CONNECTIVITY OF THE BRAIN

FROM MAGNETIC RESONANCE IMAGING

Thesis by

Dirk Neumann

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY

Pasadena, California

2010

(Defended December 10, 2009)

© 2010

Dirk Neumann

All Rights Reserved

ACKNOWLEDGEMENTS

I have much to thank my graduate advisor, Ralph Adolphs. 1) Working for the last five years in his lab has been a deeply rewarding adventure. Experiencing the lab grow, over the years, into such a diverse, creative, talented, knowledgeable and yet coherent group of scientists has been a great pleasure. 2) His passion, not only for science, but for the joys of academic life in general and his zenlike mastery of it, his guidance and countless stories have provided much inspiration. 3) I am thankful for his enthusiasm for experimenting with challenging techniques, 4) his availability to meet and talk about my research and 5) for the freedom to sometimes go off on a tangent and to explore my ideas in depth, 6) the help that I received with the preparation of this (and previous) manuscripts, and 7) his legendary (and mostly enumerated) support and insights.

I appreciate the support and generous advise I have received from my committee: Christof Koch (thanks, also, for my lab rotation), Doris Tsao, John O'Doherty, and Mike Tyszka (special thanks for introducing me to diffusion imaging). I am also grateful to the Sloan-Swartz Center of Theoretical Neurobiology at Caltech for my first year of financial support.

I would like to thank the *SURF* students for their valuable work: Wendy Zhang for tracking the Probst bundles and AnaFun work, Matt Thill for his work on a sampling fiber tracking approach and Joe Funke for his hunt for amygdala connections, Arjun Bansal for the enjoyable midnight work on your visualization project, and the very fascinating trip through India's impressive science, technology and culinary landscape.

The experimental work for this thesis took place at Caltech's Brain Imaging Center. It was a pleasure to interact with Steve Flaherty, Ralph Lee, Daniel Procissi and Hargun Sohi. I'd like to thank John Allmann and Jason Kaufman for sharing their technical experience with diffusion tensor imaging and David Anderson and Walter Lerchner for providing me the opportunity to determine the genotype of the participants of an fMRI experiment. I am indebted to Gaby Mosconi and Monica Martinez for the trust and support in teaching me how to use a pipetter and run a PCR; to Jessica Levine for helping collect all those bubbles, her cheer and entertainment; to Remya Nair for taking the time to fix topological errors; to Catherine Holcomb for her organizational skills; and for the fun times running up rainy mountains. Lynn Paul for the collaboration on the callosal project; I am glad you settled on 4 quadrants. My special thanks go to Josef Parvizi, for kindly providing the brain samples that enabled this work.

As much as anything else, what has made my experience unforgettable are all the wonderful interactions with so many smart, thoughtful and knowledgeable people in the neuroscience community at Caltech. Over the years, I have greatly benefited from discussions with Marco Andretto, Ulrik Beierholm, Signe Bray, Steph Chow, Alex Huth, Benedetto de Martino, Pierre Moreels, Alex Holub, Wei Ji Ma, Mili Milosavljevic, Zoltan Nadasdy, Leila Reddy, and Melissa Saenz. I am thankful to Farshad Moradi for his experimental and analytical advise, and sharing Boddingtons at the Siemens programming course and to Alan Hampton for the numerous discussions about MRI and regularization. Thanks also to Oana Tudusciuc for her bubbles help, and the opportunity to climb a rock with an eager 2-year old partner; to Dan Kennedy, Hackjin Kim, Jan Gläscher, and Naotsugu Tsuchiya for the many insightful discussions, and to everyone who rotated through or worked in the lab.

Thanks to my fellow class members: Anusha Narayan, Ueli Rutishauser, Ming Gu, Casimir Wierzynski, Hilary Gliden, Matthew Nelson, Jason Rolfe and Will Ford; to Jonathan Harel for the insightful Jamba juice discussions; to Kerstin Preuschoff for early advice and the very necessary helmet during the snowboarding trip; to Lavanya Reddy for her cheer while studying for the qualifying exam and for her superb cooking skills; to Constanze Hipp her enthusiasm for science and for passionately introducing me to climbing; to Julien Dubois and Christine for the pleasurable kayaking trips; to Fatma Imamoğlu for the discussions about neuroscience and hüzün, and the after-hours explorations of L.A.; to Sanook for the joyful discourses about probabilistic models, for the generous hosting me during my dissertation writing in Boston, and for the fun of discovering hidden states in everything (including my working habits).

I had a bit of a journey before coming to Caltech. I would especially like to thank my previous mentor Karl Gegenfurtner, for his generous support over the years prior to coming to Caltech. Jochem Rieger for answering all my questions at the Max Planck Institute, and the memorable discussions when hosting me later; to Thorsten Schmidt, Kathrin Scheidemann and Michael Beetz from the Munich robotics soccer team and the AI group of Andreas Zell in Tübingen, places where the Bayesian part of my brain developed. Tom Albright, Laura Abavare and Bart Krekelberg from the Salk Institute for introducing me to the scientific (and not so scientific) aspects of life in Southern California.

I am grateful beyond words (and genes) to my sister, Silke, and to my parents, Regine and Hans, who always encouraged me — at any age — to explore the world. Thanks for your passionate and educational stories when I was young, and for listening to mine when I was growing up.

Caltech has provided me with the opportunity to sample from such diverse fields as genetics (my minor) and white water kayaking (a newfound passion), and provided access to all the MR machinery and recipes to look at water in so many different ways. Thanks to everyone in the lab for many memorable experiences, like sampling spicy or otherwise exotic food (thanks, Ralph & Carol) and eccentric beer (thanks, David & Ian!).

Caltech has been a truly palatable experience !

Abstract

How do different parts of the brain work? The naive and somewhat ill-posed question nonetheless admits of a serious answer. Different regions of the brain carry out their function principally through two components: the pattern of inputs and outputs that connect a region with the rest of the brain, and the computational transformations implemented by neurons within the region itself. Here we focus on the former problem and study the connectivity of the primate brain, with an emphasis on neocortex.

We develop a novel set of algorithms for modeling anatomical connectivity based on diffusion-weighted magnetic resonance (MR) imaging. The approach is novel in several respects: it utilizes a new way of deriving a globally optimal solution from local message passing; it can be applied to the whole-brain level in a computationally tractable fashion; and it can flexibly incorporate much other information, such as constraints about the geometry of white matter tracts and high-resolution anatomical MR images. The algorithm is first described as a hierarchical Bayesian model, and then applied to the diffusion MRI data obtained from two perfusion-fixed brains of macaque monkeys.

Based on the connectivity output provided from applying our novel algorithm to high-angular resolution MR data, we next derive several new insights about the connectivity of the macaque brain. We compare our results against those from published tracer studies, and we derive the relative weights of connections known from such prior studies. We also demonstrate the ability of the algorithm to generate entirely novel connectivity data, both at the level of specific anatomical regions that are queries, and also at the whole-brain level. The latter permits new insights into whole-brain connectivity and its architecture.

In addition to this focus on the structural connectivity of the macaque brain, we also analyze an extant set of public data of BOLD-fMRI from the macaque brain. This data set yields information regarding the functional connectivity of the macaque brain that we put together with our new connectivity results in order to relate structural and functional connectivity, with several new discoveries about their relationship.

In the final chapter, we apply these methods to MR data we collected from the live human brain. We provide an overview of structural and functional connectivity results obtained from this data set, and we apply the investigation to the brains of rare patients with agenesis of the corpus callosum, who lack the normal connection between the left and right hemispheres. We close by illustrating the power of the approach to ask questions that integrate functional questions with connectivity information on which function must ultimately be based: using connectivity profiles in order to segment cortical regions based on their pattern of inputs and outputs, with the aim of then querying these segmented regions using fMRI in cognitive activation studies. The description of our algorithm, the demonstration of its reliability, validity, and application to yield new data, together with the extensive software libraries on which the work is based, will provide cognitive neuroscientists with an array of new tools to investigate brain function in both health and disease.

TABLE OF CONTENTS

List of Developed Software	xi
List of Data Sets	xii
Nomenclature	xiii
List of Figures	xiv
List of Tables	xix
1. Introduction	1
2. Background	7
Neuronal Microstructure and MRI	7
Classical Methods to Study Connectivity	18
Tractography	21
Functional Connectivity and Ultraslow Oscillations	27
Current Knowledge of Connectivity	33
3. Methods	
Graphical Models, Factor Graphs, Belief Networks, and Markov Randon	n Fields37
Path Analysis and Dimensional Embedding	
Anatomical and Diffusion Imaging	44
Preparation and MR imaging of Monkey Brains	45
Anatomical Scans and Cortical Parcellation	47
Diffusion Imaging	51
Probabilistic Fiber Tracking	52
Combining Functional Imaging and Fiber Tracking in Humans	55
Multitaper Analysis of BOLD Time Series	56
4. Belief Propagation of Connectivity	57

	Model-Based Fiber Tracking	58
	Recursive Estimation of Fiber Tracks	62
	Bayesian Model of Fractional Anisotropic Diffusion	65
	Bayesian Model of Fiber Orientation	71
	Increasing the Effective Resolution by Including Anatomical Information	74
	Belief Propagation of Connectivity	76
	Message Passing Implementation	77
	Results	82
5.	What One Can Do and Cannot Do with Tractography	91
	Reliability of Fiber Tracking	94
	Tracing Anatomical Connection in the Absence of the Corpus Callosum	100
	Validity of Tracking Fibers	106
	Weight Estimation	115
6.	Connectivity of the Macaque Brain	120
	Strength of Selected Cortical Connections	120
	Strength-Based Analysis for the Visual Motion Pathway	132
	Strength-Based Pathway Analysis of the Visual Stream	137
	Strength-Based Pathway Analysis of the Motor System	142
	Identification of Anatomical Subnetworks Using Locally Linear Embedding	145
	Functional Subnetworks Based on Ultraslow Oscillations of Blood Oxygenation	on 157
	Importance of Strength Information for Functional Networks	162
7.	Exploring Connectivity of the Human Brain	168
	Large-Scale Coherence Networks in Humans	169
	Weight Information and Functional Connectivity	177
	Connectivity of the Fusiform Gyrus and Face-Selective Activation	179
	Bibliography	185

х

LIST OF DEVELOPED SOFTWARE

Software Package Name	Number of Functions	Lines of Code (Thousand)
AnaFun (Anatomy and Function) Scripts Integrates cortical parcellation results from FreeSurfer and Caret with statistical parametric mapping and surface-based statistics.	65	3
Connectivity Belief Propagation Toolbox Implementation of belief propagation and message passing networks for fiber tracking (see chapter 4).	20	3
Brain Graph Toolbox Graph theory and path-analysis methods for anatomically and functionally defined networks. Uses data from FreeSurfer, fsl, caret and spm.	220	13
Coherence Toolbox BOLD coherence computations and artifact removal for functional time series. Uses information from FreeSurfer, caret and spm5; based on chronux.	100	5
EEG-fMRI Toolbox Removal of MRI gradient artifacts from concurrently recorded EEG signals. Adaptive noise cancellation and various filtering approaches.	30	2.5
Tensor View Toolbox Visualization functions for 3-dimensional orientation tensor fields. Supports anatomical labeling.	10	0.5
Tract View OpenGL 3-dimensional visualization of fiber tracts and tensor field line integration. Written by Arjun Bansal and D.N.	2	7
Vector Tractography G5 / AltiFiber Very fast, pipeline-optimized assembler implementation of probabilistic fiber tracking using the PowerPC/G5 vector 128-bit engine.	5	5

LIST OF DATA SETS

Data Set	MRI Sequence	Individuals	Size (GB)
AgCC anatomy acquired by D.N. in collaboration with Lynn K. Paul, Nao Tsuchiya, Mike Tyszka	DW-EPI MP-RAGE	6	10
Attention to faces EPI	DW-EPI MP-RAGE	8	35
CB57BL/6 provided by Mike Tyszka	DW gradient echo		
Macaque DWI acquired by D.N. in collaboration with Mike Tyszka and Josef Parvizi	DW spin echo FLASH	2	250
Macaque EPI publicly available (Vincent et al. 2007)	EPI T1	8	30
Human Anatomy	DW-EPI MP-RAGE	50	50

Software versions used: Caret 5.51/5.6 (van Essen 2002), Chronux 2.00 (Mitra and Bokil, 2008), FreeSurfer 3 / 4.3, FMRIB Software Library (FSL) 4.0/4.1, FMRIB Diffusion Toolbox (FDT) v2.0, Statistical Parametric Mapping (SPM) 5.

Online databases used: Collations of Connectivity Data on the Macaque Brain: CoCoMac (http://cocomac.org), Surface Management System Database: SumsDB (http://sumsdb.wustl.edu/sums/)

NOMENCLATURE

- ACC anterior cingulate cortex
- AgCC agenesis of the corpus callosum
- BOLD signal blood oxygenation level dependent signal
- BPN belief propagation network
- DLPFC dorso-lateral prefrontal cortex
- dMRI diffusion-weighted magnetic resonance imaging
- DTI diffusion tensor imaging
- DW imaging diffusion-weighted imaging
- EEG signal electroencephalography signal
- EPI echo planar imaging
- fMRI functional magnetic resonance imaging
- HMM hidden Markov model
- HRF hemodynamic response function
- KL divergence Kullback-Leibler divergence
- LFP local field potential
- LLE locally linear embedding
- MCMC sampling Markov Chain Monte Carlo sampling
- MR signal magnetic resonance signal
- MRF Markov random field
- ODF orientation density function
- PFC prefrontal cortex
- PFT probabilistic fiber tracking
- STS superior temporal sulcus

LIST OF FIGURES

Figure 2.1. Morphology of a neuron.	10
Figure 2.2. Electron microscopic image of a synapse.	11
Figure 2.3. Confocal image of a profile through a peripheral nerve fascicle.	13
Figure 2.4. White matter dissection of the human brain.	16
Figure 2.5. Hierarchical Bayes model used to constrain pathway geometry.	23
Figure 2.6. Individual normalized hemodynamic response.	28
Figure 2.7. Grand mean power spectrum and coherence of the raw LFP signal duri entire recording session.	ing an 31
Figure 2.8. The ratio of neocortical white matter relative to the volume of neocortic	cal
gray matter in insectivora and primates.	36
Figure 3.1. Overview of the Anatomical and Diffusion Processing Steps.	45
Figure 3.2. Reconstruction of Cortical Surface and Localization of Histologically De	efined
Areas.	48
Figure 3.3. Lateral view of probabilistic parcellation map of Lewis and	
van Essen (2000b).	50
Figure 3.4. Landmarks used for alignment of areas on surface map.	51
Figure 3.5. Outline of the Human Structural and Functional Processing Steps.	55
Figure 4.1. Segmentation of the model space.	61
Figure 4.2. Factor graph of the message passing algorithm.	63

Figure 4.3. Factor graph for a hierarchical local measurement model.	64
Figure 4.4. A simple hierarchical Bayes white matter model.	66
Figure 4.5. Illustration of fractional volume terms along a pathway.	68
Figure 4.6. Connectivity likelihood function. Probability density functions of empiriod	ically 69
Figure 4.7. Diffusion-based white matter classification.	70
Figure 4.8. White matter tracking using only fractional anisotropy values.	71
Figure 4.9. Hierarchical Bayesian model with curvature constraints.	73
Figure 4.10. Hierarchical Bayes model incorporating anatomical information.	75
Figure 4.11. Distribution of partial volume coefficients as a function of T1-weighted signal intensity.	MR 76
Figure 4.12. Propagation of Likelihood Information using Message Passing.	83
Figure 4.13. Tract-based statistics: Distance along the path.	84
Figure 4.14. Contribution of fractional anisotropy and curvature information.	85
Figure 4.15. Comparison between probabilistic tractography and belief propagation	. 87
Figure 4.16. Comparison between the resolution of fractional anisotropy estimates (500 μ m, left) and white matter likelihood that was computed from an anatomical scan (110 μ m, right).	88
Figure 4.17. Message passing tractography identifies voxels not found with a sample approach.	ing 89
Figure 5.1. Comparison of probabilistic and deterministic fiber tracking.	95

Figure 5.2. Example of a probability density function of fiber estimates from the p	rimary
motor cortices in a mouse brain.	98
Figure 5.3. Tractography improvement when combining both interpolation	
techniques.	100
Figure 5.4. Number of cortico- and subcortico-cortical connections found in six su	bjects
with an agenesis of the corpus callosum.	104
Figure 5.5. Projections from area MT: Identified with Fast Blue and with diffusion tractography.	-based 107
Figure 5.6. Variability of tracing results for injections in area MT.	108
Figure 5.7. Cortical parcellation atlas used for seeding the fiber tracking at the like	ely
location of more than 70 distinct cortical areas.	110
Figure 5.8. ROC when comparing the diffusion-based fiber tracking with tracer baresults.	ised 111
Figure 5.9. Odds ratio of correctness of fiber tracking as a function of connectivity strength.	113
Figure 5.10. Comparison of average and single-case fiber tracking with prediction distance.	s from 114
Figure 5.11. Weights of known connections.	119
Figure 6.1. Strength of connections with area MT.	124
Figure 6.2. Strength of connections from V3.	128
Figure 6.3. Strength of connections of secondary somatosensory cortex.	131
Figure 6.4. Transforming a weight network into a distance network.	133
Figure 6.5. Assessing multiple pathways between two brain regions.	134

Figure 6.6. Routing distance for connections through potential intermediate areas f	rom
primary visual cortex to area MT.	136
Figure 6.7. Routing distance from primary visual cortex.	140
Figure 6.8. Multinode path distance of visual areas to V1 correlates with the hierarc	chy
level of Fellman and van Essen (1991).	142
Figure 6.9. Routing distance from primary motor cortex.	144
Figure 6.10. Connectivity strength matrix.	146
Figure 6.11. Similarity between areas based on the connectivity profile.	148
Figure 6.12. Similarity between areas based on the connectivity profile for the left ar	nd
right hemisphere separately.	150
Figure 6.13. Tractography-based visual hierarchy.	152
Figure 6.14. Tractography-based hierarchy of cortical areas based on their multinod	e
connectivity to V1.	154
Figure 6.15. Tractography-based hierarchy of cortical areas based on their multinod	e
connectivity to M1.	156
Figure 6.17. BOLD clusters of similarly, functionally connected brain regions.	159
Figure 6.18. Coherence between areas of different clusters determined by BOLD	
coherence in anesthetized macaques.	161
Figure 6.19. Prediction of BOLD coherence as a function of anatomical connectivity	
strength.	163
Figure 6.20. Binary connectivity does not predict coherence.	165
Figure 6.21. Tracer results do not predict coherence.	166

xviii

Figure 7.1. BOLD coherence between cortical and subcortical regions.	171
Figure 7.2. Comparison of BOLD coherence between corresponding and	
noncorresponding areas in the left and right hemisphere.	173
Figure 7.3. Subcortical interactions cannot explain callosal coherence.	175
Figure 7.4. Cortical coherence pattern can in part be explained by subcortical BOLD	
fluctuations.	176
Figure 7.5. Monotonic relation between structural and functional connectivity.	178

LIST OF TABLES

Table 2.1. Classification of peripheral axonal fibers from Erlanger and Gasser (1937)).14
Table 2.2. Overview of main forebrain fiber tracts.	20
Table 2.3. Commonly used neuronal tracers.	20
Table 2.4: Overview of fiber tracking algorithms.	26
Table 5.1: Reliability (Kullback Leibler divergence) of interhemispheric connections	from
the mouse brain.	98
Table 5.2: Spatial noise equivalent to KL divergence of table 5.1.	98
Table 5.3: Reliability (Kullback Leibler divergence) of interhemispheric connections	from
the human brain.	99
Table 5.4. IDs of subcortical regions.	99
Table 5.5. Cortical IDs from FreeSurfer.	99
Table 6.1. Strength of largest connections with area MT.	125
Table 6.2. Cortical areas, their acronyms and Brodmann numbers.	126
Table 6.3. Strongest connections with area V3.	129
Table 6.4. Strength of connections to area SII.	130
Table 6.5. Path distance from area V1 to MT.	135

Chapter 1

INTRODUCTION

The neocortex of primates consists of hundreds, if not thousands, of densely interconnected regions. Functional specialization within these regions is reflected in a complex profile of connections. In the adult human cortex, the entire ensemble of cortical neurons are connected through approximately 60 trillion synapses (Shepherd et al., 1998). The total length of myelinated axons in the human cortex reaches 150,000 - 180,000 km (Pakkenberg and Gundersen, 1997; Pakkenberg et al., 2003) and axons occupy a large portion of our brains, with gray-to-white matter ratios ranging from 1.1 to 1.5 throughout our lifespan (Miller et al., 1980). The degree of connectivity varies considerably among different brain regions and shares many architectural similarities with man-made transportation and communication networks. Similar to the small-world property of most social networks, brain networks contain a few highly connected hubs that ensure short communication routes between any two brain areas (Bullmore and Sporns, 2009).

Although the brain resembles randomly generated scale-free and small-world networks in some respects (such as short average path length), the development of connectivity is highly organized in space and time and anything but random. Myelination of white matter extends well into adolescence and coincides with the development of cognitive abilities, language skills and increased proficiency in executive decision making: connections to primary sensory and motor cortices are the first to myelinate, followed by association cortices and last by higher-order regions such as prefrontal cortex. Differences in white matter structure predict individual performance differences for working memory, reading tasks and musical abilities. White matter differences are also a hallmark of psychopathology, and constitute a current focus of investigation in regard to several psychiatric illnesses such as depression (e.g., Pezawas et al., 2005). Neuronal activity in white matter stimulates myelination *in vitro* (Ishibashi et al., 2006), and new white matter connections are likely to contribute to reorganization and plasticity following brain injury also in neurological patients (Bridge et al., 2008). There is even evidence of white matter plasticity in healthy adults: extensive practice of piano play has been found to alter music-related connectivity (Bengtsson et al., 2005). Unfortunately, the human brain loses about 45% of its myelinated connections between the ages of 20 and 80 years, at a rate of about 10% per decade (Marner et al., 2003). An acute loss of myelinated fibers, for instance in multiple sclerosis, can result in severe impairments in domains such as cognitive abilities, attentional span, reasoning and problem solving skills. The connectivity between neurons is essential for high-level cognitive functions and memory that is ultimately encoded in the synaptic connections between nerve cells.

Historically, connectivity has been primarily studied in rodents, cats and macaque monkeys. By injecting chemicals and dyes into the cortex *in vivo*, the origins of fibers targeting an area can be traced with great precision. From such studies, we have learned that the brains of both species are highly interconnected and are organized in partially hierarchical relationships (Young et al., 1995; Felleman and Van Essen, 1991; Hegde and Felleman, 2007). Such ordered relationships between brain areas is based on the inferred information flow between locally interconnected areas, whereby higher-level areas

receive input from lower-level areas. Yet, this ordering is only locally valid and a globally consistent, total hierarchy relationship does not exists (Hegde and Felleman, 2007). Compared to macaque connectivity, much less is known about networks in the human brain. Most tracers require active transport mechanisms, necessitating injections into live animals with subsequent sacrifice and histological examination, and thus cannot be deployed for studying human connectivity.

Over the last decade, magnetic resonance imaging sequences have been extended to measuring microstructural aspects of white matter tissue (Basser and Pierpaoli, 1996; Basser et al., 2000; Le Bihan et al., 2001; Tuch et al., 2003). By quantifying the effects that microscopic barriers have on locally diffusing water, diffusion-weighted MRI provides a unique way to image the parallel microstructural organization of axonal bundles. Diffusion-based tractography falls considerably short of matching the precision and reliability of classical tracing studies. Yet, it allows measurements of large-scale connectivity of the entire brain noninvasively.

The capability of collecting high angular resolution images on a whole-brain level has so far not been matched with mathematical tools to trace fibers globally. Here, we will introduce, validate, and apply a novel method for studying cortical connectivity on a whole-brain level.

First, the current knowledge and various techniques for studying connectivity will be reviewed in chapter 2. It will provide a quick overview of the neuroanatomy relevant to connectivity, of tracer-based methods for studying connectivity *in vivo* and we will review principles of diffusion-weighted MRI and extant tractography algorithms. The chapter concludes with an overview of the current state of our knowledge of cortical connectivity and its evolutionary implications.

The next chapter provides an introduction to probabilistic methods relevant to tractography. It reviews graphical models such as factor graphs, Bayesian belief networks and Markov random fields, and discusses inference algorithms such as message passing, belief propagation and Markov chain Monte Carlo sampling. Additionally, it provides details about the methods used to prepare the brain samples, the MRI diffusion sequences used and the anatomical and functional preprocessing pipeline.

Chapter 4 describes a novel, optimal whole-brain connectivity algorithm. It first introduces ideas of global, model-based tractography in a Bayesian framework and then takes these concepts to the systems level by extending them to belief propagation and message-passing networks. The algorithm finds all optimal pathways to all locations in the brain. Unlike previous methods with exponential complexity, this algorithm converges to the optimal solution in a fixed number of steps and exhibits a low, polynomial complexity. This new algorithm forms the core of the dissertation, and global tractography is the basis on which the subsequent results were obtained in the later chapters.

The limitations and advantages of tractography are explored in chapter 5. We first study the reliability of fiber tracking by tracing callosal connections, both in mice imaged with a high field 11.7T MR system, and in humans imaged at 3T. We then quantify the intrahemispheric agreement with information-theoretic measures. Next, we compare the tractography results with a large number of published tracer studies and quantify the influence of fiber strength and distance on reliability. Most importantly, we show that diffusion-based tractography can predict connectivity strength and is in close agreement with classical tracer results. These results constitute an important validation of tractography and open the door for discovering likely new connections with tractography alone, research findings that we present in later chapters.

Measuring the connectivity weight parametrically with tracers is a laborious and technically challenging undertaking (as compared with just finding whether or not there is a connection at all, also not a trivial undertaking). The practical hurdles have so far limited our knowledge of quantitative aspects of cortical information exchange. Diffusion-weighted magnetic resonance imaging (dMRI) techniques have made it easy to image efficiently the effects of axonal composition within larger cortical fiber tracts, and we will use this method to estimate the approximate connectivity strength of cortical tracts. In chapter 6 we will report such weight estimates for neocortical connections of the macaque brain, highlighting some new discoveries with this approach.

MRI can measure anatomical and functional connectivity at the whole-brain level. To gain further understanding into the global function of the distributed cortical network, it is necessary to look beyond just pairwise interactions and to identify functionally relevant network properties and processing pathways. Previously, distributed information processing has been analyzed in terms of interacting neuronal pathways and partial hierarchies of cortical processing steps. In the second part of chapter 6 we introduce analytical techniques to obtain such pathways and partial hierarchy information with objective, unsupervised clustering approaches and network analysis techniques.

Although the activity of cortical connections cannot be directly imaged with current MR techniques, the influence that an area has on a distant connected tissue can be inferred using functional MRI. Functional MRI techniques are sensitive to the blood oxygenation level in the cortical capillaries and thus indirectly measure metabolic activity in surrounding gray matter (Logothetis, 2008). They are frequently used to measure the statistical coupling between metabolic fluctuations in distant cortical areas (Fox and Raichle, 2007), so-called BOLD coherence. This statistical coupling is thought to reflect functional connectivity between the regions, and thus is presumably mediated by the underlying structural connectivity, although this relationship has not been investigated. We will show, in chapter 6 for macaques and chapter 7 for humans, that BOLD coherence between brain areas is primarily due to the anatomical strength of the connecting fiber tract. A strong BOLD coherence was observed between the thalami and the corresponding ipsilateral cortical hemispheres and between homotopic cortical areas of both hemispheres.

Taken together, the chapters in this dissertation provide a mathematical foundation, an application, validation, and derivation of new data on the connectivity of the primate brain. The tools developed here can be used by other investigators to explore a number of specific research questions about differences in connectivity between different species, about connectivity between particular brain regions of interest, and about alterations in connectivity in neurological and psychiatric illness.

Chapter 2

BACKGROUND

This chapter will give a brief overview of neuronal microstructure and its relevance to structural and functional MR measurements. Classical methods to study connectivity will be reviewed briefly, and an overview of principles of diffusion MRI and tractography will be provided. Cortical function can be studied with functional MRI sequences by imaging blood flow and blood oxygenation changes. The functional connectivity between brain regions can be quantified by a variety of methods, including spectral coherence methods of BOLD fluctuations. The chapter concludes with an overview of the current knowledge about the evolution and global organization of cortical connectivity.

Neuronal Microstructure and MRI

Information is processed in the brain in a highly distributed way, relying on constant communication between spatially separated neuronal tissue. At each length scale, the cortical tissue is organized differently and various physical properties of the microstructure of cortical connectivity can be imaged with MRI techniques. The main communication mechanism of neurons relies on changes in electrical potential across the cell membrane. To communicate, neurons are connected by extrusions of the cellular membrane, called axons at the sending end and dendrites on the receiving end (although there are some exceptions to this polarity as well) (see figure 2.1). These neurites form treelike structures that enable the cell to exchange information with typically thousands to tens of thousands of other nerve cells. Dendrites propagate the incoming information from neighboring neurons to the soma, the cell body of the neuron. Outgoing electrical pulses are transmitted by the axon, often over long distances. In addition to replenishing ionic gradients necessary for the transmission of electrical signals within the cells, a large fraction of a neuron's energy budget is spent on interneuronal communication at synapses, the information-transmitting contact points between axons and, for the most part, dendrites. The metabolic processes necessary to sustain synaptic function, such as vesicle turnover and neurotransmitter packaging for exocytosis, have a significant effect on local blood oxygenation and blood flow. Changes in blood flow or oxygenation are measured with functional MRI. Consequently, the fMRI signal reflects primarily synaptic activity, i.e., inputs to a neural region and intrinsic processing, rather than spiking outputs (Logothetis and Wandell, 2004).

Microstructural imaging, on the other hand, can be used to generate contrast between different types of tissue, based on a large number of different possible imaging parameters. Diffusion-weighted imaging, the focus of this dissertation, exploits the directional anisotropy in the diffusion of water within tissue to infer the microanatomy of spatially oriented tissue (Basser and Pierpaoli, 1996). The measured orientation field can be used to reconstruct the pathway of axonal bundles (Basser et al., 2000; Conturo et al., 1999; Mori et al., 2002; Behrens et al., 2003a). Although the structural elements relevant to MRI, such as synapses and axonal segments, are tiny, many neocortical tissue types are homogeneous and can be imaged unambiguously due to their bulk properties (Basser and Pierpaoli, 1996). More microanatomically complex tissue types result in mixtures of signals and more sophisticated algorithms are necessary to unmix the resonance signal (Basser, 2002; Jones and Basser, 2004; Tournier et al., 2004; Jian et al., 2007) and evaluate structural (Behrens et al., 2003b) or functional hypotheses (Kriegeskorte et al., 2006). The diffusion MRI signal is sensitive to the microstructure of tissue and hence not specific to axons nor neurons. Ultimately, the limiting factors for brain imaging are thus not technical constraints of MR imaging, but the biological granularity of neuronal representations and the organization of neurons into larger functional units. Currently, the major limitation of diffusion MRI is the low amplitude of the magnetic resonance signal. The low signal-to-noise (SNR) efficiency requires long imaging times and effectively limits the spatial resolution of MRI.



Figure 2.1. Morphology of a neuron. The main functional compartments of a neuron are the (1) *cell soma* (shown cut open to reveal internal compartments) containing the nucleus and various organelles involved in metabolism and protein synthesis, (2) *dendrites* receiving chemically encoded information at the synapses and (3) *axons* transmitting action potentials (inset shows internal organization). (4) *Synapses* are the contact points between axons and other cells.¹

At the synaptic level, information is relayed chemically across the 20 nm wide synaptic cleft (see figure 2.2). In agreement with the direction of action potential propagation, the axonal, transmitter-emitting site of the synaptic protuberance is referred to as the presynaptic site. At the postsynaptic side, the neurotransmitter signal is received by membrane-bound receptors and encoded as an ion flux or as changes in

¹ Illustration by Mariana Ruiz Villarreal, 2007, public domain.

intracellular so-called second messenger molecules. The synaptically generated currents change the membrane potential over small distances and thereby convey information from the dendrites to the soma. The synaptic chemical and electrical processes are energy intensive and influence the blood oxygenation in the local environment of the neuron. Such oxygenation changes, when occurring concurrently on millions of synapses, are detectable with fMRI. At the soma, the voltage signal on the membrane is amplified by active ion channels and the electrical action potential is propagated along the axon.



Figure 2.2. Electron microscopic image of a synapse. Presynaptic vesicles contain neurotransmitter that can be secreted into the synaptic cleft. At the postsynaptic side, the neurotransmitter binds to membrane-bound receptors.²

² Image reproduced from http://txtwriter.com/Backgrounders/Drugaddiction/synapticcleft.jpg

Neocortical axonal fibers are typically just 1 μ m wide and can span the entire extent of the brain. Most of our current knowledge about axonal processes stems from studies of peripheral nerves. The morphology and functional variance of peripheral nerves are larger than in the cortex and provide insight into the trade-offs of size, speed and energy expenditure of action potential propagation. The majority of cortical axonal connections are local and the neuronal (and axonal) cell membranes are in contact with glial cells (astrocytes and oligodendrocytes) to supply oxygen, nutrition, and various auxiliary functions. For long-distance connections, axons form hierarchical structures. In the peripheral nervous system, a nerve bundle encompasses multiple fascicles that each bundle axons and blood vessels (see figure 2.3). Structurally, this hierarchical organization is supported by connective tissue. Axons in the peripheral nervous system are enclosed by Schwann cells and interconnected by the collagen-rich endoneurium to form clusters of axons. The cellular perineurium forms multiple concentric layers around each fascicle and is enclosed by the thicker epineurium that acts as a protective shell around peripheral nerve tracts. It is important to note that the extra-axonal structures show a parallel organization in-line with axonal direction and that many intracellular diffusion barriers mirror this arrangement. The fine, elongated microstructure of the axonal membrane, as well as of parallel aligned microfilaments and microtubules, are anisotropic diffusion barriers for the enclosed water. The resulting local apparent diffusion anisotropy influences can be measured with diffusionweighted MRI (dMRI; Basser et al., 1994).



Figure 2.3. Confocal image of a profile through a peripheral nerve fascicle. The axons show a parallel organization and are joined by connective tissue. Diffusion MRI measures the joined effects of diffusion barriers such as cellular and intracellular membranes. It is, however, not specific to axonal neurons and contains contributions from connective tissue and potentially smaller blood vessels.³

Peripherally and centrally, distal axons can be enclosed by myelin sheets that electrically insulate the axons and minimize leak currents. The alternation of mostly passive, electrically insulated stretches of membrane, together with actively amplifying zones containing voltage-gated channels (nodes of Ranvier) enable a very fast, saltatory form of action potential propagation. The transmission speed of axons is primarily a function of axonal diameter and myelination, which specify the longitudinal resistivity and capacitance, respectively. Fast axons, for instance in the motor and sensory systems, are typically larger and myelinated, whereas autonomic and pain fibers can be unmyelinated and as thin as cortical axons. The pioneering work of Erlanger and

³ Image reproduced from http://www.mhhe.com/biosci/ap/histology_mh/nervxsl.jpg

Gasser (1937) resulted in a velocity-based classification scheme for peripheral fibers that is summarized in table 2.1.

Table 2.1. Classification of peripheral axonal fibers from Erlanger and Gasser (1937).The table shows the nomenclature for the figure type, together with its characteristicproperties and an example of where in the nervous system it might be found.*Table adapted from* Longstaff, 2005, *p. 12.*

Fiber Type (Erlanger/Gasser)	Mean Diameter (µm)	Mean Velocity (m s ⁻¹)	Sample Location
Αα	15	100	Motor system
Αβ	8	50	Touch afferents
Αγ	5	20	Muscle spindles
Αδ	4	15	Skin temperature afferents
В	3	7	Unmyelinated pain afferents
С	1	1	Autonomic neurons

If axons were to randomly connect cortex, diffusion-weighted MRI could not resolve connectivity, since it would constitute a thin and highly distributed web of axons below the level of spatial resolution of the technique. During brain development, however, axons follow chemical gradients and adhesive cues in a more orderly and often topographic manner. As a result, the neurites travel in bundles of similar neuronal connectivity and with a highly parallel alignment. Multiple bundles of axons form larger fascicles. These larger fascicles connect adjacent and distant brain areas and are directly visible in anatomical dissections. This high degree of organization makes it possible to study the microstructure of white matter with the relatively coarse imaging resolution afforded by diffusion MRI in many parts of the brain. Some of the major fiber tracts in the brain are summarized in table 2.2, together with their putative function and the symptoms that can arise when they are lesioned. This emphasizes the point, often neglected in lesion studies, that damage to white matter can have profound consequences due to disconnection of target structures. White matter diseases, such as multiple sclerosis and other demyelinating diseases constitute an important and major set of clinical presentations in neurology. It should be noted that most of the brain is white matter in terms of volume, as can be seen in blunt dissection of a view of the medial aspect of a hemisphere shown in figure 2.4.

The fiber bundles seen in figure 2.4 suggest homogeneity within these macroscopic bundles. However, this assumption of microstructural homogeneity does not always hold. In many cases, the unit of imaging resolution, called a "voxel" and typically one to a few cubic millimeters in volume, contains multiple neuronal tissue types. The issue is especially vexing when the organization of fibers is microscopically complex, combining multiple types or composed of axons that cross in several directions. Below, we will discuss techniques aiming at decomposing the contribution from tissue types and orientation information in the case of crossing and branching fiber tracts.



Figure 2.4. White matter dissection of the human brain. Dissection of the medial aspect of the right hemisphere of a human brain reveals white matter. Large white matter tracts, occupying up to 45% of the neocortical volume, connect the different parts of the cortex and provide connectivity to subcortical structures, sensory and motor systems.

Table 2.2: Overview of main forebrain fiber tracts. The table includes putative functions based on Schmahmann et al., 2007, as well as possible clinical syndromes arising from lesion. *Table next page.*

Fiber Tract	Connections	Putative Function	Lesion Symptom
Uncinate fasciculus	Rostral temporal, Orbital and medial prefrontal cortex	Interaction between emotion and cognition; novel information, emotional response regulation for sound	Impaired learning for visual scenes and object- reward associations, emotional memory
Cingulum bundle	Hippocampus and parahippocampal, Prefrontal areas 9 and 46, rostral cingulate	Putative: motivational and emotional aspects of memory	Disruption of psychosocial behavior, obsessive- compulsive behavior, impaired memory.
Fronto- occipital fasciculus	Dorsal and medial parastriate and caudal posterior parietal, DLPFC, area 8	Spatial attention related to peripheral vision and motion	_
Superior longitudinal fasciculus I	Superior parietal, Premotor and supplemental motor areas	Higher order control and initiation of body- centered action	_
Superior longitudinal fasciculus II	Caudal inferior parietal and IPS, Posterior prefrontal cortex	Perception and awareness	Hemi-inattention and hemineglect
Superior longitudinal fasciculus III	Rostral inferior parietal, Ventral premotor and prefrontal area 46	Gestural components of language and orofacial working memory	Cortical dysarthria, aphemia, oral and buccal apraxias
Extreme capsule	Superior temporal and insula, Area 45, 47 in orbital and dorsolateral prefrontal areas	Linguistic aspects of language	Putative: associative aphasia
Arcuate fasciculus	Caudal superior temporal, Caudal dorsal prefrontal cortex	Spatial processing and attention related to sound localization; language	Conduction aphasia
Middle longitudinal fasciculus	Superior temporal, Paralimbic regions	Spatial organization, memory and motivational valence	_
Inferior longitudinal fasciculus	Occipital, Temporal lobe	Object recognition and discrimination, memory	Putative: prosapagnosia; facial emotion recognition
Corpus callosum	Corresponding areas between hemispheres	Synchronization and transfer between homotopic areas	Lateralized functionality (split brain), Cognitive impairments (AgCC)
Classical Methods to Study Connectivity

All anatomical tracing methods require some form of transport mechanism to image axons, either relying on passive diffusion or on active cellular transport mechanisms. Passive diffusion mixes water only over short length scales. Within 100 ms, water diffuses within a 10 µm radius at room temperature, but larger distances require a disproportionally longer mixing time. The diffusion is inversely proportional to distance and passive transport is extremely slow with increasing length. Some lipophilic dyes have been used to trace short-distance local connections. These carbocyanine dyes (such as DiI and DiO) have the advantage that they can be used in postmortem tissue, but they have the disadvantage that they are not directionally specific (they diffuse anywhere along the axonal membrane) and that due to the considerations just noted they only diffuse rather short distances (millimeters).

To quickly transport vesicles and proteins, cells developed a fast, active transport mechanism. Such particles are transported from the soma (where all protein translation occurs in neurons) to distant synapses (anterograde transport), as well as from axonal terminals inward (retrograde transport) via an actin-dependent mechanism. The fastest components of vesicular transport can achieve speeds of about 50-400 mm/day. Actively transported tracers first passively diffuse at the injection site in the extracellular space and are then taken up by the surrounding neurons via endocytosis or through transporters. Once inside the cells, the transport mechanism typically propels the vesicle-encapsulated molecules along the neurites. These tracers only work *in vivo* since they require ATP-dependent active transport, but they are considerably faster, more

specific, and often directionally transported. Various retrograde and anterograde tracers have been developed for tracing cortical connectivity, and some of the more commonly used ones are summarized in table 2.3 below.

The tracing of connectivity in the human brain has relied primarily on passive diffusion mechanisms in postmortem tissue. *Ex vivo*, lipophilic tracer molecules have been used to trace larger fiber bundles. Their passive diffusion rates are very slow and the obtained results, typically only for short distances of a few millimeters, are much less reliable than those obtained with modern, actively transported dyes.

Table 2.3: Commonly used neuronal tracers. The specificity of many is still unclear;several of those listed under retrograde are also transported anterogradely.*Table adapted in part from* Oztas, 2003.

Direction	Mechanism	Tracer	Туре
Anterograde		Phaseolus vulgaris leucoagglutinin (PHA-L)	Monosynaptic
	Radiolabelled amino acid	Tritiated glycine Tritiated proline	Polysynaptic
Retrograde		Horseradish peroxidase (HRP) Wheat germ agglutinin	Mostly Monosynaptic
	Nissl reaction	Methylene blue Toluidine blue Thionin Cresyl violet	Stains neurons
	Fluorescence	Lucifer yellow, Fast blue, Nuclear yellow	
	Viruses	Herpes simplex Adeno virus Pseudo rabbies	Polysynaptic
	Toxins	Cholera toxin-B Tetanus toxin-fragment	
Isotropic diffusion	Carbocyanine	Dil, DiO DiAsp	Membrane-bound (polysynaptic)

Due to the reliance of most tracers to be actively transported within the neurites, and the ethical infeasibility of injecting tracers into live humans, our knowledge of human brain connectivity is comparatively sparse and generally inferred from tracer studies in other species. Tracer studies in human brains have been restricted to passively diffusing tracers that have revealed some information about intrinsic, local connections in the brain, as noted above. Response of the brain to injury can reveal connectivity information too, through processes such as Wallerian degeneration. After cortical injuries to a brain area, e.g., after a stroke, there is both anterograde, and in some cases also retrograde, degeneration in distal processes and cells. The degeneration in a formerly connected area can be quantified postmortem, and such observations have been used historically to make some very macroscopic inferences about connectivity.

Tractography

To reconstruct fiber tracts, physical measurements about apparent water diffusion over the length of several micrometers are imaged at resolutions several orders of magnitude more coarse. Mathematical models are necessary to bridge this gap between the relatively poor spatial resolution of imaging techniques and the underlying microphysical processes. The most promising of the current algorithms aim not to resolve, but to model the uncertainties resulting from the scale mismatch (Behrens et al., 2003b; Jbabdi et al., 2007; Sherbondy, 2008). Many heuristics and, more recently, explicit probabilistic models have been proposed to tackle the problems. The next paragraphs provide a short overview of such fiber tracking algorithms and earlier, simpler approaches that do not account for the scale mismatch.

For simple white matter architectures without crossing or splitting fibers, simpler line integration approaches can be suitable and are computationally very efficient (Mori et al., 1999). Such algorithms identify the primary or most likely orientation in each voxel and extend the fiber tract estimate according to the estimated orientation. Mathematically this estimated orientation is a tensor, and therefore these earliest methods gave the name diffusion tensor imaging, or DTI, that is still commonly applied to all tractography methods today (even those not actually using a tensor model). The precision of such single-fiber tensor estimates can be increased by adapting high order filters, such as Euler's approximation and Runge-Kutta methods (Basser et al., 2000). Uncertainty bounds on such estimates can be estimated with bootstrap methods. They differ fundamentally from more sophisticated methods, such as the method of Jbabdi et al. (2007) described below, in that they do not compute a probability distribution, but deterministically operate on the highest or most probable values alone.

To bridge the gap in length scale, one would like to construct models of how the shape and location of anatomical fiber tracts change bulk molecular diffusion measurements. An elegant method has been proposed by Jbabdi et al., 2007, detailed in figure 4.5. The authors start with a geometric model of a fiber tract. They incorporate prior knowledge about the smoothness and the typical shape of cortical connections into a Bayesian model, and then construct a hierarchical probabilistic model to directly link the anatomy of a fiber to the diffusion-weighted MR measurements. The levels of the hierarchical model correspond to physical realities at decreasing length scales. At the

highest level, the approach begins with the assumption that two selected regions of interest are connected by a fiber tract, although it is important to keep in mind that this need not be monosynaptic. The model then incrementally refines its resolution at every hierarchy level, carefully bridging the decreases in length scale with probabilistic models describing the increase in uncertainty of the physical processes. Geometrically, the fiber locations are splines that are controlled by 2 control points in the source regions and a set of 3 additional points controlling the shape. At the next level, the spline-based pathway implicates anisotropic diffusion and orientation measurements that are consistent with the tangents along the curve specified by the spline. The implied constraints on the orientation distributions then predict the differences in magnetic resonance signal between diffusion-weighted and unweighted images.



Figure 2.5. Hierarchical Bayes model used to constrain pathway geometry.(a) Bayesian model of connectivity (c), spline (K, L, F) and partial volume model parameters. (b) Potential connections between two regions are modeled by a planar Catmull-Rom spline with 5 control points. The random variables of the model are

shown in circles and the dependencies between the variables are indicated by arrows. *Figures are reproduced from* Jbabdi et al., 2007.

Given a particular spline-based model, one would expect a diffusion-related attenuation along the pathway in the direction of the spline's tangents. If two areas are not connected, then a white matter tract should not be found and potential spline pathways should travel through a few voxels that do not contain white matter, or contain orientations inconsistent with the pathways. These two possibilities specify two alternate prior hypotheses: either the regions queried are connected, or they are not. Bayesian inference techniques can be used to compute from the posterior observations the likelihood for both of these prior conditions and the ratio between the loglikelihoods provides an estimate of how strongly the measured data favor one or the other hypothesis. The basic idea of Bayesian inference is to estimate the likelihood distribution over hidden parameters by applying Bayes rule.

In the case of tractography, the number of observable data points along the pathway and the model search space can be very large. To limit the effect of overfitting, additional prior constraints need to be added to the model. In the model of Jbabdi et al., 2007, the position of control points is restricted to a planar space and the central control point is fixed at the halfway plane between the source and the target regions. The algorithm is initialized with a distribution of control points describing potential pathways and the likelihood of those control points and the associated pathways is computed. A Monte Carlo sampling procedure is deployed to generate likely pathways similar to ones with a good fit and further iterations result in a convergence against a stable, ergodic distribution of control points. Distributed inference methods have been applied to the tractography problem before, although not in a Bayesian framework. Campbell et al., 2005, used information about the reliability of local orientation information to expand the wave front of potentially connected regions. Highly likely connections are thereby identified by the algorithm and connectivity information is propagated to neighboring voxels in parallel.

Many recent tractography algorithms combine elements from several of the above methods. Table 2.4 shows an overview of the historically relevant algorithms. In chapter 4 we will extend the Bayesian approach to fiber tracking to a global belief propagation network. Similar to previous methods, the pathway cost function is Bayesian and includes global smoothness constraints similar to spline-based fitting procedures. Unlike previous methods, the algorithm does not use a sampling approach and it finds all globally optimal pathways to each gray and white matter location.

	Orientation distribution	Interpolation method	Global fiber model	Whole brain estimation	Guaranteed optimal solution
Mori et al. 1999.	Tensor	-	-	-	yes
Basser et al. 2000.	Tensor	Euler, Runge-Kutta	-	-	yes
Lazar et al. 2003.	Tensor	Deflection	-	-	yes
Behrens et al. 2003.	ODF	Euler	-	-	-
Campbell et al. 2005.	ODF	-	Wave front	-	-
Jbabdi et al. 2007.	ODF	Splines	Bayesian	-	-

	Orientation distribution	Interpolation method	Global fiber model	Whole brain estimation	Guaranteed optimal solution
Belief propagation	ODF	1st and 2nd order smoothness priors	Bayesian	yes	yes

Table 2.4: Overview of fiber tracking algorithms. Tractography algorithms differ by their assumptions about the measurements and either model a single orientation ("tensor") or multiple orientations ("orientation density function", ODF) per voxel. Global pathways are generated either by line integration methods, or hierarchical Bayesian models that constrain the geometry of the tracts. For probabilistic tissue models ("ODF"), information can be propagated by heuristics or according to Bayesian models. Previous probabilistic methods only sample from the most likely connections and do not optimally solve the tractography problem on a whole-brain level. The connectivity belief propagation algorithm, in contrast, finds all Bayesian optimal solutions to every voxel in the brain.

Reconstructing the path of cortical fibers based on apparent diffusion speed measurements is challenging due to the limited spatial and angular resolution of MR scanners as noted above. The fiber reconstruction step is prone to accidental mislabeling of close-by fibers and it is limited to bigger bundles that are visible at imaging resolution. The large number of parallel axons within large fiber bundles induces a strong isotropy of water diffusion and thus MR measurements. Smaller bundles produce less profound changes in the magnetic resonance signal and are hence more difficult to trace. A major challenge is posed by crossing and neighboring pathways. Similar to first generation chemical tracers that could bleed into adjacent fibers, the mathematical reconstruction of possible fiber locations can occasionally connect segments of adjacent, but distinct pathways. To maximize the sensitivity of the imagebased reconstruction, and reduce its error rate, the brains comprising our data set were scanned with high spatial resolution in high-field MR systems (9.4T and 7T). To alleviate the risk of joining unconnected pathways, we used a reconstruction method that examines all potential connections and assigned a likelihood value to each possibility, akin to the approach by Behrens et al. (2007) but at the whole-brain level. By adjusting the threshold for the connection certainties we can restrict our analysis to the most probable connections, and trade off the algorithm's reliability against its sensitivity.

Functional Connectivity and Ultraslow Oscillations

Each cortical connection contributes to aspects of cognitive function. Impairments in connections, after stroke, as a result of surgery, following white matter disease, or due to genetic variations, can result in profound changes in cognitive ability. Most cortical areas are either directly connected or communicate through at most 2 synaptic transmission steps. The resulting interaction is thus very fast and below the time resolution of magnetic resonance imaging. Assuming a synaptic transmission time of tens of milliseconds, one would expect interactions below 100 ms, or about 10 Hz. However, due to the slowness of haemodynamic responses (Figure 2.6), only slower changes below approximately 0.5 Hz can be detected using functional MRI. One might thus have expected it to be impossible to image functional connectivity using MRI, yet surprisingly there appears to be a low-frequency correlate at the level of temporal resolution of BOLD fMRI.



Figure 2.6. Individual normalized hemodynamic response. This hemodynamic response curve is normalized to demonstrate that there are no differences in the rise and fall slopes. In the nonnormalized curves, the nonprimary areas show a much lower level of response and a lack of poststimulus undershoot (see below).⁴

Although the coupling of brain areas cannot be observed at the causal frequency band, the electrophysiological interactions on this fast timescale induce systematic coupling also at considerably slower frequencies. These lower frequency couplings are evident in the multispectral nature of the electroencephalogram (EEG), but they also result in haemodynamic correlates. The phenomenon has been mathematically modeled by Watts and Strogatz (1998). They studied weakly coupled physical and biological systems and found that above a critical coupling threshold, the systems start to synchronize and that this synchronization is a stable property of weakly coupled oscillators. The systems synchronize at their intrinsic oscillation frequencies that can be

⁴ Image reproduced from http://www.radiology.northwestern.edu/images/nr/investigate-hrf-2.gif

orders of magnitudes slower than the speed of physical interaction between the oscillators. An illustrative examples are a species of fireflies that developed light sensors to synchronize their bioluminescent activity at a frequency of several seconds (Strogatz, 2001). The speed of information transmission via the lights is almost 10 orders of magnitudes faster than the actual synchronized blinking frequency of the individual fireflies. Cortical columns exhibit intrinsic rhythms of seconds and minutes (Buzsáki, 2006) and the activity between distant cortical, but connected areas is weakly coupled.

The slow frequency coupling can be directly observed with BOLD imaging. In chapter 7 we show that the degree of coupling is not due simply to a direct effect of distance, but is primarily related to the strength of anatomical connectivity between two brain areas, relatively independent of their distance. Coherence of slow frequency oscillations has been studied electrophysiologically by Leopold et al. (2003) in monkeys. The authors placed an electrode grid across the lunate sulcus and measured the crosselectrode coherence as a function of the distance between selected electrode pairs. They found that gamma frequency coherence (in the range of 40-100 Hz) changes as a function of distance and that the fluctuations in the power of the gamma frequency band predicts the coherence in slower frequency bands. One important technical detail is the frequency dependence of the electrode-brain interface. Due to various electrolytes, one might expect a dampening of signals in higher frequencies. Logothetis et al. (2007) measured the impedance of their electrode system across a wide range of frequency ranges and found no effect of the electrode transfer function on the measured coherence. Given these electrophysiological observations, BOLD imaging is hence a powerful tool to measure the very slow frequency changes in the 0.1-0.01 Hz range.

The functional significance of the low frequency components detectable with BOLDfMRI is, for the most part, unexplored. This is mainly due to technical limitations, since most electrophysiological recording equipment, even if run in "DC" mode, have intrinsic high-pass filters, typically filtering out signal below 0.5 Hz. Figure 2.7 shows the power of a local field potential recorded with a Eckhorn multielectrode array (Leopold et al., 2003). Although on first sight, the 1/f power law seems to stop any measurable signal at around 0.5-1 Hz, with the appropriate amplifier system, no such attenuation in the ultraslow frequency bands is observed (Vanhatalo et al., 2002; Miller et al., 2007). Some validation studies have been conducted to distinguish the ultraslow EEG frequencies from respiration and motion artifacts that overlap in the same frequency range, demonstrating the presence of real signal. Little is know about the functional role of this frequency range; power changes within it have been linked to vigilance and self-initiated motor behavior (Gilden et al., 1995). In chapters 6 and 7 we show that the ultraslow oscillations are related to the underlying anatomical connectivity and that the coherence observed between two BOLD signals can be predicted by the strength of connectivity between two areas.



Figure 2.7. Grand mean power spectrum and coherence of the raw LFP signal during an entire recording session. Data were acquired from 15 electrodes in a monkey. For the data shown here, the monkey was awake and seated, but not engaged in any particular task. His scanning movements around the dim recording room determined the structure of his visual input. (a) Power spectral density of the LFP, averaged over all 15 electrodes. **(b)** Mean pairwise cross coherence between all electrode pairs in the array. The differently colored lines are different recording conditions, but all show substantial residual power and coherence even at low frequencies.

Figure reproduced from Leopold et al., 2003.

It is important to remember that magnetic resonance imaging cannot directly observe chemical or electrical neuronal activity.⁵ The most commonly used techniques are sensitive to change in either blood flow (rCBF) or the oxygenation of hemoglobin

⁵ Magnetic resonance spectroscopy can, however, measure concentrations of specific chemicals, such as intracellular osmolites and metabolites, typically with a time resolution of 10-15 minutes.

(BOLD), two measures often but not necessarily correlated. To measure functional changes and functional connectivity, both techniques utilize the fact that the blood supply to neuronal tissues is actively coupled to the metabolic activity of cortical tissue. If a cortical area increases its activity, which occurs mostly through the metabolic demands of synaptic events such as vesicle cycling, this metabolic activity will trigger an increase in local capillary blood flow and an increase in blood oxygenation. Both effects affect MR imaging sequences and can be used to determine whether a brain area changes its metabolic activity as a result of experimental manipulations.

In chapters 6 and 7, we are interested in using statistical dependencies between distant oxygenation responses as an indicator of a functional coupling between two brain areas. One difficulty of this approach is that many other global signals, such as blood pressure regulation, respiration and head motion, can affect the measured MR signal in the same frequency band. Various experimental controls have been conducted to rule out such signals as the sole explanation of observed long-distance correlations. However, it is not clear whether anatomical connectivity directly predicts ultraslow synchronization between distant brain areas. In chapter 6 we show that the predicted strength of anatomical connectivity obtained by diffusion-weighted tractography predicts the degree of ultraslow functional synchronization. Moreover, this correlation between the two measures is substantially better than the association between such functional coherence and the binary connectivity information of chemical tracer studies.

One aspect of measuring functional synchronization indirectly with BOLD-fMRI is that the BOLD response acts as an effective low-pass filter. Synchronization in frequency bands above about 0.5 Hz is thus not measurable, and the analysis needs to be restricted to ultraslow oscillations. In this thesis, we restricted ourselves to largely stationary background functional coupling that can be observed even at rest, and did not analyze dynamic changes in cortical functional connectivity, but EEG and LFP studies show that different frequency bands react differently, and in some cases, in opposite direction, to task demands and stimulus changes.

In the functional connectivity and resting-state literature, ultraslow coherence is assumed to be a uniform process. Due to the low-pass filtering properties of BOLDfMRI together with the slow drift components caused by technical aspects of the MR measuring equipment, the analysis is pragmatically restricted to the band-pass filtered signal in the 0.1-0.01 Hz frequency band. BOLD signals from two spatially separate regions are then correlated and the correlation coefficient is either taken as a direct indicator of functional connectivity, or explained in a vector autoregression or nonlinear framework. In either case, since the power in these ultraslow frequency ranges is disproportionally higher than power in the somewhat higher frequency ranges measured for standard functional BOLD-fMRI, one needs to deploy a spectral decomposition to analyze coherence in the ultraslow frequency spectrum.

Current Knowledge of Connectivity

During evolution, and in particular the evolution of mammals and their unique neocortex over the past several 100 million years, brain organization and notably white matter has undergone dramatic changes. Many species have explored new habitats and developed complex forms of interactive behavior that required new cognitive inventions, and often relied both on more processing power (larger brains) and more integrated processing (more connectivity). Brain anatomy and development are tightly related and feature specialized cortical architectures enabling new behavioral and cognitive abilities. The new capabilities are supported by enlarged neuronal white and gray matter, particularly in the prefrontal cortex during the recent phylogeny of mammals, and by a further specialization of already existing neocortical areas (Stevens, 2001; Cherniak et al., 2004; Sereno and Tootell, 2005). The expansion and specialization of gray matter has been accompanied by a disproportionately large expansion of white matter (Figure 2.8). Across species, neocortical white matter expanded at an exponential rate of about 4/3 the power of the neocortical gray matter volume (Hasenstaub in Allman, 2000, p. 165). The highest degree of connectivity can be found in the most highly encephalized mammals: primates. Relatively little is known about the changes in white matter organization itself and how the cortical fiber network has evolved to support new functionality. However, some general themes can be surmised. As brains evolved to be larger and larger, this required not only exponential volumes of white matter to achieve the same degree of interconnectivity, but also posed a problem with respect to the speed of communication. Mammals had evolved myelin to subserve rapid axonal conduction, and therefore fast communication between distal brain regions could use myelinated axons. But such myelination in turn increased the size of the axon and contributed to the ever-growing volume of white matter in brains that were already very large. This trade-off between communication and sheer volume requirements has likely been a limitation on brain expansion, and may have resulted in the increased apparent modularity of large brains (cf. Ringo et al., 1994).

Across species, increases in gray matter volume are predictably linked to increases in white matter by the power law relationship (Zhang and Sejnowski, 2000) mentioned above. The log-log ratio is constant across primate species (and about 4/3) with white matter volume increasing faster than the gray matter, as one would expect from expansion of surface and volume within a growing sphere. At first sight, this constant white-to-gray matter relationship might suggest a simple organization of white matter mirroring the expanding gray matter. However, white matter connectivity is highly specific and forms very structured communication networks that themselves constitute an important topic for evolutionary investigation. In their seminal study, Felleman and van Essen (1991) analyzed neocortical connectivity between visual areas of the macaque brain. The study used anatomical criteria about the direction of fibers, such as the originating and terminating cortical layer, and the direction of tracer transport, to classify each visual connection as feedforward or feedback. If two areas are connected, they are, in most cases, bidirectionally linked. The overall connectivity pattern shows an interesting structure. The organization of visual feedforward connections is consistent with a hierarchical processing architecture (see Chapter 6). This means there are distinct "levels" of connectivity, such that a "lower" area can route information to a "higher" area only through regions intermediate in the hierarchy. This anatomical finding mirrors our current understanding of visual information processing: visual features are extracted from the retinal image and are then iteratively refined. At the bottom of the hierarchy, in the primary visual cortex, basic local orientation information is encoded and higherlevel areas represent increasingly complex, global features. At the top of the hierarchy,

visual information leads to high-level memory representations in the hippocampus (Kreiman et al., 2000, 2002).



Figure 2.8. The ratio of neocortical white matter relative to the volume of neocortical gray matter in insectivora and primates. The red line indicates a log-log ratio of 1, the black line corresponds to the observed power law relationship with an exponent of 1.318. The white matter volume increases disproportionally when the gray matter volume expands.

Figure reproduced from Allman, 2000, p. 116.

Chapter 3

METHODS

This chapter will provide an overview of the mathematical and imaging techniques used throughout the thesis. It will first introduce computational networks, namely factor graphs, Bayesian networks and Markov random fields. These methods are used in the subsequent chapter to develop new fiber tracking algorithms. The anatomical and functional structure of networks will be analyzed with graph theoretical and lowdimensional embedding techniques. Two examples, the Dijkstra shortest path algorithm, and locally linear embedding, are discussed below. Additionally, an overview of the MR image acquisition and processing steps is provided. A description of the probabilistic fiber tracking used for parts of the thesis and the application of the spectral multitaper analysis conclude this chapter.

Graphical Models, Factor Graphs, Belief Networks, and Markov Random Fields

Extracting meaning by fitting complex models to high resolution data sets can be challenging. Frequent applications include the recognition of objects in images or the search for characteristic behavioral patterns in dense social networks. The key to solving these high dimensional problems is often to take advantage of local neighborhood relationships in the data. Graph theory is a powerful mathematical tool to capture the local structures of such problems and for developing inference algorithms on top of the network structure. A very useful computational class of models are *factor graphs* (Kschischang et al., 2001). Many complex computational and probabilistic problems can be simplified by factoring them into smaller, more tractable subfunctions. Similar to prime factorization, where any natural number can be expressed as a product of prime numbers, many functions can be written as a product of simpler factor functions. A function *g* that is defined over three variables could for example be expressed as the product of three simpler functions (f_1 , f_2 , f_3).

$$g(X_1, X_2, X_3) = f_1(X_1, X_2) f_2(X_2, X_3) f_3(X_1).$$

In general:

$$g(X_1,\ldots,X_n) = \prod_{i=1}^m f_i(S_j) \quad S_j \subseteq \{X_1,\ldots,X_n\}$$

A particularly practical property of factor graphs is that a recursive partitioning of a function into factors results in a hierarchical, treelike graphical model. Statistical inference methods on such hierarchical, loop-free models are exact, whereas inference on circular graphs, such as loopy belief propagation, may yield only approximate results.

In practice, the functions specified by graphical models can span thousands, and in the case of fiber tracking millions of locally interdependent variables. To visualize complex relationships in coding theory, a special graphical notation for factor graphs has been developed by Tanner (1981). The Tanner graphs shows the relationship between variables and factors of a network. Variables appear as a circles and factors are shown as squares (e.g., Figure 4.2 for a later example). The domain of a factor, the set of variables a functions acts on, are the edges of the graph. Solving large problems with thousands or millions of variables can be computationally challenging. By exploiting the locality of relationships through graphs, it is often surprisingly simple to approximate functions locally and to solve the global function by propagating information through the network. The exchange of local variable information in factor graphs is called *message passing*. The values that a factor function receives from its neighboring variables are called messages (Frey and Dueck, 2007).

For probability networks, the local computations are often simple arithmetic operations. Common inference mechanisms for factor graphs include sumproduct algorithms for computing conditional probabilities, max-product to obtain maximum likelihood estimates and min-sum algorithms for decoding problems. For finding pathways in diffusion images, we will for instance use a version of a maxproduct algorithm to identify all maximum likelihood solutions.

Factor graphs only specify the functional relationships between variables, without constraining the actual function. In contrast, Bayesian and Markov networks describe conditional probabilities and are spanned by parent-child edge relationships. Factor graphs are well suited for networks with many overlapping relationships. They are most useful for studying problems with complex, but local connectivity. Typical applications include low-level machine vision and identifying interaction patterns in gene and social networks.

An interesting subset of computational networks are *directional factor graphs*. Similar to Bayesian networks, they restrict edges to parent-child relationships. Directed factor graphs are closely related to Markov random fields and Bayesian networks (see Jensen

and Nielsen, 2007) and can express most instances of such models. They offer an alternative view of the information flow in such networks and they will be used in later chapters in conjunction with Bayesian formulations to describe the message propagation aspects of the fiber tracking algorithm.

Markov random fields emphasize the locality of interactions in large graphs. Nodes are typically only locally connected. The Markov property between variables follows directly from the connectivity structure of the graph: two nodes are conditionally independent if they do not share a common edge. The local neighborhood relationships are a complete model of the statistical dependencies of a node; a variable is conditionally independent of all graph nodes, given its neighboring variables. The classical example of a Markov random field is the Ising model that describes spin-spin interactions in physical chemistry. It turns out that fiber tracking shares many similarities with the Ising model. Neighboring orientation distributions are coupled in a similar way as spin states are in the simpler Ising model.

Frequently deployed inference methods in Markov random fields include Markov chain Monte Carlo (MCMC) sampling and loopy (recursive) belief propagation. With growing complexity, such methods become computationally expensive and can often only approximate results. For nonloopy Markov random fields and equivalent hierarchical Bayesian networks, exact and fast inference methods exist.

Bayesian networks share many commonalities with Markov fields. They capture conditional probability relationships between variables. An important application of Bayesian networks is the estimation of complex hidden variables, such as identifying class membership of visually distinct objects or the estimation of global shape properties

of fiber tracts through a diffusion MRI volume. Bayes networks offer an intuitive way of expressing complex interactions between various variables, such as local diffusivity of tissue, partial volume effects and noise sources.

Although not identical, the expressiveness of the various directed graphical models is quite comparable. When describing and analyzing the properties of the whole brain fiber tracking algorithm, the different notations can provide distinct insights into the various aspects of statistical inference. Markov random fields best explain the locality of interactions between variables, whereas hierarchical Bayesian networks will be used to link the algorithm to existing Bayesian tractography models. The computational aspects and the information flow through the inference network can be best highlighted using factor graphs and message passing analogies.

Previous Bayes algorithms applied to fiber tracking often estimate the complete posterior distribution over the network of hidden variables. Such fiber tracking approaches produce maps of the distributions of many possible pathways. For many practical applications, it is sufficient to identify the best paths and compute the exact probability of all optimal solutions, for example using maximum likelihood solutions.

The maximum likelihood solution in a Markov field can be obtained with the *Viterbi algorithm*. A network can be thought of as a series of local states with a finite, at times very large, number of states. Multiple sequences of states (e.g., pathways through the brain) through the network lead to a final state (e.g., a voxel in a distant brain region whose connectivity has been identified). To identify the best path, a cumulative metric is deployed when evaluating different possibilities at every node. Only the most likely pathways are kept for each state and all other possible, but less likely paths are

discarded. The search history leading to a state can be stored and used to backtrace the sequence of local events.

Path Analysis and Dimensional Embedding

The cortical network of specialized, interconnected areas, on first sight, appears very complex and random. Its topology has been compared to the properties of many realworld networks, such as national transportation networks and social graphs (descriptions of social relationships and interactions among people). Although it is quite densely connected, many characteristics of its connectivity, such as node degree distributions, clustering and centrality coefficients make it similar to randomly generated graphs where new nodes preferentially attach to existing nodes with high connectivity. Depending on the exact generative rules, the resulting networks are scalefree, with a power law distribution of the node degree, and show small world properties such as short pathways between nodes. In many respects the cortical network shows similar properties, indicating that generative rules such as preferential attachment are important during evolution and brain development.

Yet, brain connectivity is not randomly generated. Instead, it has evolved to selectively process information and support ecologically relevant cognitive and behavioral abilities. Similar to gene networks that need to constantly adapt enzyme production, the cortical network structure needs to be able to adaptively route information from primary sensory areas to high-level networks and integrate higherlevel representations with lower-level information processing. Assuming that the amount of information transfer is correlated with the anatomical size of a fiber pathway, path-analysis methods can be applied to identify the supposedly optimal routes of information through the densely interconnected cortical architecture. Optimal routes can be identified with the *Dijkstra shortest path* algorithm (Dijkstra, 1959). Starting from a start node, it identifies the distance to all connected nodes and updates their distance value. Subsequently, the most closely connected node of the start node is visited and the algorithm is repeated. Each node is only visited once and the algorithm efficiently identifies the shortest distance to each node from the start node. This algorithm will be later used to study the weight-based network distance of cortical areas and to compare this measure to the known hierarchical ordering of areas, such as the Felleman and van Essen hierarchy of visual areas mentioned earlier.

Such path-analysis methods can only estimate a multinode connectivity value between brain areas. The algorithm would, however, not be able to identify distinct parallel pathways of similar length originating in the same primary area. In such a case, for example when comparing the dorsal and ventral visual processing streams in the parietal and temporal cortex, it assigns similar distance values to areas in distinct pathways that have a similar hierarchy level. To identify separate pathways, we will use unsupervised clustering methods to segregate pathways that show distinct connectivity profiles. Our method of choice is locally linear embedding (Roweis and Saul, 2000), an algorithm that first compares feature values, that is the connectivity profile of areas, to identify similar nodes. In a second step it reduces the number of features to a lower dimensional manifold while maintaining the neighborhood relationship as far as possible. The resulting embedding is an excellent representation of the original similarity relationship and can identify pathways by identifying clusters of distinctly connected areas.

Anatomical and Diffusion Imaging

Studying brain connectivity for the entire brain is a computationally challenging task. A single in vitro scan of the macaque brain generates several gigabytes of data. Just to reconstruct the diffusion signal from the MR signal attenuation requires multiple days of computation. Reconstructing the likely location of originally histologically or functionally defined cortical areas is equally challenging and is at present only partly automated. The main step of then combining the diffusion signal with the cortical maps to reconstruct all likely fiber pathways and their probabilities could take weeks or months on conventional software and hardware. We developed our own, highly optimized implementation that efficiently estimates the probabilities with high reliability for the entire brain.

Figure 3.1 gives an overview of our anatomical processing pipeline. We first identified gray and white matter in the T1-weighted images using automated and manual segmentation tools. Using this tissue classification, 2-dimensional manifolds matching the gray-white matter boundary were reconstructed, anatomical landmarks identified, and histological parcellation maps were projected onto the surface. These reconstruction steps, mostly carried out with the software Caret 5.5, are documented in detail below.



Figure 3.1. Overview of the Anatomical and Diffusion Processing Steps. Each set of steps in the pipeline is indicated, together with the primary software used (90-degree rotated text to the right). Anatomical and diffusion processing are ultimately integrated in order to derive connectivity between segmented cortical regions of the brain.

Preparation and MR imaging of Monkey Brains

Two brains from two young adult cynomolgus monkeys (*Macaqua fascicularis*) were obtained as described earlier (Parvizi et al., 2006); both had been used in prior studies by Professor Josef Parvizi (Department. of Neurology, Stanford University Medical School), who kindly provided them for our study. Briefly, the monkeys were exsanguinted and transcardially perfused with 4% paraformaldehyde in 0.1 M phosphate buffer. Brains were then stored refrigerated after fixation. A few weeks before scanning, the brains were immersed in multiple exchanges of 1 mM sodium azide phosphate buffer to wash out the paraformaldehyde. The samples were subsequently immersed for at least 3 weeks in a gadolinium solution (ProHance, Bracco Diagnostics), a paramagnetic contrast agent that shortens T1 relaxation time. For imaging, the whole brains were mounted in perfluoropolyether (Fomblin PFPE, Solvay Solexis, USA) and imaged in 9.4T and 7T magnetic resonance systems (Bruker Biospin, Rheinstetten, Germany) at the Caltech Brain Imaging Center. To minimize temperature-dependent changes in the magnetic resonance signal, the gradient coils were insulated from the sample by a layer of self-adhesive polychloroprene.

T2 relaxation time of gray matter was determined in an initial series of spin-echo scans using different echo times. We first obtained 9 high resolution anatomical images. The images were subsequently averaged to improve the signal-to-noise ratio. Apparent diffusion speed was then obtained from the average of a custom high-angular resolution diffusion-weighted spin-echo sequence (Tyszka et al., 2006). Diffusion-weighting was applied in 72 directions with a nominal b-value of 3000 s/mm (actual values ranged from 2905 to 3104). In addition, 7 nondiffusion-weighted scans were obtained (b = 1.96 s/mm^2).

Anatomical Scans and Cortical Parcellation

High resolution anatomical scans were obtained using a Fast Low-Angle Single sHot (FLASH) gradient echo sequence. The images were taken with an in-plane resolution of 100 µm and a slice thickness of 150 µm. To increase the signal-to-noise ratio, we scanned 9 volumes that were coregistered using FSL to obtain an averaged volume that we then used to reconstruct the individual brain topologies. At higher field strength, T1 tissue contrast is inverted compared to lower field strength: white matter has a faster T1 decay constant than gray matter and thus appears darker in the images. Most software for reconstructing brain tissue anatomy is optimized for low field scans. To enable automatic reconstruction with current software, we smoothed the averaged volumes with an anisotropic diffusion filter in Amira (Mercury Computer Systems, Chelmsford, MA) and used a custom Bayesian classifier to segment white and gray matter. Subsequently, we inverted the class intensity means – by increasing the average intensity of the white matter voxels and reducing intensity of gray matter voxels - to resemble a T1 contrast at lower field strength. Using these images, the gray-white matter boundary was reconstructed in Caret 5.5 (Van Essen, 2002). The software computes the class statistics for the tissue classes, reconstructs a 2-dimensional Euclidian manifold along the gray-white-matter boundary and automatically corrects many topological errors of the manifold reconstruction. We converted the output manifold back into a volume format and inspected each hemisphere in detail. Misclassified tissue was manually corrected with Caret and the itkSnap 1.4.1 editor (Yushkevich et al., 2006). The improved masks where again reconstructed into 2-dimensional manifolds with Caret. The process was repeated a few times until the surface resembled the gray-whitematter boundary in sufficient detail (Figure 3.2). The resulting 2-dimensional cortical surfaces were deformed into a spherical representation that forms the canonical basis for Caret's realignment procedure.



Figure 3.2. Reconstruction of Cortical Surface and Localization of Histologically Defined Areas. Left: the high resolution anatomical image of a macaque brain obtained on a 9.4T Bruker MR system. Middle image shows the surface reconstruction of a standard atlas of the macaque cortex. Right image: Probabilistic parcellation map of cortical areas. Each color shows the boundaries of a histologically defined area (Lewis and van Essen, 2000b). The intensity and saturation of a color indicates the percentage agreement among different individual brains.

Cortical areas can be defined by anatomical, histological, genetic, connectivity, or functional boundaries (Toga et al., 2006). Characterizing such boundaries in a single hemisphere would be an enormous undertaking. Hence, it would be very beneficial if whole-brain fiber tracking could be performed using atlas-based regional boundaries. To test whether such an approximation would be sufficient to measure connectivity, we mapped the histologically and electrophysiologically defined cortical parcellation maps onto the reconstructed surfaces. We used a computerized registration procedure (Van Essen, 2002) to align the cortical surface of 4 parcellation atlases and the scanned brain.

The realignment procedure uses location, shape and size of sulci and gyri to adjust for the differences in folding pattern and area size between individual brains. For each hemisphere, we identified 26 anatomical landmarks on the reconstructed surface (Van Essen et al., 2001) and aligned the topologies of the individual cortices with a standard macaque brain template (F99UA1 template, van Essen, 2000). The alignment for every hemisphere was inspected manually and adjusted if necessary. Two parcellation atlases were used for the fiber tracking reported here, one for visual areas (Lewis and van Essen 2000b) and one covering the whole brain (Paxinos, 2000) (Figure 3.2). The atlas of visual areas maps functionally and histologically defined boundaries of about 90 visually responsive areas (Lewis and van Essen, 2000b) in the occipital, parietal, and temporal lobes. For studying full-brain connectivity, a more comprehensive map with the approximate location of about 150 areas covering all cortical lobes was used (Paxinos, 2000). Although such a fine-grained parcellation with precise borders is certainly beyond the specificity of diffusion imaging and fiber tracking, we found, on average, a better agreement between the fiber tracking results and tracer-based results when using finer compared to coarser cortical classification schemes. Seed regions for the probabilistic fiber tracking were identified by inverting the surface-based template registration procedure. The parcellation boundaries along the cortical surfaces were used to create 1 mm-thick gray matter sheets, located just above the white matter border. These masks were used as seed regions for the subsequent fiber tracking step.



Figure 3.3. Lateral view of probabilistic parcellation map of Lewis and

van Essen (2000b). Shown here is the probability, among four monkey brains, of finding a particular region of cortex within a specified anatomical parcel (indicated by the different colors for the different parcels, and by their saturation for the probability). The surface is a smoothed reconstruction of the gray-white matter boundary.



Figure 3.4. Landmarks used for alignment of areas on surface map. The same probabilistic parcellation atlas as shown in figure 3.3 was inflated so as to visualize all of cortex; the dotted lines indicate major sulci and gyri that were used for alignment between brains.

Diffusion Imaging

Apparent diffusion speed maps were obtained using a diffusion-weighted spin-echo sequence with a 500 μ m isotropic voxel size and 72 distinct directions. The nominal b-value of the diffusion scans was 3000 s/mm² (actual values ranged from 2905 to 3104). 7 nonweighted spin echo images (b = 1.96 s/mm²) were interspersed between blocks of

12 diffusion-weighted images to monitor temperature- and tuning-related changes in MR resonance. No significant changes in the baseline MR signal were detected and the nonweighted images were collectively used to estimate the apparent diffusion coefficient maps. The entire imaging procedure took approximately 20 hours per brain.

After the anatomical reconstruction and parcellation of cortex described above, we had obtained the approximate location of most cortical regions. The surface-based boundaries of the regions were used to label the corresponding voxels in the anatomical scans. The anatomical scans were normalized for nonuniform intensity biases inherent to the imaging equipment with N3 (Sled, 1997) and aligned to the reference frame of the diffusion-weighted images with the FSL software. Using this affine transformation, voxel labels were transfered from the anatomical to the diffusion image space using a nearest-neighbor interpolation. We then used the diffusion-weighted MR data and seeded the fiber tracking algorithm in these regions to compute the likelihood that any two are connected. For this step, we deployed our fast probabilistic fiber tracking implementation applied to the reconstructed mixture of diffusion orientations obtained from the diffusion toolbox of FSL 4.0.

Probabilistic Fiber Tracking

To compute the path of a global fiber, the tractography algorithms rely on local, voxel-based estimates of diffusion anisotropy. It is possible to model the observed direction-dependent diffusion-based MR signal attenuation as the net effect of local diffusion barriers. The FSL software computes an approximation of the likelihood of oriented diffusion barriers. The nonlinear model is fitted with a Markov chain Monte Carlo (MCMC) procedure (Behrens et al., 2003a, Behrens et al., 2007).

After modeling the fiber orientation in every voxel, the fiber tracking algorithm generates potential pathways by joining likely orientations in adjacent voxels (Behrens et al., 2003b). Due to its probabilistic nature, the algorithm can estimate fiber tracks through regions of uncertain connectivity, such as crossing fiber bundles and gray matter, better than previous, deterministic approaches. However, such a tracking through neural tissue with complex connectivity can result in erroneous connections. The approach of probabilistic fiber tracking is to trace all possible connections and compute the likelihood for each of them. Our implementation of this approach computes the connectivity for the complete brain between every parcellated cortical area and each voxel within it. Unlike previous approaches, this whole-brain analysis requires a substantially higher number of samples to achieve stable estimates of the posterior distribution for the high resolution diffusion scans. Our software was developed to be optimized for the AltiVec vector processing units available on modern PowerPC microprocessors and achieves substantially higher sampling speeds than comparable implementations on other processors. The optimization enabled us to complete the tractography for an entire brain within hours, thus producing a whole-brain "connectome" map.

Mathematically, probabilistic fiber tracking algorithms compute the likelihood of a fiber pathway. The likelihood is computed by sampling from the posterior distribution using a Markov chain Monte Carlo procedure. Each sampling step is referred to as a particle and sampling from the Markov chain corresponds to propagating the virtual
particle through the conditional orientation models in each voxel. Here, we were interested to estimate whether two areas are likely to be connected, given the likelihood of potential pathway locations. We used the density of virtual particles in the target region to judge whether two areas might be connected. To compute the particle density, we seeded a source area with a density of 1000 particles per voxel, and then counted, for every other cortical area, how many particles reached the target region. For each target area, we divided the total number of particles reaching the target area by the size of the size of the area. We found that this density results in the best agreement with tracer studies.

Combining Functional Imaging and Fiber Tracking in Humans



Figure 3.5. Outline of the Human Structural and Functional Processing Steps. This processing flowchart adds to Figure 3.1 also the functional (BOLD) pipeline and shows how functional (BOLD), connectivity (diffusion) and anatomical (structural MR) information are combined to explore both structural and functional connectivity. The particular pipeline shown here is for the human brain data, but similar flowcharts apply to the data from the monkey brain, where the functional data was obtained from a public archive.

Multitaper Analysis of BOLD Time Series

Functional connectivity is frequently studied by assessing the correlation of the BOLD signal between two areas in a broad frequency band. Due to the 1/f power spectrum of the BOLD signal, this overemphasizes the very low frequency components. To study the BOLD coherence in different, narrower frequency bands, it would be possible to band-pass filter this signal multiple times and estimate correlation or coherence coefficients in each case. As the number of frequency bands increases, the filter would need to be temporally widened to decrease the leakage between different frequency bands. To reduce leakage across frequency, then, optimal filter weights can be determined with a Slepian filter. The nonobvious solution is an eigenvector that represents the weights with the least leakage. The solution of this leakage problem has multiple eigenvalues that capture orthogonal weights. Multitaper methods use the first few solutions or "tapers" to compute spectral power with minimal leakage (Pesaran et al., 2002, Mitra and Pesaran, 2008). This optimal noise filtering in the temporal and frequency domain is very desirable for high-noise signals such as BOLD time series. In chapter 6 and 7 we show the application of using the complex BOLD spectra of the MRI signal to efficiently compute the cross-coherence across brain regions. The BOLD coherence so measured can then be combined with structural and diffusion MRI (Figure 3.5) to provide a highly multimodal imaging approach to investigating both brain structure and brain function. We present some examples of such integration for both the monkey and the human brain in the chapters below.

Chapter 4

BELIEF PROPAGATION OF CONNECTIVITY

Diffusion-weighted magnetic resonance imaging provides relatively high resolution information about the microstructure of white matter. A key challenge is how to integrate this information, often sampled at thousands to millions of different locations, in order to derive global connectivity estimates. Current approaches choose seed locations and then trace potential pathways independently. An exhaustive search with such methods quickly becomes computationally intractable, and hence all current model-based approaches avoid this problem either by estimating a single likely pathway (deterministic algorithms) or sparsely sampling the theoretical distribution of all pathways (probabilistic algorithms); neither guarantees an optimal solution and both suffer from false negative and false positive results.

Here we present a novel algorithm predicated on message passing that rapidly converges to a globally optimal solution among all possible pathways through the entire brain network. The algorithm implements a routing algorithm based solely on local decisions yet provides a solution at the level of the whole brain. Each node takes into account connectivity information that it receives from its adjacent neighbors. At each voxel location, global information, for example about the curvature of the pathway, and local diffusion information are combined into a globally optimal connectivity estimate. The effective resolution of the algorithm can be further enhanced by supplementing the diffusion data with high resolution white matter priors from T1-weighted, anatomical scans.

Model-Based Fiber Tracking

Many previous methods solve the fiber tracking problem either by starting at a seed location and tracing potential connections deterministically through the orientation field, or by evaluating many randomly generated pathway candidates independently, one by one. Algorithms thereby need to, either implicitly or explicitly, judge the reliability and goodness of fit of their fiber estimates. Formally, models of potential connections (*M*) are compared against the acquired imaging data (*y*) and a goodness of fit between the model parameters and the data is computed. In the case of Bayesian models, the likelihood distribution over the model space of different pathways *P* (*M*|*y*) is typically approximated by the tractography algorithms.⁶

$$P(M \mid y) = \frac{P(y \mid M) \cdot P(M)}{P(y)}$$

For most models, the likelihood function P(M|y) is complex and cannot be solved analytically. Instead, sampling techniques, such as Metropolis-Hastings, Gibbs, or Markov chain Monte Carlo methods (MCMC), are applied to estimate the likelihood distributions from the data. If the likelihood function is sufficiently complex, sampling quickly becomes computationally expensive when applied on a whole-brain level. The sampling procedures need to evaluate each and every pathway candidate independently and thus become very slow on a global scale.

⁶ Many papers separate a theoretical model space M from the parameter set s that instantiates a particular pathway. For clarity, M here refers to the space of model instances regardless of a chosen parameterization of model space.

A typical model evaluates pathways by scoring the agreement between the observations at multiple voxels along a pathway $(y_1, ..., y_n)$. In a Bayesian framework, the conditional probability of local diffusion signals is calculated from a global model. The observation model P(y|M) for instance estimates the probability that anisotropic diffusion components would be observed along a particularly shaped pathway. Observations that are consistent with the curve's shape are assigned a high probability, while discordant orientations are less likely, but might occasionally be observed due to thermal and other equipment noise.

$$P(M | (y_1, \dots, y_n)) = \frac{P((y_1, \dots, y_n) | M) \cdot P(M)}{P((y_1, \dots, y_n))}$$

The actual estimation of a model's likelihood value is often straightforward. The MR observations in different voxels are independent and the conditional probabilities can simply be multiplied. The observations y_i can thereby correspond to adjacent voxels chosen by a sampling procedure (Behrens et al., 2003b, Sherbondy, 2008) or represent selected points along a globally defined curve (Jbabdi et al., 2007).

$$P(M \mid (y_1, \dots, y_n)) = \prod_i \frac{P(y_i \mid M)}{P(y_i)} P(M)$$

For complex likelihood functions P(M | y), an exponentially large number of random samples are required to reliably identify the most likely pathway shape. The main challenge with current models is that small local variations of the pathway description (M) can results in drastic changes of the global pathway shape. For example, if a pathway is modeled by a Catmull-Rom spline, a change in position of one of the control points alters the path and thereby the sampled voxels for the adjacent spline segments. Each variation of the model, however small, hence needs to be evaluated independently. To increase the speed of convergence, the model complexity is often artificially limited, for instance by restricting the analysis to only planar splines with only 2 effectively free control points (Jbabdi et al., 2007). Nonplanar curves or curves that cannot be expressed by varying the location of only 2 control points, thus cannot be found by such algorithms.

An otherwise intractable model space can become computationally tractable if the global pathway can be divided into local segments. When estimating the model, one could then take advantage of the fact that many similar pathway candidates will share common segments.

$$P((m_1,...,m_l) | y) = \prod_i \frac{P(y_i | f(m_1,...,m_l))}{P(y_i)} P(f(m_1,...,m_l)).$$

For global pathway models, local changes still influence the global shape of the pathway. For Catmull-Rom splines, for example, a change of a control point influences the pathway in the adjacent segments. To leverage the existence of commonly shared pathways, the mathematical representation of initial segments needs to be conditionally independent of subsequent continuations of the track.

$$P(y | f(m_1, ..., m_n)) = P(y | f(m_1, ..., m_n, m_{n+1}, ..., m_l)).$$

If the above condition is met, the model space up to length n can be evaluated independently of further continuations of pathways. Subsequent extensions are then efficiently computed by extending the precomputed model space of length n to length n + 1. Figure 4.1 aims to illustrate the pathway fitting for the case of a splitting fiber.

Different pathway candidates for the first (m_1 , m_2) segment would be evaluated and the associated likelihood values stored. In subsequent steps, different continuations of the tract, for example to the (m_3 , m_4) or the (m'_3 , m'_4) segments, can be evaluated based not only on their agreement with local measurements, but also on how well the new global pathway resembles a smooth curve. Different continuations will extend different initial segments, thereby performing a truly global search.



Figure 4.1. Segmentation of the model space. The algorithm evaluates possible pathway candidates for the start of a fiber and assigns a likelihood value to each candidate. Subsequently, the pathways are extended, and multiple continuations of the initial segments are possible. Although a local fiber tract candidate might initially fit better, subsequent continuations can result in the selection of locally suboptimal, but globally optimal pathway estimates.

Recursive Estimation of Fiber Tracks

To estimate the likelihood function efficiently, it is necessary to find a recursive definition of *M*, such that the relevant properties of $f(M) = f(m_1,...,m_l)$ can be described by the combination of local pathway information (m_l) with one descriptor summarizing the prior history of the pathway (m^*_{l-1}) . Local curvature and expected fractional anisotropy information can be represented by the local vector m_l , and likelihood information about the curvature, length and shape of the pathway are stored in m^*_{l-1} .

$$P((m_1,...,m_l) | y) = \prod_{i} \frac{P(y_i | g(m_l, m_{l-1}^*))}{P(y_i)} P(g(m_l, m_{l-1}^*))$$

The information dependency of the above formalism can be visualized by a factor graph (figure 4.2). In the graph, circles represent variables and squares symbolize functions. Arrows indicate the dependency relationships between the variables and functions. Information about the geometry of the tract and the agreement with the observed data along the prior pathway is all conveyed in m^*_i . A global pathway likelihood and shape descriptor m^*_i is computed based on prior pathway segments m^*_{i-1} , a local shape and a tissue model m_i together with the local measurements y_i . The messages m^*_i are exchanged between adjacent nodes during the updating steps. When a voxel receives information about a prior pathway $(m^*_{i-1})_i$ it evaluates this information in light of possible local extensions m_i and of the voxel's measurements y_i .



Figure 4.2. Factor graph of the message passing algorithm. Each voxel summarizes local pathway information (m_i) and diffusion measurements (y_i), as well as global prior information (m_{i-1}^*) into global pathway descriptors (m_i^*). The pathway descriptors are exchanged between neighboring nodes during the update or message passing process (i.e., each iteration of the algorithm). The global pathway descriptors in a voxel contain the necessary information to evaluate the model fit of the associated fiber tracts. The dashed lines show the boundaries of a voxel. The arrows indicate the dependency of the random variables and information travels against the arrows during the updating steps.

The general factor graph in figure 4.2 can be further simplified into a hierarchical form. Most global models can be rewritten such that the observed voxel data y_i only depends on the local segment descriptors m_i . The algorithm is still capable of estimating global functions, since m_i is conditioned on the global pathway descriptor m_i^* .



Figure 4.3. Factor graph for a hierarchical local measurement model. In a simplified model, global pathway descriptors (m^*_i) combine global prior (m^*_{i-1}) and local pathway information (m_i) . The diffusion measurements in each voxel (y_i) are purely a function of the local properties (such as local shape; m_i) that are consistent with the global pathway estimate (m^*_i) .

The graphical models in figure 4.2 and 4.3 do not yet specify the functions implemented by the updating rule. It would for instance be possible, although impractical, to pass all local information along the path and concatenate all segment information into a complete vector $m^*_n = (m_1, ..., m_n)$. For such lossless algorithms, the size of the message to be passed would grow exponentially with each iteration and the resulting memory requirements would render it impractical. In the following paragraphs we will present updating rules with small, constant message sizes that

capture the relevant global attributes of a pathway. We first introduce a simple model that is similar to current wavefront and other flood-filling techniques, but that does not incorporate global shape constraints. We then extend this model to find pathways with low curvature changes and finally add anatomical, T1-weighted information to increase the effective resolution of the algorithm.

The models will be presented in the form of a hierarchical Bayesian model. The family of message passing algorithms is not limited to probability models and could handle consistent non-Bayesian cost or scoring functions equally well. The propagation of Bayesian information in a graphical model is called belief propagation. The measurement model of the factor graph can then be written as a conditional probability:

$$P((m_1,...,m_n) | y) = \prod_{i} \frac{P(y_i | m_i, m_{i-1}^*)}{P(y_i)} P(m_0^*)$$

Bayesian Model of Fractional Anisotropic Diffusion

The major challenge of fiber tracking is the length scale discrepancy between the microstructure of the imaged tissue and the much coarser sampling of the measurements: axons are on the order of only a micron wide and a single voxel can contain millions of axons. Diffusion imaging measures the diffusion barriers inside cells due in good part to axonal and other membranes. High-angular diffusion imaging can quantify the fraction of different anisotropic diffusion barriers within a single voxel (Tuch et al., 2003, Behrens et al., 2003b). High fractional values for an anisotropic diffusion component are a good indicator of the presence of white matter uniformly oriented in the direction of that component. Figure 4.4 shows a hierarchical Bayesian

model estimating the fractional value f within a voxel (the voxel boundary is indicated by a dashed line). In this algorithm, the observed diffusion measurements are modeled as a sum of anisotropic diffusion tensors and each component has a certain direction in space. The expected magnitude of the resonance signal is a function of a number of variables, such as diffusion speed, two noise terms, and the direction of the diffusionweightings (Behrens et al., 2003a).



Figure 4.4. A simple hierarchical Bayes white matter model. The diffusion measurements in a single voxel (*Y*) are explained by a mixture of directional anisotropic diffusion components. Each component has a direction and a fractional volume that it occupies. The voxel model is fitted for each voxel (outlined by the dashed line).Globally, connectivity information (*C*) is modeled by a global pathway variable (*F*). *See text for a detailed description.*

We extended the model of Behrens et al. (2007), to include a global pathway parameter (F). This global parameter describes the expected fractional value of a fiber

track. The global fractional value term (*F*) is then used to estimate the likelihood that two areas are connected (*C*). Strong connections with a high density of aligned diffusion barriers will have high fractional values throughout the entirety of the tract. For possibly unconnected or weakly connected regions the expected global fractional value term will be low, but larger than zero, since various other anisotropic diffusion barriers exist in nonwhite matter. The connectivity term (*C*) captures the relationship between the global fractional anisotropic term and the connectivity likelihood *P* (*C*1*F*). The likelihood function will be obtained from the diffusion measurements in later paragraphs (see figure 4.6).

To speed up computation, we estimate the hierarchical model in two separate steps. First, the posterior distributions over possible orientations and fractional values is estimated using the FSL diffusion toolbox (Behrens et al., 2007). Second, the likelihood distributions P(f|Y) are loaded into the program and are used to estimate the connectivity likelihood P(C|Y) using belief propagation.

Building on the fractional volume model, we first need to relate the global pathway term (F) to the observed fractional terms (f_i) along the path. Like streets and highways and many other man-made networks, the cortical communication network has a hierarchical organization where local denser connections are interconnected by fewer strong pathways. Axons originating in the cortical gray matter first form smaller white matter tracts that can then join larger cortical fascicles, before exiting the bundles and connecting with the target tissue. Figure 4.5 illustrates the theoretical measurement of fractional anisotropy values along a joining pathway configuration. While the

underlying fractional contribution of the pathway (*F*) is constant, the observed fractional estimates (f_1 , f_2) will increase after a second pathway joins (f_3 , f_4).



Figure 4.5. **Illustration of fractional volume terms along a pathway.** The actual contribution of a pathway (*F*) is the lower limit of the observed fractional terms ($f_1,...,f_4$). If a pathway joins a major tract, the observed fractional terms increase (f_3, f_4).

It is possible to recover the actual contribution of the traced pathway (*F*) from the individual measurements ($f_1,...,f_n$). The expected value of the fractional terms (f_i) are always larger than the global pathway value. We can hence write the conditional probability model for the fractional terms in the following form:

$$P(f_i \mid F) = \begin{cases} 1 & f_i \ge F \\ 0 & f_i < F \end{cases}$$

To complete the hierarchical model in figure 4.4, we estimated the connectivity likelihood function for the two macaque brains. We first aligned the anatomical T1 scans

to the diffusion-weighted images and constructed a Bayesian classifier to segment white from gray matter. The distribution of fractional values for both tissue classes are shown in figure 4.6 (blue and green lines). Panel A shows the data from a brain scanned at 9.4T (white matter in blue, and marginal distribution in green), and the distributions from a second brain scanned at 7T are shown in panel B.



Figure 4.6. **Connectivity likelihood function**. Probability density functions of empirically obtained connectivity models. **(A)** Data from a macaque brain scanned at 9.4T and **(B)** scanned at a 7 T. Anatomical white/gray matter classification was used to obtain the distribution of fractional *f* values in the brain *P* (*f*) (shown in green) and in

white matter P(f|c=1) (blue). The likelihood functions for connected tissue P(c=1|f) are shown in red and in turquoise for unconnected regions P(c=0|f).

The actual connectivity likelihood function P(c = 1|f) can be obtained from the conditional and marginal distribution via Bayes rule (shown in red). A prior connectivity probability of 33% was assumed and used as the Bayesian prior P(c=1).

$$P(c = 1 | f) = \frac{P(f | c = 1) \cdot P(c = 1)}{P(f)}$$

The connectivity likelihood function can be used as a Bayesian tissue classifier. Figure 4.7 shows the probabilistic classification of voxels using the maximum fractional anisotropy terms in each voxel to probabilistically classify a voxel as white matter. Voxels in large white matter tracts are reliably classified, whereas smaller tracts have smaller white matter likelihood values.



Figure 4.7. **Diffusion-based white matter classification**. Bayes-optimal classification of a diffusion-weighted scan using high-angular resolution imaging. The tissue model of figures 4.4 and 4.6 is used to estimate the white mater likelihood for each voxel. The colors indicates the likelihood that a voxel contains white matter, high values (red colors) label regions where the fractional volume terms unambiguously support the existence of white matter bundles.

The upper bound of white matter connectivity strength can be traced with the hierarchical partial volume model of figure 4.4. Figure 4.8 shows the global connectivity likelihood from a seed in the corpus callosum to each voxel in the brain. It is important to note that this version of the algorithm is blind to the directionality of diffusion and thus traces through white matter regardless of directional information. The proper fiber tracking model will be introduced in the next section.



Figure 4.8. White matter tracking using only fractional anisotropy values.

Connectivity likelihood for a region in the corpus callosum. A seed was placed in the corpus callosum (in the third slice shown, white arrow) and the global white matter connectivity model of figures 4.4 and 4.6 was used to compute an upper connectivity bound. Each voxel's value corresponds to the maximum likelihood that it could be connected through continuous white matter tract to the seed regions. Unlike a fiber tracking model, orientation information is disregarded by this algorithm.

Bayesian Model of Fiber Orientation

The previous model traced continuous white matter regardless of global fiber orientation. Fiber tracking algorithms can be greatly improved if global orientation information is integrated. A particularly interesting approach is to model cortical tracts by splines. Splines are curves with continuous derivatives and the tangent and curvature along the path changes in a smooth manner. When tracing connectivity to each location in the brain, it is not feasible to find a spline representation for each voxel. Instead, it is advantageous to take splines as a prototypical model and asses the divergence from the prototypic path in a hierarchical Bayesian model. Rather than modeling splines directly, we chose to evaluate the divergence from the curvature smoothness constraint. The overall continuity in orientation of a path can be quantified by the derivative along the curve. To obtain a splinelike preference of the fitting procedure, the aggregated L1 norm of curvature changes can be computed.

$$G = 1 - \int_{u=0}^{1} \left| \frac{d^2}{d u^2} \vec{t} (u) \right| du$$

Sharp turns, as well as curved pathways, are typically not found in the brain. Mathematically, a very curved or curled pathway might fit the measurements slightly better than a straight or smooth connection. Yet, the latter is more likely. The G termscaptures the product of curvature changes along the pathway t(u). We implemented descriptors for cumulative curvature and orientation changes, and higher order or other constraints can be easily incorporated into the model. Small changes in curvature will result in a large G values. Bayesian priors for the straightness and curvature terms result in preferentially smooth and splinelike pathway fits (T).



Figure 4.9. Hierarchical Bayesian model with curvature constraints. The diffusion measurements in a single voxel (Y) are explained by a mixture of directional anisotropic diffusion components. Each component has a direction and a fractional volume that it occupies. The voxel model is fitted for each voxel (outlined by the dashed line). Globally, connectivity information (C) is modeled by a global pathway variable (F) and a term constraining the curvature change along each pathway (G). *See text for a detailed description.*

Diffusion imaging measures orientation discretely. The conditional probability is then a product of the orientation or curvature changes between neighboring nodes along the pathway.

$$P\left(g \mid G\right) = \prod_{i} \frac{P\left(g_{i} \mid G\right)}{P\left(g_{i}\right)} P\left(G\right)$$

When estimating the connectivity likelihood, the global straightness and curvature terms can be integrated out analytically.

$$P(g) = \int_{G} \prod_{i} \frac{P(g_i \mid G)}{P(g_i)} P(G)$$

Increasing the Effective Resolution by Including Anatomical Information

One of the disadvantages of diffusion imaging is its lower resolution, compared with other modern MR sequences typically used for structural imaging. Anatomical sequences, optimized to enhance the tissue contrast between white and gray matter, can provide high-resolution pictures of the brain with relatively short acquisition times. The effective resolution of fiber tracking can be enhanced by incorporating such higher resolution scans into the belief propagation algorithm. The fiber pathway model (*T*) can be extended to generate predictions about the T1 resonance signal. By comparing pathway candidates with the T1 images, pathways that might be consistent with the measurements at a lower resolution could be rejected at the high resolution. The effective resolution of the fiber tracking algorithm can thereby be improved. Figure 4.10 shows the hierarchical Bayes model of the extended model. In the enhanced model, the tract estimates (*T*) the expected T1 tissue contrast in addition to modeling the diffusion signal. The MR resonance signal is modeled as function of that expectation (*A*) and an additional noise term (λ).



Figure 4.10. Hierarchical Bayes model incorporating anatomical information. The global pathway strength variable (F) is extended to include predictions about the T1-weighted resonance signal. The pathway model (T) predicts the T1-weighted intensity signal (A) in addition to the orientation and fractional volume parameters of the diffusion model. The T1-weighted signal intensity is a function of the global pathway strength (F). For strong pathways, a high partial volume is expected throughout the tract, and the appropriate intensity values are expected for the voxels along the pathway. Through Bayesian inference, the T1-weighted signal thus increases the effective spatial resolution by eliminating global pathway candidates that pass through T1 voxels that are inconsistent with a white matter tract.

To simplify the implementation of the algorithm, the T1-weighted signal intensity can be combined with the fractional anisotropy volumes (f). The relationship between the two measurements is shown in figure 4.11. The graph plots the mean of the fractional anisotropy values for voxels of different T1-weighted intensities. Intensities around the numerical value of 780 correspond to tissue with a large proportion of

oriented white matter. The conditional distribution of T1 intensities for white matter can be modeled by a Gaussian distribution and used to compute the white matter (and thus pathway) likelihood for each voxel.



Figure 4.11. Distribution of partial volume coefficients as a function of T1-weighted MR signal intensity. Plotted against the mean fractional anisotropy in a voxel (on the yaxis) is the T1 value from structural MRI (on the x-axis). T1-weighted intensities at the numeric value of 780 correspond to white matter tracts. A Gaussian distribution was used as a Bayesian likelihood function to quantify the T1-weighted evidence for a passing fiber tract.

Belief Propagation of Connectivity

Fiber tracking can be described as a problem of finding a sequence of states, whereby each state represents a local pathway segment. The states describe local orientations and partial anisotropy volumes in each voxel. The path then traverses from its source to multiple destinations in the brain as a sequence of state transitions. At each voxel location, the path could take various routes and the MR measurements in a single voxel might often be insufficient to resolve uncertainty about the precise location or orientation of a pathway. The maximum likelihood solutions to such state transition problems can be efficiently computed by dynamic programming. Probabilistic algorithms, in contrast, model the uncertainty by estimating or sampling from all possible combinations, as they estimate the posterior distribution across all potential connections.

In practice, it is often sufficient to find the best pathway and to estimate the likelihood that it exists. It then becomes practical to evaluate all pathways to all possible locations and to compute the exact likelihood of each solution. If the model can be formulated in a recursive way and the Markov assumption is fulfilled (i.e., the likelihood of future segments only depend on the local descriptor), then the most likely connections can be easily determined with a variant of the Viterbi algorithm. The basic idea of the algorithm is that the local distribution of the model descriptors is sufficient for finding the most likely sequence of hidden states. The distribution over all local states at step t can be used to compute the distribution of states for step t+1.

Message Passing Implementation

To find all optimal pathways, the above Bayesian model is implemented as part of a large message passing network. A typical network contains hundreds of thousands of nodes, one for each voxel, and it takes a few hours to find all optimal solutions for a seed region. The nodes form a multidimensional lattice (L), spanning the 3-dimensional spatial domain of voxel locations (x). At each of the voxel locations (x), the lattice tensor

contains multiple possible fiber orientations, enumerated by (j). Each node additionally stores a discrete representation of the hidden variables of the Bayesian fiber tracking model.

The distribution of possible orientations (θ , ϕ) is encoded in a unit vector (t) and this is associated with the fractional values (f) that are stored alongside, at each node. The global strength and curvature values (F) and (G) are sampled with the same spatial resolution (x) and angular resolution (j) as the orientation measurements.

$$L_{x}^{j} = \left(\vec{t}_{x}^{j}, f_{x}^{j}, F_{x}^{j}, G_{x}^{j}\right)_{.}$$

The posterior distribution of possible orientations (*t*) is obtained from the polar representation of the FSL diffusion toolbox. Representing all possible orientations would require vast amount of memory. For example, a directional grid with a 10 degree spacing would result in over 600 directions. We restrict the analysis to the 50 most likely orientations. For high-angular resolution diffusion imaging, the discrete orientations are obtained by sampling from the theoretical orientation density function (ODF). The attenuation of the magnetic resonance signal for each gradient direction is modeled as the effect of thousands of anisotropic diffusion components that can exist within the volume of a single voxel (Behrens et al., 2003; Jbabdi et al., 2007). We used the estimator that is part of the FSL package to choose 50 orientations according to the computed ODFs. Bidirectionality was obtained by adding the mirror-symmetric orientations of the 50 direction samples. At each voxel location, there are thus 100 directional samples.

$$\vec{t} \vec{t}_x^j \qquad \vec{t} \in \mathbb{R}^3, j \in [1, 100],$$
$$x \in \mathbb{R}^3, s \in \mathbb{N}$$

To encode multiple seed regions, an additional source or seed dimension (*s*) is introduced. In the current implementation, the algorithm is distributed along the seed dimension over the nodes in a computational cluster. The lattice is evaluated in parallel for all source regions:

$${}^{s}L_{x}^{j} = \left(\vec{t}_{x}^{j}, f_{x}^{j}, {}^{s}F_{x}^{j}, {}^{s}G_{x}^{j}\right)_{\cdot}$$

The probability that a node at location (x) and orientation (j) is connected to a seed region is computed from two sources. A fractional value term (F) describes the degree to which the pathway to (x) fits with the obtained measurements. The second, curvature term (G) captures the plausibility of a pathway shape. Less curved pathways are more likely than pathways with a high degree of curvature changes. The likelihood for a pathway is calculated by combing those two independent information sources.

$$P\left({}^{s}c_{x}^{j}\right) = P\left({}^{s}F_{x}^{j}\right) \cdot P\left({}^{s}G_{x}^{j}\right).$$

Neighboring nodes exchange those two terms during the message exchange step. Since the connectivity likelihood can be directly computed from F and G, the lattice does not need to explicitly store the connectivity values (c). The F term captures the quality of the estimated fiber track. Here, we use the minimum value of all fractional values along the candidate pathway to describe its fit with the measured volumes. A pathway The fractional value and curvature information is exchanged and updated between neighbors during the message exchange step. In addition to the minimum fractional value and cumulative curvature change along the pathway, neighboring nodes also exchange information about the measured orientations (t). First, the sending node (y) identifies a neighbor (x), by selecting one of the possible orientations (t).

$$y = \left\lfloor \vec{t}_y^j \cdot \delta + x \right\rfloor$$

Second, the receiving nodes compare their fractional value to the received global term and update an intermediate value (F). To incorporate geometric constraints, the node compares its own orientation with that of the neighboring nodes and adds a penalty term (g) to an intermediate cumulative curvature product (G). G increases with increasing discrepancy between the local curvature. The state values for (F) and (G) are updated, if the associated likelihood for the intermediate values are larger than the previously found likelihood value.

$$F_{y}^{j} \leftarrow F_{xy}^{ij}$$

$$G_{y}^{j} \leftarrow G_{xy}^{ij} \quad \text{if } P\left(F_{y}^{j}\right) \cdot P\left(G_{y}^{j}\right) < P\left(F_{xy}^{ij}\right) P\left(G_{xy}^{ij}\right).$$

By implementing the updating rules, the lattice implements a distributed maximum likelihood computation of possible pathways. To show the equivalence of this message-passing algorithm with other approaches, it is helpful to think about a pathway as a series of connected locations (x). The pathway can be fully described by the tangent (t)

at each step and the associated quality or fractional value score associated with each orientation measurement.

$$P\left(f_{x_{1}}^{j_{1}}, \vec{t}_{x_{1}}^{j_{1}}, \dots, f_{x_{n}}^{j_{n}}, \vec{t}_{x_{n}}^{j_{n}}\right)$$

As outlined on the preceding pages, the above probability can be computed from the global (F) and (G) terms that are associated with each pathway. Since the minimum and product functions used to compute (F) and (G) are separable, they can be computed iteratively. This is achieved by introducing variables that represents the goodness-of-fit and curvature term of the pathway up to the second-to-last step (n-1). The likelihood of the last step can be computed by integrating the fractional and curvature values of the last step (n). The lattice represents exactly those (n-1) and (n) terms at each voxel location. The (n-1) terms are stored in the sending node, and the (n) terms are calculated in the receiving node.

Next, we would like to show that the algorithm will find the globally optimal solution. This is achieved by iterating over all possible pathway combinations up to a length n. The number of possible pathways of length n through the lattice is limited. A pathway is completely defined by 1) its seed location x_0 , 2) the direction at each step and 3) the step length. If we consider the seed location and the step length as fixed, the number of possible pathways is limited by the number of possible orientations at each step. For the current implementation, this number is 100.

$$x_n = \left\lfloor \vec{t}_{x_n}^{j_n} \cdot \delta + x_{n-1} \right\rfloor$$

The number of all possible pathways is thus **100**ⁿ. However, since the cost functions are separable, it is possible to optimize them locally to obtain the globally optimal solution. For a cubic volume of size *m*, the number of necessary steps is $m^3 \times 100 \times n$. At each iteration, nodes at every location exchange information with their neighbors. One iteration of the algorithm thus implements m^3 propagation steps. To iterate over all possible pathways, $100 \times n$ global iterations are necessary. Although the theoretically longest loop-free path is one that touches every voxel, white matter fibers typically consist of either very direct or U-shaped segments. An exact solution can hence be obtained by scaling the number of iterations with the resolution of the imaging volume. For the results presented here, we use the length of the volume's diagonal. All optimal fiber pathways shorter than the diagonal will be guaranteed to be found by the algorithm. For pathways that are longer than the volume's diagonal, the computed likelihoods can underestimate the optimal solution. In practice however, the algorithm converges very quickly and the numerical difference between early estimates and the exact solution are often negligible.

Results

To evaluate the performance of the algorithm we collected high-resolution diffusionweighted MRI images of two macaque brains (see methods). First, we tested out the algorithm for a seed location in the corpus callosum. Starting from the midline seed region, interconnections to both hemispheres are identified. Figure 4.12 shows the propagation of likelihood information from the seed region to each voxel. Starting with a connectivity likelihood value of 1 at the seed location (panel A), the lateral callosal fibers are quickly traced in the first 50 (panel B) steps. At each location, the algorithm iteratively tries possible orientations and continuously updates the neighboring voxels about newly found connections. In this particular run, lateral connections are found first. within 100 iterations (panel C), the more superior and inferior pathway candidates are found within the first 200 iterations (panel D). The updates are monotonic and converge quickly against the final, optimal solution. Panel E and F show the likelihood after 500 and 1000 steps.



Figure 4.12. Propagation of Likelihood Information using Message Passing. The figure provides an example application of the message passing algorithm, using a seed voxel in the corpus callosum (white arrow in first panel). Subsequent panels show the results after iterations of 50, 100, 200, 500 and 1000 steps. Each voxel estimates the likelihood that the particular location is connected to the seed region in the corpus callosum. Unlike many previous algorithms that use simple voxels scores for fiber tracking, each location contains 100 instantiations of the hierarchical Bayesian model in Fig. 4.10. The best connectivity likelihood value of the hundred models is shown for each voxel. Neighboring voxels exchange their information about the models'

parameters through message passing in each iteration, efficiently propagating connectivity information through the volume. At any time, the algorithm provides a lower bound of connectivity likelihood. It converges very quickly and within a fixed number of steps against the true likelihood distribution.

Each voxel contains information about the optimal pathway from the selected seed region to each location. This information can be used to compute tract-based statistics. Figure 4.13 shows a trivial application that measures the length of the optimal pathway to each location of the brain. Other statistics, such as the average fractional anisotropy along the path or the membership of a path to one or multiple larger tracts can be easily computed. Particularly in a clinical and surgical context, such data would be invaluable and can now be computed efficiently for individual brains.



Figure 4.13. Tract-based statistics: Distance along the path. An interesting application of tractography is to compute tract-based spatial statistics (Smith et al., 2006). Previously, this method involved a multistep procedure of estimating the uncertainty boundary of tracts and finding the pathway centers for every location along a tract. Due to our algorithm's ability to find the optimal solutions to all locations, any tract-based spatial statistic can be computed efficiently to every target location. To illustrate this ability, we here computed the simplest tract statistic: its length. Each voxel contains 100 hierarchical Bayesian model instances (see Fig. 4.9) with pathway information to the seed location in the corpus callosum (indicated by the white arrow). The length of the optimal pathway to each location is encoded by the color scale.

The lattice presentation allows for an easy customization of the fiber tracking algorithm. Different information maps that are used to identify fiber tracts are shown in figure 4.14. Panel A shows the overall likelihood of connectivity with the midline corpus callosum seed after just 200 iteration steps. Panel B shows the minimum fractional value along the best pathway to each voxel location. The cumulative path orientation change field is shown in panel C. The expected value of the fractional anisotropy term along the pathways (panel D) offers a quick global connectivity map after a few initial iterations. Overall, this mean value is less sensitive than the orientation and minimum fractional value cost functions, its addition to a model can accelerate the convergence by guiding the routing during early estimation steps.



A Connectivity likelihood **B** Strength of best path

Figure 4.14. Contribution of fractional anisotropy and curvature information. The panels show: **(A)** the integrated connectivity likelihood (C), **(B)** the fractional term F providing information about the minimum fraction value from the seed region to each

location, **(C)** the curvature bias term G computing the sum of curvature changes along the pathways **(D)** the mean fractional term along the path from the seed to each shown voxel for comparison. See Fig. 4.9 for the Bayesian hierarchical model.

The overall aim of probabilistic fiber tracking methods is to integrate local orientation information throughout the volume. The Bayesian models used in our implementation resemble the previous algorithm of Behrens et al., 2003b. However, our way of optimally estimating connectivity globally provides a whole brain assessment of connectivity and results in more sensitive estimates for complex fiber architectures, such as those that are for example found for medial callosal fibers. One of the main advantages of the fiber tracking algorithm is the systematic evaluation of all locations. Figure 4.15 shows the comparison between results for the same seed in the corpus callosum, when connectivity information is evaluated with a message passing (panel A) versus a sampling approach (panel B). The results of both panels were obtained on the identical orientation flow field and using similar models for the pathway estimation. Due to the optimality of the estimation procedure, the message passing algorithm provides a more complete picture of the callosal pathways.



Figure 4.15. Comparison between probabilistic tractography and belief propagation. Comparison between **(A)** message passing and **(B)** results from a probabilistic tractography algorithm (Behrens et al., 2007) using identical measurements and seed location in the corpus callosum on a macaque brain.

The precision of the algorithm can be further improved by including high resolution anatomical information. Although the anatomical tissue contrast does not provide orientation information, it can eliminate inconsistent pathway estimates, thus potentially increasing the specificity of the global estimation and improving the location estimates of the maximum likelihood solutions. Figure 4.16 illustrates the differences in resolution of a diffusion scan (panel A) and the sample section when scanned with a high resolution anatomical sequence (panel B).



Figure 4.16. Comparison between the resolution of fractional anisotropy estimates (500 μ m, left) and white matter likelihood that was computed from an anatomical scan (110 μ m, right).

Sampling solutions sidestep the computational complexity of the fiber tracking by selectively estimating the most likely connectivity candidates only. As a result, many, particularly medium-strength, pathways are not evaluated. Figure 4.17A shows a normalized histogram of the distribution of evaluated (blue) and unevaluated (red) voxel locations. Even for the most likely locations with connection likelihood of >0.30, the message passing algorithm identifies 4 times the number of voxels that are evaluated by the sampling approach.




The belief propagation algorithm finds the optimal pathway to every voxel in a finite number of steps. The number of steps needed for convergence is a linear function of the maximum path length, typically a few hundred voxels. In practice, we observe a fast convergence within a few steps toward the optimal solution.

Summary

To study global cortical organization, we introduced a novel connectivity algorithm based on message passing. It optimally estimates connectivity on a whole-brain level in polynomial runtime. Compared to previous methods, its effective resolution is increased by incorporating probabilistic tissue information from high-resolution anatomical scans. Additionally, we have shown that diffusion-based fiber tracking can be used to estimate connectivity strength and that this information can be used to identify functionally relevant pathways and hierarchical information globally.

Chapter 5

WHAT ONE CAN DO AND CANNOT DO WITH TRACTOGRAPHY

Historically, cortical connectivity has been viewed as a binary property. Most anatomical studies explore whether two selected regions are directly connected or not. The implicit assumption of these studies is that if two regions share a common direct connection, then they are part of a tightly coupled network, whereas unconnected regions are less likely to work together. Yet, relevant for the functional importance of connectivity is not only whether two regions are connected by a fiber, but for example how strong a connection is and what importance a direct link (or nonconnection) plays in the context of the surrounding network. Two regions could, for instance ,be part of the same processing pathway and engage, tightly coupled, in the processing of similar stimulus properties, even though they do not share a direct connection but communicate through massive fiber tracts with an intermediate region.

This chapter aims to explore to what extent fiber tracking -- when applied on a whole brain level -- can reveal such nonbinary information. We first aim to quantify the reliability of fiber tracking by looking at two data sets: a 80 µm high-resolution scan of a mouse brain obtained in a 11.7T Bruker system and 2 mm resolution human scans acquired at 3T. Since the reliability of fiber tracking is not only a property of the image quality and resolution, but more importantly of the complexity and size of the underlying fiber tracks, we first aim to quantify the difficulty of the fiber tracking problem in an information theoretic framework. We bilaterally track callosal fibers from homotopic areas in both hemispheres of mice and men and estimate the two resulting posterior connectivity distributions at the midsection of the corpus callosum. The first measure we obtain for both estimates is the entropy or computational complexity of the fiber track estimate. Second, we estimate the discrepancy between the left and right density functions using Kullback-Leibler divergence, a form of mutual information. In a third step, we successively downsample our scans to quantify the dependence of fiber tract estimates on the resolution of the data. Fourth, we introduce resolution-enhancing interpolation techniques to improve the reliability at all resolutions.

Fiber tracking is frequently deployed to study connectivity noninvasively. It is often the only choice for studying connectivity in human brains. We hence explored to what extent, and under what conditions, diffusion tractography could be used to discover novel connections. The best model animal for human connectivity are macaques. We acquired two data sets of macaque brains at a 9.4T and at 7T and identified the likely locations of more than a 100 distinct cortical areas. We then compared the fiber tracking results on a whole brain level, for all pairs of possible interconnections, with results from hundreds of published tracer studies. We found that fiber tracking produces largely valid predictions, but that they cannot be proved to constitute true connectivity in all cases. However, we show that fiber tracking is an excellent tool for studying larger connections and that it provides excellent statistical predictions of connectivity.

The noninvasive nature of fiber tracking allows one to explore connectivity in clinical populations with altered brian connectivity. We were interested in tracing fibers from a patient group with a congenital absence of the corpus callosum. The brains of this patient groups exhibit a strong variability in the size and shape of ventricles and cortical midline structures. We show that fiber tracking together with a surface-based reconstruction of cortical areas provide a good way to trace many connections reliably.

The most interesting aspect of diffusion-weighted imaging is the ability to directly assess the microstructure of neurons and axons. Its primary application has been to image the ischemic effects of stroke. Here we show that it can directly measure the strength of cortical connections by comparing its results with quantifications from tracer studies.

In the previous chapter a Bayesian belief propagation approach was introduced that identifies all optimal tracts. This algorithm is the result of a long series of small, evolutionary improvements implemented by the author. The algorithm and this chapter are the fruits of this search for the limitations and capabilities of fiber tracking.

Initially, we used the probabilistic fiber tracking method of Behrens et al. (2003). To adapt it to a probabilistic model for whole-brain connectivity and high-resolution applications, we implemented the algorithm in assembler language to take advantage of the fast 128-bit vector engine of the G5 / PowerPC processor architecture. Subsequently, better models for estimating the underlying orientation distribution from diffusionweighted data became available (Behrens et al., 2007). In the process of implementing the algorithm we incorporated various smaller improvements such as trilinear data interpolation and higher-order Runge Kutta pathway interpolation. Finally, we formulated the main ideas in a global, belief propagation network framework. The various approaches do have their respective strengths. The G5/AltiVec implementation is very fast, well suited to trace larger fibers and incorporates flow field and pathway interpolation. It has been used for tracking pathways in the high-resolution mouse and macaque data sets. To facilitate comparisons, the human tract tracing was conducted with the same configuration. Results for the memory-intensive belief propagation implementation are reported separately for the macaque data set.

Reliability of Fiber Tracking

Two challenges in tracking axonal connections in the brain are to track fibers that are small compared to the imaging resolution, and to track through voxels that contain multiple and crossing fibers. Diffusion-weighted imaging with high angular resolution, together with probabilistic models of apparent diffusion, can help surmount these difficulties. However, the relationship between the uncertainty in the apparent diffusion measurement and the uncertainty of the fiber track estimates remains poorly understood.

Probabilistic methods can greatly increase the sensitivity of fiber tracking. Figure 5.1 shows the tract estimates of olfactory connections in a C57BL/6 mouse brain when traced with a probabilistic method (A) and a deterministic algorithm (B). Probabilistic methods can trace through tissue with a more complex microstructure when deterministic methods fail. The deterministic tracts follow the tract initially in the olfactory bulb, but cannot continue the pathway through the complex lateral tissue adjacent to the seed region.





Figure 5.1. Comparison of probabilistic and deterministic fiber tracking. The fiber tracking was seeded in the volume of the right olfactory bulb (orange arrow) in a 80 μm resolution, 40-directional diffusion-weighted scan of a C57BL/6 mouse brain acquired on a 11.7 T Bruker MR microscope. The position of the seed region with respect the mouse brain is shown in the inset on the bottom left of the top image. Two algorithms were deployed: (A) a custom implementation of a probabilistic tracking procedure; (B) deterministic fiber tracking with identical curvature constraints. The stream lines represent a random subset of pathways traced with both methods. The color of the

streamline represents the orientation of the tangent along the path (blue: superiorinferior, red: rostral-caudal, green: left-right). The probabilistic model finds multiple pathways including connections across the anterior commissure (red arrow). The deterministic algorithm only traced the initial anterior-superior orientations within the olfactory bulb, but fails to trace through the more complicated, converging fiber architecture lateral to the olfactory bulb.

We studied the overall reliability of a probabilistic fiber-tracking (PFT) algorithm by looking at the agreement of fiber estimates for interhemispheric connections seeded in contralateral voxels. We were particularly interested in how the relatively low spatial resolution of diffusion-weighted (DW) EPI scans influences the reliability of PFT results. To address this question, we acquired high-resolution DW images of fixed mouse brains, downsampled the data sets, and compared the reliability of the fiber tracking results at the original and the reduced resolutions for different variants of the probabilistic algorithm.

A parameter-free way to quantify the agreement between the two probability matrices is to compute the relative entropy between the two distributions. We used a symmetric version of the Kullback-Leibler (KL) divergence for this purpose. KL divergence measures the disagreement between two distributions as the increase in entropy when using information for encoding samples from one distribution with the other. KL divergence is an asymmetric measure and we computed the mean to compute a symmetric KL disagreement. The fiber tracking algorithm computes, for every region, how likely it is connected to another region in the brain. We used this likelihood distribution (Figure 5.2) to compute the KL distance for every brain area and then averaged the KL distance across all regions. KL divergence is an entropy measure and the disagreement can be quantified in bits. KL divergence does not make any assumption about the shape of the probability distribution to quantify the difference in the probability distribution and expresses the number of bits that are wasted when encoding the events from one distribution with the optimal code from the other distribution. Thus, if the two distributions were identical, they would have a KL divergence of zero, and the larger the KL divergence, the larger the nonoverlap between the two distributions (i.e., the lower the reliability of the fiber tracking algorithm).

We next related the KL divergence back to spatial reliability in the brain. Our approach was to ask how much spatial noise would result in an increase in KL divergence that was equivalent to a given number of bits (Table 5.1 and 5.2).

C57BL/6 mouse brain images were acquired from paraformaldehyde perfusion fixed samples. MR diffusion microscopy was performed with a 11.7T Bruker system using a conventional pulsed-gradient spin-echo sequence ($256x150x130 \mu m$ matrix, $80 \mu m$ resolution, 40 directions). The human data was acquired with a 3T Siemens system using a diffusion-weighted EPI sequence (64x64x34 mm matrix, 3 mm resolution, 72 directions) (Table 5.3).

The orientation density distribution was estimated using the FMRIB diffusion toolbox. Seed regions for fiber tracking were selected anatomically on a digital version of a Paxinos mouse brain atlas and the ICBM brain template that had been coregistered to the human brain. Custom software was used for the probabilistic sampling and Runge Kutta line integration.



Figure 5.2. Example of a probability density function of fiber estimates from the primary motor cortices in a mouse brain. Probabilistic tractography was seeded in the (A) left and (B) right primary motor cortex of a C57BL/6 mouse brain volume. Callosal motor connections were traced from both sides and the probability density function of the left and right estimate at the identical midline section are shown. The color map indicates the proportion of particles reaching each midline voxel. In this example, the estimates from the left are more extended compared to the right, but the most likely location of crossing is found in either case (yellow colors).

	80 µm		160 µm		320 µm	
Seed region	PFT	PFT+	PFT	PFT+	PFT	PFT+
Primary auditory cortex	5.80 bit	3.86 bit	3.96 bit	3.38 bit	1.72 bit	0.61 bit
Primary motor cortex	Inf	1.7 bit	1.62 bit	2.12 bit	2.05 bit	0.55 bit
Primary sensory cortex	2.47 bit	2.49 bit	2.84 bit	2.28 bit	1.65 bit	0.55 bit
Basal amygdaloid nucleus	3.63 bit	2.46 bit	4.58 bit	3.14 bit	1.44 bit	0.95 bit
Central amygdaloid nucleus	5.26 bit	2.84 bit	5.27 bit	2.41 bit	1.56 bit	0.67 bit

Table 5.1: Reliability (Kullback Leibler divergence) of interhemispheric connections
from the mouse brain.

	80 µm		160 µm		320 µm	
	PFT	PFT+	PFT	PFT+	PFT	PFT+
Primary Auditory cortex	5.25 voxel	4.11 voxel	3.80 voxel	2.90 voxel	3.99 voxel	2.90 voxel
Primary Motor cortex	-	5.06 voxel	2.69 voxel	3.78 voxel	3.80 voxel	1.81 voxel
Primary Sensory cortex	2.85 voxel	4.53 voxel	3.74 voxel	3.61 voxel	2.69 voxel	1.83 voxel
Basal amygdaloid nucl.	3.55 voxel	2.81 voxel	3.02 voxel	2.99 voxel	2.35 voxel	2.44 voxel
Central amygdaloid nucl.	3.71 voxel	1.90 voxel	3.76 voxel	2.68 voxel	2.17 voxel	1.88 voxel

Table 5.2: Spatial noise equivalent to KL divergence of table 5.1.

	PFT	PFT+	PFT	PFT+
Precentral gyrus	3.60 bit	2.17 bit	4.79 voxel	3.91 voxel
Postcentral gyrus	3.26 bit	2.00 bit	3.76 voxel	3.02 voxel
Superior occipital gyrus	3.83 bit	2.86 bit	3.76 voxel	3.81 voxel

Table 5.3: Reliability (Kullback Leibler divergence) of interhemispheric connectionsfrom the human brain.

Initial probabilistic fiber tracking estimates from the mouse brain were quite reliable at high as well as low imaging resolution. An improved PFT algorithm (PFT+) that used regularization of the orientation distribution function (trilinear interpolation, Runge Kutta integration) resulted in a significant improvement in reliability (p<0.001). The KL divergence provides a quantitative approach to study the reliability of probabilistic fiber tracking results.



Figure 5.3. Tractography improvement when combining both interpolation

techniques. The average KL divergence between left and right probabilistic estimates of callosal fibers is shown when different interpolation methods are used by the algorithm. The higher-order pathway interpolation provides slightly better results than the trilinear flow field interpolation. The estimates can be further improved by interpolating both the orientation flow field and fiber path.

Tracing Anatomical Connection in the Absence of the Corpus Callosum

Tractography reconstructs fiber pathways by following the local orientation field. If a large amount of noise is present, and when the size of a fiber is small compared to the imaging resolution, erroneous estimates can be produced. With current diffusionweighted MRI methods, it is, for example, not possible to detect local synaptic connections, or the splitting or joining of finer fibers. Similar to multisynaptic tracers, the diffusion-based tractography traces irrespective of synaptic connectivity and through intermediate areas. For practical application, it is necessary to establish the specificity of the tracking and to what extent erroneous connections could be generated.

Here we trace connections in a group of rare patients with a congenital agenesis of the corpus callosum (AgCC). The only interhemispheric connections are the anterior and posterior commissures, while all callosal fibers are completely absent. The failure to develop the corpus callosum results in a rearrangement of cortical midline structures such as the anterior cingulate cortex. Most prominently, the ventricles are also enlarged and developing fiber tracts that would have otherwise formed the corpus callosum can form additional connectivity. Identifying and tracking brain connectivity in an abnormally organized brain is challenging and we try to solve this problem by first applying a surface-based automatic parcellation technique to the cortex, and second by identifying pairs of interconnected brain regions that can be consistently tracked. By combining computation anatomy with automatic full-brain fiber tracking, we produce a full-brain connectivity matrix, from hundreds of cortical and selected subcortical regions in both hemispheres.

A group of 6 individuals with a complete agenesis of the corpus callosum were scanned in a 3T Siemens scanner. First, we obtained 2-3 sets of anatomical T1-weighted images. The T1 images were realigned and a white-gray matter tissue segmentation was performed. The cortical surface was reconstructed with FreeSurfer and a probabilistic surface map was aligned to the individual anatomy of each patient. Subsequently, the likely outline of anatomically defined regions was identified on the cortex map and the location of these areas in the original T1 volume was determined. Second, a 70-directional diffusion-weighted EPI volume was obtained for each person (2 mm isotropic, 64x64x64 matrix). A nonlinear partial volume model of local orientations was fitted to each voxel and the resulting orientation density function was used in a custom implementation of fiber tracking. We traced the likely tracts from each of the anatomically defined seed regions identified from the anatomical scans. For every target region, we determined the number and density of the tract estimates. We classified two brain regions as connected if the particle density reached at least 1 particle per voxel (i.e., the number of traced particles matched the size of a region). Regions with zero or only a few estimated paths were counted as likely not being connected.

Figure 5.4 summarizes the reliability of fiber tracking between all regions. Fiber tracking was conducted in all 6 brains and the color map shows in how many subjects a particular connection has been found. If a connection is found in all 6 brains, it is marked by a red dot, if a connection is consistently not found in all cases, it is shown in a dark blue color. Light blue and yellow colors indicate connections with inconsistent results (found in some individuals but not in others). Overall, a high proportion of connections that are either consistently found (red) or consistently flagged as nonexistent (dark blue) indicate a good reliability of the anatomical parcellation and fiber tracking. As expected, connections between regions of the left and right hemispheres are typically not found at all in these acallosal brains. The only exception are a few connections between homotopic regions in the two hemispheres (found on the

diagonal of the plot). These pathways are likely candidates for regions being connected by the anterior and posterior commissures, or possibly through intermediate subcortical structures.

Unlike after an acute surgical resection of the corpus callosum, patients with a congenital agenesis of the corpus callosum retain many functions that are otherwise affected by acute corpus callosum removal. One hypothesis is that the remaining interhemispheric connections, subcortical and cortical, fulfill the compensatory role. We aimed to identify the likely location of possible thalamo-cortical loops that could underly such compensatory abilities. A number of ipsilateral thalamo-cortical connections were consistently found in all subjects, and in both hemispheres: to the calcarine suclus, the superior frontal gyrus and sulcus and to cingulate cortex. These regions are candidates for the cross-hemispheric interactions too: in typically 3 out of 6 brains, connections from the left thalamus to these right cortical regions and from the right thalamus to the left cortical regions were found. The observed variability of results is, in part, a result of the limited resolution of fiber tracking, and of the difficulty of reliably identifying homologous cortical areas in the sometimes reorganized brains.



Figure 5.4. Number of cortico- and subcortico-cortical connections found in six subjects with an agenesis of the corpus callosum. The different colors encode the discrete (N=0-6) number of brains in which a connection was found; the symmetrical x- and y-axes plot anatomical regions (see Table 5.5 for detailed labels).

Subcortical region	Left ID	Right ID
Thalamus	10	49
Caudate	11	50
Putamen	12	51
Hippocampus	17	53
Nucleus Accumbens	26	58

Table 5.4: IDs of subcortical regions (from FreeSurfer).

Cortical region	ID	Cortical region	ID
Cingulate Cortex	103 104	Calcarine sulcus	144
Cuneus / Gyrus	105	Central sulcus	145
Inferior frontal gyrus	106-108	Central insula	146
Middle frontal gyrus	109	Cingulate sulcus	147-148
Superior frontal gyrus	110	Anterior insula	149-151
Frontomarginal gyrus	111	Collateral traverse	152-153
Insula	112-113	Inferior frontal sulcus	154
Inferior occipital areas	114	Middle frontal sulcus	155
Middle occipital gyrus	115	Superior frontal sulcus	156
Superior occipital gyrus	116	Frontomarginal sulcus	157
Fusiform gyrus	117	Intermedius primus Jensen	158
Lingual gyrus	118	Intraparietal sulcus	159
Parahippocampal area	119	Anterior occipital sulcus	160
Orbital gryus	120	Lunate / middle occipital sulcus	161
Paracentral gyrus	121	Superior occipital sulcus	162
Supramarginal gyrus	123	Occipital-temporal sulcus	163
Superior parietal gyrus	124	Lingual sulcus	164
Postcentral gyrus	125	Orbital sulcus	165-167
Precentral gyrus	126	Paracentral sulcus	168
Precuneus	127	Occipital-parieto sulcus	169
Gyrus rectus	128	Pericallosal sulcus	170
Subcallosl gyrus	129	Postcentral sulcus	171
Subcentral gyrus	130	Precentral sulcus	172-173
Inferior temporal gyrus	131	Subcentral sulcus	174-175
Middle temporal gyrus	132	Suborbital sulcus	176
Superior temporal gyrus	133	Subparietal sulcus	177
Superior temporal gyrus	134-136	Temporal sulcus	179-181
Frontopolar tranverse	137		
Anterior lateral fissure	138-140		
Medial wall	141		
Occipital pole	142		
Temporal pole	143		

Table 5.5: Cortical IDs from FreeSurfer.

Validity of Tracking Fibers

The predictive power of tractography is not only determined by the spatial and angular resolution of the measurements, but -- and maybe more importantly -- by the overall organization and complexity of the fiber tracts themselves. The different cortical fiber tracts have distinct functions and often group axons from similar cortical regions.

Here we compare connectivity obtained from high-resolution diffusion MRI with results from a large number of published tracer studies. First we show examples of tracer and tractography results for the well-studied visual area MT. We then extend the fiber tracking to the whole brain and determine the sensitivity and specificity to predict connectivity. We further try to distinguish between two classes of connections: larger fascicles that can be reliably traced, and smaller ones that are harder to assess. We quantify the validity of fiber tracking for connections of different strength.

Area MT is central to visual motion processing and it is highly connected in the visual domain. It receives strong connections from early visual areas and projects to areas in the parietal lobe. We located the likely location of area MT by mapping a probabilistic histological atlas onto the reconstructed cortical surfaces (see methods chapter). The fiber tracking was seeded in a uniform 1 mm gray matter sheet, just above the white matter, covering the likely location of area MT. The fiber tracking was performed with a fast, custom implementation of a probabilistic tractography algorithm.

Figure 5.6b shows the most likely connections in the early visual cortex. Area MT projects to primary visual cortex V1, area V2, and we found an overlap of the projection with area V3. Panel A shows the projections traced with Fast Blue (FB) for comparison. Similar to diffusion imaging, strong projections to the early visual areas V1, V2 and a weaker overlap with parts of V3 were found. Unlike the tracer study, no connections to the ventral parts of area V4 nor to V4t was found by tractography at the chosen threshold in this hemisphere.



Figure 5.5. Projections from area MT: Identified with Fast Blue and with diffusionbased tractography. (A) Fast Blue (FB) labelled projections from area MT (blue dots, Figure from Lewis and van Essen, 2000) and **(B)** projections traced with probabilistic tractography from MT and are similar to A in early visual cortices (V1, V2, V3, V3A).

Although modern tracers have a high specificity and for example do not easily crosscontaminate adjacent axons, their sensitivity and reliability can vary from study to study. Different tracers can introduce variance, and individual anatomical variability can make it difficult to reliably inject a precise amount of tracer into the same cortical layer in corresponding areas of different individuals. Figure 5.6 reproduced from Lewis and van Essen (2000) shows results from tracer injections in 3 monkey hemispheres. Connections to dorsal area V4d were found in 2 out of 3 cases, connectivity to ventral V4v has not been studied and projections to V4ta have been found in one case and but not in another. The strength estimate to V4tp varies from very strong (+++) in case 17 to weak (+) in case 18. Unlike most tracer studies that occasionally report such variability, our fiber tracking results are more heterogenous. The projection to area V4t not found in the above analysis were, for example, found in the other 3 hemispheres not shown here.



Figure 5.6. Variability of tracing results for injections in area MT. Projections from area MT are traced in 3 hemispheres (subject 17, 18, 19). Connections are quantified on a

scale from "-" (nonexistent) to "+++" (strong projections). For multiple areas, results
vary from case to case (see text for a discussion)
Figure is reproduced from Lewis and van Essen, 2000.

Anecdotally, fiber tracking result do (and sometimes do not) agree with tracer results, but very little is know about how valid tractography predictions are. To assess the validity of the fiber tracking results, we compared them directly against the tracer results obtained from the CoCoMac database. Although we currently have no quantitative estimate of the variability of tracer applications, the comparison is nevertheless useful and provides a lower bound on the validity of the fiber tracking method. We identified the likely location of mostly visual, histologically defined cortical areas using the surface-based approach of Caret 5.5. Fiber tracking was seeded in a 1 mm gray matter sheet of each area. Figure 5.7 shows the location of about 70 different locations of this atlas. A large number of samples was drawn uniformly from each seed region and the location of likely pathways was estimated. The intersections of pathways from a seed region with the other areas was used to compute the full-brain connectivity matrix. The results from the 4 scanned hemispheres were averaged.



Figure 5.7. Cortical parcellation atlas used for seeding the fiber tracking at the likely location of more than 70 distinct cortical areas.

For the mostly visual areas present in the Lewis and van Essen (2000) map, connectivity information was obtained from the CoCoMac database (Stephan et al., 2001). Unlike most tracers, fiber tracking does not differentiate between anterograde and retrograde connections from an area. If either a retrograde or an anterograde connection was reported in the database, the connection was expected to be found by tractography. Probabilistic fiber tracking does not classify connections into existing and nonexisting. Instead, a connectivity likelihood value is computed for each pair of cortical areas. Higher thresholds for the likelihood value produce more conservative estimates, whereas lower thresholds increase the sensitivity of the method at the cost of more false-positive mistakes. Figure 5.8 shows the resulting ROC curve when the

sensitivity (y-axis) is plotted against the false positive rate (x-axis). For high thresholds, with an error rate below 10%, a surprisingly large number of tracts can be identified (approximately 38% of all connections in the database).



Figure 5.8. ROC when comparing the diffusion-based fiber tracking with tracer based results. For every pairing between the cortical areas, we estimated the diffusion-based evidence that a fiber connects the pair. If the evidence is strong, then it is very likely that the connection is indeed there and can be identified with a chemical tracing technique. By increasing the cutoff for the diffusion-estimate strength, we can raise the reliability of our algorithm. By decreasing the cutoff we can increase its sensitivity at the expense of increasing the number of wrongly identified connections. We compared our fiber tracking algorithm against some known cortical connections of the macaque brain (Young, 1993; Sporns et al., 2002) in order to generate the blue ROC curve shown. A' = 0.72 (ROC curve area).

Another method of quantifying the agreement with the CoCoMac data is to compute the odds ratios: the likelihood that a prediction is correct vs. the likelihood that a predicted connectivity does not exist (Figure 5.9). Similar to the ROC analysis above, the odds ratio can be computed for different fibers of different strengths. The ROC curve plots the hit rate against the global false positive rate, the proportion of all unconnected cortical pairs that are wrongly flagged when lowering the threshold. In comparison, the odds ratio indicates how many of the identified connections above a chosen threshold do indeed exist divided by the proportion of chosen connections that were found to be unconnected in tracer studies. The estimates are more volatile than the ROC curve, but may provide a better intuition into how reliable global tractography is. At low thresholds (>0.05) the odds ratio it is about 1:5 and it rises to 1:10 (and above) for stronger connections (>0.2). Fiber tracking, especially of stronger connections, is hence a good method to identify potential pathways and can guide further anatomical studies.

Fiber Tracking





Figure 5.9. Odds ratio of correctness of fiber tracking as a function of connectivity strength. Larger connections are easier to detect and result in more reliable predictions of connectivity. We compared the whole brain connectivity matrix of the macaque cortex with a large number of tracer studies (obtained through the CoCoMac database, Stephan et al., 2001). The probabilistic tractography estimates a weight or likelihood value for each connection. Stronger connections are more likely to exist. We identified potential connections by selecting all tractography results with weights larger than 0, 0.05, ..., 0.35. for each of the sets, we counted the number of true (positive) and false positive entries in the CoCoMac database. This true-to-false positive or odds ratio is shown for each threshold value.

Distance is an important predictor of cortical connectivity. Cortical areas are arranged to minimize wiring length and close-by regions are more likely to be connected than distant pairs. To quantify the influence of distance, we computed the Euclidian distance between the centers of gravity of each region. The resulting distance matrices were computed for each of the four hemispheres and then averaged. We thresholded the inverse of the distance matrix to test to what extent close regions are more likely found to be connected. The resulting ROC curve can be directly compared to the validity results from the fiber tracking (figure 5.8). For visual areas, we found that distance alone is a predictor of connectivity. More distant visual areas are less likely to be connected and the fiber tracking cannot predict whether distant areas are connected with high reliability.



Figure 5.10. Comparison of average and single-case fiber tracking with predictions from distance. For every pairing between the cortical areas, we estimated the diffusionbased evidence that a fiber connects the pair. If the evidence is strong, then it is very likely that the connection is indeed there and can be identified with a chemical tracing technique. By increasing the cutoff for the diffusion-estimate strength, we can raise the reliability of our algorithm. By decreasing the cutoff we can increase its sensitivity at the expense of increasing the number of wrongly identified connections. Spatially close

areas are more likely being connected (red curve) and fiber tracking improves if results from multiple hemispheres are combined (compare blue and black curve).

Weight Estimation

To understand the functional role an area has in a larger network, it is not only important to know to which other nodes it projects, but the different strength can give an insight about the potential information flow. Two areas that are, for example, not directly connected are likely to be part of the same pathway if they are strongly connected through an intermediate region. In contrast, two distant areas that are directly connected could be part of functionally distinct networks if their connection is weak. The strength of fiber pathways greatly varies in different parts of the network. Early sensory areas are, for example, connected by massive fiber bundles, whereas many higher-level areas are connected by smaller bundles. Fibers that contain millions of axons are indicative of high bandwidth transfer between the connecting brain areas, whereas smaller fibers might suggest a less central role of a particular connection in the distributed processing architecture. Similarly, myelinated or unmyelinated axons conduct information at very different speeds and would be suited to different kinds of transmission. Due to practical and technical challenges, explicit tracer studies of connectivity strength are rare and can be difficult to compare.

Diffusion imaging directly measures the net effect of anisotropic diffusion barriers and has been successfully deployed to image the effect of myelination and degeneration of axons. Currently, little is know about how such measurements of fractional anisotropy is related to the actual strength of a larger fiber bundle. We here compare the strength of cortical projections from area V4, as measured by fiber tracking, with weight estimates obtained from a published tracer study (Ungerleider et al., 2007).

To estimate the reliability of our weight estimates, we determined the connectivity strength by fiber tracking from each cortical region. Each region was seeded with a large number of particles proportional to the volume of the cortical area. For each target region, the density of particles from each other area was counted. The correlation between the weight estimates of the 4 hemispheres ranged from r=0.67 to r=0.76, with an expected value of <r>=0.72. The reliability of our weight estimates could be further improved by averaging the results from two hemispheres. We first averaged the results from the left and right hemispheres of each macaque. The aggregated weight estimates correlated with r=0.70 (p < 0.01) between the two brains. We then combined the weight matrices from the 2 left and the 2 right hemispheres and found a left-right agreement of r=0.75 (p < 0.01). The two correlation coefficients are not significantly different and we finally combined the results from all 4 hemispheres for comparison with the tracer data.

Ungerleider et al. (2008) traced and quantified projections from the visual area V4 of the macaque in a total of 20 hemispheres. The strength of each connection was classified as either light (1), moderate (2) or heavy (3). We encoded the strength of each case numerically and computed the expected value for each projection. When comparing the results to the tracer quantification we observed a nonlinear relationship between the raw tractography likelihood and the strength values. To correct for the propensity to weight stronger connections disproportional higher, we applied a gamma correction x^{γ} (γ =0.33) to the fiber tracking estimates. This transformation maps the tractography-based density into a more linear weight space. The resulting numerical values were linearly mapped into a range from 0 to 3 typically used in tracer studies to ease comparison (figure 5.11A). The strength estimates vary between the different cases and tracer molecules used. The average correlation coefficient between one case and the study's mean was <r>=0.75.

We then transformed the particle density matrix obtained from fiber tracking into a similar scale ranging from 0 to 3 (figure 5.11B) and correlated the diffusion weights with the tracer average. Overall, we found an excellent agreement between the tractography estimates and a single tracer quantification of r=0.66. The weight of cortical connections can hence be quite reliably measured with tractography and we will use diffusion-derived weight estimates to study the organization of the cerebral cortex in later sections.



A) Tracer-based connection strength

B) Diffusion-based connection strengths



Figure 5.11. Weights of known connections. A) Strength of connection to area V4 estimated by quantifying dye staining (data from Ungerleider et al., 2007). **B)** Strength estimates obtained by the diffusion-based fiber tracking algorithm. The correlation between our estimates and the tracer based strength (r = 0.66, p < 0.01) is comparable to the reliability of the tracer injections (< r > = 0.75).

Chapter 6

CONNECTIVITY OF THE MACAQUE BRAIN

Understanding the function of cortical regions requires knowledge of their inputs and outputs. Yet many areas are so densely connected that their connectivity profiles reveal very little about their particular functional role in a network. Tractography offers a novel way of measuring the strength of connections (see previous chapter). Here, we analyze the strength of connections that are already known and make predictions about new connectivity. Using path analysis, we use the pairwise strength information between areas to identify likely multinode pathways in the visual and motor domain. We show that such a strength-weighted multinode path distance correlates with anatomically defined hierarchy levels for visual areas. Using dimensionality reduction methods, we visualize global connectivity of cortical regions and demonstrate an architecture that likely reflects functional modularity. The results generate new hypotheses about connectivity as well as provide criteria for the application of similar methods to the study of the human brain.

Strength of Selected Cortical Connections

An important aspect of cortical information processing is the strength of connectivity. An area might, for example, receive input from dozens of cortical areas and project to many parts of the cortex, but not every connection will be of equal importance for understanding the functional role of that area. Here, we deploy tractography to quantify the strength of ipsilateral cortical projections between 91 mostly

visual neocortical regions of the macaque brain. We report detailed results from three cortical areas: the middle temporal area (MT), the visual area 3 (V3) and the secondary somatosensory cortex (SII).

Area MT is central to visual motion processing and is densely connected with the majority of visual areas. It received retinotopic input from all early visual areas: V1, V2, V3, V3A, V4 and V4t. Its projections reach the dorsal parts of the medial superior temporal area (MST), the floor of the superior temporal sulcus (FST) and it also projects to the ventral intraparietal area (VIP) and connects with the parieto-occipital area (PO/V6) (Maunsell and Vanessen, 1983, Ungerleider and Desimone, 1986, Lewis and Van Essen, 2000a). The strength of most MT projections has been quantified with tracers and we were interested to apply the tractography method to quantify these well-studied connections of the visual stream.

Two perfusion-fixed, Gadolinium-treated macaque brains were scanned for about 20 h in a 9.4T and 7T Bruker Biospin MR system. Diffusion-weighted images (b-value: 3000 s/mm^2) with 500 µm spatial resolution and 72-directional angular resolution were acquired as detailed in the earlier chapters. Nine high-resolution anatomical volumes (ca. 100 µm in-plane resolution, 150 µm slice thickness) were obtained and used to locate the likely location of 91 histologically defined areas (see method chapter for details). A custom implementation of a probabilistic fiber tracking algorithm (Behrens et al., 2003b) was seeded in each of the cortical regions and the average connectivity likelihood to every other target region was determined (see previous chapter for details). A gamma correction x^{γ} (γ =0.33) was applied to transform the tractography-based

density into a weight estimate (see previous chapter). The strength was remapped to a numerical range from 0 to 3 that is typically used in tracer studies.

We first assessed the reliability of strength estimates for the connections to and from the middle temporal area (MT) and compared the reliability of the weight estimates measured between the four hemispheres. The singleton estimates were quite reliable with a mean correlation coefficient of $\langle r \rangle_{hemisphere} = 0.73$ (expected value computed using Fisher's transformation; Fisher, 1915). Correlation coefficients between the different hemispheres ranged from r = 0.58 to r = 0.85 (p < 0.01). By combining the weights from the 2 left and 2 right hemispheres the reliability of the estimates could be further improved (r left,right = 0.80, p < 0.01). For the final analysis reported below, results of all 4 hemispheres were combined into a single estimate.

Probabilistic tractography identifies likely pathways, but cannot trace single axons nor show monosynaptic connectivity. To test the validity of the predictions, we compared the tractography results from area MT against results from published tracer studies. The data was obtained from the CoCoMac database (Stephan et al., 2001, see previous chapter for details). The tractography results predict MT connections well with an ROC value of A' = 0.71, similarly to the validity observed on a whole-brain level (A' = 0.73; see previous section).

The tractography weights of projections with area MT are shown in Fig. 6.1. The graph shows a subset of all cortical regions, selected from the visual hierarchy diagram in Felleman and Van Essen, 1991, and the surface atlas of cortical regions used to locate the area boundaries in each hemisphere (Lewis and Van Essen, 2000b). The strongest connections were found with the medial superior temporal (MST) area, the floor of the

superior temporal sulcus (FST), visual area 4 (V4) and area TE. The numerical strength values of the largest tracts (with a weight larger than 1) are shown in table 6.1. Strength values are computed based on the complete cortical map of 91 regions (Lewis and Van Essen, 2000b), including many subregions of larger cortical areas. The atlas, for example, differentiates between the dorsal and ventral parts of many visual areas. These subregions often do have different connectivity patterns and distinct roles in the dorsal and ventral processing stream. Information about such smaller regions with the densest projections are reported in table 6.1. Connections from area MT to area MST are, for example, the strongest in the medial and dorsal part of MST, and MT projections are stronger in area V4t than in V4 proper.

The strength estimates are consistent with tracer results for area MT (Ungerleider and Desimone, 1986, Lewis and Van Essen, 2000a). The size of fibers to most areas in early visual cortex and in the temporal and parietal streams are moderate to light (Fig. 6.1). Interestingly, connections from area V1 and V2 are found to be of moderate strength. The strength quantifications in tracer studies differ between the dorsal and ventral parts of V1 and V2, generally reporting stronger projections from the dorsal areas. Our tractography method defines the strength as density per surface area and could potentially underestimate the quantification of tracer studies, particularly for large areas such as V1 and V2. At the end of this chapter we will compare the density estimates from tractography and tracer quantifications with the strength of functional BOLD coherence.



Figure 6.1. Strength of connections with area MT. The strength of projections to or from area MT were quantified with a custom implementation of probabilistic tractography (see previous chapter for details). The strongest projections were found to the medial superior temporal area (MST), the floor of the superior temporal areas (FST), and visual area 4 (V4) (see Table 6.1 below for numeric values). Fiber tracts to early visual areas V1, V2, V3 and parietal areas are of moderate strength. Cortical areas form functional streams that can be identified from connectivity strength information with unsupervised dimensionality reduction techniques (see later in this chapter). The color of the areas indicate the membership in the different functional clusters (see Fig. 6.11). Negative results, connections with a tractography weight of 0 are indicated by a dotted line.

Strongest	Streng	gt Highest	Tracer	Area
projections	h	density	Strength	
FST	2.02		++	FST
MST	1.84	MSTm/d	+++	MSTdp
V4	1.66	V4t	?	V4tp

Table 6.1. Strength of largest connections with area MT. The strongest connections were found with areas upstream of MT: to superior temporal areas FST and MST, V4 and temporally, to TE. Within these areas, the strongest projections were found to the medial and dorsal subregions (MSTm/d, TE am/d, V4t). Tracer studies report moderate (FST) to strong retrograde projections (MSTdp, V4tp). Tractography results were obtained for 91 cortical areas in 4 macaque hemispheres using a custom implementation of probabilistic tractography.
Acronym	Name	Brodmann number
Occipital		
V1	Visual area 1	17
V2	Visual area 2	18
V3	Visual area 3	
VP	Ventral posterior area	
V3A	Visual area 3A	
V4	Visual area V4	
VOT	Ventral occipitotemporal	
MT	Middle temporal	21
Parietal		
MST	Medial superior temporal	
РО	Parieto-occipital	
IPa		
MDP	Medial dorsal parietal	
MIP	Middle intraparietal	
VIP	Ventral intraparietal	
LIP	Lateral intraparietal	
7a	7a	7
7b	7b	7
DP	Dorsal prelunate	
Temporal	<u></u>	
FST	Floor of superior temporal	
TF/TH	TF/TH	
TE		
Tpt		
TPO		
ТАа		
Somatosensory		
SI	Primary somatosensory	1, 2, 3
SII	Secondary somatosensory	5
Cingulate		
23	Ventral posterior cingulate area	23
24	Ventral anterior cingulate area	24
31	Dorsal posterior cingulate area	31
Frontal		
FEF	Frontal eye fields	8
9	Part of DLPFC	9
45	Inferior frontal / Broca	45
46	Middle frontal / DLPFC	46
11	OFC	11
13	Posterior insular cortex	13

Table 6.2. Cortical areas, their acronyms and Brodmann numbers. The likely location of 91 cortical areas was identified by mapping an atlas (Lewis and van Essen, 2000) to

the cortical surface of each scanned hemisphere. The acronyms of the areas shown in Fig. 6.1 are listed.

We next measured the strength of connections from area V3. The visual area 3 is the highest early visual area before visual processing separates into the dorsal "where" and ventral "what" pathway (Milner and Goodale, 1993). Similar to neurons in area MT, V3 neurons prefers low spatial and high temporal frequencies and many are colorselective (Gegenfurtner et al., 1997, Tootell et al., 1997). The V3 weight measurements found by tractography are reliable (r left,right = 0.88, <r> hemisphere = 0.88, p < 0.01) and the tractography agrees with tracer studies (A' = 0.73). The strongest V3 connections were found to area LOP and to V4 (Fig. 6.2). Projections to V1, V2 and area PIP are overall less dense than to LOP and V4, but denser projections were found to dorsal V2 (Table 6.3). Unlike tracer results that are difficult to compare across studies, tractography provides globally consistent strength estimates for the whole brain. In comparison with area MT, the projections from early visual cortex to V3 are, for instance, stronger than the respective MT connections.



Figure 6.2. Strength of connections from V3. The strength of projections to or from area V3 were quantified with a custom implementation of probabilistic tractography (see previous chapter for details). Strong connections were found to early visual areas. The densest V3 projections are to the lateral occipital area (LOP), dorsal V2 (V2d) and area V4 (see Table 6.3 below). Cortical areas form functional streams that can be identified from connectivity strength information with unsupervised dimensionality embedding technique (see later in this chapter). The color of the areas indicate the membership to the different functional clusters.

Strength	Highest density	Tracer
2.05		
1.09	MSTdp	+++
1.56		х
1.07		x
1.29		х
3.44	V2d	
2.65		х
	Strength 2.05 1.09 1.56 1.07 1.29 3.44 2.65	Strength Highest density 2.05

Table 6.3. Strongest connections with area V3. Connections are strong to all of early visual cortex (V1-V4, PIP, PO). Dense projections are found to the lateral occipital area (LOP) and medial superior temporal area (MST). Tractography results were obtained for 91 cortical areas in 4 macaque hemispheres using a custom implementation of probabilistic tractography and connections with a strength value larger than 1 are reported.

Area SII has historically is a higher-level somatosensory area. SII has connections with primary somatosensory cortex (Brodmann areas 3,2,1), area 7b and the insular fields (Friedman et al., 1986). We estimated the strength of connectivity from Brodmann area 5 to visual, parietal, temporal and cingulate areas. The combined estimates and weights computed for single hemispheres are reliable (r _{left,right} = 0.81, <r> hemisphere = 0.76, p < 0.01). In agreement with previous publications, strong connections to area 7b were found. Additionally, connections to the granular and dysgranular insular fields (Ig/Id), the gustatory precentral opercular area (PrCO), and superior temporal areas (Ts) were quite strong, as were connections to Brodmann area 6 containing premotor cortex (PM) and the supplemental motor area (SMA).

Strongest projections SII	Strength	Highest density	Tracer
Ig/Id	2.33		
PrCo	2.36		
Ri	1.91		
S1	1.85		
Ts	2.33		
6	1.20	6Vb	
7	1.64	7b	+++

Table 6.4. Strength of connections to area SII. SII is strongly connected to the granular and dysgranular insular fields (Ig/Id), the gustatory precentral opercular area (PrCO), and superior temporal areas (Ts), and there are also connections to Brodmann area 6 containing premotor cortex (PM) and the supplemental motor area (SMA). Tractography results were obtained for 91 cortical areas in 4 macaque hemispheres using a custom implementation of probabilistic tractography and connections with a strength value larger than 1 are shown.



Figure 6.3. Strength of connections of secondary somatosensory cortex. The strength of projections to or from area SII were quantified with a custom implementation of probabilistic tractography (see previous chapter for details). Strong tracts were found to parietal areas. The connections to cingulate cortex (Brodmann areas 23, 31) and areas in the temporal parietal junction are of medium strength. The color of the areas indicate the membership to the different functional clusters, and the width of connections as well as their yellow-orange-red color indicates their strength.

Strength-Based Analysis for the Visual Motion Pathway

Strength information can provide insight into the relative importance of the direct input or output connections of an area. Information processing, however, is highly distributed and to understand the function of a network, it is necessary to develop methods that go beyond analyzing pairwise connectivity. Area MT, for example, receives its visual input via multiple routes: Directly, from area V1, and indirectly from V1 via V2, and V3 (Ungerleider and Desimone, 1986). The strength of these routes, from V1 to MT, can provide some insight about their relative importance.

Here we extend the concept of pairwise connectivity strength to quantify the strength of multinode pathways. The problem of optimal information routing has been extensively studied in telecommunication networks. To find an optimal path, routing algorithms assign each direct connection a cost or distance value. If the cost functions are nonnegative, it is possible to find the best or shortest route efficiently with dynamic programming (Dijkstra, 1959). For the analysis presented here, we chose the inverse of the weight matrix as the cost function of pairwise connectivity. If two cortical areas are strongly connected, they are considered to be close in the routing space, whereas weakly connected areas will be placed at a greater distance. Fig. 6.4 illustrates this placement of areas in an Euclidian routing space (B) based on their weights (A).



Figure 6.4. Transforming a weight network into a distance network. Possible pathways of information flow through the cortical network can be identified by quantifying the strength of connectivity. A strong direct connection between two areas suggests that they are involved in joint processing of stimuli. Similarly, two areas that are strongly, but indirectly, connected through an intermediate area are likely to be part of the same pathway. Globally, the connectivity similarity between two areas can be defined as the shortest multinode path between them. Here, we invert the **(A)** strength network obtained from tractography and obtain **(B)** the cortical connectivity distance network. The shortest path for each area from primary visual cortex (Fig. 6.14) and from motor cortex (Fig. 6.15) were computed by the Dijkstra shortest path algorithm. For the visual network, we found a good agreement of this connectivity distance with the hierarchy levels of visual areas when defined using anatomical criteria (Fig. 6.8).

In the previous chapter we have shown that the relative weights of connections can be determined with tractography. Here we extend this concept to multinode routes and evaluate all possible pathways from primary sensory and motor areas to many higherlevel regions. Fig. 6.5 illustrates two hypothetical pathways from a low level (at the bottom) to a high-level area (at the top). For the illustrated case, the 2-node route on the left is more strongly connected than the more direct, weakly connected 1-step route at the center. In the neocortex, most regions could theoretically exchange information often through direct connections or through at most one intermediate region. Functionally however, information processing is typically organized in functional streams and requires multiple, often sequential, processing steps.



Figure 6.5. Assessing multiple pathways between two brain regions. Cortical processing is distributed in parallel pathways. Information is sent from primary visual cortex to higher-level areas via multiple routes. To visualize such routes of likely information flow, we transformed the strength information obtained by tractography into a distance network. For a selected pair of regions, the length of the shortest path through those regions can be computed and provides a quantification of how central an area is for the information exchange between the selected area pair. This analysis does not account for the direction of information flow.

We first studied the visual input paths to area MT. MT receives input from most low-level visual areas: V1, V2, V3, VP and PIP (Maunsell and Vanessen, 1983, Ungerleider and Desimone, 1986). The direct input from V1 and an indirect pathway from V1 via V2 to MT are the most important. To compare the relative importance of the indirect input paths to MT, we estimated all 15 pairwise connectivity weights between the visual areas V1, V2, V3, VP, PIP and MT as described earlier. We next computed the shortest paths from V1 and from MT to each of the four possible intermediate areas. Table 6.5 lists the scores for direct and indirect pathways to the intermediate areas. Area V2 has the strongest connection with V1 (inverted weight score of d=0.24), whereas the direct connection to area MT (d=6.21) is much weaker than the V1 to V3 connection (d=1.83) or the V1 to PIP connection (d=1.83). The strongest direct inputs to MT are from area V2 (d=3.77) and V3 (d=4.27). The best multinode path thus passes through area V2 (d=4.01), and is, according to our cost function, stronger than the direct connection (d=6.21) and the pathway through V3 (d=5.19).

Path distance measure	V1	V2	V3	VP	PIP	MT
Direct from V1	0	0.24	1.83	(-)	1.83	6.21
Direct from MT	6.21	3.77	4.27	24.73	9.81	0
Path from V1	0	0.24	0.92	0.88	1.45	4.01
Path from MT	4.01	3.77	4.27	4.40	4.97	0
Path score V1-MT	(6.21)	4.01	5.19	5.28	6.43	(6.21)

Table 6.5. Path distance from area V1 to MT. Motion information is computed in multiple, parallel processing steps from V1 to MT. Tractography can identify possible pathways by scoring the strength of such multinode pathways. The existence of a strong multinode pathway makes it likely that a large amount of information is exchanged between the areas along such a path. This table shows the computation of the pathway score (bottom row), based on the strength of direct connections (first two rows) and pathway length obtained by the Dijkstra shortest path algorithm (intermediate rows). Area MT received its input from primary visual cortex (V1) and from areas V2 and V3.

The connectivity distance values, from V1 and from area MT are shown in the top rows. The path distance to each intermediate area from V1 and MT are computed below. The direct distance was defined as the inverse of the strength of the respective fibers. The distance of a path through an intermediate area is the sum of the shortest path to that area from V1 plus the path length from that area to area MT. The pathway length through each intermediate area is shown in the bottom row. The distance score for direct connection from V1 and MT is given in parentheses.



Figure 6.6. Routing distance for connections through potential intermediate areas from primary visual cortex to area MT. Hypothetical routing distance for pathways between primary visual cortex and area MT through different intermediate areas. Shown are the tractography-derived route distances of the shortest multinode pathways from V1 to MT that pass through the individual cortical areas. The inverted weight matrix was used to quantify the strength of each cortical connection and Dijkstra's shortest path algorithm was used to identify the shortest route from V1 to each intermediate region and from each intermediate region to MT (see Table 6.5 and text for details). The length of the shortest route through each region is shown above: yellow and orange hues indicate a strong multinode connectivity from one area to both V1 and MT. This analysis treats feedforward and feedback connections equivalently, since the directionality of connections cannot be obtained with tractography.

Strength-Based Pathway Analysis of the Visual Stream

Cortical information processing is highly distributed and the anatomical connectivity pattern shows aspects of small-world and scale-free networks (Bullmore and Sporns, 2009). In the visual domain, the processing of stimuli proceeds through multiple levels that are characterized by increasing receptive field sizes and by an increase in complexity of the neuronal representations. The organization of feedforward connections is consistent with a partial order of visual areas, whereby higher-level areas, by definition, receive inputs from lower-level areas. Initially, this ordered organization was interpreted as a total hierarchy, whereby each area could be assigned one, globally consistent hierarchy level (Felleman and Van Essen, 1991; Scannell et al., 1995; Hilgetag et al., 2000). However, cortical processing is organized in multiple, parallel streams, suggesting a partial hierarchical rather than a total ordering of areas. Feedback connections play an important and often neglected role in stimulus processing. Feedback connection are, for example, thought to be responsible for the complex shape sensitivity in primary visual cortex (see Hegde and Felleman, 2007).

In the previous section we identified the shortest path from primary visual cortex to area MT and assigned a length score to each pathway. Here, we apply this shortest path algorithm to the complete weight graph between 91 cortical areas. The idea of the graph-based algorithm is to identify the likely paths of information flow from primary visual cortex by following the strongest connections from area V1. We use the Dijkstra shortest path algorithm (Dijkstra, 1959) to find the best network path to each area. The length score of the path is a measure of multinode connectivity and it correlates, as we will show later, with the hierarchy level of an area.

The algorithm assigns a distance scores to each area that quantifies the multinode connectivity to area V1. The order implied by this distance score, for example for early visual areas (V1 < V2 < V3 and V4), is in good agreement with the known functional organization of these areas and we will compare these scores with anatomically defined hierarchy levels later (see Fig. 6.8). The relatively coarse resolution of the diffusion imaging does not allow the assignment of an exact hierarchy level and cannot replicate the precision of anatomical studies. The strength of the tractography method is its whole-brain coverage and the possibility to compute approximate hierarchy levels for different parts of the cortex. According to the visual pathway scores, clusters of highlevel areas, areas that do not receive strong inputs from early visual areas, are found in the intraparietal sulcus, on the ventral surface of the temporal lobe, medially for anterior regions of cingulate cortex and frontally in the dorsolateral prefrontal cortex (Figure 6.7). These cluster locations agree with known centers for action planning (intraparietal sulcus), high-level object and face recognition (temporal cortex), behavioral monitoring (anterior cingulate) and behavioral inhibition (prefrontal cortex).

<image><page-header><page-header>

B Graph distance from primary visual cortex (medial view)



A Graph distance from primary visual cortex (lateral view)



C Graph distance from primary visual cortex (flatmap view)

Figure 6.7. Routing distance from primary visual cortex. Length of the best network route from primary visual cortex to each cortical area: **(A)** lateral view, **(B)** medial view and **(C)** flattened cortical map. Connectivity strength between cortical areas were obtained with probabilistic tractography. The direct connectivity distance between two areas was defined as the inverse of the connectivity strength. Predicted from this distance metric, many weakly connected areas have their shortest connectivity path through one or multiple intermediate areas. The length of this best multinode route from primary visual cortex is shown for each cortical area. Clusters of high-level areas in the intraparietal sulcus, areas in the cingulate cortex and at the ventral surface of the temporal lobe are most distant from V1. These clusters contain areas for high-level motor planning, and are involved in monitoring and correcting behavior. Single areas can appear distant from V1 if the tractography did not identify strong pathways to these areas.

We next wanted to know to what extent the weight-based pathway distance agrees with anatomically defined hierarchy levels. Hierarchy relationships are exclusively defined by the input an area receives, irrespective of the strength of such connections: if an area receives input from areas with hierarchy levels of at most *n* and does not provide input to any areas of or below level n, it will be assigned the level n+1. Fiber tracking by itself does not provide information about the originating or terminating layer of a connection and hence cannot discriminate between feedforward and feedback connections. Yet, if weight information is available, it is possible to use pathway analysis in order to identify the likely flow information through the cortical network. Areas with similar hierarchy levels are likely to be more strongly connected than areas on distant levels of the hierarchy. Felleman and van Essen (1991) defined the hierarchy level of an area as the node distance from primary visual area V1. We defined the length of the shortest multinode pathway from primary visual cortex as an indicator of the hierarchy of an area in visual cortex. The length of a direct connection between two areas was defined as the inverse of the diffusion-based strength estimate. The multinode distance between two areas is the sum of the length scores of the shortest path between two areas (see previous section). We then compared the multinode connectivity distance with the hierarchy relationships determined from tracer studies (Felleman and Van Essen, 1991). Fig. 6.8 shows the correlation between the graph-based distance and the hierarchy of an area (defined by the y-position in figure 4 of Felleman and Van Essen, 1991). We found a good agreement (r = 0.75, p < 0.01) between the anatomical hierarchy levels and weight-based path distance (Fig. 6.8). Global quantitative tractography is thus a good tool for studying (partial) hierarchical structures of the cortex.



Figure 6.8. Multinode path distance of visual areas to V1 correlates with the hierarchy level of Fellman and van Essen (1991). The multinode path distance of visual areas was computed using tractography-based connectivity weights (see Fig. 6.7). This routing distance quantifies the connectivity strength of the strongest multinode pathway from primary visual cortex to a cortical area. It correlates with anatomically defined hierarchy levels (r = 0.83, p < 0.001). Each dot indicates the hierarchy level of an area (y-axis) and the path distance from V1 (x-axis).

Strength-Based Pathway Analysis of the Motor System

A large proportion of the macaque cortex is dedicated to the planning, preparation, and control of motor behavior (Rizzolatti et al., 1998). We here study the importance of cortical areas for motor behavior by calculating the path distance to primary motor cortex. Similar to the analysis in the visual stream, the inverse of the pairwise connectivity strength was used as a distance function. Dijkstra's shortest path algorithm was used to compute the best multinode route from each area to primary motor cortex. In contrast to the analysis of multinode connectivity from primary visual cortex, the information flow in the motor domain is inverted, and information is routed towards the primary motor cortex. Fig. 6.9 shows the path distance from each area towards M1. Many areas along the central sulcus are closely connected to the primary motor area. Interestingly, most medial areas in the cingulate cortex have stronger multinode pathways to M1 than to V1, in line with the importance of cingulate areas in high-level behavioral control (compare Fig. 6.7 and Fig. 6.9). Clusters in the intraparietal sulcus and at the ventral surface of the temporal lobe are distant to the motor system.



Figure 6.9. Routing distance from primary motor cortex. The length of the multinode distance from a primary sensory area is a measure of hierarchy level of an area (see Fig. 6.8 and 6.7). The multinode route distance from primary motor cortex is shown on a flattened cortical map. Connectivity strength between all cortical areas was derived with a custom implementation of probabilistic tractography. The inverted pairwise connectivity strength values were used to define a direct path distance between two areas. The shortest multinode route between M1 and each cortical area was computed and the length of this shortest path is shown above. Clusters in the intraparietal sulcus and on the surface of the ventral temporal lobe are the most distant from M1. Areas in the cingulate cortex are more closely connected to primary motor cortex than to the visual network.

Identification of Anatomical Subnetworks Using Locally Linear Embedding

The connectivity profile of cortical areas are signatures of their functional role in the cortical network. Functionally similar areas receive similar input and project to similar output regions. In the previous sections we have shown that the strength of connections provides information about their relative importance. Here, we deploy a dimensionality reduction technique, locally linear embedding (Roweis and Saul, 2000) - for the first time with anatomical data — to visualize the global organization of cortical connectivity. The algorithm fits a low-dimensional manifold — a 2-dimensional surface in our case to the high dimensional data by analyzing local changes of the data distribution. The algorithm first identifies, for each area, the most similarly connected regions ("neighbors"). The number of neighbors k used for the local embedding is the only free parameter of the LLE algorithm. We observed reliable embedding results for the range of k = 6 to k = 9, and chose k = 8 for all subsequent analyses. For each neighborhood clique, the main axes of variance are determined by principal component analysis and then used to construct the low-dimensional embedding manifold. Globally, the embedding preserved the local neighborhood relationships. Areas that are similar in connectivity space will be close in the embedding map. Local variations on the map signify systematic changes in connectivity between areas.

We first selected a subset of areas with the most reliable weight estimates for the embedding procedure. The reliability of weight estimates depends on a number of circumstances: the anatomical localization of area boundaries requires proximate anatomical landmarks that are not always present and complex white matter organization can negatively affect the weight estimation. The 59 most reliable regions were identified based on the observed variance of the weight estimates across the 4 hemispheres. Fig. 6.10 shows the connectivity strength between these best 59 regions. The main axis of variance of this 59 \times 59 matrix was determined with principal component analysis (PCA). Regions are ordered according to the first principal component: roughly from somatosensory to motor areas.



Figure 6.10. Connectivity strength matrix. The matrix shows the connectivity weights between 59 cortical areas with the most reliable estimates. Regions are sorted according

to the first principle component of this connectivity matrix: from somatosensory to motor areas. The weights are a direct transformation of the tractography density to a numerical range from 0 to 3.

This subset of most reliable regions was analyzed by locally linear embedding (LLE). Fig. 6.10 shows the original weight estimates of the tractography algorithm. For all further analyses, known false positive results – defined from cases where no connectivity had been found in the CoCoMac database (see Fig 5.9 and Fig. 5.10) – were remove from the weight matrix. Fig. 6.11 shows the results of the LLE analysis. Multiple clusters are identified: visual areas (on the left, in green), somatosensory areas (at the bottom, in red), motor areas (on the right, in purple) and executive areas (on the top, in blue).



Figure 6.11. Similarity between areas based on the connectivity profile. The connectivity profiles of 59 cortical regions (Fig. 6.8) were embedded in a 2-dimensional manifold by locally linear embedding (Roweis and Saul, 2000). The dimensionality reduction technique preserves local neighborhood relationships: areas with similar connectivity will appear close on the 2-dimensional map. Systematic differences across regions appear as local changes in position and generate the stringlike appearance of the embedding. The similarity in connectivity mirrors the know functional organization. Multiple clusters are found: visual areas (in green on the left), somatosensory areas (red), motor and premotor areas (purple) and prefrontal areas (blue). The hue was assigned

according to the angular position of an area on the map (e.g., green to the top left, blue to top right, etc.).

A Similarity between areas of the left hemisphere





B Similarity between areas of the right hemisphere

Figure 6.12. Similarity between areas based on the connectivity profile for the left and right hemisphere separately. To assess the split-half reliability of the embedding technique we performed the locally linear embedding separately for the data of the two left hemispheres (A) and for the right hemispheres (B). The connectivity profiles of the 59 most reliable cortical regions were embedded in a 2-dimensional manifold by locally linear embedding (Roweis and Saul, 2000). The dimensionality reduction technique preserves local neighborhood relationships: areas with similar connectivity will appear close on the 2-dimensional map. Systematic differences across regions appear as local changes in position and generate the stringlike appearance of the embedding. The similarity in connectivity mirrors the know functional organization. Multiple clusters are found: the two visual streams, somatosensory areas, motor and premotor areas and prefrontal areas. The hue was assigned according to the angular position of an area on the map.

In the previous sections we have shown that weights can identify the most relevant input or output connections of an area and that they can be used to trace likely multinode pathways to an area. Clusters of similarly connected regions can be identified with locally linear embedding in an unsupervised fashion. Each of these methods characterizes a different aspect of the anatomical organization. We next tried to visualize the network properties revealed by these analysis methods in a single diagram. Fig. 6.13 summarizes the results for visual areas: the weight estimates, the graph-based path analysis and the clustering using locally linear embedding. LLE assigns each node a 2-dimensional coordinate that characterizes its membership with similarly connected regions ("clusters"). The angular position in the LLE embedding space identifies the membership in a cluster. In figure 6.13, we have taken this LLE angular position to define the x-position of each area, grouping the mostly visual areas into separate processing streams. The routing distance from primary visual cortex is a good predictor of the hierarchy level of visual areas (see Fig. 6.8). The path distance defines the y-axis of each region in the plot: low-level areas appear at the bottom of the plot and higher-level areas are at the top. The weights of the 100 strongest connections are shown and the color of these connections reflects their relative strength. The hierarchy or path distance score is most reliable for the strongly connected low level areas. It for example identifies the hierarchical relationship between V1, V2, V3 and V4. Areas involved in higher0level processing, such as area 7a, 7b, the perirhinal cortex (Brodmann 36) and the dorsolateral prefrontal cortex (Brodmann 45) received high hierarchy scores (Fig. 6.13). The xposition of a region (corresponding to the angular position in the 2-dimensional LLE)

correlates with the membership in functional streams and was used to select the hues for coloring the areas. The LLE embedding procedure separates areas of the 'what' and 'where' pathways. Within the dorsal stream, the embedding separates motion-selective areas (e.g., MT and LIP) from areas involved in planning reach and eye movements (e.g., MDP and MIP).



Figure 6.13. Tractography-based visual hierarchy. Visual information processing is organized into parallel streams. Areas within each stream have similar connectivity profiles. This organization can be revealed by locally linear embedding (LLE) (Roweis and Saul, 2000), an unsupervised dimensionality reduction technique (see Fig. 6.11).

Within each stream, areas at adjacent hierarchy levels are strongly connected. The most likely pathways from primary visual cortex (V1) to each area are obtained by graphbased pathway analysis (Fig. 6.7). The connectivity strength of a pathway correlates with anatomically defined hierarchy levels (Fig. 6.8). The x-coordinate of each area in the plot shows the cluster membership in the LLE analysis (the angle in a polar coordinate system in embedding space). Similarly, color indicates the cluster membership; gray areas have a more ambiguous cluster membership. The y-position corresponds to the multinode pathway distance from V1. The weights of the 100 strongest connections are shown to illustrate important connectivity within and between pathways. The width and color of connections indicates their strength (stronger connections are red, lighter connections are shown in yellow).

Fig. 6.13 shows the embedding and hierarchy information for only a few selected visual areas. Fig. 6.14 provides an overview of the connectivity information of all cortical areas of the Lewis and van Essen (2000) atlas. In additional to the ventral and dorsal visual stream, areas are grouped into a premotor stream (including the posterior cingulate), a frontal stream (including the orbitofrontal cortex, Brodmann areas 12 and 11) and the dorsolateral prefrontal cortex (DLPFC, Brodmann 46). The orbitofrontal cortex is involved in valence processing and goal-directed action selection (Schoenbaum et al., 1998; O'Doherty et al., 2001; Gottfried et al., 2003; Anderson et al., 2003). The dorsolateral prefrontal cortex (Brodmann 46), and the insular cortex (Brodmann 13) have been implicated in emotional awareness and decision making (Craig, 2009).



Figure 6.14. Tractography-based hierarchy of cortical areas based on their multinode connectivity to V1. Information processing is organized into parallel streams. Areas within each stream have similar connectivity profiles and this organization can be revealed by an unsupervised dimensionality reduction technique. Within each stream, areas at adjacent hierarchy levels are strongly connected and the likely pathways from

primary visual cortex (V1) to each area can be revealed by a graph-based pathway analysis. The x-position of each area was determined with locally linear embedding (see Fig. 6.11) and was used to assign colors to each area too. The y-position corresponds to the multinode pathway distance from V1. The 100 strongest connections are shown to illustrate important connectivity within and between pathways. The width and color of connections indicates their strength (stronger connections are red, lighter connections are shown in yellow).

The previous analyses have shown the cortical organization based on the path distance to the primary visual cortex. In the macaque, many areas are involved in highlevel motor control and the hierarchy levels of such areas can be estimated with graphbased methods (see Fig. 6.9). Fig. 6.15 summarizes the connectivity information for the motor hierarchy. The motor hierarchy score is calculated as the path distance from M1 that determines the y-position of each area. The algorithm is otherwise identical to the analysis of the visual hierarchy (see Fig. 6.14) and locally linear embedding is used to group all areas according to their similarity of their respective connectivity profiles (see Fig. 6.11). Lower level motor areas, such as area 4C and premotor area 6, are placed close to primary motor cortex by the path analysis. Similarly, the posterior cingulate cortex (Brodmann 24) is classified as being relatively low in the motor hierarchy. Part of the motor hierarchy are subregions of Brodmann area 8 that includes the frontal eye fields. In agreement with the involvement of Brodmann area 5 (SII) in motor control, the connectivity algorithm places this area low in the hierarchy. Most visual areas are not directly involved in motor control. The pathway scores of many purely visual areas likely does not provide useful information about their involvement in behavioral and motor control.



Figure 6.15. Tractography-based hierarchy of cortical areas based on their multinode connectivity to M1. Information processing in organized into parallel streams. Areas within each stream have similar connectivity profiles and this organization can be revealed by an unsupervised dimensionality reduction technique. Within each stream,

areas at adjacent hierarchy levels are strongly connected and the likely pathways to primary motor cortex (M1) from each area can be revealed by a graph-based pathway analysis. The x-position of each area was determined with locally linear embedding (see Fig. 6.11) and was used to assign colors to each areas too. The y-position corresponds to the multinode pathway distance from M1. The 100 strongest connections are shown to illustrate important connectivity within and between pathways. The width and color of connections indicates their strength (stronger connections are red, lighter connections are shown in yellow).

Functional Subnetworks Based on Ultraslow Oscillations of Blood Oxygenation

Most of the brain's energy expenditure is related to ongoing activity (Raichle and Mintun, 2006). This "default mode" of spontaneous brain synchronization has been thought to reflect ongoing conscious mental activity (Morcom and Fletcher, 2007), yet temporally coherent activity of BOLD activity in the oculomotor, somatomotor and visual systems has however been shown to persist during deep anesthesia in macaques (Vincent et al., 2007), at isoflurane levels known to induce the complete loss of consciousness (0.8%-1.5%). Under deep anesthesia, the BOLD correlations have been found to mirror the awake synchronous activity in humans and macaques. Anecdotally, the reported correlations agree with anatomical connections between the frontal eye fields and intraparietal areas. Anatomical similarities were also found for the precuneus / posterior cingulate cortex and prefrontal BOLD correlations. This BOLD

activity under anesthesia could thus reflect baseline communication between connected regions regardless of the level of consciousness (Vincent et al., 2007).

To test this hypothesis globally, we obtained this publicly available data set (Vincent et al., 2007), reanalyzed the BOLD coherence pattern and compared the strength of pairwise cortical coherence with the strength of the anatomical connections. The data consists of multiple 15 min gradient-echo EPI runs that were acquired on a 3T Siemens Allegra scanner (TR = 3.02 s, TE = 25 ms, 90° flip angle, 1.5 mm in-plane resolution and 1.6 mm slice thickness) in 8 macaques. The BOLD images were transformed to atlas space, resampled to 1.5 mm isotropic resolution and temporally filtered to retain ultraslow BOLD oscillations above 0.0025 Hz (Vincent et al., 2007). We reconstructed the cortical surface of the atlas brain and located 91 cortical regions as described previously. The mean BOLD signal across all gray matter voxels in each of the 91 regions were extracted for each run. The temporal coherence between the time courses for all cortical regions was estimated with multitaper analysis (Mitra and Pesaran, 1999; Mitra and Bokil, 2008). The resulting 91x91 coherence matrices were combined across the 84 available runs.

Traditionally, ongoing neocortical activation is analyzed by either correlating signals to a default mode region (such as the posterior cingulate cortex or the precuneus) or by identifying independent combinations of signals with principal or independent component analysis. Here, we chose a different approach to group distributed activity into similar clusters. Rather than comparing the signal with prototypical signals, we clustered regions based on their full coherence profile using locally linear embedding (LLE). Each region was projected into a high-dimensional feature space, whereby the coherence values to all cortical regions define the coordinates of each region in this 91-dimensional feature space. A 2-dimensional manifold is fitted to distribution of areas by LLE. The resulting 2-dimensional map is shown in Fig. 6.17. We found three clusters of regions that exhibit similar coherence profiles. The biggest cluster spans the visual domain and includes somatosensory and primary motor cortex (blue areas in Fig. 6.17). The two other clusters have their most characteristic regions in the prefrontal cortex (red and green in Fig. 6.17). Regions in the intraparietal sulcus belong to either of the two frontal clusters, consistent with the previous findings of frontal-parietal correlations (Vincent et al., 2007).



Figure 6.17. BOLD clusters of similarly, functionally connected brain regions. The color code shows the cluster membership (the angular position in the embedding plot), projected onto the cortical surface of the template brain for reference in identifying regions. Three clusters were identified and are shown in blue, red and green colors. Similarly colored brain regions have a more similar functional connectivity profiles,

whereas distinctly colored regions fall into separate clusters. The coherence between areas of each cluster are shown in Fig. 6.18.

Unlike previous analyses, locally linear embedding groups regions based on the similarity in their coherence profile, rather than based on the similarity of the signal per se. For further analysis of the connectivity profiles we selected the most characteristic regions for each cluster. The embedding technique assigns a 2-dimensional coordinate to every profile. This 2-dimensional Cartesian coordinate can be converted into a polar system, whereby the angle corresponds to the cluster membership and the radius indicates the distinctiveness of a region. Fig. 6.18 shows the coherence profile for the most characteristic regions of each cluster. Interestingly, only the visual cluster shows a strong widespread within-cluster coherence. Regions belonging to the frontal clusters 1 and 3 have relatively sparse coherence profiles.




Importance of Strength Information for Functional Networks

Under anesthesia, the temporally coherent activity overlaps with the anatomical projection pattern in the oculomotor and visual system (Vincent et al., 2007). It has been suggested that this background, 'always-on' activity reflects anatomical connectivity irrespective of the level of consciousness or wakefulness. Here we test this hypothesis directly, by comparing the degree of temporally coherent activity with results from a large number of published tracer studies. The BOLD coherence between 176 cortical regions (Paxinos atlas) was computed from a total of 21 h of BOLD recordings in 8 anesthetized macaques (Vincent et al., 2007).

The neocortical connections form a very dense network of interconnected regions, and not every connection might contribute equally to the observed coherence pattern in the brain. In the previous section we have shown that the strength of connections can identify the most important input and output of a region, and that the modified strength measure is a good function to identify functionally relevant pathways through the cortical network. In humans, tractography-based connectivity strength correlates with BOLD signal coherence between cortical regions (Hagmann et al., 2007). Yet, little is known about this relationship under anesthesia. Here, we correlated the 91x91 anatomical weight matrix obtained with tractography for high angular resolution diffusion-weighted scans of 4 macaque hemispheres with the coherence matrix computed from 15 min BOLD runs in 8 macaque monkeys under deep anesthesia. We first identified the percentiles of connections based on their relative strength. For each connectivity strength segment, we computed the average complex BOLD coherence value. Fig. 6.19 shows the mean absolute value of the coherence. The stronger two area are connected, the higher is their observed BOLD coherence. The coherence mean and the connectivity strength show a strong correlation relationship (r = 0.79, p < 0.01). This relationship is independent of distance and the correlation is unaltered if distance is removed from either the coherence or strength distribution with linear regression.



Figure 6.19. Prediction of BOLD coherence as a function of anatomical connectivity strength. Ultraslow BOLD coherence (0.01-0.25 Hz) between different brain regions depends on the strength of anatomical connectivity between the regions (r = 0.79, p < 0.01). Weakly connected regions show a small increase in coherence and the largest coherence values were observed for areas with strong connectivity. The strength of cortical connections was obtained with a custom implementation of probabilistic tractography. The blue line shows the mean coherence values and the red lines the

standard error of the estimates. The coherence values were computed by multitaper analysis (Mitra and Bokil, 2008) of the BOLD signal in anatomically defined cortical regions of eight isoflurane-anesthetized macaques (Vincent et al., 2007).

Given the strong correlation between the strength of anatomical connectivity from tractography and BOLD coherence, we next tested whether binary connectivity per se, regardless of its strength, influences coherence. We obtained the published results of tracer injections for each of the possible cortical pairs from an online database (CoCoMac, Stephan et al., 2001). Only information from tracer studies that used similar definition of cortical areas as the Paxinos atlas were included in the analysis. Information about 1454 area pairs were retrieved from the database: 703 connections were found to be connected (48%) and for 751 pairings negative results are reported in the database (52%). We then computed the average BOLD coherence for connected and unconnected cortical pairs. Only a small difference in BOLD coherence between regions that are directly connected was found, compared to regions that do not share a direct connection (p=0.06; t=1.87) (see Fig. 6.20). Connections as defined from tracer studies thus do not bear a strong relationship to the ultraslow BOLD coherence between regions.



Figure 6.20. Binary connectivity does not predict coherence. The coherence between connected areas ("x") and unconnected areas ("-") is not significantly different (p=0.06, t=1.87). The mean coherence values across all connected and unconnected pairs are shown above. The error bars show the standard error of the mean estimates. Anatomical connectivity information was obtained from the CoCoMac database. The coherence values were computed by multitaper analysis (Mitra and Bokil, 2008) of the BOLD signal in anatomically defined cortical regions of eight isoflurane-anesthetized macaques (Vincent et al., 2007).

Tractography provides globally consistent strength estimates. In contrast, the quantification in tracer studies depends on many technical details. It, for example, varies with the amount of injected tracer, and the location and depth of the injection site. When we compared connections with weak, medium and heavy tracer results, an analysis of variance (ANOVA) found no difference between the average BOLD coherence values (p = 0.16, F = 1.74).



Figure 6.21. Tracer results do not predict coherence. More strongly connected regions ("++" / "+++") do not show a higher degree of coherence than weakly connected regions ("+") (p = 0.16, F=1.74). Anatomical connectivity information was obtained from the CoCoMac database. The coherence values were computed by multitaper analysis (Mitra and Bokil, 2008) of the BOLD signal in anatomically defined cortical regions of eight isoflurane-anesthetized macaques (Vincent et al., 2007).

Summary

Tractography can be used to reliably measure the strength of cortical connections. The strength is a good indicator of the relative importance of the different input and output connections of an area. More importantly, tractography allows for the estimation of weights globally. Using path analysis, it is then possible to quantify and compare the strength of parallel multinode pathways to an area. The strength or distance of such multinode pathways correlates with anatomically defined hierarchy levels in the visual system. Similar regions have similar connectivity profiles and locally linear embedding (LLE), an unsupervised dimensionality reduction technique, can be used to identify functionally similar regions. Functionally, the strength of BOLD coherence between cortical areas is strongly correlated with the strength of connections, but not with the (binary) connectivity obtained from tracer studies.

Chapter 7

EXPLORING CONNECTIVITY OF THE HUMAN BRAIN

Neurophysiological signals such as BOLD fMRI are highly variable. Spontaneous BOLD fluctuations in the anesthetized macaque are synchronized via large intrahemispheric fiber tracts (see previous chapter). Less is known about the influence that subcortical and interhemispheric loops have on ultraslow BOLD oscillations. To explore these effects, we studied the global pattern of ultraslow BOLD coherence in the awake human brain, while participants performed a visual attention and discrimination task. Similar to the results in the anesthetized macaque, we find strong correlations between the anatomical strength of intrahemispheric cortical connections and the degree of BOLD coherence between two cortical areas. Across hemispheres, only the corresponding areas are strongly coupled. Subcortical areas have a strong effect on the apparent cortical coherence pattern and can explain much of the similarity in the apparent spontaneous BOLD synchronization between many areas.

In the second part of this chapter we study how fiber tracking can be used to reveal functional specialization in the fusiform gyrus. We show, in a small pilot study, that tractography predicts the site of the face patch within the fusiform gyrus. The location of the FFA coincides with the subregion with strong connectivity to the lingual gyrus.

To explore the relationship between the cortical coherence patterns of ultraslow BOLD oscillations (<0.5 Hz) and the anatomical connectivity, we conducted a visual attention and discrimination experiment. Grayscale photographs of faces and houses were presented to five participants. In each trial, one pair of houses and one pair of face photographs was shown. The participants paid attention to either the horizontal or vertical image pair, while maintaing fixation at a cross in the center of the screen. The five participants were instructed to report whether the two images were identical or different by pressing one of two buttons (see Vuilleumier et al., 2001 for detailed description of the task). Two additional 'localizer' runs and 2-3 anatomical scans were obtained for each subjects. During the localizer blocks, faces and houses were presented centrally while subjects who performed a simple 1-back memory task for centrally presented digits. Scanning was performed in a 3T Siemens scanner with an 8-channel head coil. Standard EPI imaging procedures were used (TR=1.5s, TE=30ms, flip angle of 75°, 3mm in-plane resolution, slice thickness of 4.5mm, 64x64x25 matrix). Anatomical ROIs were automatically identified by the "FreeSurfer" software on the individual cortical surfaces. The anatomical scans were aligned to the functional BOLD images using FSL. Prior to the alignment procedure, spatial MR intensity biases were removed with a nonuniformity-correction algorithm (Sled, 1997) and the spatial alignment was manually verified. The anatomical surface-based labels were then used to label the respective voxels of the functional scans. Preprocessing of the EPI images was performed in SPM5. The acquired images were corrected for slice timing and head motion. No spatial smoothing nor spatial atlas normalization was performed.

Correlates of translational and rotational motion and any correlation with a full eventrelated design matrix were removed from the BOLD signal in each voxel. The residuals of this regression were extracted from SPM.

For each anatomical region, the average BOLD signal across all voxels of an ROI was extracted. To correct for global signal drifts inherent to the scanning equipment, respiration artifacts, residual effects of head motion and other artifacts such as global blood pressure changes, we extracted the MR time courses from metabolically inactive regions in each subjects' ventricles and a subcortical white matter location. Any correlates of these signals in the extracted signals of each region were removed.

The coherence between the different BOLD time courses were computed by multitaper analysis (Mitra and Bokil, 2008). The estimation bias inherent to the multitaper coherence analysis was corrected with the CMTM procedure⁷. Coherence values were obtained for an ultraslow frequencies range (0.01-0.33 Hz) for all possible pairs of anatomical regions. The resulting coherence value is a complex number that quantifies the phase difference and similarity ("correlation") between two signals at a particular frequency. We computed the mean of the complex value first across all sessions, then across subjects and the relevant frequency range. This analysis identifies pairs of brain regions that show coherent signals with a stable phase difference across sessions and subjects. The amplitude of the BOLD coherence between all regions in the 0.01 Hz frequency band is shown in Fig. 7.1.

⁷ http://www.people.fas.harvard.edu/~phuybers/Mfiles/cmtm.m





Ultraslow fluctuations of BOLD-sensitive MR signals (0.01 Hz) between subcortical and cortical areas are highly correlated. The amplitude of the average coherence for each possible pairing between anatomically defined brain regions is shown here. Strong coherence was observed between **1**) the subcortical ROIs and the ipsilateral cortical regions, **2**) many cortical areas within a hemisphere and **3**) between corresponding areas in the left and right hemispheres. Participants were performing a visual discrimination and attention task (*see text for details*). Any correlates of motion and of stimulus- and task-evoked activity, physiological and background fluctuations of the BOLD signal were removed via multiple regression prior to the coherence analysis.

Our exploratory analysis of the 22 BOLD sessions reveals strong coherence between subcortical regions and many ipsilateral cortical spontaneous BOLD fluctuations (see top rows in Fig. 7.1). Slow thalamo-cortical neuronal oscillations (<15 Hz) have been observed in the awake cat (Steriade, 1997) and during sleep (<1 Hz) (Steriade and McCarley, 1990; Steriade, 2000). Two of the participants suffer from a rare, autosomal recessive, genetic disease, Urbach-Wiethe disease (Hofer, 1973), which resulted in calcification of the amygdala. No significant or otherwise apparent differences in the coherence pattern compared to the healthy subjects was observed. For the final analysis reported here, results from all participants were combined. The subcortical-cortical synchronization of oscillations is assumed to be caused by weak neuronal coupling between neurons. It is conceivable that this weak coupling is not specific to a single frequency and similarly synchronizes the ultraslow intrinsic oscillations in subcortical and cortical areas (see Leopold et al., 2003). Cortical BOLD coherence is stronger between areas within the same hemisphere than between (nonhomolog) areas in different hemispheres. No significant difference between the left and right hemisphere was found. Interestingly, we also found strong coherence between the corresponding regions of the left and right hemisphere. Homolog areas are strongly connected through the corpus callosum and the high degree interhemispheric connectivity could explain the observed coherence pattern (Fig. 7.2).





If spontaneous activity is caused by weak loops between the subcortical and cortical regions (see Steriade and McCarley, 1990, Crick and Koch, 1998), then it should be possible to improve the cortical coherence analysis by removing the influence of the subcortical-cortical loops. To obtain coherence information that exclusively reflects cortical interactions, we performed a multiple regression analysis on the artifactcorrected BOLD time course of each cortical region. We entered the signal from each subcortical region, the motion and design matrix as regressors into the model. Subsequently, the multitaper coherence analysis was performed between the residual BOLD time series in each cortical region. The absolute value of the coherence is a measure of similarity of the BOLD signal largely due to cortico-cortical interactions (Fig. 7.4). The mean absolute coherence value within the left and right hemisphere (r left = 0.0950, r right = 0.0959) are very small and not significantly different on the left and right (p = 0.19, paired t-test). The interhemispheric coherence between corresponding areas is strong (r homolog = 0.25) and significantly higher that the coherence between nonhomolog areas of the left and right hemisphere (r non_homolog = 0.08, p difference < 0.01; t-test). Intrahemispheric regions are slightly more coherent than interhemispheric, noncorresponding areas (r = 0.10 vs. 0.08, p < 0.01; t-test).



Figure 7.3. Subcortical interactions cannot explain callosal coherence. The coherence of spontaneous ultraslow (0.01-0.33 Hz) BOLD fluctuations are stronger between corresponding regions in the left and right hemisphere. Interhemispheric correlations between noncorresponding regions are lower than the average coherence observed between areas of the left or right hemisphere. The coherence values, when corrected for subcortical influences, are lower than the apparent coherence values (compare to Fig. 7.2).



Figure 7.4. Cortical coherence pattern can in part be explained by subcortical BOLD fluctuations. Spontaneous BOLD fluctuations are coherent between cortical and subcortical regions (Fig. 7.1). The subcortical-cortical interaction produce apparent cortical coherence patterns that obscure cortico-cortical interactions. We removed correlates of all subcortical signals with multiple regression and computed the coherence of the residuals of this analysis in the ultraslow frequency band (0.01-0.33 Hz) using multitaper analysis (Mitra and Pesaran, 2008). The absolute value of the mean coherence between cortical regions is shown. Strong coherence values are observed between corresponding regions in the left and right hemisphere (see Fig. 7.3). The estimates of the coherence pattern are very reliable and highly correlated between the left and right hemisphere (r = 0.88, p < 0.001).

The *intrahemispheric* coherence patterns between corresponding regions in the left and right hemispheres are very similar ($\mathbf{r} = 0.88$, $\mathbf{p} < 0.001$). The similarity between the two hemispheres suggests that our method can reliably discriminate between more strongly correlated and less coherent regions. Such strongly correlated regions, for example in the 'default mode' network, have been extensively studied. Less is known about the relationship between intra- and interhemispheric BOLD coherence for such regions. We tested this hypothesis, that cortical hubs – areas with strong coherence within a hemisphere – are also more strongly coupled across the corpus callosum. We compared the mean coherence of an area within one hemisphere to the coherence value of its corresponding area in the other hemisphere. The intrahemispheric coherence strongly correlates the interhemispheric coherence ($\mathbf{r} = 0.53$, $\mathbf{p} < 0.001$), suggesting that strongly coherent regions typically synchronize cortical regions within a hemisphere and with the corresponding regions in the opposite hemisphere.

Weight Information and Functional Connectivity

To test the hypothesis that strong anatomical connectivity implies strong coherence in the awake human brain, we performed global probabilistic fiber tracking. The tractography was seeded in each anatomical region as labelled by the "FreeSurfer" software. The connectivity likelihood to each other anatomical regions was measured with probabilistic tractography (see methods) and computed the BOLD coherence between all regions. To control for the effect of distance, we computed the Euclidian distance between the center of mass of all areas. We then identified sets of brain region pairs that are either not connected, or that are weakly, medium or heavily connected. We further split these sets of pairings into regions with similar distances. The absolute coherence for each pair in a set was converted to a z-score using Fisher's transform and the mean z-score is shown in Fig. 7.5. The anatomical connection strength correlates with coherence. Strongly connected regions, whether they are distant and close-by, show strong coherence in their ultraslow BOLD oscillations. The degree of coherence is monotonically related to the strength of connectivity between two brain regions.



Figure 7.5. Monotonic relation between structural and functional connectivity. BOLD coherence between pairs of cortical regions is primarily related to the strength of fibers connecting them. Strongly connected regions show strong BOLD coupling,

178

independent of the distance between the cortical regions. The observed correlations between connectivity strength and ultraslow BOLD coherence suggests that ultraslow intrinsic oscillations in cortical tissue is synchronized by long range connections and that the degree of coherence is related to the strength of connectivity. The finding cannot be explained by nonfunctional artifacts, such as respiration and blood pressure regulation. Leopold et al., 2003 have shown that ultraslow LFP oscillations correlate with the power envelope of neuronal gamma activity. It is hence possible that the anatomical connectivity synchronizes the overall activity of neurons in distant cortical areas and that the degree of synchronization is a primarily a function of the anatomical connectivity between them.

Connectivity of the Fusiform Gyrus and Face-Selective Activation

Face processing is distributed over multiple face patches in the temporal lobe (Tsao et al., 2003, Moeller et al., 2008). In humans, the face patch in the fusiform gyrus has been extensively studied (Kanwisher, 2000, Kanwisher et al., 1997). The anatomical location of these fusiform face patches varies from person to person and little is known about the relevance of its position. Functional specialization implies a specialization of anatomical connectivity (Passingham et al., 2002) and in order to process facial information, the face patches receive visual input from visual cortex and project to upstream face patches (Moeller et al., 2008). Here, we were interested to use tractography to identify subregions of the fusiform gyrus with distinct connectivity profiles. We then correlated the location of such connectivity-based anatomical

landmarks with the location of face-selective BOLD activity. This analysis is exploratory in nature. For this pilot experiment we analyzed 6 fMRI sessions from 3 subjects.

For each subjects, a minimum of 2 face localizer runs, 2-5 T1-weighted anatomical images and a 72-directional diffusion-weighted data sets were acquired. Scanning was performed in a 3T Siemens scanner with an 8-channel head coil. Standard EPI imaging procedures were used (TR=2.0s, TE=30ms, flip angle of 80°, 3mm in-plane resolution, slice thickness of 3.5mm, 64x64x32 matrix). Anatomical regions of interest were identified for each individual brain using "FreeSurfer". The fusiform gyrus and its adjacent anatomic regions were used as seed regions for probabilistic tractography. Fiber tracking was performed from the occipital pole, the inferior occipital cortex, the lingual gyrus, the parahippocampal gyrus, the the inferior, the middle and the superior temporal gyrus. The connectivity likelihood from each of these seven region to each voxel in the fusiform gyrus was estimated. The connectivity density maps were projected onto the individual cortical surfaces and smoothed with a Gaussian kernel to Subsequently, the region with the strongest improve the signal-to-noise ratio. connection was determined for each fusiform voxel. This maximum likelihood classification for an individual brain is shown in Fig. 7.6. The plot shows, for each voxel, the identity of the most likely connected region. It should be noted that the participant shown here had a bilateral calcification of the amygdala. Up to the present, we have not found any anatomical (or functional) differences between the amygdala lesion participants and the normal subjects in the fusiform gyrus and the data from all subjects were combined for the following analysis.



Figure 7.6. Connectivity-based segmentation of the fusiform gyrus.

The anatomical connectivity strength of each voxel in the fusiform gyrus was determined with a probabilistic tractography algorithm. The identity of the seed region that is most strongly connected to each fusiform voxel is shown in its respective color. Seven anatomically defined regions adjacent to the fusiform gyrus were automatically identified with "FreeSurfer". Fiber tracking was started in each regions and traced with a fast custom fiber tracking implementation to the fusiform gyrus. The density of connectivity samples was projected on the cortical surface. The figure shows the parcellation for a participant (AP) with calcification of the amygdalae.

To assess the correspondence between the anatomical connectivity of the fusiform gyrus and its function specialization, we identified face-selective voxels by centrally presenting faces and houses to the subjects. The EPI images were preprocessed in SPM and corrected for slice timing and head motion. The BOLD signal was correlated with the presentation of faces and houses and with the participants' key presses related to a numerical 1-back memory task. The beta regression coefficients of the main contrast between face and house presentations were extracted from SPM. The fusiform gyrus was parcellated into multiple subregions based on their connectivity to the the anatomical regions adjacent to the fusiform gyrus. For each of the subregions, the mean of the beta values of the face-house contrast was estimated. This number provides a good estimate of the relative face selectivity of a subregion, but its absolute value differs from subject to subject. To correct for individual differences, we normalized the mean scores across the seven subregions to yield a cumulative score of 1. This procedure gives every subject an equal weight in the analysis, regardless of differences in the absolute size of the beta regression weights. The resulting normalized coefficients were compared in an analysis of variance in R (implemented via the more general linear mixed-effect model; see Pinheiro and Bates, 2000). The normalized activation is significantly different between subregions (F=5.11, p < 0.001) and no laterality effect was found. That largest activation was found

in the subregion of the fusiform gyrus that is most strongly connected to the lingual gyrus (Fig. 7.7). The lingual gyrus has been found to be the likely source of the N170 face-selected EEG component (Scott et al., 2006) and shows face-selective changes in regional cerebral blood flow for various face tasks in PET (Nakamura et al., 2000).



Figure 7.7. Face selectivity of connectivity-defined subregions of the fusiform gyrus. Largest than expected face selectivity was found for a subregion of the fusiform gyrus that is most strongly connected to the lingual gyrus. Normalized beta values of the BOLD contrast for face vs. house presentations are shown for each subregion. (*see text for details*)



Summary

Spontaneous cortical BOLD fluctuations are synchronized through cortical tracts. The fiber strengths determines the coherence in the ultraslow frequency band: the stronger a connecting tract, the larger is the observed coherence between two brain areas, within and across hemispheres. A substantial fraction of apparent intrahemispheric cortical coherence is due to subcortical-cortical interactions.

Many functionally defined cortical areas show a high degree of individual variance. Such functional differences could reflect an underlying variability in anatomical connectivity. Tractography could reliable identify landmarks based on connectivity and we showed in a pilot experiment how such an algorithm can be applied to study faceselective patches in the temporal lobe.

BIBLIOGRAPHY

ALLMAN, J. 2000. Evolving Brains, New York: W. H. Freeman.

- ANDERSON, A. K., K. CHRISTOFF, I. STAPPEN, D. PANITZ, D. G. GHAHREMANI,
 G. GLOVER, J. D. E. GABRIELI and N. SOBEL 2003. Dissociated neural representations of intensity and valence in human olfaction. *Nature Neuroscience*, 6, 196–202.
- BASSER, P. J. 2002. Relationships between diffusion tensor and q-space MRI. *Magnetic resonance in medicine*, 47, 392–397.
- BASSER, P. J., J. MATTIELLO and D. LEBIHAN 1994. MR diffusion tensor spectroscopy and imaging. *Biophys J*, 66, 259–267.
- BASSER, P. J., S. PAJEVIC, C. PIERPAOLI, J. DUDA and A. ALDROUBI 2000. In vivo fiber tractography using DT-MRI data. *Magn Reson Med*, 44, 625–632.
- BASSER, P. J. and C. PIERPAOLI 1996. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J Magn Reson B*, 111, 209–219.
- BEHRENS, T. E. J., H. JOHANSEN-BERG, S. JBABDI, M. F. S. RUSHWORTH and M.
 W. WOOLRICH 2007. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *Neuroimage*, 34, 144–155.
- BEHRENS, T. E. J., M. W. WOOLRICH, M. JENKINSON, H. JOHANSEN-BERG, R. G. NUNES, S. CLARE, P.M. MATTHEWS, J. M. BRADY and S. M. SMITH 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med*, 50, 1077–1088.
- BENGTSSON, S. L., Z. N. NAGY, S. SKARE, L. FORSMAN, H. FORSSBERG and F. ULLÈN 2005. Extensive piano practicing has regionally specific effects on white matter development. *Nat Neurosci*, 8, 1148–1150.
- BRIDGE, H., O. THOMAS, S. JBABDI and A. COWEY 2008. Changes in connectivity after visual cortical brain damage underlie altered visual function. *Brain*, 131, 1433–1444.

- BULLMORE, E. and O. SPORNS 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10, 186–198.
- BUZSÁKI, G. 2006. Rhythms of the Brain, Oxford; New York: Oxford University Press.
- CAMPBELL, J. S. W., K. SIDDIQI, V. V. RYMAR, A. F. SADIKOT and G. B. PIKE 2005.
 Flow-based fiber tracking with diffusion tensor and q-ball data: Validation and comparison to principal diffusion direction techniques. *NeuroImage*, 27, 725–36.
- CHERNIAK, C., Z. MOKHTARZADA, R. RODRIGUEZ-ESTEBAN and K. CHANGIZI. 2004. Global optimization of cerebral cortex layout. *Proc Natl Acad Sci USA*, 101, 1081–1086.
- CONTURO, T. E., N. F. LORI, T. S. CULL, E. AKBUDAK, A. Z. SNYDER, J. S. SHIMONY, R. C. MCKINSTRY, H. BURTON and M. E. RAICHLE 1999. Tracking neuronal fiber pathways in the living human brain. *Proc Natl Acad Sci USA*, 96, 10422–10427.
- CRAIG, A. D. 2009. How do you feel—now? The anterior insula and human awareness. *Nature Reviews Neuroscience*, 10, 59–70.
- CRICK, F. and C. KOCH 1998. Constraints on cortical and thalamic projections: The nostrong-loops hypothesis. *Nature*, 391, 245–250.
- DIJKSTRA, E. 1959. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1, 269–271.
- ERLANGER, J. and H. S. GASSER 1937. *Electrical signs of nervous activity,* London, H. Milford, Oxford University Press.
- FELLEMAN, D. J. and D. C. VAN ESSEN 1991. Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex*, 1, 1–47.
- FISHER, R. A. 1915. Frequency distribution of the values of the correlation coefficient in samples of an indefinitely large population. *Biometrika*, 10, 507–521.
- FOX, M. D. and M. E. RAICHLE 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8, 700–711.

- FREY, B. J. and D. DUECK 2007. Clustering by passing messages between data points. *Science*, 315, 972–976.
- FRIEDMAN, D., E. MURRAY, J. O'NEILL and M. MISHKIN 1986. Cortical connections of the somatosensory fields of the lateral sulcus of macaques: Evidence for a corticolimbic pathway for touch. *The Journal of Comparative Neurology*, 252, 323–347.
- GEGENFURTNER, K., D. C. KIPER, and J. B. LEVITT. Functional properties of neurons in macaque area V3. *J Neurophys*, 77 (4). 1906-1923.
- GILDEN, D. L., T. THORNTON and M.W. MALLON 1995. 1/F Noise in Human Cognition. *Science*, 267, 1837–1839.
- GOTTFRIED, J. A., J. O'DOHERTY and R. J. DOLAN 2003. Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science*, 301, 1104–1107.
- HAGMANN, P., M. KURANT, X. GIGANDET, P. THIRAN, V. J. WEDEEN, R. MEULI, and J. P. THIRAN. Mapping human whole-brain structural networks with diffusion MRI. *PloS One*, 2 (7), 597.
- HEGDE, J. and D. J. FELLEMAN 2007. Reappraising the functional implications of the primate visual anatomical hierarchy. *Neuroscientist*, 13, 416–421.
- HILGETAG, C. C., M. A. O'NEILL and M. P. YOUNG 2000. Hierarchical organization of macaque and cat cortical sensory systems explored with a novel network processor. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 355, 71–89.
- ISHIBASHI, T., K.A. DAKIN, B. STEVENS, P. R. LEE, S. V. KOZLOV, C. L. STEWART and R. D. FIELDS 2006. Astrocytes promote myelination in response to electrical impulses. *Neuron*, 49, 823–832.
- JBABDI, S., M. W. WOOLRICH, J. L. R. ANDERSSON and T. E. J. BEHRENS 2007. A Bayesian framework for global tractography. *Neuroimage*, 37, 116–129.
- JENSEN, F. and T. NIELSEN 2007. *Bayesian Networks and Decision Graphs*, Springer Verlag.

- JIAN, B., B. C. VEMURI, E. OZARSLAN, P. R. CARNEY and T. H. MARECI 2007. A novel tensor distribution model for the diffusion-weighted MR signal. *Neuroimage*, 37, 164–176.
- JONES, D. K. and P. J. BASSER 2004. "Squashing peanuts and smashing pumpkins": how noise distorts diffusion-weighted MR data. *Magn Reson Med*, 52, 979–993.
- KANWISHER, N. 2000. Domain specificity in face perception. *Nature Neuroscience*, 3, 759–763.
- KANWISHER, N., J. MCDERMOTT and M. M. CHUN 1997. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci*, 17, 4302–4311.
- KREIMAN, G., I. FRIED and C. KOCH 2002. Single-neuron correlates of subjective vision in the human medial temporal lobe. *Proc Natl Acad Sci USA*, 99, 8378–83.
- KREIMAN, G., C. KOCH and I. FRIED 2000. Category-specific visual responses of single neurons in the human medial temporal lobe. *Nat Neurosci*, 3, 946–953.
- KRIEGESKORTE, N., R. GOEBEL and P. BANDETTINI. 2006. Information-based functional brain mapping. *Proc Natl Acad Sci USA*, 103, 3863–3868.
- KSCHISCHANG, F., B. FREY and H. LOELIGER 2001. Factor graphs and the sumproduct algorithm. *IEEE Transactions on information theory*, 47, 498–519.
- LAZAR, M., D. M. WEINSTEIN, J. S. TSURUDA, K. M. HASAN, K. ARFANAKIS, M. E. MEYERAND, B. BADIE, H. A. ROWLEY, V. HAUGHTON, A. FIELD, and A. L. ALEXANDER 2003. White matter tractography using diffusion tensor deflection. *Hum Brain Mapp*, 18 (4), 306–321.
- LE BIHAN, D., J. F. MANGIN, C. POUPON, C. A. CLARK, S. PAPPATA, N. MOLKO and H. CHABRIAT 2001. Diffusion tensor imaging: Concepts and applications. *J Magn Reson Imaging*, 13, 534–546.
- LEOPOLD, D. A., Y. MURAYAMA and N. K. LOGOTHETIS 2003. Very slow activity fluctuations in monkey visual cortex: Implications for functional brain imaging. *Cereb Cortex*, 13, 422–433.

- LEWIS, J. W. and D. C. VAN ESSEN 2000a. Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *Journal of Comparative Neurology*, 428, 112–137.
- LEWIS, J. W. and D. C. VAN ESSEN 2000b. Mapping of architectonic subdivisions in the macaque monkey, with emphasis on parieto-occipital cortex. *Journal of Comparative Neurology*, 428, 79–111.
- LOGOTHETIS, N. K. 2008. What we can do and what we cannot do with fMRI. *Nature*, 453, 869–878.
- LOGOTHETIS, N. K., C. KAYSER and A. OELTERMANN 2007. In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron*, 55, 809–823.
- LOGOTHETIS, N. K. and B. A. WANDELL 2004. Interpreting the BOLD signal. *Annual Review of Physiology*, 66, 735–769.
- LONGSTAFF, A. 2005. *Neuroscience*, Taylor and Francis.
- MARNER, L., J. R. NYENGAARD, Y. TANG and B. PAKKENBERG 2003. Marked loss of myelinated nerve fibers in the human brain with age. *Journal of Comparative Neurology*, 462, 144–152.
- MAUNSELL, J. H. R. and D. C. VANESSEN 1983. The Connections of the Middle Temporal Visual Area (Mt) and Their Relationship to a Cortical Hierarchy in the Macaque Monkey. *Journal of Neuroscience*, 3, 2563–2586.
- MILLER, J. W., W. KIM, M. D. HOLMES and S. VANHATALO 2007. Ictal localization by source analysis of infraslow activity in DC-coupled scalp EEG recordings. *Neuroimage*, 35, 583–597.
- MILNER, A. D. and M. A. GOODALE 1993. Visual Pathways to Perception and Action. *Progress in Brain Research*, 95, 317–337.
- MITRA, P. P. and H. BOKIL 2008. *Observed Brain Dynamics,* Oxford; New York: Oxford University Press.
- MITRA, P. P. and B. PESARAN 1999. Analysis of dynamic brain imaging data. *Biophysical Journal*, 76, 691–708.

- MOELLER, S., W. FREIWALD and D. TSAO 2008. Patches with links: a unified system for processing faces in the macaque temporal lobe. *Science*, 320, 1355.
- MORCOM, A. M. and P. C. FLETCHER 2007. Does the brain have a baseline? Why we should be resisting a rest. *Neuroimage*, 37, 1073–1082.
- MORI, S., and P. B. BARKER 1999. Diffusion magnetic resonance imaging: its principles and applications. *Anat Rec*, 257 (3), 102–109.
- MORI, S., W. E. KAUFMANN, C. DAVATZIKOS, B. STIELTJES, L. AMODEI, K. FREDERICKSEN, G. D. PEARLSON, E. R. MELHEM, M. SOLAIYAPPAN, G. V. RAYMOND, H. W. MOSER and P. C. M. VAN ZIJL 2002. Imaging cortical association tracts in the human brain using diffusion-tensor-based axonal tracking. *Magn Reson Med*, 47, 215–223.
- O'DOHERTY, J., M. L. KRINGELBACH, E. T. ROLLS, J. HORNAK and C. ANDREWS 2001. Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, 4, 95–102.
- OZTAS, E. 2003. Neuronal tracing. Neuroanatomy, 2, 2.
- PAKKENBERG, B. and H. J. G. GUNDERSEN 1997. Neocortical neuron number in humans: Effect of sex and age. *Journal of Comparative Neurology*, 384, 312– 320.
- PAKKENBERG, B., D. PELVIG, L. MARNER, M. J. BUNDGAARD, H. J. G. GUNDERSEN, J. R. NYENGAARD and L. REGEUR 2003. Aging and the human neocortex. *Experimental Gerontology*, 38, 95–99.
- PARVIZI, J., G. W. VAN HOESEN, J. BUCKWALTER and A. DAMASIO 2006. Neural connections of the posteromedial cortex in the macaque. *Proc Natl Acad Sci USA*, 103, 1563–1568.
- PASSINGHAM, R., K. STEPHAN and R. KOTTER 2002. The anatomical basis of functional localization in the cortex. *Nature Reviews Neuroscience*, 3, 606–616.
- PAXINOS, G. 2000. The Rhesus Monkey Brain in Stereotaxic Coordinates. San Diego: Academic Press.

- PESARAN, B., J. S. PEZARIS, M. SAHANI, P. P. MITRA and R. A. ANDERSEN 2002. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci*, 5, 805–811.
- PEZAWAS, L., A. MEYER-LINDENBERG, E. M. DRABANT, B. A. VERCHINSKI, K. E. MUNOZ, B. S. KOLACHANA, M. F. EGAN, V. S. MATTAY, A. R. HARIRI and D. R. WEINBERGER 2005. 5-HTTLPR polymorphism impacts human cingulateamygdala interactions: A genetic susceptibility mechanism for depression. *Nat Neurosci*, 8, 828–834.
- PINHEIRO, J. and D. BATES 2000. *Mixed-effects models in S and S-PLUS*, Springer Verlag.
- RAICHLE, M. E. and M. A. MINTUN 2006. Brain work and brain imaging. *Annual Review of Neuroscience*, 29, 449–476.
- RINGO, J., R. DOTY, S. DEMETER and P. SIMARD 1994. Time is of the essence: a conjecture that hemispheric specialization arises from interhemispheric conduction delay. *Cerebral Cortex*, 4, 331.
- RIZZOLATTI, G., G. LUPPINO and M. MATELLI 1998. The organization of the cortical motor system: New concepts. *Electroencephalography and Clinical Neurophysiology*, 106, 283–296.
- ROWEIS, S. T. and L. K. SAUL 2000. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290, 2323–2326.
- SCANNELL, J. W., C. BLAKEMORE and M. P. YOUNG 1995. Analysis of connectivity in the cat cerebral-cortex. *Journal of Neuroscience*, 15, 1463–1483.
- SCHMAHMANN, J. D., D. N. PANDYA, R. WANG, G. DAI, H. D'ARCEUIL, A. DE CRESPIGNY and V. WEDEEN 2007. Association fibre pathways of the brain: parallel observations from diffusion spectrum imaging and autoradiography. *Brain*, 130, 630–653.
- SCHOENBAUM, G., A. A. CHIBA and M. GALLAGHER 1998. Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature Neuroscience*, 1, 155–159.

- SERENO, M. I. and R. B. H. TOOTELL 2005. From monkeys to humans: What do we now know about brain homologies? *Curr Opin Neurobiol,* 15, 135–144.
- SHEPHERD, G., J. MIRSKY, M. HEALY, M. SINGER, E. SKOUFOS, M. HINES, P. NADKARNI and P. MILLER 1998. The Human Brain Project: Neuroinformatics tools for integrating, searching and modeling multidisciplinary neuroscience data. *Trends in Neurosciences*, 21, 460–468.
- SHERBONDY 2008. ConTrack: Finding the most likely pathways between brain regions using diffusion tractography. *JOV*, 8, 1–16.
- SLED, J. 1997. A non-parametric method for automatic correction of intensity nonuniformity in MRI data. *Department of Biomedical Engineering, McGill University.*
- SMITH, S., M. JENKINSON, H. JOHANSEN-BERG, D. RUECKERT, T. NICHOLS, C. MACKAY, K. WATKINS, O. CICCARELLI, M. CADER and P. MATTHEWS 2006. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *Neuroimage*, 31, 1487–1505.
- SPORNS, O., G. TONONI and G. EDELMAN 2002. Theoretical neuroanatomy and the connectivity of the cerebral cortex. *Behavioural Brain Research*, 135, 69–74.
- STEPHAN, K. E., L. KAMPER, A. BOZKURT, G. A. BURNS, M. P. YOUNG and R. KOETTER 2001. Advanced database methodology for the collation of connectivity data on the macaque brain (CoCoMac). *Philos Trans R Soc Lond B Biol Sci*, 356, 1159–1186.
- STERIADE, M. 1997. Synchronized activities of coupled oscillators in the cerebral cortex and thalamus at different levels of vigilance. *Cereb Cortex*, 7, 583–604.
- STERIADE, M. 2000. Corticothalamic resonance, states of vigilance and mentation. *Neuroscience*, 101, 243–276.
- STERIADE, M. and R. W. MCCARLEY 1990. *Brainstem Control of Wakefulness and Sleep,* New York: Plenum Press.
- STEVENS, C. F. 2001. An evolutionary scaling law for the primate visual system and its basis in cortical function. *Nature*, 411, 193–195.
- STROGATZ, S. H. 2001. Exploring complex networks. Nature, 410, 268–276.

- TANNER, R. 1981. A recursive approach to low complexity codes. *IEEE Trans Inf Theory*,
- TOGA, A. W., P. M. THOMPSON, S. MORI, K. AMUNTS and K. ZILLES 2006. Towards multimodal atlases of the human brain. *Nature Reviews Neuroscience*, 7, 952– 966.
- TOOTELL, R. B. H., J. D. MENDOLA, N. K. HADJIKHANI, P. J. LEDDEN, A. K. LIU, J. B. REPPAS, M. I. SERENO and A. M. DALE 1997. Functional analysis of V3A and related areas in human visual cortex. *Journal of Neuroscience*, 17, 7060–7078.
- TOURNIER, J. D., F. CALAMANTE, D. G. GADIAN and A. CONNELLY 2004. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *Neuroimage*, 23, 1176–1185.
- TSAO, D. Y., W. A. FREIWALD, T. A. KNUTSEN, J. B. MANDEVILLE and R. B. H. TOOTELL 2003. Faces and objects in macaque cerebral cortex. *Nat Neurosci*, 6, 989–995.
- TUCH, D. S., T. G. REESE, M. R. WIEGELL and V. J. WEDEEN 2003. Diffusion MRI of complex neural architecture. *Neuron*, 40, 885–895.
- TYSZKA, J., C. READHEAD, E. BEARER, R. PAUTLER and R. JACOBS 2006. Statistical diffusion tensor histology reveals regional dysmyelination effects in the shiverer mouse mutant. *Neuroimage*, 29, 1058–1065.
- UNGERLEIDER, L. G., T. W. GALKIN, R. DESIMONE and R. GATTASS 2008. Cortical connections of area V4 in the macaque. *Cereb Cortex*, 18, 477-499.
- UNGERLEIDER, L. G. and R. DESIMONE 1986. Cortical connections of visual area MT in the cacaque. *Journal of Comparative Neurology*, 248, 190–222.
- VAN ESSEN, D. C., H.A. DRURY, J. DICKSON and J. HARWELL 2001. An integrated software suite for surface-based analyses of cerebral cortex. *Journal of the American Medical Informatics Association*, 8, 443–459.
- VAN ESSEN, D. C. 2002. Windows on the brain: the emerging role of atlases and databases in neuroscience. *Current Opinion in Neurobiology*, 12, 574-579.

- VANHATALO, S. K., M. D. HOLMES, P. TALLGREN, J. VOIPIO, K. KAILA and J. W. MILLER 2002. Direct current EEG is needed for accurate recording of slow EEG waveforms during sleep. *Sleep*, 25, A44–A44.
- VINCENT, J. L., G. H. PATEL, M. D. FOX, A. Z. SNYDER, J. T. BAKER, D. C. VAN ESSEN, J. M. ZEMPEL, L. H. SNYDER, M. CORBETTA and M. E. RAICHLE 2007. Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, 447, 83–84.
- VUILLEUMIER, P., J. L. ARMONY, J. DRIVER and R. J. DOLAN 2001. Effects of attention and emotion on face processing in the human brain: An event-related fMRI study. *Neuron*, 30, 829–841.
- WATTS, D. J. and S. H. STROGATZ 1998. Collective dynamics of 'small-world' networks. *Nature*, 393, 440–442.
- YOUNG, M. 1993. The organization of neural systems in the primate cerebral cortex. *Proceedings R Soc Lond B*, 252, 13–18.
- YOUNG, M. P., J. W. SCANNELL, M. A. ONEILL, C. C. HILGETAG, G. BURNS and C. BLAKEMORE 1995. Nonmetric multidimensional-scaling in the analysis of neuroanatomical connection data and the organization of the primate cortical visual system. *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 348, 281–308.
- YUSHKEVICH, P. A., J. PIVEN, H. C. HAZLETT, R. G. SMITH, S. HO, J. C. GEE and G. GERIG 2006. User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *Neuroimage*, 31, 1116–1128.
- ZHANG, K. and T. J. SEJNOWSKI 2000. A universal scaling law between gray matter and white matter of cerebral cortex. *Proc Natl Acad Sci USA*, 97, 5621–5626.