of cortical depth and change when moving along the cortical surface. These results support the idea that hMT+ makes up part of the content-specific NCC for visual motion and, more importantly, extend this idea by demonstrating involvement of specific horizontal and vertical columnar clusters in tracking conscious experience of a particular motion axis.

While this study imposed important constraints on the location of content-specific NCC, it did not offer a strong constraint on mechanisms that link cortical activity to conscious visual perception. In Chapter 4, we therefore explore the ability of the population receptive field (pRF) mapping paradigm to connect responses in early and mid-level visual cortex to the perceived position of visual stimuli. Using psychophysics, we first identify stimulus contrast as an important factor that influences the strength of motion-induced position shifts. We observe that at high stimulus contrast, when sufficient sensory evidence is available, little to no perceptual displacements occur, while at low stimulus contrast, when signal uncertainty is increased, substantial perceptual displacements occur. Correspondingly, we find that pRF centers are substantially shifted against the direction of motion for low-contrast, but not high-contrast, stimuli. Based on these observations, we propose a model that links pRF estimates to perceived stimulus position. This model is based on the idea that early visual areas send shifted population codes in response to motion stimuli that bias the conscious percept of stimulus position in the direction of motion.

Neuroimaging methods in humans have often been considered ill-suited to investigate content-specific NCC, given difficulties to reliably map and track representational content for neural systems in humans. This thesis can be understood as a demonstration that, with the introduction of UHF scanners and the development of appropriate analysis techniques, it is now possible with fMRI to map (sub-categorical) representational content in human cortical systems and to formulate and test mechanisms that link these representations to conscious experience.

Knowledge Valorization

The practice of science can be divided into basic and applied science. While basic science is concerned with the investigation and discovery of fundamental principles, applied science is practiced with the aim to solve a specific practical problem, often of a technical nature. The work in this thesis can be understood as basic science. The aim here was to study the principles that describe the relationship between the activity of cortical systems in humans and the contents of visual consciousness.

Valorization is defined as "the process of creating value from knowledge, by making knowledge suitable and/or available for social (and/or economic) use and by making it suitable for translation into competing products, services, processes and new activities" (report of the National Valorisation Committee, Indicators for valorisation, 2011, p. 8). Given the fundamental nature of the research presented in this thesis, the goal was not to create value from knowledge but to create knowledge in the first place. This thesis therefore does not make knowledge directly suitable for social and/or economic use. However, I believe that the work presented here can still be valuable to society in at least two different ways. First, it can bring pleasure to its readers by presenting knowledge about something that is extremely central to our existence - the internal and underlying workings of our brain and mind. Knowing and understanding can be very satisfactory in its own right. Second by organizing, maintaining and creating knowledge and making this knowledge openly available to society, this thesis contributes to the foundation that is absolutely necessary to solving any specific, applied problem. In what follows, these two aspects are elaborated on.

First, humans are innately curious beings, showing an immense drive to explore the environment that surrounds them and often displaying pleasure when discovering new things. Human explorations probably start in mother's womb and if it was not for the drive to explore, none of us would be here today; for the existence of each of us presupposes a long chain of exclusively successful reproductions. These reproductions were successful only because our ancestors followed their drive to forage their

environment for appropriate food and mating resources. The human drive to explore, navigate, manipulate and understand our environment is so deeply ingrained that it let us to discover the laws that underlie our immediate and visible surrounding (why does the apple fall from the tree?), discover the laws of particles invisible to the eye and let us to understand the movement of the stars.

A more recent trend in human history is the exploration of the universe inside of our heads rather than of things and processes external to us. The significance of our mental lives for our existence was already pointed out by Rene Descartes. He noted that our internal universe, i.e. our experiences, feelings and thoughts, is the only universe whose existence we can take for granted and, in fact, is the only proof of our very existence. These insights are famously captured by his "Cogito, ergo sum" ("I think, therefore I am"). Given that our mental life is so central to our existence, most humans are curious to learn more about it and its relationship with the brain. By presenting knowledge on this relationship, I think the current thesis has societal value: it affords the pleasure, I hope, of finding things out on a topic that is very central to everyone of us.

Second, given the fundamental nature of this thesis, applied science is required to translate the insights of this thesis into applications. However, it is important to emphasize that basic science, in turn, is a necessary condition for applied science since, in many ways, it relies on the knowledge that is gained in the process of basic science. Imagine we were interested in solving the particular problem of restoring vision in congenitally blind people. These people have been blind from birth - is it in any way possible to give them back vision, even if in a rudimentary form? Tackling such questions is only possible when a rich body of knowledge exists since it requires knowledge in many different aspects: What is a photo-receptor? How does neuronal processing work? How is the light that hits the retina translated into neuronal signals? What are the brain areas involved in visual processing? How are the neuronal signals processed and transformed to give rise to conscious experience of a visual scene? Which neuronal processes render a visual stimulus conscious rather than unconscious? We need to have a basic understanding of all these aspects first before we can address practical questions about restoring vision in the blind. This thesis aims to contribute to such an understanding.

Furthermore, it is important to emphasize that translating fundamental to applied science, is usually not a straightforward process. This is because - by design - most attempts at this translation will fail. Imagine, for example, we were interested in increasing the number of grains that we can harvest from a crop plant - a specific problem of a technical nature. Based on previous fundamental research, we can formulate several hypotheses about potentially effective genetic modifications. Problematically, we do not know which of the modifications will work out in practice, given that it is unclear how the suggested genetic modifications will play out in a complex environment. Furthermore, we do not know, which of the hypothesised modifications, if any, will yield the highest amount of harvest. The hope is that, if we try out ten different modifications, maybe two will work in practice and one of them might increase the harvest by an amount that makes the research efforts worth their while (assuming an exclusively utilitarian perspective on knowledge).

Given these difficulties that are present by design already, it is important not to introduce any additional, unnecessary obstacles to translating knowledge into applications. Such obstacles are difficulties to access the knowledge because it is hidden behind a journal paywall, difficulties to reproduce the process of gaining knowledge because either the custom analysis code or the data are unavailable, or difficulties to understand the knowledge in the first place because the authors did not attempt to use plain language or did not make an effort to visualize what they did in a comprehensible manner. All of these obstacles are still surprisingly common in current scientific practice.

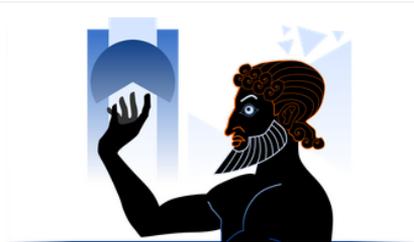
Being aware of these obstacles, we tried to eliminate them whenever possible. This thesis and all of the chapters in it will be made available to the public and free to download for everyone in the world. Two out of three empirical chapters have already been published in journals where articles are openly accessible. Both of these publications were accompanied by the release of code and data that are necessary to perform the described analyses. Furthermore, throughout the thesis, I have attempted to use clear and accessible English and to present comprehensible visualizations of our analysis steps and results. In order raise awareness of the research with a non-expert audience, we also described our research in media outlets such as the in-mind magazine and the Maastricht University news.

Most of the research presented in this thesis required the development of new software tools. Whenever possible and appropriate, we turned these software tools into open-source packages and made them openly available on GitHub. These packages include the [Segmentator package](#), the [PyPrf Motion package](#), the [PyPrf Feature package](#) and the [Hrf Optimize package](#). As illustrated in Figure V1 for the example of the Segmentator package, all of the packages are accompanied by user documentation.

Segmentator

Segmentator is a free and open-source package for multi-dimensional data exploration and segmentation for 3D images. This application is mainly developed and tested using ultra-high field magnetic resonance imaging (MRI) brain data.

The goal is to provide a complementary tool to the already available brain tissue segmentation methods (to the best of our knowledge) in other software packages (FSL, CBS-Tools, ITK-SNAP, Freesurfer, SPM, Brainvoyager, etc.).



Citation:

- Our paper can be accessed from [this link](#).
- Released versions of this package can be cited by using our [Zenodo DOI](#).

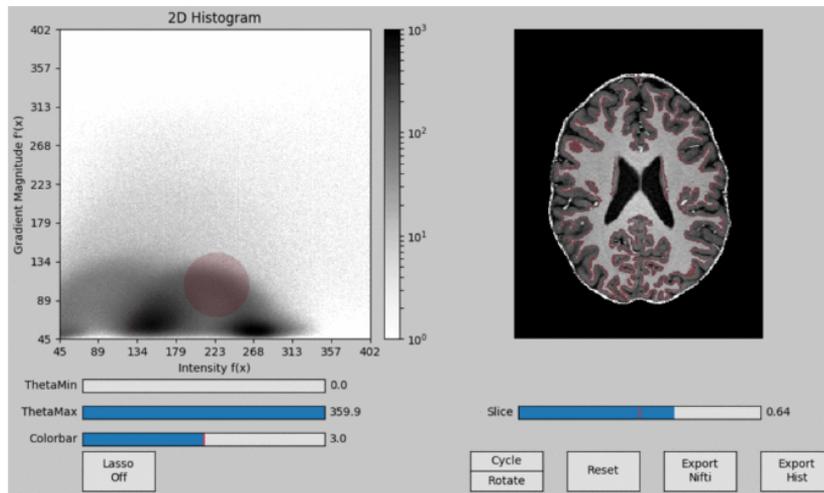


FIGURE V1: Example of open-source software package. As part of this PhD thesis, several software tools were developed that were published on GitHub and accompanied by user instructions.

Acknowledgments

This thesis would not have been possible without the help and continuous support of many people, who I hold very dear.

I would like to thank my supervisory team. **Federico**, I know you always read this section of the thesis with great attention. So be assured that these lines have been chosen carefully for you. I am grateful for many things that you have done for me but two of them in particular. First, that you did not run away when I asked you to join my supervisory team. You are the best academic closer in town: I don't think there is a single piece of research that goes unpublished in your hands and it was this push towards completing work that enormously benefited this PhD project and, I know, many PhD projects before that. Second, I am grateful that no matter what happened, no matter how desperate the situation appeared (*to me*), you kept on reminding me that the most important thing in science (and life) is to be a good fellow human. This mindset came in many different guises and it meant that I did not just feel like another student of yours but like a family friend, always welcome in your home. May the cores of your computer always be busy.

Rainer, I am extremely grateful that, almost from the first time we met (I was giving a presentation on 3D GRASE in your Advanced fMRI course), you placed your trust in my abilities. I am not sure what it was that you saw in me but I am glad that it kicked off a successful cooperation over many years. No matter how high the target, you always seemed to think that I was up to the job. I am also very thankful that this trust, together with your tireless efforts to remove every formal obstacle in my way, provided for an environment where I had the luxury to dedicate almost all my time and mental resources to research. Finally, I would like to thank you for informal discussions and ideas that turned out to be very influential for my thinking on the future of (artificial) consciousness. May your optimism prevail in times of light and in times of darkness.

Nikolaus, thank you for welcoming me to your lab, for introducing me to basic (down to the tone!) and more advanced topics of MRI, for supervising my first steps in academia, for being such a patient and

encouraging teacher, and for inspiring professionalism wherever you go. May your kindness and professionalism define the future of academic research.

I would like to thank my brothers in arms, **Ingo** and **Faruk**. When I was deciding where to do the PhD, there were many aspects that weighed into my decision. One thing that I totally underestimated was the influence of peers. As the ancient German proverb has it "even a blind chicken sometimes finds a grain to pick" - indicating that even undeserving people sometimes get lucky. I guess, in this case, the chicken (me) was not only blind but also headless and got more than lucky. I am tremendously indebted to our (meta) journeys to foreign lands, during which we saw blunt swords and shiny armor, profound abysses and mountains of hope, monkey parades and acts of great human humbleness (or did we?), exotic acronym animals, mind palaces and gipsy houses. May your thirst for exploration never die and your fountains of knowledge never run dry.

Yawen, *in general*, you are a good person. In particular, I am thankful that life gave me a new friend (and such a loyal one, at that!) in the very last moments of my Maastricht life. Thank you for introducing me to the sinosphere (especially the mooncakes) and for tirelessly tearing down the cultural barriers between us. May your cosmopolitan spirit take over the world.

Mario, dr. dr. amigo, my most competent friend and breaker of the chain of causation that keeps your free will unfree. I am not being ironic here - or am I? I guess it doesn't matter as long as it is "just" in my mind. Thank you for all the exciting discussions about our bionic future, extraordinary life advice and the many private GP sessions. May your spirit remain unconditionally free.

My PhD sisters, **Geraldine** and **Linda**, thank you for all the dinners, cocktails and a few memorable nights of dancing in the Sway. **Geraldine**, amiga, may your lust for life brighten a many peoples' mood. **Linda**, may your pointed remarks make many more people smile and may the dancing continue.

My PhD brothers, **Franc** and **Sri**, thank you for keeping me company. **Franc**, may your moustache make a reappearance. **Sri**, may you never run out of "chill pill" catchphrases.

Thank you to the senior generation of PhD students - **Valentin, Thomas, Mario** and **Arko** - all of them honorable gentlemen for being living and breathing examples of the future. May your good manners inspire generations to follow.

Jessica, I am grateful for all the time spent together and the discussions we had - in university buildings and on pique-nique blankets in the park. May your critical mind inconvenience the establishment.

Tyler, the friendly visitor from Glasgow. Thank you for all the balanced discussions. May your voice of reason be heard.

Dimo, thank you for all the help and support at the scanner. May your sequences continue to run flawlessly.

Joel - even though our paths only crossed shortly, I am thankful for the things that you taught me. May your example lead the way.

Thank you to **the Dutch people and state**, for providing near-Utopian conditions for living and doing research. May your model of the welfare state and open exchange of goods and ideas outlast right-wing populism.

Thanks to the **European Union** for giving me the freedom of movement. May you be populated with citizens that appreciate you.

Danke an meine Familie, **Claudia, Norbert** und **Sarah**, für Eure bedingungslose Unterstützung und unermüdliche Fürsorge. Wenn man in einer bildungsnahen Umgebung wie der Euren aufwächst, könnte man all die erfolgreichen Schul- und Uniabschlüsse fast als Selbstverständlichkeit erachten. Das sind sie natürlich nicht und ich bin für den Rückhalt, den Ihr mir über all die Jahre gegeben habt, ausgesprochen dankbar. Möge die Selbstverständlichkeit, mit der Ihr Bildung unterstützt und ermöglicht habt, die normalste Sache der Welt werden.

Undine, thank you for all your love and sparkles. It has been a tremendous pleasure to evolve together with you over all those years, from the kids we were at UCM to what almost feels like proper adults. I am thankful to have in you an honest friend, a supportive partner and a mobile family. May you always be my home in the world, no matter where we are.

Marian Schneider
Maastricht
09-10-2019

About the author

Marian Schneider was born on October 10th 1988 in Oldenburg, Germany. He graduated from high school at Gymnasium Eversten Oldenburg in 2008. He subsequently was admitted to the Bachelor program of University College Maastricht, the Netherlands, where he completed courses in psychology, economics, philosophy and neuroscience. In 2011, he spent an Erasmus semester at the Institut d'Études Politiques in Toulouse, France. He received a Bachelor of Arts (summa cum laude) in 2012. Next, he enrolled in the Research Master program in Cognitive Neuroscience at Maastricht University. In 2013, he was awarded a scholarship by the German Academic Scholarship Foundation ('Studienstiftung des deutschen Volkes'). For his Master thesis, he investigated layer-specific attention effects in human primary visual cortex using 7 Tesla magnetic resonance imaging. This research was conducted from 2013 to 2014 under the supervision of Prof. Dr. Nikolaus Weiskopf at the Wellcome Trust Centre for Neuroimaging, University College London, United Kingdom. He obtained a Master of Science (cum laude) and his master thesis won the prize for the best thesis of the Faculty of Psychology and Neuroscience in 2014. Funded by an NWO Top Talent grant, he conducted his doctoral research from 2014 to 2019 under supervision of Dr. Federico De Martino, Prof. Dr. Nikolaus Weiskopf, and Prof. Dr. Rainer Goebel at the Department of Cognitive Neuroscience, Faculty of Psychology and Neuroscience, Maastricht University. Marian currently lives in Tallinn, Estonia, where he founded a one-man business and does freelance work for BrainInnovation to develop deep neural net architectures for tissue class segmentation of the human brain.