



Learning from tractography: reconstructing and analysing structural connections

Niklas Kasenburg

This thesis has been submitted to the PhD School of The Faculty of Science,
University of Copenhagen.

February 1, 2016

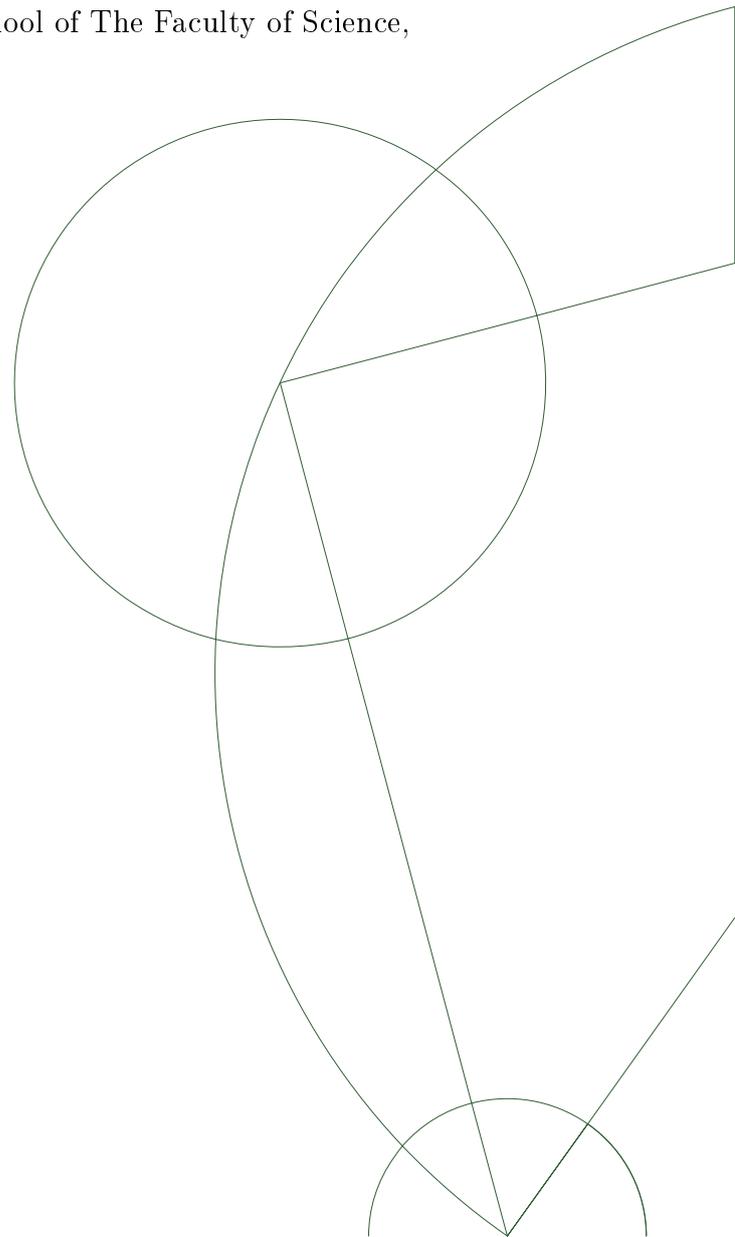
Principal supervisor

Aasa Feragen

Co-supervisors

Matthew Liptrot

Mads Nielsen



Acknowledgements

I want to thank all the people that have helped or supported me during my PhD and all that contributed to a productive working atmosphere. I thank my supervisor Aasa Feragen who always had an open ear and gave valuable feedback even while being on maternity leave. I further thank Matthew Liptrot from whom I learned (almost) all I know about diffusion weighted imaging, who improved the language of most of my publications and because of whom I write all my papers in British English. I also thank Mads Nielsen who always had good advice and Karsten Borgwardt who allowed me to stay a year longer in Germany and sparked my interest in graph analysis.

I further thank all collaborators and co-authors I worked with on the projects and research presented in this thesis. I appreciated your valuable inputs and feedback while writing the papers, without you it would not have been possible.

Furthermore, I thank all of my colleagues both from Tübingen and at the Department of Computer Science at the University of Copenhagen. You made research and working at the office much more fun. I thank my parents, who although not understanding most of what I work on, always support me. Last but not least I thank the committee and you for reading this thesis. I hope you will enjoy it!

Summary

In this thesis I will cover three topics concerned with reconstruction and analysis of structural brain connections, which will be briefly introduced in Chapter 1:

- Tractography: Reconstructing structural connections
- Statistical significance of connections and application to connectivity-based parcellation
- Supervised dimensionality reduction of networks of structural connections

In Chapter 2 I will give an overview of the current state-of-the-art research for reconstructing and analysing structural connections with a focus on the aforementioned topics. Chapter 3 briefly summarises the motivation, approach and results of the four papers included in this thesis. This chapter is organized according to the topics of the papers:

- Shortest-path tractography with spatial priors
- Statistical significance of connections applied to parcellation of the thalamus
- Supervised hub-detection for structural brain networks

In Chapter 4 I will conclude the thesis with an outlook towards further extension and applications of the presented methods. Chapter 5 contains the discussed papers in the following order:

1. Training shortest-path tractography: Automatic learning of spatial priors ([Kasenburg et al., 2016c](#))
2. Finding significantly connected voxels based on histograms of connection strengths ([Kasenburg et al., 2016d](#))
3. Structural parcellation of the thalamus using shortest-path tractography ([Kasenburg et al., 2016a](#))
4. Supervised hub-detection for brain connectivity ([Kasenburg et al., 2016b](#))

I also participated in other projects that are not the focus of my thesis, but which are related to the research presented herein. Two of these projects are concerned with machine learning on networks. In the first we developed a new graph kernel based on shortest paths to measure similarity between attributed graphs ([Feragen et al., 2013](#)). In the second we proposed a new approach to find common sub-networks in a dataset that significantly distinguish graphs of two different classes ([Sugiyama et al., 2015](#)). The last project also involves tractography, but utilises simpler DTI models allowing tractography output as probability distributions over tracts ([Schober et al., 2014](#)).

List of papers and abstracts

The following list contains all papers and abstracts published during my PhD in chronological order:

1. A. Feragen, **Niklas Kasenburg**, J. Petersen, M. de Bruijne and K. Borgwardt, “Scalable kernels for graphs with continuous attributes”, *Advances in Neural Information Processing Systems 26*, 216–224 (2013)
2. **Niklas Kasenburg**, M. Liptrot, K. M. Borgwardt, S. N. Ørting, M. Nielsen and A. Feragen, “Graph-based fibre tractography computing shortest paths between regions of interest”, In: *22nd annual meeting of International Society for Magnetic Resonance in Medicine (ISMRM)*, 4512 (May 2014)
3. M. Schober, **Niklas Kasenburg**, A. Feragen, P. Hennig and S. Hauberg, “Probabilistic shortest path tractography in DTI using Gaussian Process ODE solvers”, *Medical Image Computing and Computer-Assisted Intervention–MICCAI 17 (Pt 3)*, 265–272 (2014)
4. M. Sugiyama, F. Llinares-López, **Niklas Kasenburg** and K. M. Borgwardt, “Significant subgraph mining with multiple testing correction”, *Proceedings of the 2015 SIAM International Conference on Data Mining*, Vol. 5 pp. 37–45 (2015)
5. **Niklas Kasenburg**, M. Liptrot, N. L. Reislev, S. N. Ørting, M. Nielsen, E. Garde and A. Feragen, “Training shortest-path tractography: automatic learning of spatial priors”, *NeuroImage* (2016)
6. **Niklas Kasenburg**, M. G. Liptrot, N. L. Reislev, E. Garde, M. Nielsen and A. Feragen, “Supervised hub-detection for brain connectivity”, to appear in *Proceedings SPIE 9784, Medical Imaging 2016: Image Processing* (chosen as oral)
7. **Niklas Kasenburg**, M. V. Pedersen and S. Darkner, “Finding significantly connected voxels based on histograms of connection strengths”, to appear in *Proceedings SPIE 9784, Medical Imaging 2016: Image Processing*
8. **Niklas Kasenburg**, S. Darkner, U. Hahn, M. Liptrot and A. Feragen, “Structural parcellation of the thalamus using shortest-path tractography”, to appear in *2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI)*

Contents

1	Introduction: What can we learn from structural connections?	3
1.1	Tractography: Reconstructing structural connections	3
1.2	Significance analysis of structural connections	4
1.3	Supervised dimensionality reduction of networks of structural connections	4
2	State-of-the-art in reconstruction and analysis of structural connections	5
2.1	Tractography	5
2.1.1	From DWI to fibre orientation	5
2.1.2	From fibre orientation distribution functions to tracts	6
2.2	Connectivity-based parcellation	8
2.3	Structural brain network analysis	9
2.3.1	From diffusion images to brain networks	9
2.3.2	From brain networks to hubs	10
3	Summary of published papers: Learning from tractography	13
3.1	Shortest-path tractography with spatial priors	13
3.1.1	What can we gain from spatial priors?	13
3.1.2	Creating learned spatial priors	14
3.1.3	Application of learned priors	15
3.2	Significance analysis of connections and application to parcellation . . .	16
3.2.1	How to determine significance?	16
3.2.2	False discovery rate	17
3.2.3	Rank-based test	17
3.2.4	Hard and soft parcellation of the thalamus	18
3.3	Supervised hub-detection for structural brain networks	20
3.3.1	Dimensionality reduction	20
3.3.2	Supervised hub-detection	20
4	Outlook: What else can we learn?	23
4.1	Weighted and iteratively applied learned prior	23
4.2	Connection significance beyond parcellation	23
4.3	Significance analysis in tractography	24
4.4	Structural networks revisited	25
	References	26

5	Papers	35
5.1	Training shortest-path tractography: Automatic learning of spatial priors	36
5.2	Finding significantly connected voxels based on histograms of connection strengths	69
5.3	Structural parcellation of the thalamus using shortest-path tractography	76
5.4	Supervised hub-detection for brain connectivity	81

Abstract

Analysis of structural connections between brain regions enables us to gain insight into the structural architecture of the human brain and into how connections are affected by age or pathology. Tractography is the standard tool for automatic delineation of structural connections or tracts. Post-processing of tractography results using expert prior knowledge is often performed to ensure a robust delineation. In this thesis, I present a shortest-path tractography (SPT) framework that can automatically incorporate any prior knowledge about the location of a tract. Furthermore, I show how such a prior can be learned from previous tractography results.

A confound common to all SPT methods is their sensitivity to finding many false-positive connections, since a path between two locations in the brain is always found. To address this issue I present two approaches to measure the statistical significance of a connection and demonstrate their application in connectivity-based parcellation.

Network models are a common way to represent structural connections of the whole brain. With supervised learning methods, features are extracted from these networks and are associated with a parameter of interest. Dimensionality reduction is often performed as pre-processing, since network analysis typically suffers from high-dimensionality low-sample-size problems. Preceding the supervised analysis with unsupervised dimensionality reduction can, however, smooth the discriminative signals, degrading predictive performance. In this thesis I present a novel supervised dimensionality reduction algorithm that clusters network nodes into hubs, which reflect common connectivity structures in the population, and that retains predictive performance of the lower dimensional features.

Chapter 1

Introduction: What can we learn from structural connections?

Structural connections are defined as physical connections between different parts of the human brain, also called tracts. Analysing these tracts enables us to understand how structural connections of the brain are organized and how tracts are affected by age or pathology. The motivation for the work in this thesis is to search for differences in structural brain connectivity within a population. This requires finding the connections, evaluating whether connections are true or false positives, and finally analysing the network of brain connections as a whole.

1.1 Tractography: Reconstructing structural connections

Reconstruction of structural connections is a required pre-processing step for tract analysis. While it is possible to delineate tracts *ex vivo* after injection of fluorescent tracers (Haber, 1988; Mufson et al., 1990), this can only give insights into the general structure of the brain and is only applicable for post mortem analysis.

Diffusion weighted imaging (DWI) measures the diffusion of water in the brain. Since diffusion of water is greater along than across tracts (Henkelman et al., 1994; Moseley et al., 1991), fibre orientation distribution functions (fODFs) can be computed from the DWI. Tractography delineates tracts *in vivo* based on fODFs computed for each voxel in the image.

Unlike tractography methods that follow the local maximum likelihood direction of the fODF of the current voxel, shortest-path tractography (SPT) is able to find the globally optimal path between two voxels over all possible paths. In SPT, the DWI is modelled as a brain graph, where nodes represent each voxel in the image, edges connect neighbouring nodes and edge weights reflect the likelihood of a tract connecting two voxels derived from the fODFs of the corresponding voxels.

Post-processing based on expert prior knowledge of a tract is often applied to tractography outputs from both local tracking and SPT to ensure a reliable delineation. However, until now it has not been possible to include *any* prior knowledge about the location of a tract into a tractography framework.

The first objective of this thesis is to formulate a new approach to shortest-

path tractography in which prior spatial location of a tract is incorporated into the tractography framework, in order to find the globally optimal path based on both the prior information and the voxel-wise fODFs.

1.2 Significance analysis of structural connections

Shortest-path tractography will always find a path for every given pair of voxels, whether there exists an actual physical connection or not. It is therefore necessary to be able to evaluate the likelihood of a path and to determine the statistical significance of a path to be a true-positive connection.

The second aim of this thesis is to measure whether a voxel is significantly connected to a target region. Moreover, analysis of the significance of the connections of a seed region to multiple target regions (ROIs) is applied in structural parcellation of the thalamus.

1.3 Supervised dimensionality reduction of networks of structural connections

Structural connections can be either analysed directly by comparing the delineated tracts or in the form of an ROI-to-ROI graph. In contrast to the voxel graph used in graph-based SPT, an ROI-to-ROI graph is a brain graph on a higher level, where the nodes are anatomical units of the brain consisting of hundreds or thousands of voxels and edges reflect the structural connectivity between these units. In these graphs, connectivity between nodes is measured based on the tracts found between the respective ROIs and metrics that can be derived from these tracts.

Analysis of these structural brain networks is commonly preceded by dimensionality reduction to account for the apparent high-dimensionality low-sample-size problem. When supervised learning is applied to brain networks to compare them with respect to a factor of interest like age or pathology, preceding unsupervised dimensionality reduction can, however, smooth the discriminative signal.

The last objective of this thesis is to incorporate the learning problem into the dimensionality reduction, thereby finding clusters of nodes, so called *hubs*, that are related to the factor of interest and represent common structures in the dataset.

Chapter 2

State-of-the-art in reconstruction and analysis of structural connections

Here, I will describe three different topics related to the objectives from Chapter 1, in the context of the international state-of-the-art. I will first discuss how to estimate a fibre orientation model from DWI, which is necessary to perform tractography. I will then give an overview of different tractography methods with a focus on shortest-path tractography. Next, I will describe current approaches to connectivity-based parcellation (CBP). Finally, I will give an introduction into how ROI-to-ROI brain graphs are constructed and describe state-of-the-art network analysis with a focus on hub-detection and network dimensionality reduction. I will also briefly outline the new ideas and approaches developed during my project in these fields but refer to Chapter 3 for a more in-depth description.

2.1 Tractography

The human brain consists mainly of nerve cells. The bodies of these cells are mostly located in the cortex of the brain, forming the grey matter. The axons of the nerve cells that connect cells with each other form the white matter. Bundles of these axons, also called fibre bundles or tracts, connect different regions of the cortex. The goal of tractography is to delineate those tracts based on voxel-wise fibre orientations derived from DWIs.

2.1.1 From DWI to fibre orientation

Tractography methods rely on voxel-wise fibre orientation models derived from DWI. In DWI the diffusion of water is measured along different directions. Since diffusion depends on the orientation of the measured tissue ([Henkelman et al., 1994](#)) and is faster along nervous fibres ([Moseley et al., 1991](#)), it is possible to estimate the direction of fibres from DWI. Different models exist to model the fibre direction from the measured diffusion.

The diffusion tensor (DT) is a 3×3 positive definite, symmetric matrix representing displacements in 3D. Assuming that the diffusion of water is distributed as a Gaussian, the DT models the covariance between diffusion along the three main axes ([Basser et al., 1994](#)). The main fibre direction lies along the eigenvector with the largest

eigenvalue, called the principal eigenvector. However, since the DT assumes that there is only one fibre bundle within a voxel, DTI is not able to model crossing, branching or fanning fibres (Pierpaoli and Basser, 1996).

One possible solution to allow for multiple fibre populations is to model the diffusion as a mixture of Gaussian functions (Alexander, 2005). Another approach is to model the diffusion outside the fibre as an isotropic Gaussian and the diffusion along a fibre as a DT with only one non-zero eigenvalue (Behrens et al., 2003b). This “ball and stick” model can be extended to multiple fibre populations by allowing for additional “sticks”. Nevertheless, the number of fibre populations in a voxel needs to be defined or estimated beforehand for all of the above models.

Other methods exist that try to estimate a fibre orientation distribution function (fODF) for each voxel in the image. The fODF is a probability function on the unit sphere. For each direction pointing out of the voxel centre it returns a probability that there is a fibre along that direction.

Non-parametric methods compute the fODF by, for example, using a funk radon transform of the measured signal (Aganj et al., 2010; Descoteaux et al., 2007; Tuch, 2004) or by deconvolution of the measurements (Tournier et al., 2007, 2004). For a complete overview of existing fODF models and reconstruction, the interested reader is referred to the reference books *Diffusion MRI: Theory, methods, and application* by Jones (2010) and *Introduction to Diffusion Tensor Imaging and Higher Order Models* by Tournier and Mori (2014).

Due to the diversity of fODF models and their variable suitability for different DWI acquisition protocols, there is also a great variety of tractography algorithms. In the next section, I will give an overview of current state-of-the-art tractography methods that either focus on a certain model or allow a free choice of the fODF.

2.1.2 From fibre orientation distribution functions to tracts

Tractography algorithms can be divided into local and global methods. The former class proceeds step-wise along the most likely direction or maximum likelihood of the fODF from the current position until reaching a termination criterion. The latter seeks the globally most optimal path between two voxels or points in the image.

Local tracking

Local tracking algorithms can be further divided into deterministic tracking (Basser et al., 2000; Garyfallidis, 2012; Mori et al., 1999) and probabilistic tracking (Behrens et al., 2007; Parker and Alexander, 2003; Parker et al., 2003). While deterministic tracking follows the direction with the largest fODF value for every step, probabilistic tracking chooses a direction sampled from the fODF (Behrens et al., 2007; Jeurissen et al., 2011) or some other kind of (un)certainly of the fibre orientation (Parker and Alexander, 2003; Parker et al., 2003).

The first deterministic methods simply follow the principal eigenvector of the DT in every step (Mori et al., 1999). Similarly, one can interpolate the direction from the underlying diffusion tensor in every step to obtain a tract curve (Basser et al., 2000). Deterministic tractography has also been extended to work with any fODF, e.g. by Garyfallidis (2012) implemented in DiPy (Garyfallidis et al., 2014).

Most probabilistic tracking methods stochastically sample the direction for the next step from the fODF. Methods based on (multi)-tensor models measure the uncertainty for each principal eigenvector, for example, using the fractional anisotropy (FA) of the tensor (Parker et al., 2003) or a noise model (Parker and Alexander, 2003). The same sampling approach can be applied to the fODFs computed from the “ball and stick” model (Behrens et al., 2007) or a bootstrap approach can be used together with the constrained spherical deconvolution model (Jeurissen et al., 2011).

While probabilistic tractography accounts for uncertainty in the measurement and direction for a single voxel (Behrens et al., 2003b; Lazar and Alexander, 2003), it suffers from path-length dependency (Liptrot et al., 2014), i.e. connection probability decreases and uncertainty accumulates with the distance to the start region (Morris et al., 2008). Path-length dependency can be avoided by considering all possible paths to find the globally optimal one, instead of a locally optimal path obtained through local tracking by following the most likely direction from the current position.

Global tractography

Global tractography methods try to find the optimal path between two points or voxels, which makes them less susceptible to local noise and removes the problem of path-length dependency. These methods can be distinguished by the function they optimize and by how the optimization is performed.

Inspired by probabilistic tracking, global probabilistic tractography estimates the maximum likelihood path between two points using a Markov Chain Monte Carlo (MCMC) algorithm to sample from a posterior distribution over paths (Jbabdi et al., 2007; Yendiki et al., 2011). The posterior distribution is based on the fODF of a fixed number of points along the path, and can include prior information about the existence of a path (Jbabdi et al., 2007) or limited types of anatomical information about the path (Yendiki et al., 2011). Due to the MCMC sampling, these global probabilistic methods are not guaranteed to find the globally optimal solution, but can instead result in a local optimum (Teeuw et al., 2015).

Alternatively, the global optimum is defined as a shortest path under a Riemannian metric (Lenglet et al., 2004; O’Donnell et al., 2002), which is defined as being inversely proportional to the diffusion tensor in every voxel. Other metrics have also been used in place of the inverse DT (Fuster et al., 2014; Hao et al., 2011; Hauberg et al., 2015; Schober et al., 2014). These continuous shortest-path methods depend on sampling and a good initialisation similar to the global probabilistic tractography methods and therefore also risk to only find locally optimal solutions.

In contrast, discrete shortest-path methods model the DWI as a graph in which nodes represent voxels and edges link neighbouring voxels in either a $3 \times 3 \times 3$ (Iturria-Medina et al., 2007; Zalesky, 2008) or $5 \times 5 \times 5$ (Sotiropoulos et al., 2010; Vorburget al., 2013) neighbourhood. The weight of an edge reflects the probability that a fibre is connecting the two corresponding voxels. This probability is computed by integrating the fODF over a cone around the corresponding direction. A path between two voxels is then computed by finding the most probable path in the graph (Iturria-Medina et al., 2007; Sotiropoulos et al., 2010) or the shortest path in a graph where edge weights were negatively log-transformed (Vorburget al., 2013; Zalesky, 2008). This guarantees a globally optimal solution, which does not require a good initial

solution, unlike the continuous shortest-path methods.

Graph-based shortest-path tractography (SPT) is, moreover, well suited for including grey and white matter probability to account for partial volume effects (Iturria-Medina et al., 2007), using the cone of uncertainty to reflect noise in the data (Vorburger et al., 2013) or defining a weight for each fibre population in a voxel to model fibre crossings (Sotiropoulos et al., 2010).

However, current graph-based SPT approaches only work with specific fODF models (Sotiropoulos et al., 2010; Vorburger et al., 2013; Zalesky, 2008) or use specialised algorithms to find the global optimum (Iturria-Medina et al., 2007; Sotiropoulos et al., 2010).

In this thesis, I will describe an extremely flexible framework for graph-based SPT which allows for any choice of fODF and makes it algorithmically and mathematically easy to include prior information, such as the spatial location of the tract (see Section 5.1). Furthermore, I will show how such a prior can be learned from previous tractography results.

2.2 Connectivity-based parcellation

It is important to determine whether a tract found by SPT is a true-positive connection, since SPT will always find a path between two voxels. In Section 3.2 I will present an approach to test the significance of connections derived from SPT and how the significance can be used in connectivity-based structural parcellation (CBP). CBP is a form of data-driven segmentation, where the aim is to divide either the whole cortex or a specific region of the brain into smaller “parcels”. While both functional and structural connectivity are often used for parcellation (Eickhoff et al., 2015), I will focus on structural connectivity, which is derived from tractography.

The first paper describing structural parcellation (Behrens et al., 2003a) segmented the thalamus, and was based on the connections of the thalamus to seven target cortical regions known as thalamic connection sites. More specifically, every thalamus voxel was assigned to the target region to which it had most connections. However, this method did not take correlations between voxels into account, since every voxel was labelled independently, and assumed that every voxel is truly connected to *exactly* one target region.

Another approach to thalamus segmentation appeared in the same year and used k -means clustering, with the distance between voxels defined as a linear combination of the Mahalanobis voxel coordinate distance and the Frobenius diffusion tensor distance (Wiegell et al., 2003). While this approach segmented the thalamus purely based on similarity in diffusivity between voxels, it was not related to their connections to the cortex.

Later approaches computed a correlation between seed voxels based on their so-called tractograms (Johansen-Berg et al., 2004), a binarized vector of connections to all other voxels in the brain. The correlation between tractograms both used connectivity information and accounted for correlations between voxels. Similar to Wiegell et al. (2003) k -means clustering was performed for both whole-brain parcellation (Anwander et al., 2007) or to segment specific brain regions like the frontal cortex (Johansen-Berg et al., 2004), parietal cortex (Mars et al., 2011) or the insula (Nanetti et al., 2009).

When the number of target parcels is known beforehand, k -means is an efficient and easy-to-use method to perform the parcellation. However, especially for whole-brain parcellation, the number of clusters is unknown and it is hard to evaluate results for different choices of k (Eickhoff et al., 2015).

Recent approaches for whole-brain parcellation use hierarchical clustering to retrieve groupings of voxels at different levels (Moreno-Dominguez et al., 2014). Hierarchical clustering has the advantage that no predefined number of clusters is required. Nevertheless, it is sensitive to noise (Eickhoff et al., 2015) and depends on the choice of the linkage criterion, i.e. how clusters are joined.

Clustering of voxels leads to a parcellation of a seed region where voxels within a parcel have a similar connectivity. When the goal is to also gain insight about to which target region(s) a seed voxel is most strongly connected to, another approach similar to the one by Behrens et al. (2003a) needs to be applied instead. However, assigning a fixed label to a seed voxel does not reflect whether the voxel is connected to multiple target regions. Moreover, such a hard parcellation neglects that a seed voxel might not be connected to any target region.

In this thesis I present an approach for performing CBP in which parcellation is based on how *significantly* a seed voxel is connected to a target region, thereby yielding a soft segmentation of the seed region that reflects the degree of connection significance to every target region (see Sections 5.2 and 5.3).

2.3 Structural brain network analysis

Structural brain networks are often modelled as graphs consisting of nodes that represent brain regions and edges between pairs of nodes that have an assigned weight (Bullmore and Sporns, 2009; Sporns, 2014). These weights typically reflect either the structural or the functional connectivity between the corresponding regions.

The focus in this thesis is on structural connectivity networks that are created using tractography algorithms. Existing studies analyse how connections are affected by age (Perry et al., 2015; Robinson et al., 2010), pathology (Fornito et al., 2015; Griffa et al., 2013) or gender (Perry et al., 2015), or how functional connections can be predicted from the underlying structural network (Deligianni et al., 2013). All of these methods differ in acquisition protocols, pre-processing of the data, the fODF model, the way region of interests (ROIs) are defined and the tractography method.

In the following, I will not comment on the acquisition of the DWI, pre-processing or choice of diffusion model and tractography method. Instead I will only briefly describe often encountered differences in construction of structural brain networks, including choice of ROIs and the way tractography is performed. Next, I will briefly survey structural networks with a focus on hub-detection and network dimensionality reduction.

2.3.1 From diffusion images to brain networks

The first step of constructing an ROI-to-ROI brain network is to define the ROIs, or nodes, of the graph. Regions of interest are most commonly defined by a structural template atlas (Zalesky et al., 2010). Instead of using the regions provided by the

template atlas directly, they can also be subdivided to obtain a finer parcellation of the brain (Hagmann et al., 2008; Perry et al., 2015).

After definition of the nodes, tractography is performed to obtain the edge weights of the graph. Here, two main strategies are used: whole-brain tractography or ROI-to-ROI tractography.

Whole-brain tractography is performed from a set of seed voxels equally distributed in the white matter (Gong et al., 2009; Hagmann et al., 2008; Zalesky et al., 2010), and tracts are assigned to a pair of regions if they connect both of them. Since tracts start and end in grey matter, found tracts are, moreover, post-processed to select those with endpoints at the grey-matter-white-matter (GM-WM) boundary. In ROI-to-ROI tractography, tracts are computed between the GM-WM boundary of the ROIs (Ghanbari et al., 2014b; Robinson et al., 2010). Whole-brain and ROI-to-ROI tractography both result in a set of tracts connecting each pair of regions. Which of the two approaches is chosen mainly depends on whether the tractography method requires endpoint regions or whether tracts are computed from a set of seed voxels.

Once the tracts connecting each pair of regions are found, the next step is to derive a measure representing each connection. One possible choice is to count the number of tracts reconstructed between two ROIs. However, tract counts are biased by the size of the regions and the physical distance between them. Tract counts are therefore either binarized (Gong et al., 2009) or normalized by the length of the tracts (Hagmann et al., 2008; Perry et al., 2015) or the size of the endpoint regions (Hagmann et al., 2008) to construct weighted graphs. The mean anisotropy along a tract (Deligianni et al., 2013; Robinson et al., 2010) presents an alternative way of defining edge weights.

While binarized graphs only represent the absence or existence of a connection, weighted networks try to estimate the density of connections or white matter integrity along the tract. Nevertheless, the anatomical interpretation of these measures remains unclear (Jones et al., 2013) and they do not account for false-positive tracts due to noise or tracts that do not reflect a physical connection.

In contrast to estimates of the white matter integrity of a tract, edge weights can be defined as the uncertainty of the connection computed from the likelihood assigned to a tract by either probabilistic tractography (Ghanbari et al., 2014b) or graph-based SPT (Iturria-Medina et al., 2007).

Since probabilistic tractography suffers from path-length dependency and is not guaranteed to find the globally optimal path between two region, as described in Section 2.1, we perform SPT to reconstruct the tracts connecting the nodes of the brain network. We further define edge weights as the likelihood of a tract similar to Iturria-Medina et al. (2007), but use the mean likelihood over all paths between two ROIs instead of the maximum (see Section 5.4).

2.3.2 From brain networks to hubs

Structural brain network analysis suffers from the high dimensionality of the networks compared to the typically small sample size, since graphs consist of thousands of edges while datasets generally contain fewer than 100 subjects. Dimensionality reduction is therefore often performed as a pre-processing step. In the case of brain networks, where the features are edges, dimensionality reduction aims to find sub-networks that describe the common structures in the population while still retaining the subject-

specific variation (Clayden et al., 2013).

Commonly, supervised learning is performed on the dimensionality reduced networks to find those sub-structures that are affected by a factor of interest like age (Robinson et al., 2010) or pathology (Ghanbari et al., 2014a,b). However, most dimensionality reduction techniques extract sub-networks without taking the learning problem into account and tend to smooth out the discriminative signal, leading to lower-dimensional brain networks with reduced predictive power (Cheplygina et al., 2014; Rubinov and Bullmore, 2013). The learning goal should therefore instead be part of the dimensionality reduction to retain the predictive power.

In this thesis I present a novel dimensionality reduction algorithm for structural brain networks where the supervised learning problem is incorporated into the algorithm (see Section 5.4). The aim is to find clusters of nodes, so called *hubs*, that form a lower dimensional representation of all networks in the dataset while also being related to the factor of interest.

Chapter 3

Summary of published papers: Learning from tractography

In Chapter 2, I introduced the current state-of-the-art research in the field of tractography, connectivity-based parcellation (CBP) and hub-detection in structural brain networks. Both CBP and hub-detection are ways to learn from data obtained through tractography.

This chapter first describes a novel framework for shortest-path tractography (SPT), where the tractography algorithm can learn from previous tractography results. I will further describe how to determine significance of connections obtained from SPT and how this contributes to CBP. Furthermore, I will summarise my contributions to hub-detection in the context of existing state-of-the-art methods.

3.1 Shortest-path tractography with spatial priors

As introduced in Section 2.1, the goal of tractography is to delineate tracts between regions of the brain. Delineated tracts are compared across different subjects based on properties derived from the corresponding tracts. These properties include, but are not limited to, the shape of the tract or different diffusivity metrics (Benedetti et al., 2011; Davis et al., 2009; Galantucci et al., 2011; Wozniak et al., 2013), such as the fractional anisotropy (FA) (McGrath et al., 2013; Whitford et al., 2015; Xia et al., 2012) or generalized FA (Tang et al., 2010). Furthermore, there exist more complex properties that estimate white matter integrity (Assaf et al., 2013; Golestani et al., 2014).

To properly compare properties derived from tracts across subjects and between populations, it is necessary that the tract delineation is robust and consistent. In this section, I present how incorporating prior knowledge into the tractography algorithm leads to robust delineation *without* the need for further, usually manual, processing of the tractography results.

3.1.1 What can we gain from spatial priors?

In tractography algorithms, FA (Iwata et al., 2011; Whitford et al., 2015) or white matter (WM) tissue probabilities (Iturria-Medina et al., 2007) are often used as a

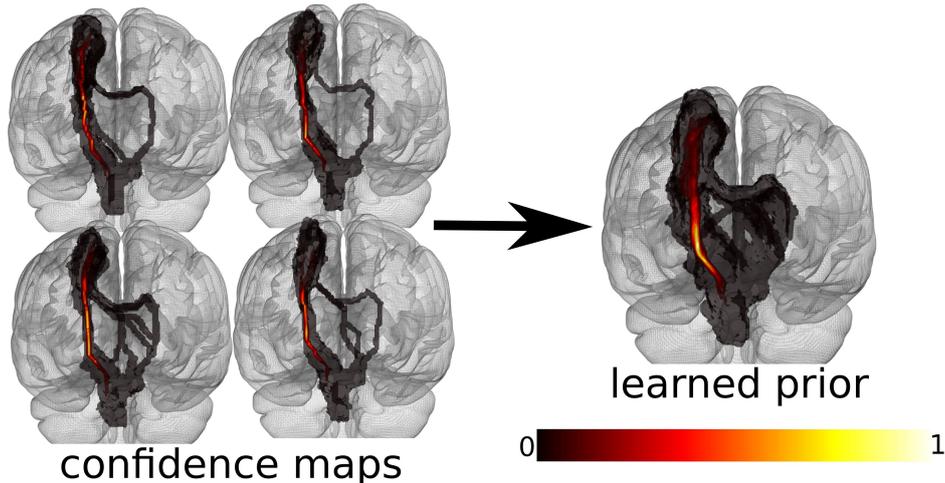


Figure 3.1: The learned prior (right) is created from a set of single-subject confidence maps (left; only a subset is shown). Both the learned prior and the confidence maps are shown in MNI standard space represented as 20 opaque isosurfaces equally distributed over the range of values. The figure was created from results for the left cortico spinal tract presented by [Kasenburg et al. \(2016c\)](#).

guide or stopping criterion for tractography. Furthermore, waypoint or exclusion regions are commonly applied as post-processing of tractography outputs to ensure a reliable delineation of tracts ([Benedetti et al., 2011](#); [Connally et al., 2014](#); [Galantucci et al., 2011](#); [Rojkova et al., 2015](#)). Incorporating such domain knowledge into the tractography framework is a special case of integrating prior information regarding tract location into the tract delineation.

Existing global probabilistic tractography approaches are able to incorporate more restrictive prior knowledge such as the existence or absence of a tract ([Jbabdi et al., 2007](#)) or prior anatomical information about a fixed number of segments along each tract ([Yendiki et al., 2011](#)). These methods are, however, limited to manually pre-defined priors.

In the first paper of this thesis, found in Section 5.1, we derived intuitive, exact and efficient algorithmic solutions to incorporate any spatial prior information about a tract from multiple sources, into a graph-based SPT framework. First, we described how common techniques to prune or post-process tractography results, such as WM probabilities or waypoint masks, can be included as a prior. In a similar fashion, we applied an independent tract atlas as a prior to guide the tractography.

Furthermore, we created *learned* spatial priors that improved the accuracy of the delineated tracts compared to the application of other prior information and results obtained without the use of a prior.

3.1.2 Creating learned spatial priors

We proposed that a learned prior can be used for aiding tractography. These learned priors are computed from single-subject tractography results over a population.

First, tractography results are summarised in so called confidence maps, where the confidence maps assign a likelihood or confidence to each voxel to lie upon the sought tract ([Kasenburg et al., 2016c](#)). The confidence value is computed as the sum of path scores from all found shortest-paths containing the voxel. The path score represents

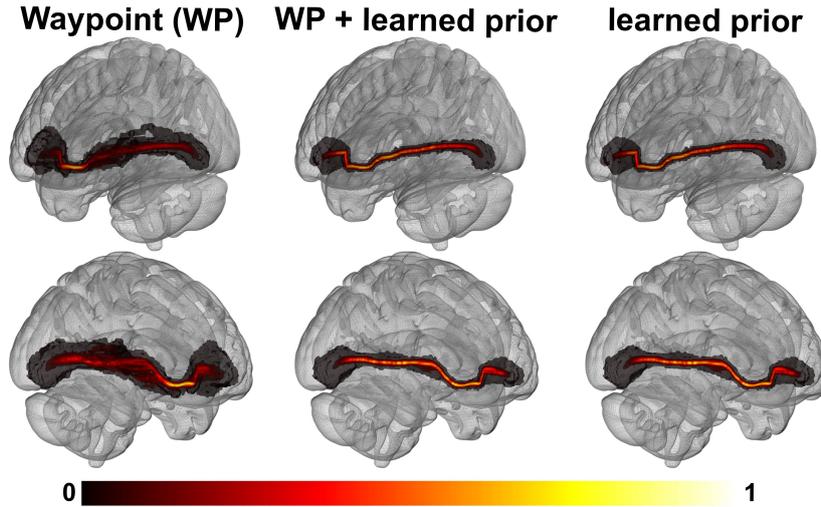


Figure 3.2: Tractography results for the left (top row) and right (bottom row) inferior fronto-occipital fasciculus using a waypoint prior (left column), a waypoint together with a learned prior (middle column) and only using the learned prior. The learned prior was created from the results using the waypoint prior (left column). Confidence maps are shown in MNI standard space represented as 20 opaque isosurfaces equally distributed over the range of values. The figure is an excerpt from results shown in the paper presented in Section 5.1 (Kasenburg et al., 2016c).

the likelihood of a path which is based on both the underlying fODFs of all voxels in the path and any prior information.

Next, each confidence map is normalised to sum up to one and warped into a common space. Finally, the learned prior is computed as the average over all warped confidence maps, see Figure 3.1 for an illustration. Such a learned prior reflects the variation of tract location in a dataset, taking the subject-specific confidence into account.

3.1.3 Application of learned priors

Learned priors can be applied to the same population from which they are created, or as independent priors on a different population. We have shown that a learned prior increases the robustness of tract delineation, focusing on the high confidence region of the prior. We further described that the accuracy of tractography can be improved when a learned prior constructed from high quality data is used in tractography on data acquired on a typical clinical scanner (Kasenburg et al., 2016c).

In particular, including a learned prior created from tractography results where a waypoint or atlas prior was used, renders the use of these priors unnecessary in future tractography (see Figure 3.2).

In summary, creation and application of a learned prior presents a new approach to robustly delineate tracts without the need of manual, often extensive post-processing of tractography results.

3.2 Significance analysis of connections and application to parcellation

Shortest-path tractography always returns a path between two given voxels. Although this path is the most likely of all possible paths between these voxels, it is not guaranteed that there exists a physical connection between the two voxels or the corresponding regions. The path score, which reflects the likelihood of the connection, is one way to account for this fact, as shown in the tractography paper presented in Section 5.1.

Another way is to determine the *significance* of a connection. In the two papers presented in Section 5.2 and 5.3, we described two approaches to measure the significance with which a voxel is connected to a target region. In this section, I will summarise the methods of these papers and show how they can be applied to connectivity-based parcellation of the thalamus.

3.2.1 How to determine significance?

When performing SPT from a given seed voxel to a target region, we get a path and its corresponding score for every voxel in the target region. We define the connection profile of a seed voxel to be the distribution of scores over the set of paths to a target region, and represent this distribution as a histogram over observed scores. Our aim is to determine whether a seed voxel is significantly connected to a target region based on its connection profile.

We assume that not every seed voxel is physically connected to a given target region. Under this assumption, most paths found between the seed and the target region do not reflect a true connection but instead reflect noise in the data. Since these false-positive paths will have a low path score, a connection profile generated from noise should be mainly distributed around low scores.

In statistics, significance is commonly measured with a p -value that reflects the probability that a given observation stems from a noise- or null-distribution. The lower the p -value is, the lower is the probability that the observation is purely generated from noise. If the p -value is lower than a predefined threshold, the null-hypothesis can be rejected and the observation is deemed significant. In terms of structural connectivity, we define an observation to be the connection profile of a seed voxel, i.e a histogram of observed scores, and the null-distribution to be a connection profile that is purely generated by noise, i.e. a histogram of low scores.

The goal is now to compute the p -value for a seed voxel by comparing the histogram of scores of the seed voxel to the null-distribution. This comparison should result in a low p -value if the voxel is significantly connected to the target region, and a high p -value if it is not. The p -value will be lower the more different a seed voxel's histogram is from the null-distribution. Since we define the null-distribution to mainly consist of low scores, the more a histogram is skewed towards high scores the more significant it will be.

Computing the p -value requires the definition of a null-distribution to describe a distribution of low scores and a method for comparing connection profiles or histograms of scores respectively. In the following Sections 3.2.2 and 3.2.3 I will explain two approaches to tackle this task.

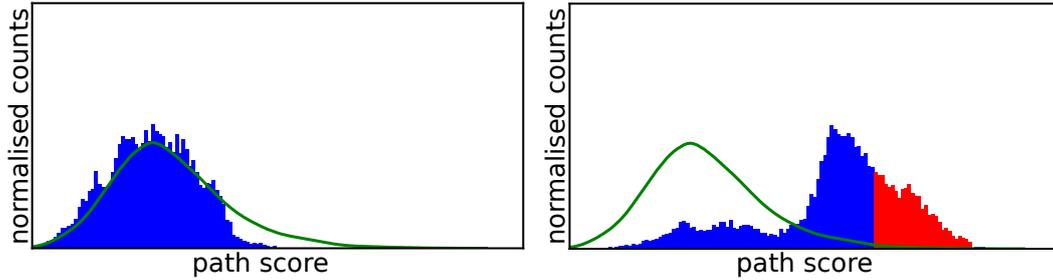


Figure 3.3: Histograms of path scores for a non-significantly connected seed voxel (left) and a significantly connected seed voxel (right) shown together with the average histogram over all seed voxels (green line). Significant bins ($FDR < 0.05$) are coloured in red. The figure is partly an excerpt from the paper presented in Section 5.2 (Kasenburg et al., 2016d).

3.2.2 False discovery rate

The first approach, presented in Section 5.2, directly performs a test of significance on the connection profile histograms and is inspired by Efron (2004). In his work, Efron presents a method to test multiple hypotheses simultaneously. Given a set of hypotheses with an assigned likelihood value for every hypothesis, the goal is to find the significant hypotheses based on their likelihood values. Efron shows that the probability of a hypothesis being generated from the null-distribution is equal to the false discovery rate (FDR) defined as (Efron, 2004):

$$FDR(z) = \frac{f_0(z)}{f(z)}, \quad z \in \mathbb{R}, \quad (3.1)$$

where f_0 is the null-distribution and f is a function fitted to the histogram of observed likelihood values and z is the likelihood of the tested hypothesis. Significant hypotheses are then determined by checking whether the FDR of the z -value assigned to a given hypothesis is below a fixed threshold.

In our case, f is the normalized histogram associated to each seed voxel, and the z -values are the scores of paths from the seed voxel to a target region (Kasenburg et al., 2016d). Since the null-distribution f_0 is unknown, we estimate it as the average histogram over all seed voxels. To determine whether a voxel is significantly connected, we check whether, for any of the bins of the corresponding histogram, the FDR is below a given threshold.

The FDR approach allows us to determine which bins or score ranges in a seed voxel’s histogram are significantly different from the null-distribution, and whether a voxel is significantly connected or not (see Figure 3.3). However, the significance is only computed bin-wise for every voxel. We have therefore proposed a second test of significance, which is able to measure significance of a single voxel with a p -value, as described in Section 3.2.3 and the paper in Section 5.3.

3.2.3 Rank-based test

The second approach found in Section 5.3 uses cumulative histograms similar to the cumulative distribution functions commonly used in statistics. When defining the noise-distribution to be centred around low scores, significantly connected voxels will have more high-scoring paths to the target region, and we expect their histograms to

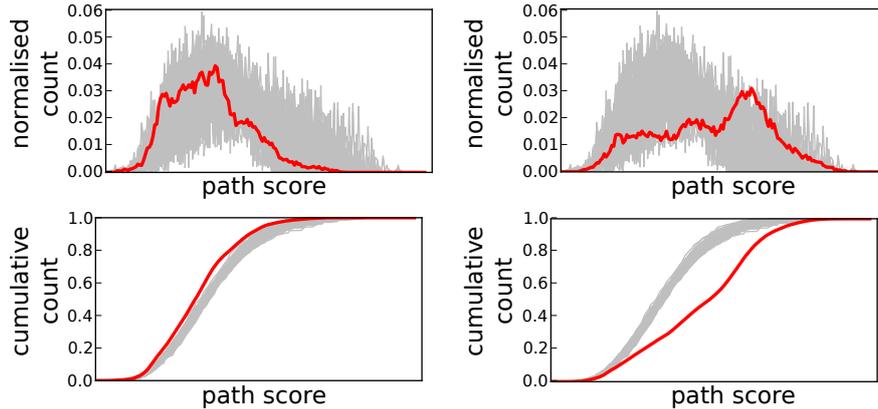


Figure 3.4: Normalised (top) and cumulative histograms (bottom) shown in red for a seed voxel with an assigned p -value close to 1 (left) and a seed voxel with an assigned p -value below 0.0005 (right) together with 100 samples from the null-distribution shown in grey. For a description of how the histograms were created, how the p -values were computed and how the sample was constructed, see Section 5.3 and the corresponding paper (Kasenburg et al., 2016a).

be skewed to the right. The corresponding cumulative histograms then lie below the cumulative histogram of the noise-distribution, because the maximum is reached later, i.e. at a higher score (see Figure 3.4).

Given the cumulative histogram C of a seed voxel, its p -value is computed using a one-sided rank test inspired by Myllymäki et al. (2013). The p -value reflects the probability that a cumulative histogram comes from the null-distribution.

The null-distribution is represented by samples. Since the null-distribution is still unknown, the samples are drawn directly from the data, assuming that most connections are false positive. Each sample histogram is created by drawing a value for each bin from the values of all seed voxel histograms at the corresponding bin. The sample histograms are then transformed into the corresponding cumulative histograms.

The p -value is computed as the fraction of samples that have a rank lower or equal to the rank of the cumulative histogram associated with the seed voxel. The rank is computed as the average rank over all bins over the set of samples including C (Kasenburg et al., 2016a).

3.2.4 Hard and soft parcellation of the thalamus

We applied both significance tests to Behrens style thalamus parcellation (Behrens et al., 2003a) on subjects of the HCP dataset (Essen et al., 2013), where the cortex was divided into four target regions. We obtained a hard parcellation using the FDR method (Kasenburg et al., 2016d), while the p -values computed with the rank-based test were used to generate a soft parcellation (Kasenburg et al., 2016a).

Hard parcellation

The hard parcellation finds significantly connected voxels that are mainly located on the surface of the anterior part of the thalamus, when applying a threshold of 0.05 for the FDR. As shown in Figure 3.5, this effect is increased when lower thresholds (0.01 and 0.005) are used to filter out more voxels.

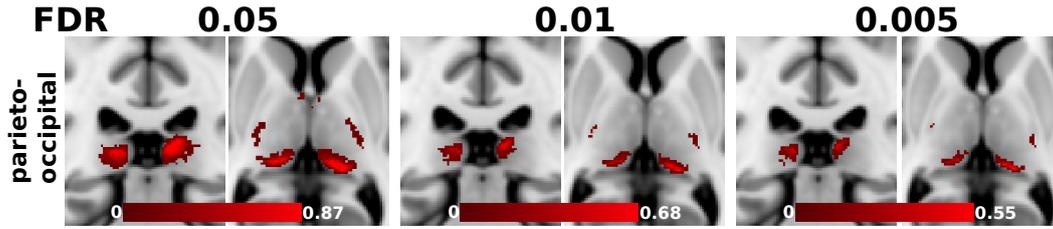


Figure 3.5: Probability overlap maps of hard parcellations of the thalamus based on connections to the parieto-occipital zone in standard MNI space (row: $x = 90$, $y = 96$, $z = 77$), zoomed in on the thalamus for different thresholds (from left to right: 0.05, 0.01 and 0.005). The probability overlap maps were created from single-subject parcellations and show the fraction of subjects in which the parcellations overlap. More details and the whole figure can be found in the paper presented in Section 5.2 (Kasenburg et al., 2016d).

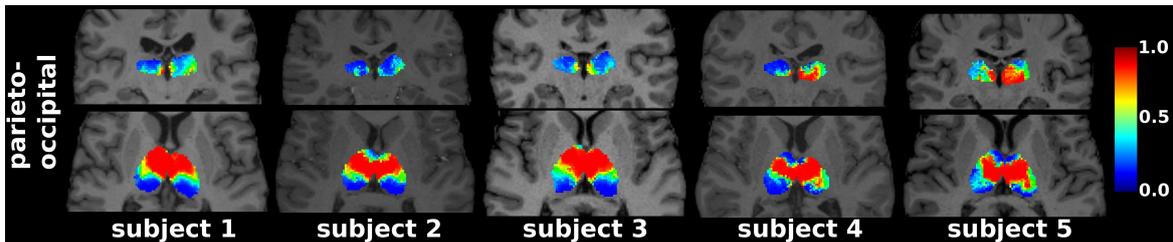


Figure 3.6: Soft parcellations of the thalamus based on connections to the parieto-occipital zone for five subjects shown in the respective subjects' space. Low p -values (blue) reflect regions that are more likely connected to the respective target region than high p -values (red). The figure is an excerpt from the paper presented in Section 5.3 (Kasenburg et al., 2016a).

Furthermore, there is a strong overlap of significant regions between different target regions, and results are consistent over the whole dataset, although there is subject-specific variation as expected.

Soft parcellation

The soft parcellation leads to similar observations: voxels with low p -values are mainly located on the surface of the anterior part of the thalamus and low p -value regions overlap between different target regions. However, this parcellation shows a gradient in the significance whose direction differs for the target regions, but usually decreases towards the centre of the thalamus (see Figure 3.6).

The soft parcellation can, moreover, determine significantly connected voxels for target regions that would not be detected by a hard segmentation due to an additional yet stronger connection to another target region.

Problems addressed by significance-based parcellation

Commonly used hard parcellation methods assign a discrete label to all voxels in the seed region based on how many tracts are found between the corresponding voxel and each target region. The following problems apparent in these parcellation methods are addressed by the parcellation using significance tests of SPT results:

1. The labels assigned to seed voxels that are not connected to any of the target regions are purely based on noise and are unstable.

2. It is not possible to reflect whether and to what degree a seed voxel is connected to multiple target regions, since a voxel is assigned to the region to which it has most connections (Gorbach et al., 2011).
3. Even if a seed voxel is connected to a specific target region, there might not be enough signal in the data to find this connection.
4. False-positive connections are known to increase with the noise in the data (Fillard et al., 2011; Neher et al., 2014).

In summary, parcellation based on significance analysis of connections found by SPT presents an alternative approach to CBP that accounts for false-positive connections and noise, and that also reflects significant connections to multiple target regions.

3.3 Supervised hub-detection for structural brain networks

Structural connectivity networks, as described in Section 2.3, are typically analysed using supervised learning methods like regression or classification. The goal of this analysis is to gain insight into which network features are related to a factor of interest, such as age or pathology. Since the number of features is often much higher than the number of subjects or samples, it is hard to robustly extract these features. Dimensionality reduction is therefore commonly performed beforehand to reduce the number of features.

3.3.1 Dimensionality reduction

One form of dimensionality reduction for brain networks, called *hub-detection*, aims to find clusters of nodes common to the population in order to represent the original network as a hub-network with a lower number of nodes (see Section 2.3.2). Since hub-detection is unrelated to the learning problem, it has the tendency to smooth out the discriminative signal. This results in lower-dimensional brain networks with reduced predictive power (Cheplygina et al., 2014; Rubinov and Bullmore, 2013).

3.3.2 Supervised hub-detection

Integrating the learning problem into the hub-detection would therefore be beneficial to robustly extract lower dimensional features that are related to the learning problem. To the best of my knowledge, this problem has not been addressed in brain network analysis before. However, in the field of genetic research, a similar approach has recently appeared, where the learning problem is integrated into a graph-based feature selection (Allahyar and de Ridder, 2015).

We proposed to integrate the learning problem into dimensionality reduction by incorporating a supervised learning term into the hub-detection from Ghanbari et al. (2014a), originally developed for functional brain networks. The supervised hub-detection (SHD) algorithm described in Section 5.4 was applied to structural networks

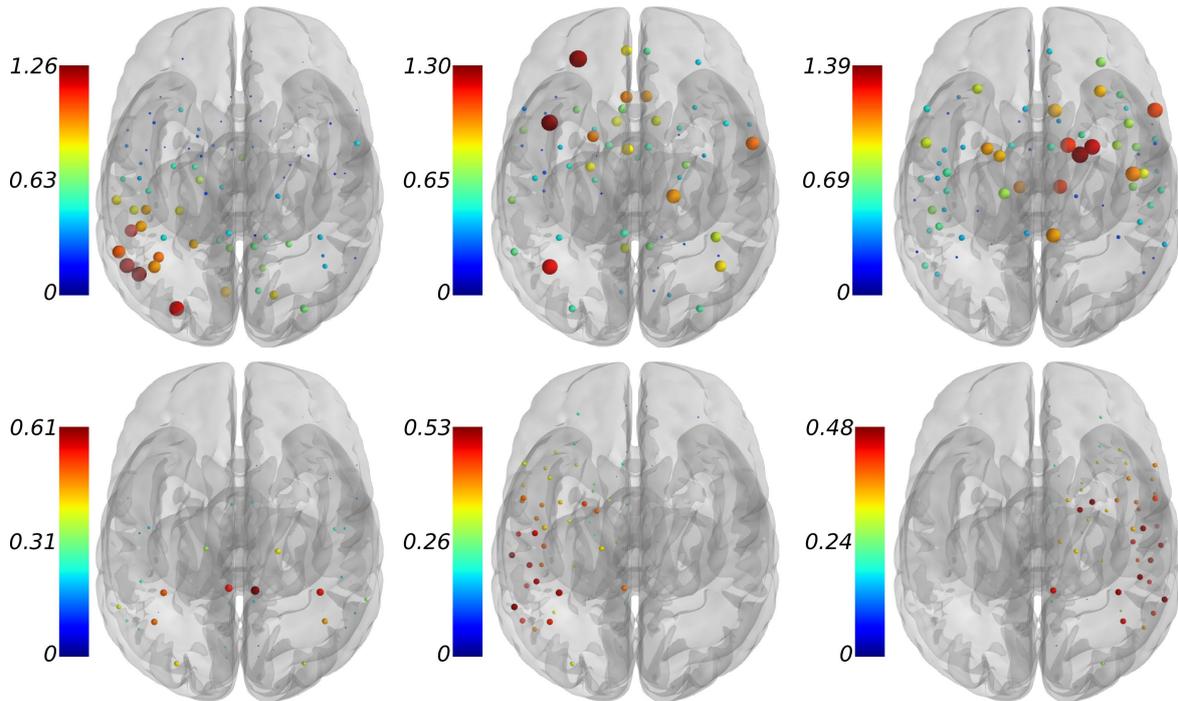


Figure 3.7: The top row shows hubs strongly contributing to the prediction of age (high regression weight), while the hubs in the lower row only have a weak contribution (low regression weight). Hubs are visualised by the hub-membership scores for each node, represented by a colour scale (within hub) and by size (across all hubs). The figure was created from hubs found by the SHD as presented in the hub-detection paper (Kasenburg et al., 2016b) that can be found in Section 5.4.

from an ageing population obtained with shortest-path tractography (Kasenburg et al., 2016b).

We showed that predictive performance is retained on the data after applying SHD, as opposed to after applying the original hub-detection algorithm, where predictive performance decreased. In addition, the hubs found by the SHD still represent the population well, and to a similar degree as the hubs found by the original unsupervised hub-detection.

We found that although no spatial coherency is enforced, hubs are often spatially coherent and are either symmetric or have symmetric counterparts. Furthermore, hubs found by the SHD that contribute more strongly to the prediction of age are more globally distributed throughout the brain than hubs with a weaker contribution (see Figure 3.7). This is in line with results on the effect of ageing found by Lebel et al. (2012).

The presented approach is not limited to structural brain networks or regression, but could also be applied to classification problems and to other types of biological networks like gene networks or protein-protein interaction networks.

Chapter 4

Outlook: What else can we learn?

In Chapter 3, I summarised how learned spatial priors can be constructed and used in shortest-path tractography (see Section 3.1); how significance of a structural connection can be evaluated and used for connectivity-based parcellation (see Section 3.2); and how a supervised learning problem can be integrated into hub-detection for structural brain networks (see Section 3.3). In this chapter I will give an outlook towards possible extensions and future applications of the presented methods.

4.1 Weighted and iteratively applied learned prior

In the presented shortest-path tractography method, a shortest path is chosen according to its length, where the length is the negative log-transformed likelihood of the path (Kasenburg et al., 2016c). The likelihood is a product of the node priors and the edge weights encountered along the path. Since the edge weights are derived from the diffusion data, the prior and the signal from the data are weighted equally. However, it would be interesting to analyse how different weighting schemes change the results.

A different weighting of the prior can algorithmically be introduced by raising the product of the node priors along a path by a factor k in the path likelihood computation. Particularly interesting is the question of how the subject-specific variance is affected for different values of $k > 1$, when the prior is weighted more than the data.

The same question could be asked when using an iteratively applied learned prior. We used a learned prior created from tractography results that included other priors, but no other learned prior (Kasenburg et al., 2016c). Creating a learned prior from those tractography results and applying it again to the data could be considered the next iteration step. Evaluating how multiple iterations affect the subject-specific variation, and to what extent the diffusion data still matters, is another problem for future investigation.

4.2 Connection significance beyond parcellation

The two approaches for testing significance of connections found by graph-based shortest-path tractography both aimed to quantify significance. We have shown how the associated p -values or FDRs can be used to perform soft (Kasenburg et al., 2016a) or hard (Kasenburg et al., 2016d) connectivity-based parcellation (CBP) of the thala-

mus. While *ex vivo* techniques looking at histological slices are better suited to gain insight into the structure of the thalamus in general (Morel et al., 1997), CBP makes it possible to compare the obtained segmentations *in vivo* across subjects.

Given a hypothesis that the connectivity of a certain region of the brain is affected by a disease or pathology, both the soft and hard segmentations could be compared between a control group and patients. While hard segmentations only provide an insight into whether different voxels are significantly connected between the two groups, soft segmentations are also able to capture the degree of difference in significance. Such a population study is therefore a clear next step.

Furthermore, it could be interesting to look at the high scoring paths of significant voxels or those paths that contribute to the significant bins of a voxel. This would, in addition to insights into the projection of the tracts, also give insights into how the actual connections are affected.

4.3 Significance analysis in tractography

While we have shown how to measure the significance of a voxel's connection to a target region, no significance analysis has been performed on single paths. One problem in this regard is the creation of a proper null-distribution. On one hand, we do not expect many false-positive paths when performing tractography between two regions that are known to be connected. On the other hand, creating a null-hypothesis from paths between a different pair of regions that are known not to be connected would be biased towards this specific "connection".

If we are, however, interested in whether a single path is significantly different from the majority of computed paths, it would be possible to create a null-hypothesis without the assumption that most paths are false-positive. Such a null-hypothesis would consist of a distribution of scores over all found paths and a significance test would then test whether the score of a single path is a significant outlier.

Assuming we are able to construct a proper null-distribution and thereby are able to assign a p -value to every path, then another problem arises if the significance of a path should be determined with a given threshold. Since the number of paths can be in the order of thousands or millions depending on the size of the seed and target regions, a proper statistical analysis would require adjustment for multiple-hypothesis testing. Such a correction would then mean that no path would be deemed significant, if the correction is too conservative.

Since paths of neighbouring seed voxels share large parts of the paths connecting them to the target region, path scores and the corresponding hypotheses are not independent. Taking the correlation between hypotheses into account, when performing multiple-hypothesis testing, would therefore be a way to adjust for the dependence between paths and to perform a less conservative correction.

Despite the challenges of defining a proper null-hypothesis and dealing with the multiple-hypothesis testing, I still think it is interesting and, more importantly, also possible to determine the significance of a single path which can then be used for further analysis of tractography results.

4.4 Structural networks revisited

Structural networks are commonly generated using either deterministic or probabilistic tractography (see Section 2.3). In this thesis, we showed several advantages of shortest-path tractography (SPT):

1. SPT finds the *globally* optimal path between two voxels.
2. Each path is endowed with a score that reflects the likelihood of the path. This score can be used in the analysis of the tract.
3. Prior information can be integrated into tractography with an efficient and algorithmically simple formulation as SPT on a modified graph (Kasenburg et al., 2016c).

This leads to the question as to why only a few studies exist where structural networks are created by graph-based SPT (Iturria-Medina et al., 2007, 2008). One reason could be the problem of how to define the edge weight of an edge connecting two regions of interest (ROIs). Another reason could be that SPT always returns a path between every pair of voxels in the two endpoint regions, whether there is a connection or not. In the following, I will address these two problems.

In the hub-detection paper (Kasenburg et al., 2016b), the edge weights between two regions were defined as the average score over all paths found between all pairs of voxels from the pair of ROIs. Iturria-Medina et al. (2007, 2008) used instead the maximum over the scores of all paths and variations thereof.

While the maximal and average score are reasonable representations of the set of paths connecting two ROIs, it is also possible to create a distribution of path scores as used in the significance analysis (see Section 3.2). Other parameters than the mean, like the standard deviation or different quantiles of the distribution, could give a more complete description to define the connection between ROIs.

Deriving the edge weight for a pair of ROIs from the distribution of scores of paths connecting them, moreover, can be tailored to account for false-positive paths. When the maximum, a specific quantile or more descriptive parameters of the score distribution are chosen, low-scoring, false-positive paths will have a smaller contribution to the edge weight.

In ROI-to-ROI graphs generated from tracts found by deterministic or probabilistic tracking, edge weights are defined by “fibre” density or metrics of white matter integrity, like mean anisotropy, derived from the found tracts. Tracts computed with tractography reflect, however, not always an actual physical fibre and in the best case are only parallel to a fibre bundle (Jones et al., 2013). The interpretation of “fibre” density and metrics of white matter integrity as edge weights is therefore unclear.

In structural brain networks generated from SPT, edges reflect instead the likelihood of the tracts, which is purely based on the voxel-wise fODF models derived from the DWI and possible prior information. Thereby, no assumption about whether a tract reflects a true physical fibre affects the definition of the edge weight, and interpretation of the likelihood is directly related to the data.

In summary, generating ROI-to-ROI graphs from SPT can be achieved, and it presents an interesting possibility to perform structural brain network analysis.

References

- Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., Harel, N., Aug 2010. Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. *Magn Reson Med* 64 (2), 554–566.
- Alexander, D. C., Dec 2005. Multiple-fiber reconstruction algorithms for diffusion MRI. *Ann NY Acad Sci* 1064, 113–133.
- Allahyar, A., de Ridder, J., 2015. FERAL: network-based classifier with application to breast cancer outcome prediction. *Bioinformatics* 31 (12), i311–i319.
- Anwander, A., Tittgemeyer, M., von Cramon, D. Y., Friederici, A. D., Knösche, T. R., Apr 2007. Connectivity-based parcellation of Broca’s area. *Cereb Cortex* 17 (4), 816–825.
- Assaf, Y., Alexander, D. C., Jones, D. K., Bizzi, A., Behrens, T. E. J., Clark, C. A., Cohen, Y., Dyrby, T. B., Huppi, P. S., Knoesche, T. R., Lebihan, D., Parker, G. J. M., Poupon, C., consortium, C. O. N. N. E. C. T., Anaby, D., Anwander, A., Bar, L., Barazany, D., Blumenfeld-Katzir, T., De-Santis, S., Duclap, D., Figini, M., Fisch, E., Guevara, P., Hubbard, P., Hofstetter, S., Jbabdi, S., Kunz, N., Lazeyras, F., Lebois, A., Liptrot, M. G., Lundell, H., Mangin, J.-F., Dominguez, D. M., Morozov, D., Schreiber, J., Seunarine, K., Nava, S., Poupon, C., Riffert, T., Sasson, E., Schmitt, B., Shemesh, N., Sotiropoulos, S. N., Tavor, I., Zhang, H. G., Zhou, F.-L., Oct 2013. The connect project: Combining macro- and micro-structure. *NeuroImage* 80, 273–282.
- Basser, P. J., Mattiello, J., LeBihan, D., Mar 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. *J Magn Reson B* 103 (3), 247–254.
- Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., Oct 2000. In vivo fiber tractography using DT-MRI data. *Magn Reson Med* 44 (4), 625–632.
- Behrens, T. E. J., Berg, H. J., Jbabdi, S., Rushworth, M. F. S., Woolrich, M. W., Jan 2007. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage* 34 (1), 144–155.
- Behrens, T. E. J., Johansen-Berg, H., Woolrich, M. W., Smith, S. M., Wheeler-Kingshott, C. A. M., Boulby, P. A., Barker, G. J., Sillery, E. L., Sheehan, K., Ciccarelli, O., Thompson, A. J., Brady, J. M., Matthews, P. M., Jul 2003a. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nat Neurosci* 6 (7), 750–757.

- Behrens, T. E. J., Woolrich, M. W., Jenkinson, M., Johansen-Berg, H., Nunes, R. G., Clare, S., Matthews, P. M., Brady, J. M., Smith, S. M., Nov 2003b. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med* 50 (5), 1077–1088.
- Benedetti, F., Absinta, M., Rocca, M. A., Radaelli, D., Poletti, S., Bernasconi, A., Dal-laspezia, S., Pagani, E., Falini, A., Copetti, M., Colombo, C., Comi, G., Smeraldi, E., Filippi, M., Jun 2011. Tract-specific white matter structural disruption in patients with bipolar disorder. *Bipolar Disord* 13 (4), 414–424.
- Bullmore, E., Sporns, O., Mar 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 10 (3), 186–198.
- Cheplygina, V., Tax, D. M. J., Loog, M., Feragen, A., 2014. Network-guided group feature selection for classification of autism spectrum disorder. In: *Machine Learning in Medical Imaging*. Vol. 8679 of *Lecture Notes in Computer Science*. pp. 190–197.
- Clayden, J. D., Dayan, M., Clark, C. A., 2013. Principal networks. *PLoS One* 8 (4), e60997.
- Connally, E. L., Ward, D., Howell, P., Watkins, K. E., Apr 2014. Disrupted white matter in language and motor tracts in developmental stuttering. *Brain Lang* 131, 25–35.
- Davis, S. W., Dennis, N. A., Buchler, N. G., White, L. E., Madden, D. J., Cabeza, R., Jun 2009. Assessing the effects of age on long white matter tracts using diffusion tensor tractography. *NeuroImage* 46 (2), 530–541.
- Deligianni, F., Varoquaux, G., Thirion, B., Sharp, D., Ledig, C., Leech, R., Rueckert, D., Dec 2013. A framework for inter-subject prediction of functional connectivity from structural networks. *IEEE Transactions on Medical Imaging* 32 (12), 2200–2214.
- Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R., Sep 2007. Regularized, fast, and robust analytical q-ball imaging. *Magn Reson Med* 58 (3), 497–510.
- Efron, B., 2004. Large-scale simultaneous hypothesis testing. *J Amer Statist Assoc* 99 (465), 96–104.
- Eickhoff, S. B., Thirion, B., Varoquaux, G., Bzdok, D., Sep 2015. Connectivity-based parcellation: Critique and implications. *Hum Brain Mapp*.
- Essen, D. C. V., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium, Oct 2013. The WU-Minn Human Connectome Project: An overview. *NeuroImage* 80, 62–79.
- Feragen, A., Kasenburg, N., Petersen, J., de Bruijne, M., Borgwardt, K., 2013. Scalable kernels for graphs with continuous attributes. In: *Advances in Neural Information Processing Systems* 26. pp. 216–224.

- Fillard, P., Descoteaux, M., Goh, A., Gouttard, S., Jeurissen, B., Malcolm, J., Ramirez-Manzanares, A., Reisert, M., Sakaie, K., Tensaouti, F., Yo, T., Mangin, J.-F., Poupon, C., May 2011. Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. *NeuroImage* 56 (1), 220–234.
- Fornito, A., Zalesky, A., Breakspear, M., 2015. The connectomics of brain disorders. *Nat Rev Neurosci* 16 (3), 159–172.
- Fuster, A., Tristan-Vega, A., Haije, T., Westin, C.-F., Florack, L., 2014. A novel Riemannian metric for geodesic tractography in DTI. In: *Computational Diffusion MRI and Brain Connectivity. Mathematics and Visualization.* pp. 97–104.
- Galantucci, S., Tartaglia, M. C., Wilson, S. M., Henry, M. L., Filippi, M., Agosta, F., Dronkers, N. F., Henry, R. G., Ogar, J. M., Miller, B. L., Gorno-Tempini, M. L., Oct 2011. White matter damage in primary progressive aphasia: a diffusion tensor tractography study. *Brain* 134 (Pt 10), 3011–3029.
- Garyfallidis, E., 2012. Towards an accurate brain tractography. Ph.D. thesis, University of Cambridge.
- Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., van der Walt, S., Descoteaux, M., Nimmo-Smith, I., Dipy Contributors, 2014. Dipy, a library for the analysis of diffusion MRI data. *Front Neuroinform* 8, 8.
- Ghanbari, Y., Bloy, L., Shankar, V., Edgar, J., Roberts, T., Schultz, R., Verma, R., 2014a. Functionally driven brain networks using multi-layer graph clustering. *MICCAI* 17 (Pt 3), 113–120.
- Ghanbari, Y., Smith, A. R., Schultz, R. T., Verma, R., Dec 2014b. Identifying group discriminative and age regressive sub-networks from DTI-based connectivity via a unified framework of non-negative matrix factorization and graph embedding. *Med Image Anal* 18 (8), 1337–1348.
- Golestani, A. M., Miles, L., Babb, J., Castellanos, F. X., Malaspina, D., Lazar, M., Nov 2014. Constrained by our connections: white matter’s key role in interindividual variability in visual working memory capacity. *J Neurosci* 34 (45), 14913–14918.
- Gong, G., He, Y., Concha, L., Lebel, C., Gross, D. W., Evans, A. C., Beaulieu, C., Mar 2009. Mapping anatomical connectivity patterns of human cerebral cortex using in vivo diffusion tensor imaging tractography. *Cereb Cortex* 19 (3), 524–536.
- Gorbach, N. S., Schütte, C., Melzer, C., Goldau, M., Sujazow, O., Jitsev, J., Douglas, T., Tittgemeyer, M., 2011. Hierarchical information-based clustering for connectivity-based cortex parcellation. *Front Neuroinform* 5, 18.
- Griffa, A., Baumann, P. S., Thiran, J.-P., Hagmann, P., Oct 2013. Structural connectomics in brain diseases. *NeuroImage* 80, 515–526.
- Haber, S., Feb 1988. Tracing intrinsic fiber connections in postmortem human brain with WGA-HRP. *J Neurosci Methods* 23 (1), 15–22.

- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., Sporns, O., Jul 2008. Mapping the structural core of human cerebral cortex. *PLoS Biol* 6 (7), e159.
- Hao, X., Whitaker, R. T., Fletcher, P. T., 2011. Adaptive Riemannian metrics for improved geodesic tracking of white matter. *Inf Process Med Imaging* 22, 13–24.
- Hauberg, S., Schober, M., Liptrot, M., Hennig, P., Feragen, A., 2015. A random Riemannian metric for probabilistic shortest-path tractography. *MICCAI 18 (Pt 1)*, 597–604.
- Henkelman, R. M., Stanisz, G. J., Kim, J. K., Bronskill, M. J., Nov 1994. Anisotropy of NMR properties of tissues. *Magn Reson Med* 32 (5), 592–601.
- Iturria-Medina, Y., Canales-Rodríguez, E. J., Melie-García, L., Valdés-Hernández, P. A., Martínez-Montes, E., Alemán-Gómez, Y., Sánchez-Bornot, J. M., Jul 2007. Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *NeuroImage* 36 (3), 645–660.
- Iturria-Medina, Y., Sotero, R. C., Canales-Rodríguez, E. J., Alemán-Gómez, Y., Melie-García, L., Apr 2008. Studying the human brain anatomical network via diffusion-weighted MRI and graph theory. *NeuroImage* 40 (3), 1064–1076.
- Iwata, N. K., Kwan, J. Y., Danielian, L. E., Butman, J. A., Tovar-Moll, F., Bayat, E., Floeter, M. K., Sep 2011. White matter alterations differ in primary lateral sclerosis and amyotrophic lateral sclerosis. *Brain* 134 (Pt 9), 2642–2655.
- Jbabdi, S., Woolrich, M. W., Andersson, J. L. R., Behrens, T. E. J., Aug 2007. A bayesian framework for global tractography. *NeuroImage* 37 (1), 116–129.
- Jeurissen, B., Leemans, A., Jones, D. K., Tournier, J.-D., Sijbers, J., Mar 2011. Probabilistic fiber tracking using the residual bootstrap with constrained spherical deconvolution. *Hum Brain Mapp* 32 (3), 461–479.
- Johansen-Berg, H., Behrens, T. E. J., Robson, M. D., Drobnjak, I., Rushworth, M. F. S., Brady, J. M., Smith, S. M., Higham, D. J., Matthews, P. M., Sep 2004. Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. *Proc Natl Acad Sci* 101 (36), 13335–13340.
- Jones, D. K., 2010. *Diffusion MRI: Theory, methods, and applications*. Oxford University Press.
- Jones, D. K., Knösche, T. R., Turner, R., Jun 2013. White matter integrity, fiber count, and other fallacies: the do’s and don’ts of diffusion MRI. *NeuroImage* 73, 239–254.
- Kasenburg, N., Darkner, S., Hahn, U., Liptrot, M., Feragen, A., 2016a. Structural parcellation of the thalamus using shortest-path tractography. In: *IEEE 13th International Symposium on Biomedical Imaging (ISBI)*.

- Kasenburg, N., Liptrot, M., Reisle, N. L., Garde, E., Nielsen, M., Feragen, A., 2016b. Supervised hub-detection for brain connectivity. Proc. SPIE 9784, Medical Imaging 2016: Image Processing.
- Kasenburg, N., Liptrot, M., Reisle, N. L., Ørting, S. N., Nielsen, M., Garde, E., Feragen, A., 2016c. Training shortest-path tractography: Automatic learning of spatial priors. NeuroImage.
- Kasenburg, N., Pedersen, M. V., Darkner, S., 2016d. Finding significantly connected voxels based on histograms of connection strengths. Proc. SPIE 9784, Medical Imaging 2016: Image Processing.
- Lazar, M., Alexander, A. L., Oct 2003. An error analysis of white matter tractography methods: synthetic diffusion tensor field simulations. NeuroImage 20 (2), 1140–1153.
- Lebel, C., Gee, M., Camicioli, R., Wieler, M., Martin, W., Beaulieu, C., 2012. Diffusion tensor imaging of white matter tract evolution over the lifespan. NeuroImage 60 (1), 340–352.
- Lenglet, C., Deriche, R., Faugeras, O., 2004. Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping. In: ECCV.
- Liptrot, M. G., Sidaros, K., Dyrby, T. B., 2014. Addressing the path-length-dependency confound in white matter tract segmentation. PLoS One 9 (5), e96247.
- Mars, R. B., Jbabdi, S., Sallet, J., O’Reilly, J. X., Crosson, P. L., Olivier, E., Noonan, M. P., Bergmann, C., Mitchell, A. S., Baxter, M. G., Behrens, T. E. J., Johansen-Berg, H., Tomassini, V., Miller, K. L., Rushworth, M. F. S., Mar 2011. Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and macaque resting-state functional connectivity. J Neurosci 31 (11), 4087–4100.
- McGrath, J., Johnson, K., O’Hanlon, E., Garavan, H., Leemans, A., Gallagher, L., 2013. Abnormal functional connectivity during visuospatial processing is associated with disrupted organisation of white matter in autism. Front Hum Neurosci 7, 434.
- Morel, A., Magnin, M., Jeanmonod, D., Nov 1997. Multiarchitectonic and stereotactic atlas of the human thalamus. J Comp Neurol 387 (4), 588–630.
- Moreno-Dominguez, D., Anwender, A., Knösche, T. R., Oct 2014. A hierarchical method for whole-brain connectivity-based parcellation. Hum Brain Mapp 35 (10), 5000–5025.
- Mori, S., Crain, B. J., Chacko, V. P., van Zijl, P. C., Feb 1999. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Ann Neurol 45 (2), 265–269.
- Morris, D. M., Embleton, K. V., Parker, G. J. M., Oct 2008. Probabilistic fibre tracking: differentiation of connections from chance events. NeuroImage 42 (4), 1329–1339.

- Moseley, M. E., Kucharczyk, J., Asgari, H. S., Norman, D., Jun 1991. Anisotropy in diffusion-weighted MRI. *Magn Reson Med* 19 (2), 321–326.
- Mufson, E. J., Brady, D. R., Kordower, J. H., 1990. Tracing neuronal connections in postmortem human hippocampal complex with the carbocyanine dye DiI. *Neurobiol Aging* 11 (6), 649–653.
- Myllymäki, M., Mrkvicka, T., Grabarnik, P., Seijo, H., Hahn, U., June 2013. Global envelope tests for spatial processes. eprint arXiv:1307.0239.
- Nanetti, L., Cerliani, L., Gazzola, V., Renken, R., Keysers, C., Oct 2009. Group analyses of connectivity-based cortical parcellation using repeated k-means clustering. *NeuroImage* 47 (4), 1666–1677.
- Neher, P. F., Laun, F. B., Stieltjes, B., Maier-Hein, K. H., Nov 2014. Fiberfox: facilitating the creation of realistic white matter software phantoms. *Magn Reson Med* 72 (5), 1460–1470.
- O’Donnell, L., Haker, S., Westin, C.-F., 2002. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. *MICCAI* 5 (Pt 1), 459–466.
- Parker, G. J. M., Alexander, D. C., Jul 2003. Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. *Inf Process Med Imaging* 18, 684–695.
- Parker, G. J. M., Haroon, H. A., Wheeler-Kingshott, C. A. M., Aug 2003. A framework for a streamline-based probabilistic index of connectivity (PICO) using a structural interpretation of MRI diffusion measurements. *J Magn Reson Imaging* 18 (2), 242–254.
- Perry, A., Wen, W., Lord, A., Thalamuthu, A., Roberts, G., Mitchell, P. B., Sachdev, P. S., Breakspear, M., Jul 2015. The organisation of the elderly connectome. *NeuroImage* 114, 414–426.
- Pierpaoli, C., Basser, P. J., Dec 1996. Toward a quantitative assessment of diffusion anisotropy. *Magn Reson Med* 36 (6), 893–906.
- Robinson, E. C., Hammers, A., Ericsson, A., Edwards, A. D., Rueckert, D., 2010. Identifying population differences in whole-brain structural networks: A machine learning approach. *NeuroImage* 50 (3), 910–919.
- Rojkova, K., Volle, E., Urbanski, M., Humbert, F., Dell’Acqua, F., de Schotten, M. T., Feb 2015. Atlasing the frontal lobe connections and their variability due to age and education: a spherical deconvolution tractography study. *Brain Struct Funct*.
- Rubinov, M., Bullmore, E., 2013. Fledgling pathoconnectomics of psychiatric disorders. *Trends in Cognitive Sciences* 17 (12), 641–647, Special Issue: The Connectome.
- Schober, M., Kasenburg, N., Feragen, A., Hennig, P., Hauberg, S., 2014. Probabilistic shortest path tractography in DTI using Gaussian Process ODE solvers. *MICCAI* 17 (Pt 3), 265–272.

- Sotiropoulos, S. N., Bai, L., Morgan, P. S., Constantinescu, C. S., Tench, C. R., Feb 2010. Brain tractography using q-ball imaging and graph theory: Improved connectivities through fibre crossings via a model-based approach. *NeuroImage* 49 (3), 2444–2456.
- Sporns, O., 2014. Contributions and challenges for network models in cognitive neuroscience. *Nat Neurosci* 17 (5), 652–660.
- Sugiyama, M., López, F. L., Kasenburg, N., Borgwardt, K. M., 2015. Significant subgraph mining with multiple testing correction. In: *Proceedings of the 2015 SIAM International Conference on Data Mining*. Vol. 5. pp. 37–45.
- Tang, P.-F., Ko, Y.-H., Luo, Z.-A., Yeh, F.-C., Chen, S.-H. A., Tseng, W.-Y. I., Jun 2010. Tract-specific and region of interest analysis of corticospinal tract integrity in subcortical ischemic stroke: reliability and correlation with motor function of affected lower extremity. *Am J Neuroradiol* 31 (6), 1023–1030.
- Teeuw, J., Caan, M. W. A., Olabarriaga, S. D., 2015. Robust automated white matter pathway reconstruction for large studies. *MICCAI* 18 (Pt 1), 101–108.
- Tournier, J.-D., Calamante, F., Connelly, A., May 2007. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35 (4), 1459–1472.
- Tournier, J.-D., Calamante, F., Gadian, D. G., Connelly, A., Nov 2004. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage* 23 (3), 1176–1185.
- Tournier, J.-D., Mori, S., 2014. *Introduction to Diffusion Tensor Imaging and Higher Order Models*. Academic Press.
- Tuch, D. S., Dec 2004. Q-ball imaging. *Magn Reson Med* 52 (6), 1358–1372.
- Vorburger, R. S., Reischauer, C., Boesiger, P., Feb 2013. BootGraph: Probabilistic fiber tractography using bootstrap algorithms and graph theory. *NeuroImage* 66, 426–435.
- Whitford, T. J., Kubicki, M., Pelavin, P. E., Lucia, D., Schneiderman, J. S., Pantelis, C., McCarley, R. W., Shenton, M. E., Jan 2015. Cingulum bundle integrity associated with delusions of control in schizophrenia: Preliminary evidence from diffusion-tensor tractography. *Schizophr Res* 161 (1), 36–41.
- Wiegell, M. R., Tuch, D. S., Larsson, H. B. W., Wedeen, V. J., Jun 2003. Automatic segmentation of thalamic nuclei from diffusion tensor magnetic resonance imaging. *NeuroImage* 19 (2 Pt 1), 391–401.
- Wozniak, J. R., Mueller, B. A., Bell, C. J., Muetzel, R. L., Lim, K. O., Day, J. W., Apr 2013. Diffusion tensor imaging reveals widespread white matter abnormalities in children and adolescents with myotonic dystrophy type 1. *J Neurol* 260 (4), 1122–1131.

- Xia, S., Li, X., Kimball, A. E., Kelly, M. S., Lesser, I., Branch, C., Nov 2012. Thalamic shape and connectivity abnormalities in children with attention-deficit/hyperactivity disorder. *Psychiatry Res* 204 (2-3), 161–167.
- Yendiki, A., Panneck, P., Srinivasan, P., Stevens, A., Zöllei, L., Augustinack, J., Wang, R., Salat, D., Ehrlich, S., Behrens, T., Jbabdi, S., Gollub, R., Fischl, B., 2011. Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy. *Front Neuroinform* 5, 23.
- Zalesky, A., Oct 2008. DT-MRI fiber tracking: a shortest paths approach. *IEEE Transactions on Medical Imaging* 27 (10), 1458–1471.
- Zalesky, A., Fornito, A., Harding, I. H., Cocchi, L., Yücel, M., Pantelis, C., Bullmore, E. T., Apr 2010. Whole-brain anatomical networks: does the choice of nodes matter? *NeuroImage* 50 (3), 970–983.

Chapter 5

Papers

This chapter includes all papers discussed in this thesis. They are included directly as the state they were in at the time of submission of the thesis.

The order of papers is not chronological with respect to publication, but according to the structure used throughout this thesis, starting with tractography, followed by the two papers on statistical analysis of tractography results and its application to connectivity-based parcellation, and finally the paper on supervised hub-detection.

5.1

Training Shortest-Path Tractography: Automatic Learning of Spatial Priors

Niklas Kasenburg^{a,*}, Matthew Liptrot^{a,b}, Nina Linde Reisleiv^c, Silas N. Ørting^a, Mads Nielsen^a, Ellen Garde^c, Aasa Feragen^a

^a*Department of Computer Science, University of Copenhagen, Denmark*

^b*DTU Compute, Technical University of Denmark, Denmark*

^c*DRCMR, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital Hvidovre, Denmark*

Abstract

Tractography is the standard tool for automatic delineation of white matter tracts from diffusion weighted images. However, the output of tractography often requires post-processing to remove false positives and ensure a robust delineation of the studied tract, and this demands expert prior knowledge. Here we demonstrate how such prior knowledge, or indeed any prior spatial information, can be automatically incorporated into a shortest-path tractography approach to produce more robust results. We describe how such a prior can be automatically generated (learned) from a population, and we demonstrate that our framework also retains support for conventional interactive constraints such as waypoint regions. We apply our approach to the open access, high quality Human Connectome Project data, as well as a dataset acquired on a typical clinical scanner. Our results show that the use of a learned prior substantially increases the overlap of tractography output with a reference atlas on both populations, and this is confirmed by visual inspection. Furthermore, we demonstrate how a prior learned on the high quality dataset significantly increases the overlap with the reference for the more typical yet lower quality data acquired on a clinical scanner. We hope that such automatic incorporation of prior knowledge and the obviation of expert interactive tract delineation on every subject, will improve the feasibility of large clinical tractography studies.

Keywords: Tractography, Diffusion MRI, Graph theory, Prior information

1. Introduction

Diffusion weighted imaging (DWI) of the human brain provides local estimates of water diffusion summarised as voxel-wise diffusion orientation distribution functions (dODFs) (Hagmann et al., 2006). These can be transformed into fibre orientation distribution functions (fODFs), representing estimates of the fibre directions within a voxel. Subsequently, tractography attempts to delineate the underlying anatomical *tracts* connecting brain regions by inferring inter-voxel connectivity from these fODFs.

*Corresponding author. DIKU North Campus, Universitetsparken 5, DK-2100 København Ø, Denmark.
Email address: niklas.kasenburg@di.ku.dk (Niklas Kasenburg)

Over the last decade it has become increasingly clear how critical the integrity of these white matter (WM) tracts is to the healthy functioning of the brain (Abhinav et al., 2014; Connally et al., 2014; Iwata et al., 2011; Steinbach et al., 2015). Therefore, techniques are needed which allow evaluation and monitoring of microstructural tissue properties to gain insight into the mechanisms underlying brain development, ageing and pathology. Metrics of such properties include those derived from DWI data, such as the simple fractional anisotropy (FA) (McGrath et al., 2013; Whitford et al., 2015; Xia et al., 2012), generalized FA (Tang et al., 2010), diffusivity measurements (Benedetti et al., 2011; Davis et al., 2009; Galantucci et al., 2011; Wozniak et al., 2013) or other more complex estimates of microstructure (Assaf et al., 2013; Golestani et al., 2014). Additionally, such tract-based integrity measures can also be derived from other magnetic resonance sequences such as T1 relaxometry or magnetisation transfer imaging (Alexander et al., 2011).

When comparing such tract-specific features across subjects, it is important that the tracts from which they originate are robustly and reliably reconstructed. Voxel-wise FA values (Iwata et al., 2011; Whitford et al., 2015) or white matter tissue probabilities (Iturria-Medina et al., 2007) are often used to either prune results or guide the tractography. Post-processing of tractography outputs using waypoint or exclusion regions (Benedetti et al., 2011; Connally et al., 2014; Galantucci et al., 2011; Rojkova et al., 2015) is also often necessary to ensure consistent and accurate tract delineation. It would therefore be beneficial, in terms of manual effort and reliability, if domain knowledge about the connection or subject could be automatically included in the tract delineation. This can be achieved by integrating prior information regarding tract location into the tractography framework.

Jbabdi et al. (2007) present a Bayesian framework for global probabilistic tractography, which aims to find the optimal tract between two regions, incorporating prior information. However, the possible priors only include knowledge about the existence (or not) of a connection and do not include any prior information about the tract location. Furthermore, due to the large complexity of the problem, the optimal solution is intractable and the tracts are estimated by heuristically sampling from the posterior distribution. Building upon this global tractography framework, Yendiki et al. (2011) include prior anatomical information about the spatial location of a fixed number of segments along each fibre for whole-brain tractography. The prior is computed from training data comprised of manually labelled and verified tracts. Although this reduces the search space and is able to guide the tractography reliably on unseen data without manual intervention, it still requires that domain experts post-process the tractography training data and it may need to be repeated for any future novel tasks or derived features.

We present an extension to a shortest-path tractography framework that can include any type of prior information about the spatial location of a tract. Such prior information could, for example, consist of per-voxel white matter probabilities, which guide tractography through white matter, or anatomical knowledge in the form of a tract atlas. Our algorithm also allows priors that are generated from expert annotation, similar to Yendiki et al. (2011). In addition, we present a method for automatically learning such spatial priors from previous tractography results. We demonstrate, in particular, how a prior learned on independent, high quality data, where tract delineation is easier and more accurate, is able to improve the performance of tractography on lower quality data.

Our framework employs a shortest-path tractography (SPT) approach, which finds the globally optimal path connecting two voxels. Like the framework described by Jbabdi et al. (2007),

SPT has the advantage of being less susceptible to local noise in the data because it evaluates all possible connections. Moreover, because the discretisation into a graph allows the use of optimal graph-based shortest-path algorithms, graph-based SPT methods (Iturria-Medina et al., 2007; Sotiropoulos et al., 2010; Vorburger et al., 2013; Zalesky, 2008) are guaranteed to find the best path connecting any pair of voxels. In contrast, their continuous counterparts (Fuster et al., 2014; Hauberg et al., 2015; Lenglet et al., 2004; O’Donnell et al., 2002; Schober et al., 2014), as well as the probabilistic approach by Jbabdi et al. (2007), require a good initialization to avoid local optima. While existing graph-based SPT algorithms often impose strict assumptions upon the form that the f/dODF may take, our framework gives full modelling flexibility by permitting any form of fODF. We obtain a Bayesian SPT algorithm by interpreting spatial priors as soft or hard constraints on tract location. As existing graph-based tractography methods do not provide algorithmic solutions to constrained tractography problems (Iturria-Medina et al., 2007; Sotiropoulos et al., 2010), we furthermore derive intuitive, exact and efficient algorithmic solutions to incorporate prior information from multiple sources into our tractography framework.

In addition to determining the most probable path for the tract connecting two voxels, our SPT algorithm also returns a confidence score which provides a quantitative measure of how well a shortest path is supported by both the underlying fODFs of all component voxels and by the prior information. This “importance” evaluation of any path provides a numerical score that permits our framework to automatically learn a tract prior from training data without requiring expert interaction.

In the next section we briefly review graph-based shortest-path tractography and how it can be applied for region to region global tractography. We then describe how we integrate prior information in Section 3. In Section 4 we describe the two datasets used throughout this study, how they were pre-processed and how the tractography experiments were performed. We also describe the reference used for evaluation of tractography results. In Section 5, we show results of the tractography, first using a simple subject-specific prior given by white matter probability, second using study-specific or independent learned priors, and finally using both simple and learned priors in combination with a binary waypoint prior. We show, in particular, results for tractography on the dataset acquired on a typical clinical scanner with a prior learned from the high quality dataset. We conclude with a discussion of the results and a brief conclusion.

2. Revisiting graph-based shortest-path tractography

In this section we review graph-based tractography and phrase its solution as a shortest-path problem, which in Section 3 will allow us to efficiently integrate spatial priors into the tractography algorithm.

From the DWI data of a brain we extract an undirected brain-graph $G = (V, E, w_E)$ whose node set V contains all DWI voxels within the brain, excluding those classified as cerebrospinal fluid (CSF) by prior tissue segmentation. Each node is connected by an edge $e \in E$ to all white matter voxels in its $3 \times 3 \times 3$ neighbourhood on the 3D image grid. Each edge e is assigned a weight $w_E(e) \in [0, 1]$ reflecting the probability of a fibre bundle connecting its two endpoint nodes; this process is described in Section 2.1.

A path $\pi_{v,v'}$ connecting nodes (or voxels) $v \in V$ and $v' \in V$ in G is defined as a sequence of nodes $\pi_{v,v'} = [v_1, v_2, \dots, v_n]$, where $v_1 = v$, $v_n = v'$ and $(v_i, v_{i+1}) \in E$ for all $i = 1, \dots, n - 1$. The cardinality $|\pi_{v,v'}| = n$ of a path is given by the number of nodes in the path. The likelihood

of the path $\pi_{v,v'}$ is defined as the product of all edge weights $w_E(e)$ encountered along the path:

$$\mathcal{L}(\pi_{v,v'}) = \prod_{i=1}^{n-1} w_E(v_i, v_{i+1}) . \quad (1)$$

2.1. From fODF to edge probability

For each voxel, assume that the diffusion information from the DWI is summarized in an fODF $f: S^2 \rightarrow R_+$ associating to any given direction θ on the unit sphere S^2 a probability that there exists a fibre along that direction. We are interested in the 26 directions θ_i with $i = 1, \dots, 26$, pointing from the centre of a voxel towards its 26 neighbouring voxels.

We model the connectivity $w(\theta_i)$ along an edge from the voxel centre in the direction $\theta_i \in S^2$ by integrating the fODF over the set C_i of all directions $\theta \in S^2$ pointing out of the voxel that are closer to θ_i than to any other of the 26 directions $\theta_j, j \neq i$. The set C_i is called a Voronoi cell (Voronoi, 1908). Since computing integrals over Voronoi cells on the sphere is computationally hard, we numerically approximate the integral through sampling. The weight $w(\theta_i)$ describes the probability of connection in the direction θ_i and is defined and approximated as follows:

$$w(\theta_i) = \int_{C_i} f(\theta) d\theta \approx \sum_{\tilde{\theta}_k \in S_i} \left(f(\tilde{\theta}_k) \cdot \frac{\text{Vol}(S^2)}{N} \right) , \quad (2)$$

where $S = \{\tilde{\theta}_k \in S^2, k = 1 : N\}$ is a uniform sample of N directions, $S_i = S \cap C_i$ is the set of direction samples belonging to C_i and $\text{Vol}(S^2)/N$ is the average volume corresponding to a sample direction $\tilde{\theta}_k$. As $w(\theta_i)$ depends on its source node, the edge weight $w_E(v, v')$ is defined as the average

$$w_E(v, v') = 1/2 \cdot (w(v \rightarrow v') + w(v' \rightarrow v)) , \quad (3)$$

where $v \rightarrow v'$ is the direction from v to v' . This yields an undirected graph with symmetric edge weights: $w_E(v, v') = w_E(v', v)$.

2.2. From most-likely path to shortest path

Given any two nodes $v, v' \in V$, tractography is defined as finding the path between v and v' in the brain graph G that maximizes the path likelihood from Equation 1, i.e. finding the maximum likelihood solution

$$\text{argmax}_{\pi_{v,v'}} \mathcal{L}(\pi_{v,v'}) . \quad (4)$$

The most-likely path problem on G can be reformulated as a shortest-path problem on a modified graph \tilde{G} . This gives access to classical efficient algorithms for computing shortest paths on \tilde{G} to obtain most-likely paths on G . The log-likelihood of the path $\pi_{v,v'}$ is equal to the sum of modified edge weights

$$\tilde{w}_E(e) = -\log w_E(e) , \quad (5)$$

corresponding to edges encountered along the path, that is:

$$\begin{aligned} -\log(\mathcal{L}(\pi_{v,v'})) &= -\log \prod_{i=1}^{n-1} w_E(v_i, v_{i+1}) \\ &= \sum_{i=1}^{n-1} \tilde{w}_E(v_i, v_{i+1}) . \end{aligned} \tag{6}$$

Defining the modified graph $\tilde{G} = (V, E, \tilde{w}_E)$, where \tilde{w}_E is the log-transformed edge probability (5), we observe that finding the path that maximizes (1) in G is equivalent to finding the path in \tilde{G} that minimizes (6). That is, the maximum likelihood path problem on G can be solved as a shortest-path problem on \tilde{G} . We solved this using Dijkstra's single-source shortest-path algorithm (Dijkstra, 1959), but any other shortest-path algorithm can be applied.

The formulation of tractography as a shortest-path problem first appeared in (Zalesky, 2008). The use of Dijkstra's shortest-path algorithm for most-likely paths is algorithmically similar to the algorithm employed in (Iturria-Medina et al., 2007; Sotiropoulos et al., 2010), where instead of translating the most-likely path to a shortest path, Dijkstra's algorithm was converted to the domain of multiplicative probabilities. In the following sections we will use the negative log-transform to turn more complex path optimization problems into shortest-path problems as well.

2.3. Region to region tractography

In Section 2.2 we defined the most-likely path between two nodes in the brain graph. We now describe how to perform shortest-path tractography (SPT) between two regions of interest (ROIs) and how to obtain confidence maps illustrating the likelihood of tract location. The SPT outlined in Algorithm 1 takes as input the brain-graph G and two sets of nodes (i.e. voxels) $R_1, R_2 \subset V$ representing the endpoint ROIs of the sought tract. Then tractography is performed by finding the set Π of all shortest paths π_{r_1, r_2} from any voxel r_1 in R_1 to any voxel r_2 in R_2 in \tilde{G} , along with their scores $s(\pi_{r_1, r_2})$. We do this using Dijkstra's single-source shortest-path algorithm for every node $r_1 \in R_1$. See lines 3–8 in Algorithm 1, where Π_{r_1} denotes the set of shortest paths in \tilde{G} that start at r_1 .

Algorithm 1 Shortest-path tractography (SPT) between ROIs R_1 and R_2

Input: $G = (V, E, w_E)$, $R_1, R_2 \subset V$
1: $\tilde{G} \leftarrow (V, E, \tilde{w}_E)$, where $\tilde{w}_E(v, v') = -\log w_E(v, v')$
 for all $(v, v') \in E$
2: $\Pi \leftarrow \emptyset$, $s(\Pi) \leftarrow \emptyset$
3: **for** $r_1 \in R_1$ **do**
4: $(\Pi_{r_1}, s(\Pi_{r_1})) \leftarrow \text{single_source_shortest_path}(r_1, \tilde{G})$
5: **for** $r_2 \in R_2$ **do**
6: $(\Pi, s(\Pi)) \leftarrow$
 $(\Pi, s(\Pi)) \cup \{(\pi_{r_1, r_2}, s(\pi_{r_1, r_2})) \mid \pi_{r_1, r_2} \in \Pi_{r_1}\}$
7: **end for**
8: **end for**
Output: Π , $s(\Pi)$

In addition to the actual paths, Dijkstra’s algorithm also outputs the length of each shortest path in Π_{r_1} , that is, the log-transformed likelihood (see Equation 6) of each path in Π_{r_1} , from which we can compute the path likelihood (see Equation 1). As a path *score* we use the normalized path likelihood, which accounts for path cardinality $n = |\pi_{v,v'}|$:

$$s(\pi_{v,v'}) = \sqrt[n]{\mathcal{L}(\pi_{v,v'})} . \quad (7)$$

The SPT algorithm outputs the shortest paths $\pi \in \Pi$, as well as their scores $s(\pi)$ for all $\pi \in \Pi$. We summarize the SPT output in a scalar confidence map M , which for each voxel $v \in V$ is the sum of all scores of the paths containing v :

$$M(v) = \sum \{s(\pi) | v \in \pi, \pi \in \Pi\} . \quad (8)$$

We describe in Section 3.5 how a population heatmap can be created from such subject-specific confidence maps.

3. Shortest-path tractography with spatial priors

The goal of this paper is to integrate spatial prior information into the shortest-path tractography algorithm. This is achieved by assigning a prior probability $p_V(v) \in [0, 1]$ to every node $v \in V$ in the graph, such that $p_V(v)$ reflects the probability that the node v lies within the sought tract(s). This prior information can for instance consist of anatomical knowledge about the location of a tract, white matter probability to account for partial volume effects or way-point masks to constrain the tractography to pass through predefined regions. In Section 3.5 we shall, moreover, see how we can automatically *learn* a prior from previous tractography results.

The node prior $p_V: V \rightarrow [0, 1]$ generates a soft or hard constraint on the shortest paths, encouraging them (or, in the case of a binary prior $p_V: V \rightarrow \{0, 1\}$, forcing them) to pass through certain regions. In Section 3.2 below we show how this constrained shortest-path problem can be reinterpreted as a shortest-path problem on a new modified graph in order to obtain an efficient algorithm.

3.1. A spatial Bayesian prior

We incorporate prior information about the spatial location of a tract in a Bayesian model of connectivity along edges in the brain-graph $G = (V, E, w_E)$. A spatial prior is most naturally formulated as a prior $p_V: V \rightarrow \mathbb{R}_+$ describing whether given nodes (or voxels) belong to the tract. As our original connectivity probabilities are defined over edges, the Bayesian model for connectivity along edges is obtained by translating the node prior to a prior over edges $e \in E$. For a single edge $e = (v, v') \in E$ we define the prior probability $P(e)$ that the sought tract runs along the edge e as the square root of the product of its end node probabilities:

$$P(e) = \sqrt{p_V(v)p_V(v')} . \quad (9)$$

Given a single edge $e = (v, v') \in E$, we define a probability density function $f_e: W \rightarrow \mathbb{R}_+$ which, given any edge weight $w_E(e)$ in the set of possible edge weights $W \subset \mathbb{R}_+$ extracted

from the DWI image, will return the likelihood of structural brain connectivity along the edge e . We take f_e to be the edge probability w_E described in Section 2.1, that is

$$f_e(w_E(e)) \propto w_E(e) = \exp(-\tilde{w}_E(e)) \quad , \quad (10)$$

which gives a likelihood $P(w_E(e)|e) = f_e(w_E(e)) = w_E(e)$ of connectivity along e . As before, large values of $w_E(e)$ correspond to small edge lengths $\tilde{w}_E(e)$, so that short edges in $\tilde{G} = (V, E, \tilde{w}_E)$ encode high likelihood of structural connectivity along e .

In order to incorporate the spatial prior probability $P(e)$, we use Bayes' theorem to obtain the posterior probability of structural brain connectivity along the edge e :

$$P(e|w_E(e)) \propto P(w_E(e)|e) P(e) = w_E(e)P(e) \quad . \quad (11)$$

This returns an edge-wise posterior probability for connectivity along each edge $e \in E$.

3.2. An algorithm for SPT with a spatial prior

From the Bayesian model for structural brain connectivity along edges, we obtain tractography with a spatial prior as a problem of finding shortest paths in a new modified graph $\tilde{G} = (V, E, \tilde{w}_E)$ where

$$\begin{aligned} \tilde{w}_E(e) &= -\log P(e|w_E(e)) \\ &= -\log(w_E(e)P(e)) \\ &= -\log\left(\sqrt{p_V(v)} \cdot w_E(v, v') \cdot \sqrt{p_V(v')}\right) \quad , \end{aligned} \quad (12)$$

for any given edge $e = (v, v')$.

Note that since the length of a path $\pi_{v, v'} = [v = v_1, v_2 \dots, v_n = v']$ in \tilde{G} can be rewritten as

$$\begin{aligned} \sum_{i=1}^{n-1} \tilde{w}_E(e_i) &= \sum_{i=1}^{n-1} -\log P(e_i|w_E(e_i)) \\ &= -\log\left(\prod_{i=1}^{n-1} P(e_i|w_E(e_i))\right) \quad , \end{aligned} \quad (13)$$

where $e_i = (v_i, v_{i+1})$ for all i , a shortest path in \tilde{G} corresponds to a path in G which maximizes the ‘‘posterior’’ probability of the path given the prior. That is, if we assume that our edges are independent, we obtain a path probability

$$\begin{aligned} P(\pi_{v, v'}|G) &= \prod_{i=1}^{n-1} P(e_i|w_E(e_i)) \\ &= \sqrt{p_V(v_1)p_V(v_n)} \prod_{i=1}^{n-1} w_E(e_i) \prod_{i=2}^{n-1} p_V(v_i) \quad , \end{aligned} \quad (14)$$

and the path $\pi_{v, v'}$ that maximises the path probability $P(\pi_{v, v'}|G)$ is the shortest path connecting v and v' in \tilde{G} :

$$\operatorname{argmin}_{\pi_{v, v'}} \sum_{i=1}^{n-1} \tilde{w}_E(v_i, v_{i+1}) = \operatorname{argmax}_{\pi_{v, v'}} P(\pi_{v, v'}|G) \quad . \quad (15)$$

Note that as in (Yendiki et al., 2011) the path prior in (14) becomes a product of prior information about voxels along the path. However, in contrast to Yendiki et al. (2011), where only a specific and limited anatomical prior (a region label) is used, we formulate the problem such that it can accept any type of spatial prior. Furthermore, and unlike Jbabdi et al. (2007), we do not need to estimate the solution by sampling a path from a posterior distribution. Instead we find the optimum, maximum probability path (see Equation 15) directly, as described in Algorithm 2. This is made possible due to the discrete graph representation, in which we know all possible paths. Note also that the reinterpretation of the prior as an edge prior allows us to phrase tractography with a prior as a shortest-path problem on a modified graph. This formulation simplifies the computational problem as it gives access to a rich family of shortest-path algorithms. In the following, we will describe explicit examples of spatial priors p_V .

Algorithm 2 Shortest-path tractography with a spatial prior

Input: $G = (V, E, w_E)$, $R_1, R_2 \subset V$, $p_V : V \rightarrow [0, 1]$
 1: $\tilde{G} \leftarrow (V, E, \tilde{w}_E)$, where
 $\tilde{w}_E(v, v') = -\log\left(\sqrt{p_V(v)} \cdot w_E(v, v') \cdot \sqrt{p_V(v')}\right)$
 for all $(v, v') \in E$
 2: $\Pi \leftarrow \emptyset$, $s(\Pi) \leftarrow \emptyset$
 3: **for** $r_1 \in R_1$ **do**
 4: $(\Pi_{r_1}, s(\Pi_{r_1})) \leftarrow \text{single_source_shortest_path}(r_1, \tilde{G})$
 5: **for** $r_2 \in R_2$ **do**
 6: $(\Pi, s(\Pi)) \leftarrow$
 $(\Pi, s(\Pi)) \cup \{(\pi_{r_1, r_2}, s(\pi_{r_1, r_2})) \mid \pi_{r_1, r_2} \in \Pi_{r_1}\}$
 7: **end for**
 8: **end for**
Output: Π , $s(\Pi)$

3.3. Subject-specific simple prior: white matter probability

In existing methods for streamline and SPT tractography, simple descriptors are often extracted from each subject’s DWI volume to guide the tractography. For instance, due to partial volume effects and the fact that CSF regions have high diffusivity, voxels that partially contain CSF tend to attract shortest paths. This can be alleviated by using white matter probabilities for each voxel v , e.g. (Iturria-Medina et al., 2007). These probabilities can be integrated into our framework as a spatial prior as described in (11). In our experiments we demonstrate the use of a white matter prior $p_V^{WM} : V \rightarrow [0, 1]$ defined by a tissue segmentation.

Note that if all voxels on the path have prior value $p_V(v) = 1$, the original path likelihood is obtained (see Equation 1). Each voxel on the path with a lower prior probability, such as a lower white matter probability in the case of partial volume effects, decreases the total likelihood of the path.

3.4. Binary priors: waypoints and exclusion masks

Other commonly used tools in tractography are exclusion regions and waypoints (Catani et al., 2002; Rojkova et al., 2015; Wakana et al., 2007) that either forbid or force the tracts

to pass through pre-defined regions. These masks are usually applied after whole-brain tractography to delineate those streamlines that correspond to the tract of interest. Binary masks can be integrated as hard constraints into our tractography framework by using a binary prior $p_V^{bin}: V \rightarrow \{0, 1\}$. Exclusion priors are then set to 0 in the voxels belonging to the exclusion region and 1 everywhere else.

We define waypoint masks in a single slice or a thin volume transversal to the tract. In order to apply the waypoint to SPT, it must provide the only possible route for all potential paths between the endpoint regions. To ensure this, all voxels in the waypoint slice (or thin volume) are set to 0, except for those of the waypoint itself. The remainder of voxels in the full volume (and excluding those within the waypoint slice mask) are set to 1. The latter is important as it enforces that no additional prior is applied to voxels outside of the slice of interest.

3.5. Learned spatial prior

We propose to learn a prior for the location of a tract based on tractography performed on training data. First tractography is performed for a set of N training subjects. A confidence map $M_i, i = 1, \dots, N$, where $M_i(v)$ is the scalar confidence value for voxel v (as defined in Equation 8), is created for each of the N subjects. In the next step the maps are first normalized to sum up to 1 and then warped into a common standard space resulting in the normalized warped maps $\tilde{M}_i, i = 1, \dots, N$.

Then the population heatmap \tilde{H} is computed for every voxel in the standard space by first adding up the corresponding values of each subject and then dividing by the maximum to normalize to values between 0 and 1

$$\tilde{H}(u) = \frac{\sum_{i=1}^N \tilde{M}_i(u)}{\max\left(\left\{\sum_{i=1}^N \tilde{M}_i(u') \mid u' \in \tilde{H}\right\}\right)}, \quad (16)$$

where u and u' are voxels in standard space. Given a new subject i in which tractography is to be performed, the learned prior probability $p_V^{learned}: V \rightarrow [0, 1]$ is obtained by warping the population heatmap \tilde{H} to its individual subject space, obtaining a warped population heatmap H_i :

$$p_V^{learned}(v) = H_i(v). \quad (17)$$

The prior $p_V^{learned}$ thus quantifies how likely a voxel is to lie on a tract between two specific ROIs, based on tractography results of the training population.

3.6. Combining multiple priors

Since all priors are represented by a probability whose values are between 0 and 1, different priors can be combined by taking their product. Given for example a white matter prior p_V^{WM} , a waypoint prior p_V^{bin} and a learned spatial prior $p_V^{learned}$ for all nodes, the combined prior for a node v can be defined as

$$p_V(v) = p_V^{WM}(v) \cdot p_V^{bin}(v) \cdot p_V^{learned}(v). \quad (18)$$

4. Data

We applied the presented method to two different datasets. Firstly, to pre-processed diffusion data of 38 subjects (21 females and 17 males, age 22–35) provided as a sub-sample from the Q3 release of the Human Connectome Project (Essen et al., 2013; Glasser et al., 2013; Sotiropoulos et al., 2013), henceforth referred to as *HCP data*. The pre-processed HCP diffusion data contains 270 diffusion directions distributed equally over 3 shells defined with b -values of 1000, 2000 and 3000 s/mm², with the following parameters: repetition time (TR)= 5500 ms; echo time (TE)= 89 ms; $1.25 \times 1.25 \times 1.25$ mm³ voxels. Further sequence details are provided by (Sotiropoulos et al., 2013).

Secondly, we used a dataset of 53 healthy subjects (23 females and 30 males) with an age range from 18 to 81 acquired on a typical clinical scanner (Siemens Magnetom Trio 3T MR scanner (Erlangen, Germany) with an 8-channel head coil (Invivo, Gainesville, FL, USA)), from here on referred to as *standard data* (Ramsøy et al., 2012). Whole brain diffusion weighted (DW) images were acquired using a twice-refocused balanced spin echo sequence that minimised eddy current distortion. Ten non-DW images ($b = 0$ s/mm²) and 61 DW images ($b = 1200$ s/mm²), encoded along independent collinear diffusion gradient orientations, were acquired (TR= 8200 ms; TE= 100 ms; field of view (FOV)= 220×220 mm, matrix = 96×96 ; GRAPPA: factor 2, 48 lines; 61 transverse slices; no gap; $2.3 \times 2.3 \times 2.3$ mm³ voxels). Additionally, a 3D whole brain T1-weighted magnetization prepared rapid acquisition gradient echo (MPRAGE) scan (voxel dimension of $1 \times 1 \times 1$ mm, FOV= $256 \times 256 \times 192$ mm, matrix = $256 \times 256 \times 192$, TR= 1540 ms; TE= 3.93 ms ; inversion time (TI)= 800 ms, and a flip-angle of 9 degrees) was also acquired.

The reference tract atlas was obtained from www.natbrainlab.com and its creation is described in detail by de Schotten et al. (2011). In brief, it is derived from a population of 40 (20 females and 20 males) healthy right-handed subjects in an age range from 18 to 22 years. For all subjects diffusion tensor images (60 directions, $2.5 \times 2.5 \times 2.5$ mm) were obtained. Deterministic tractography was performed from manually defined seed regions and delineated using either a one-ROI approach (arcuate fasciculus and fornix) or two-ROI approach (cortico spinal tract and inferior fronto-occipital fasciculus). The atlas was then created for each tract individually by summing binarised visitation maps warped to MNI space over all subjects. The final atlas represents a percentage overlap map (de Schotten et al., 2011) reflecting for each voxel in MNI space the percentage of subjects where a streamline of the corresponding tract passes through that voxel. Note that the percentages are discretised in four intervals (0.25, 0.5, 0.75 and 1). Figure 1 shows a 3D visualisation of the reference for all tracts used in this study.

A second white matter tract atlas from the IIT Human Brain Atlas (Varentsova et al., 2014) was used as a prior for the fornix. In contrast to the reference atlas, this atlas was created by performing tractography directly on a high angular resolution diffusion image (HARDI) template constructed from DWI data of 72 subjects. The confidence for this atlas in each voxel is defined as the ratio of the number of streamlines that belong to a certain tract and go through that voxel divided by the total number of streamlines that belong to that tract.

4.1. Pre-processing

The T1-weighted images of the standard data were registered to MNI standard space using first a linear registration with 12 degrees of freedom (FLIRT) (Jenkinson et al., 2002) followed

by a non-linear registration (FNIRT) (Andersson et al., 2010). For the HCP data we used the provided non-linear transformations to MNI standard space. Tissue segmentations were obtained for both datasets from the T1-weighted image using FAST (Zhang et al., 2001) after brain extraction using BET (Jenkinson et al., 2005).

No further pre-processing was performed on the already extensively pre-processed HCP diffusion data (Glasser et al., 2013). Pre-processing of the DWIs of the standard data to reduce distortions due to eddy currents, motion and susceptibility differences, was performed using ExploreDTI (Leemans et al., 2009) and the gradient directions were adjusted accordingly. The structural T1-weighted image was registered to the mean b0-image using a two stage linear registration: first a linear registration and then a boundary-based registration (BBR) (Greve and Fischl, 2009). Both registrations were performed with FLIRT (Jenkinson et al., 2002) using 6 degrees of freedom.

For every voxel, a fibre orientation distribution function (fODF) based on constrained spherical deconvolution (CSD) (Tournier et al., 2007) with order 8 was computed with the DiPy package (Garyfallidis et al., 2014); on the HCP DWIs using all three shells, and on the pre-processed DWIs from the standard dataset.

4.2. Tractography

As a first step, the graphs required for the SPT were constructed based on the given fODFs and tissue segmentation for each subject as described in Section 2. Then SPT was performed between two given regions of interest using different priors.

We tested the SPT on four different tracts and compared the results to the tract atlas (de Schotten et al., 2011) described above: the arcuate fasciculus (AF), the cortico spinal tract (CST), the inferior fronto-occipital fasciculus (IFOF) and the fornix. Seed regions were constructed in MNI152 “template space” by overlapping the tract atlas with regions defined by the Harvard-Oxford atlas (Desikan et al., 2006), Juelich atlas (Eickhoff et al., 2007), MNI structural atlas (Mazziotta et al., 2001) or Talairach atlas (Lancaster et al., 2007). The overlap regions for the AF were chosen according to Rojkova et al. (2015): 1) anterior middle and inferior temporal gyrus and 2) Broca’s area 44. Both seed regions were dilated with a $3 \times 3 \times 3$ morphological kernel after the overlap was computed. The overlap regions for the CST included: 1) brain stem, hippocampus and amygdala 2) superior frontal gyrus, precentral and postcentral gyrus. The following seed regions were used for the IFOF, where the atlas was dilated with a $3 \times 3 \times 3$ morphological kernel: 1) occipital pole and 2) frontal pole. Waypoint masks for the IFOF were created by overlapping the tract with the occipital lobe and external capsule, similar to the masks used by Rojkova et al. (2015). The seed regions for the fornix were 1) the hypothalamus and 2) the temporal pole. Here, only the second region was overlapped with the reference, while the first was taken in full. Waypoint masks were constructed by drawing regions on the FA-MNI152 template that overlap with the superior right and left branch of the reference.

The waypoint priors for the IFOF and fornix were specified as binary masks covering the entire volume. All voxels were set to 1, with the exception of the slice containing the waypoint region. In this slice, which was aligned transversally to the tract, only those voxels within the waypoint region were set to 1 and all other voxels were set to 0. Any path crossing this slice was thereby forced to pass through the waypoint region.

We used the warp fields provided in the HCP data or computed as described above (see Section 4.1) to warp the constructed seed regions from the “template space” to individual

“subject space”. Tractography was then performed for every subject in its respective space.

For comparison, TRACULA (Yendiki et al., 2011) was performed only on the CST for the HCP data, as the other three tracts are not available in TRACULA. We run TRACULA with standard parameters on the already pre-processed HCP data, and fODFs computed with FSL’s bedpostx (Behrens et al., 2007) using the zeppelin model and 3 fibres per voxel.

4.3. Implementation and runtime of shortest-path tractography with spatial priors

The SPT method including prior information as described in Section 3 and the corresponding graph construction and transformation have been implemented in Python and C++ and are available [download link will be available upon acceptance].

All experiments were performed on a desktop computer running Ubuntu 14.04.3 with twelve 3.2 GHz Intel Core i7 CPUs and 64 GB of memory. Creating the graph for one subject takes, on average, 18.6 ± 2.3 min for the HCP data and 187 ± 23 s for the standard data. Integrating prior and ROI information into the graph for one subject takes, on average, 3.5 ± 0.5 min for the HCP data and 43 ± 5 s for the standard data.

The runtime of SPT for individual tracts depends on the size of the ROIs. Runtimes for all four tracts and information about the corresponding region sizes can be found in Supplementary Table 1 and 2.

4.4. Quantitative evaluation of tractography results

Confidence maps were generated from the tractography results as defined in Section 2.3 and then warped into standard space. To perform a quantitative comparison between those maps obtained by SPT using different prior information, we use a “true positive” score (TP). This score consists of the sum of confidence values over all voxels inside a subject’s normalized map \tilde{M} warped into the reference space (see Section 3.5), weighted by the respective value of the reference R . That is:

$$\text{TP}(\tilde{M}, R) = \sum_{v \in \tilde{M}} R(v) \tilde{M}(v) . \quad (19)$$

The TP reflects how many voxels are contained in the high probability regions of the reference. If we define the reference for being outside the tract as $(1 - R)$, the false positive score can be defined as

$$\begin{aligned} \text{FP}(\tilde{M}, R) &= \sum_{v \in \tilde{M}} (1 - R(v)) \tilde{M}(v) \\ &= \sum_{v \in \tilde{M}} \tilde{M}(v) - \sum_{v \in \tilde{M}} R(v) \tilde{M}(v) \\ &= 1 - \text{TP}(v \in \tilde{M}, R) , \end{aligned} \quad (20)$$

since each warped map is normalized to sum up to 1. It is therefore clear that the TP score is punished by high confidence voxels outside the reference whilst it benefits from high confidence voxels occurring within the reference. Note that in the case of a binarized reference the optimal TP is 1, which is achieved when all voxels of the predicted tract lie within the reference.

simple prior	learned prior	HCP		standard	
		left AF	right AF	left AF	right AF
	no prior	0.548±0.051	0.449±0.052	0.149±0.044	0.152±0.058
Subject-specific white matter prior					
WM		0.553±0.048	0.451±0.050	0.196±0.053	0.175±0.063
Learned prior including WM					
WM	standard _{WM}	0.246±0.029	0.223±0.038	0.196±0.035	0.178±0.047
WM	HCP _{WM}	0.633±0.023	0.501±0.033	0.569±0.049	0.472±0.051
Reference atlas as prior					
reference atlas		0.692±0.022	0.640±0.031	0.547±0.074	0.511±0.048

Table 1: Mean and standard deviation of TP scores over all subjects of the HCP (left) and standard data (right) for different priors (WM: white matter probability, HCP_{WM}: learned prior from the HCP data using WM as prior, standard_{WM}: learned prior from the standard data using WM as prior) for the arcuate fasciculus (AF).

5. Tractography results

We performed SPT for the arcuate fasciculus (AF), the cortico spinal tract (CST), the inferior fronto-occipital fasciculus (IFOF) and the fornix on both the HCP dataset and the standard dataset. On the AF and CST we used both the subject-specific priors described in Section 3.3 and the learned priors described in Section 3.5. On the IFOF and fornix we additionally included a waypoint prior, as described in Section 3.4. Furthermore, the IIT atlas was included as a prior for the fornix. Visualizations of population heatmaps for all tracts are shown in Figures 2, 3, 4 and 5, where each heatmap is represented by 20 opaque isosurfaces equally distributed over the range of heatmap values.

5.1. Tractography with a subject-specific simple prior

On the AF and CST we performed SPT using the white matter probability (WM), as defined in Section 3.3. This was compared to SPT without a prior. The results are found in lines 1–2 in Tables 1 and 2. The use of subject-specific WM prior generally gives a moderate increase in the TP score.

5.2. Learned priors

For the AF and CST, learned priors were created from both the standard and HCP datasets using SPT with the WM prior as defined in Section 3.5. These priors are hereafter denoted standard_{WM} and HCP_{WM}, respectively. Tractography was then performed upon both datasets using HCP_{WM} and standard_{WM} combined with the corresponding WM prior. The use of the HCP_{WM} resulted in the highest TP score for both tracts (see Tables 1 and 2) on both datasets, except on the left CST where the standard_{WM} prior achieves a higher score on the HCP data. For the right CST the standard_{WM} achieved the second best TP score. On the AF, it resulted in a similar TP score to that generated using only the white matter priors for the standard data, but has the lowest score on the HCP data.

simple prior	learned prior	HCP		standard	
		left CST	right CST	left CST	right CST
	no prior	0.503±0.030	0.416±0.038	0.434±0.052	0.323±0.039
Subject-specific white matter prior					
WM		0.503±0.030	0.416±0.038	0.437±0.055	0.334±0.036
Learned prior including WM					
WM	standard _{WM}	0.563±0.024	0.419±0.019	0.505±0.048	0.381±0.030
WM	HCP _{WM}	0.549±0.019	0.480±0.016	0.508±0.044	0.420±0.033
Reference atlas as prior					
reference atlas		0.656±0.014	0.561±0.016	0.552±0.040	0.471±0.036

Table 2: Mean and standard deviation of TP scores over all subjects of the HCP (left) and standard data (right) for different priors (WM: white matter probability, HCP_{WM}: learned prior from the HCP data using WM as prior, standard_{WM}: learned prior from the standard data using WM as prior) for the cortico spinal tract (CST).

5.3. Tractography with a waypoint prior

On the IFOF, SPT was initially performed with no prior, with the subject-specific WM prior and with the learned HCP prior (HCP_{WM}), see lines 1–3 in Table 3. Use of the WM prior resulted in a lower score than with no prior on the HCP data, and a similar score on the standard data (see Table 3). The learned HCP prior decreased the TP score on the HCP dataset. Note, however, that on the standard dataset, the learned HCP prior significantly increased the TP score on the left IFOF.

To guide the tractography, we introduced a waypoint (WP) prior as described in Section 4.2. Use of the WP prior alone (see fourth line of Table 3) achieved a higher TP score than the previously used priors without a waypoint, with the exception of that for the left IFOF in the standard dataset. Including the subject-specific WM prior together with the WP prior gave a similar small improvement in TP scores as observed in the previous tracts, see line 5 in Table 3.

Next we derived new learned priors for each dataset, obtained using the combination of WM and WP priors. These are termed standard_{WM,WP} and HCP_{WM,WP}. Once again we saw an improvement in the TP score analogous to that observed on the previous tracts (see lines 6–7 in Table 3). Finally, we performed tractography using the HCP_{WM,WP} prior, but this time *without* the subject-specific waypoint prior, and note that the performance was retained on both datasets, see line 8 in Table 3.

5.4. Tractography with an independent atlas prior

Similar to learned priors that reflect a population heatmap of the tract location, it is also possible to use a white matter tract atlas that reflects a probability. Here, we used the IIT atlas (see Section 4), together with a waypoint prior and the WM probability, as a prior for the fornix. The combination of all three priors achieved a higher TP score on both datasets than any of the following alternatives: no prior, a subject-specific WM prior, or the learned HCP prior (HCP_{WM}). See lines 1–4 in Table 4.

simple prior	learned prior	HCP		standard	
		left IFOF	right IFOF	left IFOF	right IFOF
	no prior	0.118±0.088	0.096±0.053	0.081±0.053	0.047±0.021
WM		0.104±0.081	0.078±0.042	0.082±0.049	0.050±0.021
WM	HCP _{WM}	0.070±0.062	0.087±0.033	0.222±0.080	0.098±0.049
		Waypoint (WP)			
WP		0.278±0.051	0.255±0.051	0.188±0.048	0.191±0.052
WM, WP		0.285±0.043	0.268±0.049	0.188±0.046	0.200±0.056
		Learned prior including WM and WP			
WM, WP	standard _{WM,WP}	0.290±0.028	0.282±0.027	0.217±0.036	0.248±0.038
WM, WP	HCP _{WM,WP}	0.313±0.015	0.300±0.014	0.296±0.028	0.289±0.025
		Learned prior without WP including WM			
WM	HCP _{WM,WP}	0.313±0.015	0.300±0.014	0.296±0.028	0.290±0.027
		Reference atlas as prior			
reference atlas		0.463±0.021	0.511±0.026	0.552±0.040	0.434±0.039

Table 3: Mean and standard deviation of TP scores over all subjects of the HCP (left) and standard data (right) for different priors (WM: white matter probability, WP: two-ROI waypoint mask, HCP_{WM}: learned prior from the HCP data using WM as prior, HCP_{WM,WP}: learned prior from the HCP data using WM and WP as prior, standard_{WM,WP}: learned prior from the standard data using WM and WP as prior) for the inferior fronto-occipital fasciculus (IFOF).

Similar to the experiments on the IFOF, new learned priors were derived for each dataset by using the combination of WM, WP and the IIT atlas as priors. These learned priors are denoted standard_{WM,WP,IIT} and HCP_{WM,WP,IIT}. While we saw an improvement in the TP score for HCP_{WM,WP,IIT} on both datasets, there was only a slight improvement on the standard data using standard_{WM,WP,IIT} (see lines 5–6 in Table 4). Finally, we performed tractography using the HCP_{WM,WP,IIT} prior, but this time *without* the subject-specific waypoint prior and the IIT atlas, and note that the performance was retained on the HCP data, but not the standard data see line 7 in Table 4.

5.5. Comparison to TRACULA

As a comparison, TRACULA (Yendiki et al., 2011) was performed on the HCP data and compared to the results for the CST with our method using the WM prior and the learned prior from the HCP data (HCP_{WM}). To construct a population heatmap and compute the TP score for the TRACULA results, the posterior distribution of the tract was treated in the same way as the confidence maps. The corresponding TP scores and visualization for the left and right CST are shown in Figure 6.

6. Discussion

In this paper we presented a shortest-path tractography (SPT) framework which is able to incorporate prior knowledge about tract location. Our framework supports frequently used

simple prior	learned prior	HCP		standard	
		left fornix	right fornix	left fornix	right fornix
	no prior	0.117±0.018	0.106±0.029	0.069±0.017	0.039±0.016
WM		0.127±0.016	0.127±0.039	0.090±0.015	0.048±0.021
WM	HCP _{WM}	0.133±0.013	0.120±0.037	0.123±0.016	0.093±0.019
		Waypoint (WP) and IIT atlas			
WM, WP, IIT atlas		0.335±0.032	0.526±0.031	0.175±0.088	0.336±0.105
		Learned prior including WM, WP and IIT atlas			
WM, WP, IIT atlas	standard _{WM,WP,IIT}	0.328±0.030	0.503±0.029	0.187±0.079	0.327±0.105
WM, WP, IIT atlas	HCP _{WM,WP,IIT}	0.339±0.025	0.545±0.026	0.270±0.069	0.417±0.128
		Learned prior without WP including WM			
WM	HCP _{WM,WP,IIT}	0.320±0.024	0.545±0.024	0.176±0.074	0.189±0.117
		Reference atlas as prior			
reference atlas		0.537±0.048	0.625±0.042	0.353±0.044	0.361±0.055

Table 4: Mean and standard deviation of TP scores over all subjects of the HCP (left) and standard data (right) for different priors (WM: white matter probability, WP: one-ROI waypoint mask, HCP_{WM}: learned prior from the HCP data using WM as prior, HCP_{WM,WP,IIT}: learned prior from the HCP data using WM, WP and the IIT atlas as prior, standard_{WM,WP,IIT}: learned prior from the standard data using WM, WP and the IIT atlas as prior) for the fornix.

prior information such as waypoints or white matter probabilities, and in particular is able to generate and use a prior learned from previously performed tractography experiments.

6.1. Simple subject-specific prior: white matter probability

We demonstrated how simple subject-specific priors, such as white matter (WM) probability, typically lead to a moderate improvement of the overlap score. Such priors have the inherent advantage of being easily determined automatically by conventional image processing software. The WM prior is similar to the WM-weighting schemes adapted by Iturria-Medina et al. (2007) in order to mollify the partial-volume effects that cause inappropriate attraction to CSF regions (Fuster et al., 2014). The WM prior improves tractography in its own right when measured by the TP score on some tracts (see Tables 1 and 2), although the improvement is modest. It is also clear from a visual inspection of the population heatmaps in Figures 2 and 3 that the inclusion of simple priors helps to clean up the tractography output. For instance, for the CST on the standard data, paths that incorrectly cross to the contralateral hemisphere are removed (see red arrows in Figure 3). For the use of subject-specific priors on the IFOF and fornix see Section 6.2 and Section 6.3 respectively.

The WM prior has most effect on the standard data as it achieves slightly better TP scores on all tracts compared to no prior (see Tables 1 and 2). On the HCP data, on the other hand, the WM prior is only slightly better for the AF. This could be explained by the lower image resolution and stronger partial volume effects of the standard data, since these are precisely the effects which are accounted for by the WM prior.

6.2. Binary prior: waypoints

For long tracts that intersect with or are close to other tracts, waypoint regions are commonly used to guide the tractography (Wakana et al., 2007). We showed that waypoints can also be included in our framework as a binary prior. We illustrated the use of a waypoint prior by performing tractography on the inferior fronto-occipital fasciculus (IFOF).

It is evident both from the TP scores in Table 3 and a visual inspection of the results in Figure 4 that the tractography fails without the use of a waypoint. The inclusion of a waypoint stops the tractography from following incorrect paths, marked by red arrows in Figure 4, except for a few, low confidence paths. Table 3 also shows that, as for the other tracts, the TP score is moderately improved when using the waypoint in combination with subject-specific priors like the WM prior.

Note that the TP score for the left IFOF on the standard data is high when the HCP prior is applied without including a waypoint (see line 3 in Table 3). Visual inspection of Figure 4 shows that the used HCP prior contains three high confidence regions, where one of them overlaps with the reference (see green arrows in Figure 4). As a result, tractography results on the standard data lie within the same region. However when applying the HCP prior, it is evident from visual inspection of the resulting population heatmap in Figure 4, that in several subjects wrong paths are also detected with substantial confidence (see red arrow). These paths are removed by including the waypoint prior.

6.3. Independent tract atlas as a prior

Since our framework allows the use of any spatial prior, a white matter tract atlas can also be used as a prior. We showed the application of such an independent atlas prior combined with a waypoint prior on the fornix (see Figure 5 and Table 4).

On the HCP dataset, the combination of these priors is enough to ensure correct delineation of the tract in almost all subjects. The addition of the learned HCP prior removes the remaining false positives, as seen in the fifth column of Figure 5. As shown in the sixth column, the HCP prior learned with the atlas and waypoint is strong enough to produce correct delineation even without using the atlas and waypoint, in a similar fashion as we observed on the IFOF.

On the standard data, however, when using the waypoint and atlas there are still incorrect paths included with a considerable confidence (see red arrows in Figure 5). As a consequence, the learned prior from the standard data does not perform as well as the learned prior from the HCP data, and the incorrect paths are not removed. The learned HCP prior improves the delineation, but it remains necessary to use the waypoint and IIT prior on the standard data, also in combination with the learned HCP prior (see fifth and sixth column in Figure 5).

6.4. Path score and heatmaps

The SPT returns a score for each path, which reflects how well the path is supported by both the underlying fODFs and the prior. Confidence maps created from these path scores then represent a subject-wise confidence of the location of the tract (see Equation 8). These maps are similar to heatmaps generated by conventional fibre-tracking methods. However, while traditional heatmaps count the fibres passing through a voxel, our confidence maps are weighted sums of paths passing through the voxel. Since the confidence map is not normalized, it is not a real probability, but after normalization it can be interpreted as a probability density function for the probability of the tract passing through any given voxel.

In addition, population heatmaps can be created from subject normalized confidence maps, thereby incorporating the variation within a dataset (see Section 3.5). Such population heatmaps are similar to the tract probability overlap maps created for the major fibre bundles in studies that aim to generate a white matter tract atlas (de Schotten et al., 2011; Varentsova et al., 2014), but here they also take subject-wise confidences into account.

6.5. Learned prior

We applied the population heatmaps as population-based learned spatial priors for tractography. In general, a learned prior can be constructed from tractography results created by any tractography method. Such priors can be generated from all investigated subjects, creating a study-specific prior, or from an independent representative population, creating an independent prior.

Note that the use of a study-specific prior is *not* a case of training on a corresponding test set, as the learned prior is not trained with the help of an annotated ground truth, but is entirely unsupervised. The application of the reference atlas as prior, when the same reference is used to compute the TP score, is a case of training on the test set. This is known, from machine learning, to lead to an overestimated performance score. The corresponding results should be therefore an upper benchmark (see also Section 6.10).

First, we observe that using a study-specific learned prior reduces the standard deviation and thereby the variation of the tractography results, both on the standard data and the HCP data, see Tables 1, 2 and 3. This effect is also evident from the visual inspection of the tractography results, see Figures 2, 3 and 4, where low confidence tracts, most of which are incorrect, are consistently removed by the learned prior. For example, on the HCP dataset, for the AF, the use of the study-specific prior concentrates the paths in the high confidence regions (see green arrows in Figure 2). For the CST the learned prior removes paths that follow the CST on the contralateral hemisphere (see red circles in Figure 3). Even though restricting the source and target ROIs to the ipsilateral hemisphere would probably have the same effect, the learned prior is generated automatically and is robust for such errors in the ROI definition.

On the standard dataset, the performance of the study-specific learned prior is variable. The TP score is increased on the CST (see Table 2) when compared to the WM prior. For the AF, however, there is no significant improvement of the TP score, see Table 1, compared to using subject-specific priors like WM. A visual inspection of Figure 2 reveals that the paths obtained by tractography with only a WM prior on the standard dataset follow white matter tracts inferior of the “true” location of the AF (see Figure 1). It is therefore to be expected that a learned prior constructed from these results would not help tractography.

However, the high quality HCP prior consistently increased the TP score for the standard dataset compared to all other results (see Tables 1, 2, 3 and 4). This is supported by visual inspection of Figures 2, 3, 4 and 5. In particular, the TP score for the AF on the standard dataset was lifted to the same level as on the HCP dataset, where inspection of Figure 2 shows that the HCP prior guides the tractography to the correct tract location even on the standard dataset. While using a waypoint to guide the tractography might have the same effect, the learned prior has the advantage that it is obtained automatically and without manual intervention.

The results on the IFOF and fornix (see Table 3 and Figure 4) show that a learned prior constructed with a waypoint or other independent prior can render the use of these priors in subsequent tractography unnecessary. This clearly illustrates the importance and potential

usefulness of a high quality learned prior. Note that it is also possible to manually define high precision waypoints for a small subset of the dataset and apply a learned prior from that subset to the whole data.

6.6. Comparison to TRACULA

We compared our best result on the CST based on a learned prior from the HCP data to TRACULA (Yendiki et al., 2011), which also uses a trained prior. While the TP scores are lower for the TRACULA results, the visual inspection looks equally plausible (see Figure 6). The TP score might be lower because different ROIs are used in TRACULA. Note that we could only compare results for the CST, as the other tracts in this study are not included in TRACULA.

6.7. Learned prior for clinical data

Tractography is often a pre-processing tool for analysis of white matter changes in patients with a disease or pathology. In this paper, we have demonstrated the creation and application of a learned prior on two healthy populations, and care should be taken when applying an independent prior created from a healthy population in order to study patients with pathology. An independent prior created from healthy subjects should only be applied if the tracts of the analysed population are *not* assumed to be drastically disrupted or remodelled, since tracts will be encouraged to stay within the typical tract location of the “healthy” prior.

In cases where disease is expected to generate disruptions or remodelling of fibre pathways, a study-specific learned prior would not enforce the same constraints. However, for a study-specific prior to be effective, it would need the pathological effect on the tract to be spatially similar across patients.

6.8. Application of anatomical or learned priors

An anatomical (or indeed any) prior comes with both benefits and downsides. Whilst they aid in the delineation of known pathways, they inherently also provide bias away from the evidence provided by the individual subject. As such, the application of priors is always a balancing-act between the discovery of possibly-unique results in the individual versus the robustness of finding group-wise commonalities. In practice, this balance could be changed by giving different weights to the prior and the individual information from the diffusion data. In this paper, we give equal weight to these terms, and as observed in Figure 7, subject-specific variability is retained, although concentrated near the high-confidence region of the prior.

6.9. Dependence on the delineation of the endpoint regions

The SPT method requires a source and a target region and thus depends on the delineation of those regions. Therefore it is important that the regions are chosen to reflect reasonable endpoints of the tracts. However, the SPT is robust to small errors or changes in the delineation of the ROIs, as can be seen in the low confidence for tracts in the contralateral hemisphere for the CST results (see Figure 3). Here, these false tracts occur due to the overly large seed region in the brain stem which incorporates fibres from both hemispheres. Furthermore, since a path is always found between any pair of voxels, the SPT also returns good results when the ROIs are small, as for the fornix (see Supplementary Table 2).

6.10. Evaluation criterion

We chose to evaluate our results through a combination of visual inspection and quantitative comparison with the reference tract atlas from de Schotten et al. (2011) described in Section 4. We visually inspected the computed shortest paths in single subjects, where we saw a concentrating effect of the use of different priors. However, visual inspection of shortest paths does not give a summarized population evaluation. We therefore generated population heatmaps (see Section 3.5) for population-level visualization of tractography results. Since the population heatmap is constructed based on the scalar path-confidence scores, it automatically down-weights low confidence paths and enhances high confidence paths. While it is evident from the heatmaps in Figures 2 and 3 that some of the followed paths were incorrect, it is also evident that these are appropriately rated as low confidence paths.

We find that when accompanied with qualitative visual inspection, a quantitative evaluation of tract co-occurrence on a population of subjects can be very valuable. For the quantitative comparison we introduced the true positive (TP) score for measuring the overlap of the population heatmaps with the reference described in Section 4.4. The TP simultaneously measures true and false positives and is able to take the probability of the reference into account. In this way, a full TP score of 1 can only be obtained by a tract which lies within the region where the reference is 1. The TP score is therefore rather conservative, as even a perfect tractography result can not be expected to generate a score of 1 on tracts where the reference has a high level of variation.

To obtain an upper benchmark TP score for every tract and dataset, we performed tractography using the reference atlas as a prior. This should lead to a near optimal tractography result, while still considering subject-specific variation. As expected, the resulting TP scores are higher than or similar to the highest scores with any other prior. Furthermore, the TP score differs for the different tracts and datasets (see last line in Tables 1, 2, 3 and 4), indicating that the TP score should not be considered an absolute measure of quality, but a relative one.

Note further that there are differences in the TP score between the left and right hemisphere, which is especially true for the AF and CST. These differences are most likely caused by asymmetry of the reference, and are also observed when applying the reference atlas as prior. This imbalance does not make a difference for comparison of different methods, as the comparison should be done on the different sides separately.

Any overlap measure depends on accurate image registration, which is known to be imperfect. Furthermore, comparing the tractography results to a reference created from a significantly different population, e.g. with respect to age or gender, can have a negative effect on the score, as the registered reference is not necessarily anatomically correct. This is especially true when the analysed population consists of patients where severe changes of the tracts are expected. One can, however, use the comparison score to quantify the difference between computed tracts and the reference and thereby detect patients where the tract differs from the reference.

6.11. Future extensions of the method

In our experiments and definition, the prior contributes equally to the path likelihood as the weight of a connection (see Equations 11 and 14). However, it would also be possible to introduce a weighting of the prior in Equation 11 resulting in $P(e|w_E(e)) \propto w_E(e)P(e)^k$, where k is the weight of the prior. Setting k to 1 would result in the same weighting used in this paper, while a $k > 1$ would weight the prior more than the edge weight and $0 < k < 1$

would weight the prior less, respectively. While an adaptation of the formulas to include k is straightforward, a detailed study of the effect of different values for k is out of the scope of this study.

It is known that fibre bundles have a direction and therefore can be distinguished by being afferent or efferent. Although the DWI method itself is undirected, it could, in theory, be possible to include anatomical (topological) based directionality information into the tractography. This would require modelling the diffusion information as a directed graph with asymmetric edges and hence reformulation of the SPT method. As such, even though it may be interesting to explore directional-dependent differences in the transition between voxels, this would need to be considered as future work.

7. Conclusion

In this paper we derived a shortest-path tractography framework which includes prior information about the spatial location of a tract. We demonstrated that the use of such prior information improves the accuracy and robustness of the delineated tracts. We especially showed how this prior information can be derived automatically, thereby permitting our tractography to learn based on previous tractography results. This allows prior results, ROIs and tract masks to be automatically applied to future studies, alleviating the need for tedious and often error-prone manual intervention. Here, we applied learned priors obtained by running the same tractography algorithm with a simple prior only. However, our framework is able to integrate a prior learned from any tractography algorithm or from previously created tractography results.

In summary, we hope that the work presented herein will enable tractography and tract-based analysis of large populations to be made more feasible by permitting prior domain knowledge to be automatically included in the tract delineation of individual subjects.

Acknowledgements

Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

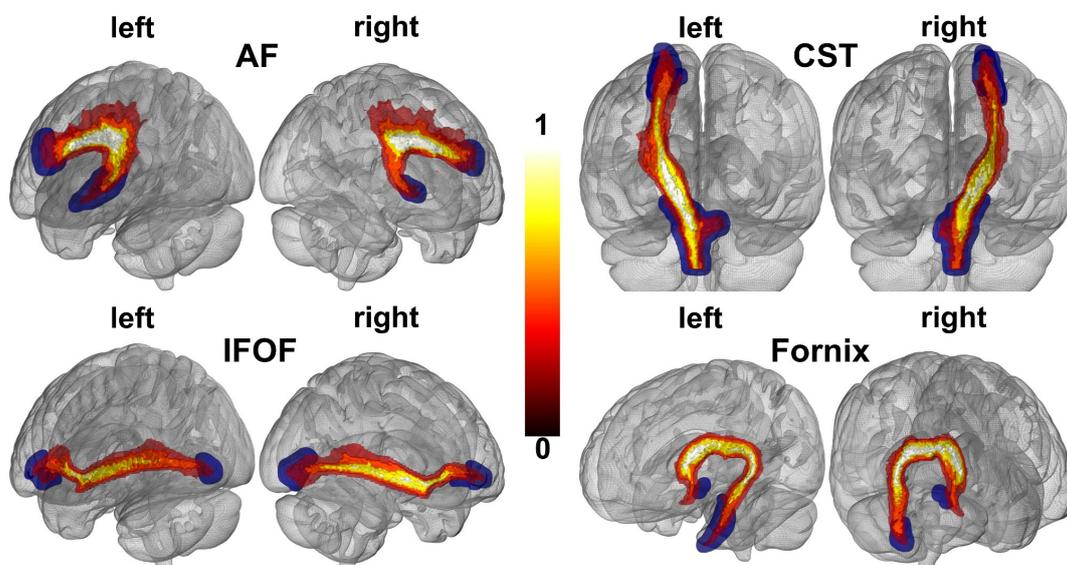


Figure 1: Reference tract atlas. The atlas obtained from de Schotten et al. (2011) is visualized as heatmap (white=high overlap, red=low overlap) together with the seed regions (blue) used for tractography. Shown are all tracts used in this study: the arcuate fasciculus (AF), the cortico spinal tract (CST), the inferior fronto-occipital fasciculus (IFOF) and the fornix. The seed regions were spatially smoothed by an isotropic Gaussian filter and are shown opaque for visualization purposes.

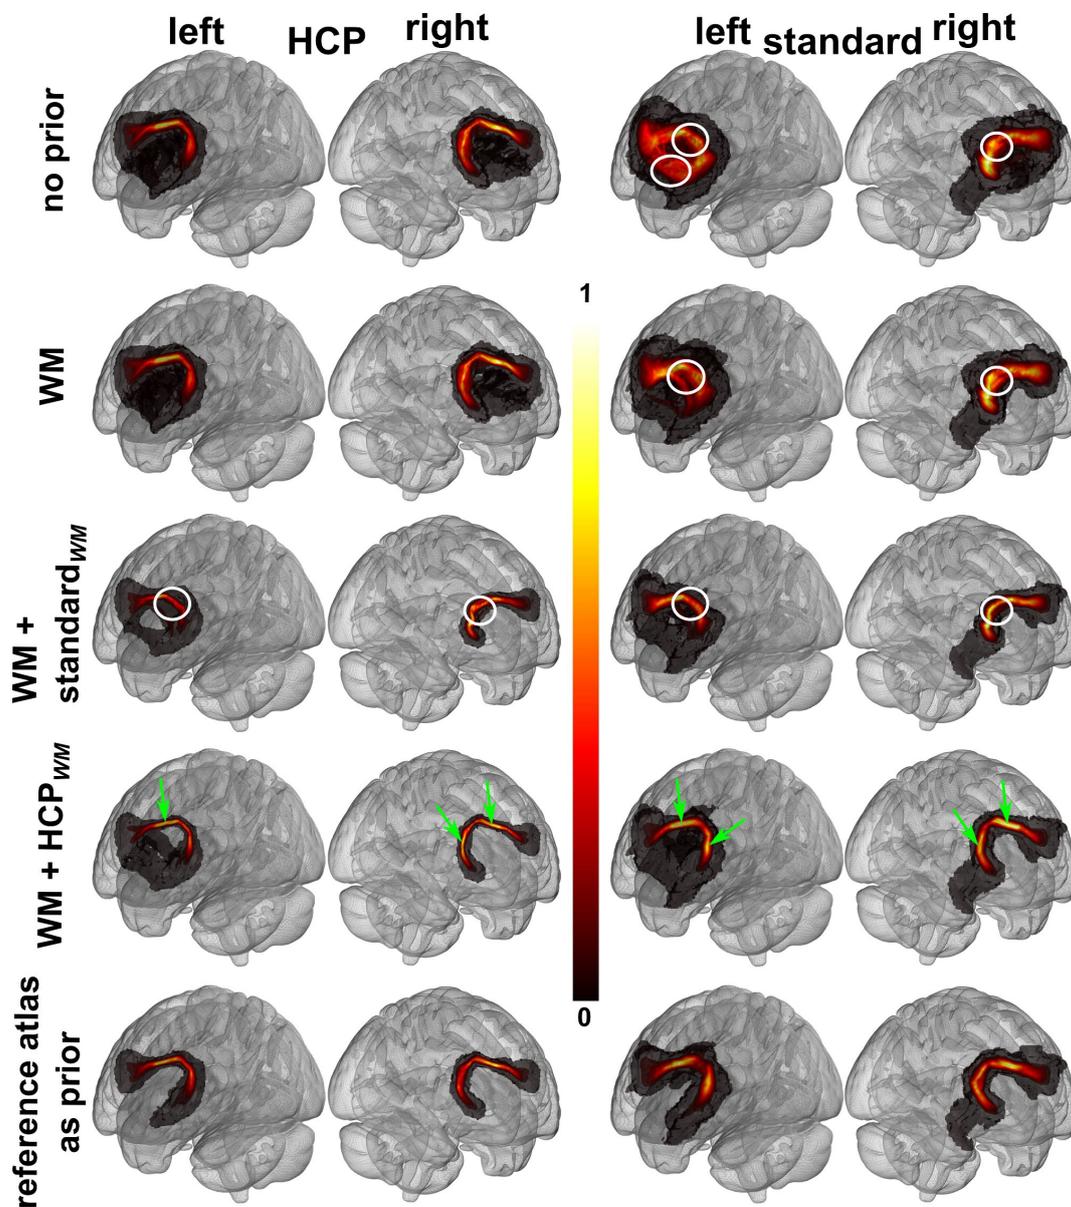


Figure 2: Population heatmaps of the arcuate fasciculus (AF). Heatmaps (white=high confidence, red=low confidence) were created as average from individual results in MNI space (see Equation 16) and are shown for results using different priors both on the HCP data (left) and the standard data (right). Learned priors (rows 3–4) were created from the results including WM (second row). White circles mark regions in the heatmaps inferior to the “true” location of the AF, while green arrows mark the high confidence areas that are enhanced by the learned HCP prior.

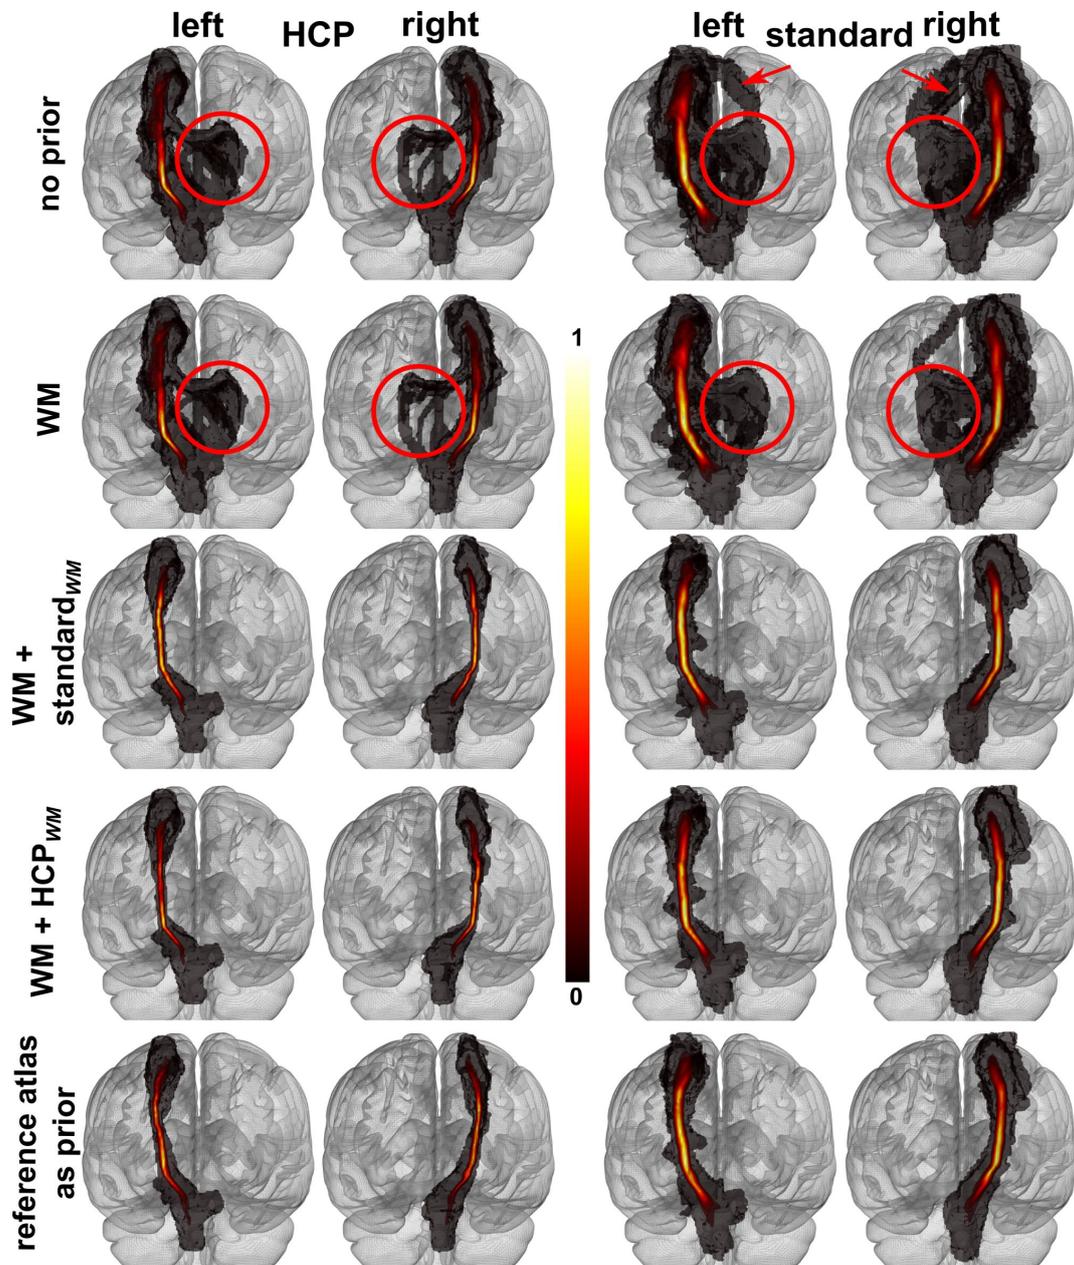


Figure 3: Population heatmaps of the cortico spinal tract (CST). Heatmaps (white=high confidence, red=low confidence) were created as average from individual results in MNI space (see Equation 16) and are shown for results using different priors both on the HCP data (left) and the standard data (right). Learned priors (rows 3–4) were created from the results including WM (second row). Red arrows mark paths that are removed with the application of subject-specific priors and red circles mark paths removed by a learned prior.

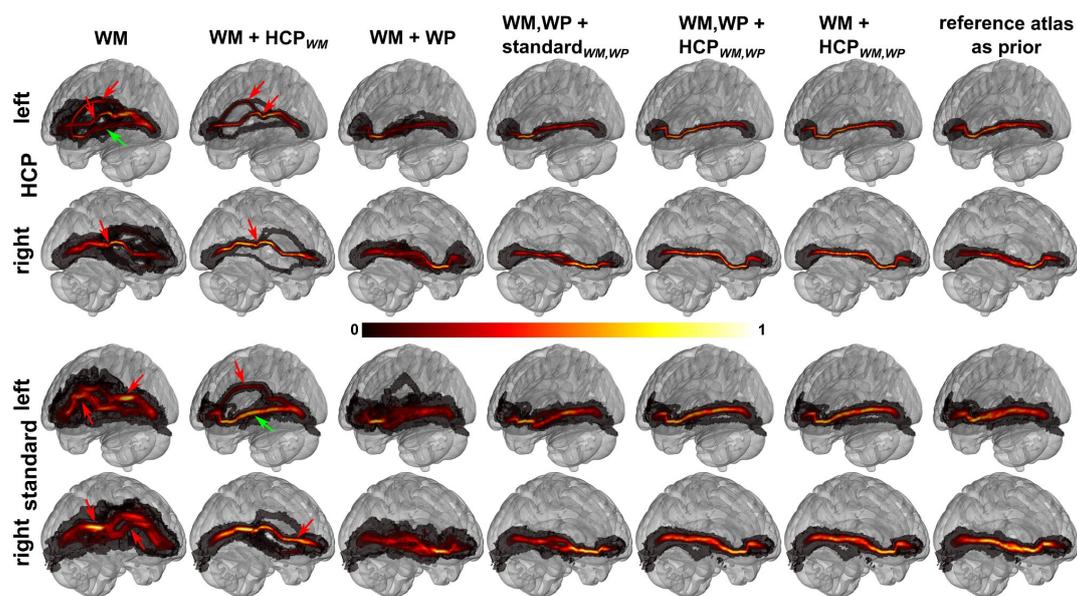


Figure 4: Population heatmaps of the inferior fronto-occipital fasciculus (IFOF). Heatmaps (white=high confidence, red=low confidence) were created as average from individual results in MNI space (see Equation 16) and are shown for results using different priors both on the HCP data (top) and the standard data (bottom). Learned priors (see columns 2, 4, 5 and 6) were created from the results including WM (first column) or from results using the waypoint and WM prior (third column). Red arrows mark wrong paths that are removed by the use of a waypoint prior, while green arrows mark “true” paths found without the WP prior.

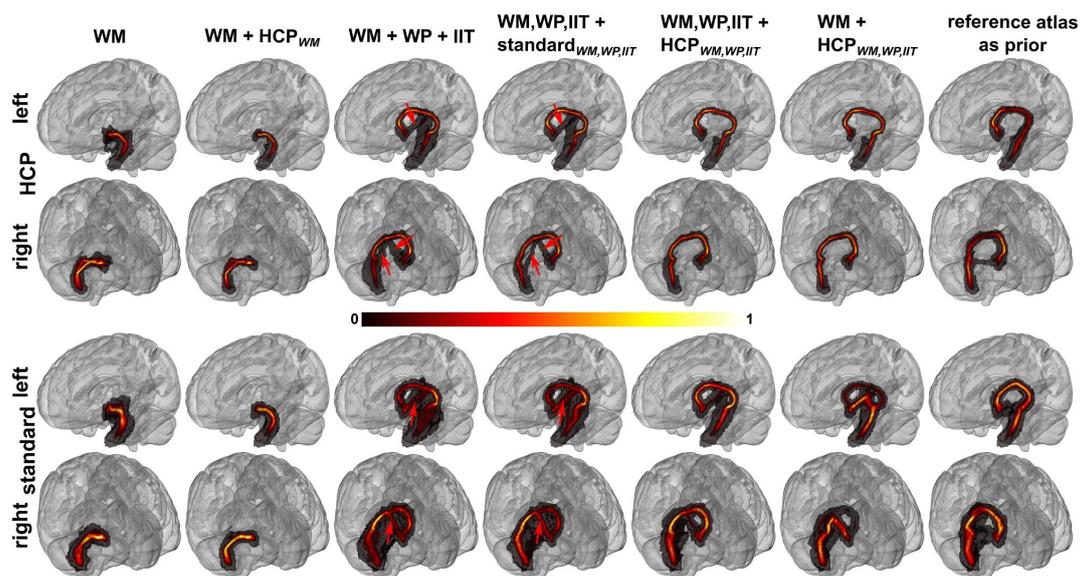


Figure 5: Population heatmaps of the fornix. Heatmaps (white=high confidence, red=low confidence) were created as average from individual results in MNI space (see Equation 16) and are shown for results using different priors both on the HCP data (top) and the standard data (bottom). Learned priors (see columns 2, 4, 5 and 6) were created from the results including WM (first column) or from results using the waypoint, WM and IIT atlas as prior (third column). Red arrows mark wrong paths that are removed or reduced by the use of the learned HCP prior (see column 5).

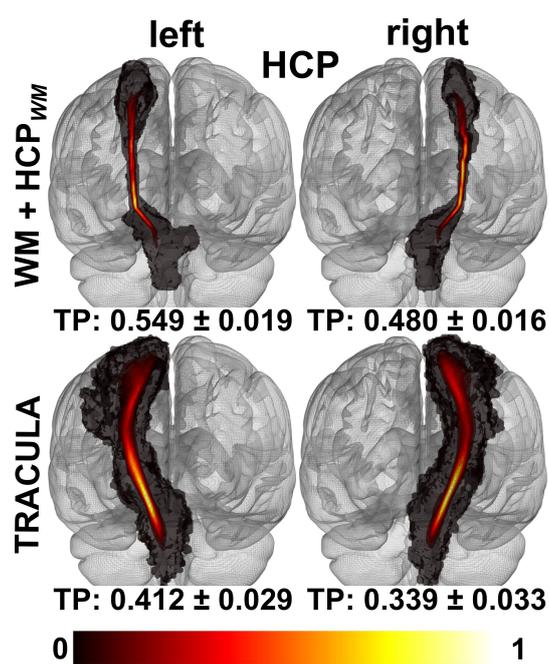


Figure 6: Comparison with TRACULA on the cortico spinal tract (CST). Shown are population heatmaps (white=high confidence, red=low confidence) created as average from individual results in MNI space (see Equation 16) for the HCP data together with the TP scores. Tractography results are compared between SPT using a learned prior from the HCP data (top) and TRACULA (Yendiki et al., 2011) (bottom).

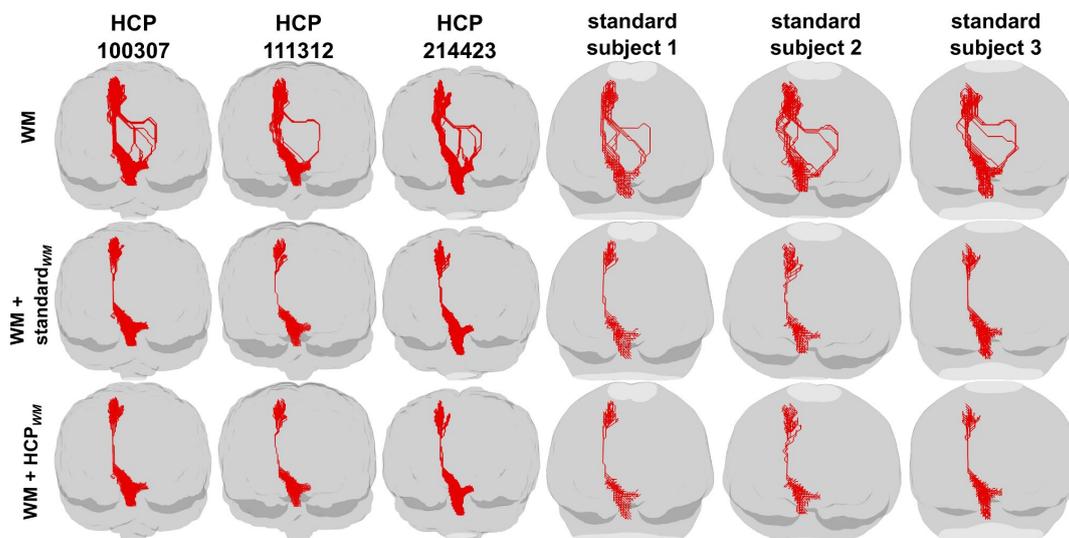


Figure 7: Single-subject trajectories for the left cortico spinal tract (CST). Trajectories are visualized three HCP subjects (left) and three subjects from the standard data (right) using different priors: the white matter (WM) prior, the WM prior with a prior learned from the HCP data and the WM prior learned from the standard data. Shown are all edges that lie on at least one of the found shortest paths. All edges are coloured equally, and hence are independent of the score of the paths that travel along that edge.

References

- Abhinav, K., Yeh, F.-C., Pathak, S., Suski, V., Lacomis, D., Friedlander, R. M., Fernandez-Miranda, J. C., Nov 2014. Advanced diffusion MRI fiber tracking in neurosurgical and neurodegenerative disorders and neuroanatomical studies: A review. *Biochim Biophys Acta* 1842 (11), 2286–2297.
- Alexander, A. L., Hurley, S. A., Samsonov, A. A., Adluru, N., Hosseinbor, A. P., Mossahebi, P., Tromp, D. P. M., Zakszewski, E., Field, A. S., 2011. Characterization of cerebral white matter properties using quantitative magnetic resonance imaging stains. *Brain Connect* 1 (6), 423–446.
- Andersson, J. L. R., Jenkinson, M., Smith, S., 2010. Non-linear registration, aka spatial normalisation. Tech. rep., FMRIB.
- Assaf, Y., Alexander, D. C., Jones, D. K., Bizzi, A., Behrens, T. E. J., Clark, C. A., Cohen, Y., Dyrby, T. B., Huppi, P. S., Knoesche, T. R., Lebihan, D., Parker, G. J. M., Poupon, C., consortium, C. O. N. N. E. C. T., Anaby, D., Anwender, A., Bar, L., Barazany, D., Blumenfeld-Katzir, T., De-Santis, S., Duclap, D., Figini, M., Fisci, E., Guevara, P., Hubbard, P., Hofstetter, S., Jbabdi, S., Kunz, N., Lazeyras, F., Lebois, A., Liptrot, M. G., Lundell, H., Mangin, J.-F., Dominguez, D. M., Morozov, D., Schreiber, J., Seunarine, K., Nava, S., Poupon, C., Riffert, T., Sasson, E., Schmitt, B., Shemesh, N., Sotiropoulos, S. N., Tavor, I., Zhang, H. G., Zhou, F.-L., Oct 2013. The connect project: Combining macro- and micro-structure. *NeuroImage* 80, 273–282.
- Behrens, T. E. J., Berg, H. J., Jbabdi, S., Rushworth, M. F. S., Woolrich, M. W., Jan 2007. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage* 34 (1), 144–155.
- Benedetti, F., Absinta, M., Rocca, M. A., Radaelli, D., Poletti, S., Bernasconi, A., Dallspezia, S., Pagani, E., Falini, A., Copetti, M., Colombo, C., Comi, G., Smeraldi, E., Filippi, M., Jun 2011. Tract-specific white matter structural disruption in patients with bipolar disorder. *Bipolar Disord* 13 (4), 414–424.
- Catani, M., Howard, R. J., Pajevic, S., Jones, D. K., Sep 2002. Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *NeuroImage* 17 (1), 77–94.
- Connally, E. L., Ward, D., Howell, P., Watkins, K. E., Apr 2014. Disrupted white matter in language and motor tracts in developmental stuttering. *Brain Lang* 131, 25–35.
- Davis, S. W., Dennis, N. A., Buchler, N. G., White, L. E., Madden, D. J., Cabeza, R., Jun 2009. Assessing the effects of age on long white matter tracts using diffusion tensor tractography. *NeuroImage* 46 (2), 530–541.
- de Schotten, M. T., Ffytche, D. H., Bizzi, A., Dell’Acqua, F., Allin, M., Walshe, M., Murray, R., Williams, S. C., Murphy, D. G. M., Catani, M., Jan 2011. Atlasing location, asymmetry and inter-subject variability of white matter tracts in the human brain with MR diffusion tractography. *NeuroImage* 54 (1), 49–59.

- Desikan, R. S., Sgonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., Killiany, R. J., Jul 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 31 (3), 968–980.
- Dijkstra, E. W., 1959. A note on two problems in connexion with graphs. *Numerische Mathematik* 1 (1), 269–271.
- Eickhoff, S. B., Paus, T., Caspers, S., Grosbras, M.-H., Evans, A. C., Zilles, K., Amunts, K., Jul 2007. Assignment of functional activations to probabilistic cytoarchitectonic areas revisited. *NeuroImage* 36 (3), 511–521.
- Essen, D. C. V., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., Ugurbil, K., WU-Minn HCP Consortium, Oct 2013. The WU-Minn Human Connectome Project: An overview. *NeuroImage* 80, 62–79.
- Fuster, A., Tristan-Vega, A., Haije, T., Westin, C.-F., Florack, L., 2014. A novel Riemannian metric for geodesic tractography in DTI. In: Schultz, T., Nedjati-Gilani, G., Venkataraman, A., O'Donnell, L., Panagiotaki, E. (Eds.), *Computational Diffusion MRI and Brain Connectivity. Mathematics and Visualization*. Springer International Publishing, pp. 97–104.
- Galantucci, S., Tartaglia, M. C., Wilson, S. M., Henry, M. L., Filippi, M., Agosta, F., Dronkers, N. F., Henry, R. G., Ogar, J. M., Miller, B. L., Gorno-Tempini, M. L., Oct 2011. White matter damage in primary progressive aphasia: a diffusion tensor tractography study. *Brain* 134 (Pt 10), 3011–3029.
- Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., van der Walt, S., Descoteaux, M., Nimmo-Smith, I., Dipy Contributors, 2014. Dipy, a library for the analysis of diffusion MRI data. *Front Neuroinform* 8, 8.
- Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Essen, D. C. V., Jenkinson, M., WU-Minn HCP Consortium, Oct 2013. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* 80, 105–124.
- Golestani, A. M., Miles, L., Babb, J., Castellanos, F. X., Malaspina, D., Lazar, M., Nov 2014. Constrained by our connections: white matter's key role in interindividual variability in visual working memory capacity. *J Neurosci* 34 (45), 14913–14918.
- Greve, D. N., Fischl, B., Oct 2009. Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* 48 (1), 63–72.
- Hagmann, P., Jonasson, L., Maeder, P., Thiran, J.-P., Wedeen, V. J., Meuli, R., Oct 2006. Understanding diffusion MR imaging techniques: from scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *Radiographics* 26 Suppl 1, S205–S223.
- Haugberg, S., Schober, M., Liptrot, M., Hennig, P., Feragen, A., 2015. A random Riemannian metric for probabilistic shortest-path tractography. *MICCAI*.

- Iturria-Medina, Y., Canales-Rodríguez, E. J., Melie-García, L., Valdés-Hernández, P. A., Martínez-Montes, E., Alemán-Gómez, Y., Sánchez-Bornot, J. M., Jul 2007. Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *NeuroImage* 36 (3), 645–660.
- Iwata, N. K., Kwan, J. Y., Danielian, L. E., Butman, J. A., Tovar-Moll, F., Bayat, E., Floeter, M. K., Sep 2011. White matter alterations differ in primary lateral sclerosis and amyotrophic lateral sclerosis. *Brain* 134 (Pt 9), 2642–2655.
- Jbabdi, S., Woolrich, M. W., Andersson, J. L. R., Behrens, T. E. J., Aug 2007. A bayesian framework for global tractography. *NeuroImage* 37 (1), 116–129.
- Jenkinson, M., Bannister, P., Brady, M., Smith, S., Oct 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17 (2), 825–841.
- Jenkinson, M., Pechaud, M., Smith, S., 2005. BET2: MR-based estimation of brain, skull and scalp surfaces. In: Eleventh annual meeting of the organization for human brain mapping. Vol. 17.
- Lancaster, J. L., Tordesillas-Gutierrez, D., Martinez, M., Salinas, F., Evans, A., Zilles, K., Mazziotta, J. C., Fox, P. T., Nov 2007. Bias between MNI and Talairach coordinates analyzed using the ICBM-152 brain template. *Hum Brain Mapp* 28 (11), 1194–1205.
- Leemans, A., Jeurissen, B., Sijbers, J., Jones, D. K., 2009. ExploreDTI: a graphical toolbox for processing, analyzing, and visualizing diffusion MR data. In: 17th Annual Meeting of Intl Soc Mag Reson Med. p. 3537.
- Lenglet, C., Deriche, R., Faugeras, O., 2004. Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping. In: ECCV.
- Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., Parsons, L., Narr, K., Kabani, N., Goualher, G. L., Boomsma, D., Cannon, T., Kawashima, R., Mazoyer, B., Aug 2001. A probabilistic atlas and reference system for the human brain: International consortium for brain mapping (ICBM). *Philos Trans R Soc Lond B Biol Sci* 356 (1412), 1293–1322.
- McGrath, J., Johnson, K., O’Hanlon, E., Garavan, H., Leemans, A., Gallagher, L., 2013. Abnormal functional connectivity during visuospatial processing is associated with disrupted organisation of white matter in autism. *Front Hum Neurosci* 7, 434.
- O’Donnell, L., Haker, S., Westin, C.-F., 2002. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. In: MICCAI. pp. 459–466.
- Ramsøy, T. Z., Liptrot, M. G., Skimminge, A., Lund, T. E., Sidaros, K., Christensen, M. S., Baar, W., Paulson, O. B., Jernigan, T. L., Siebner, H. R., Sep 2012. Healthy aging attenuates task-related specialization in the human medial temporal lobe. *Neurobiol Aging* 33 (9), 1874–1889.

- Rojkova, K., Volle, E., Urbanski, M., Humbert, F., Dell'Acqua, F., de Schotten, M. T., Feb 2015. Atlasing the frontal lobe connections and their variability due to age and education: a spherical deconvolution tractography study. *Brain Struct Funct*.
- Schober, M., Kasenburg, N., Feragen, A., Hennig, P., Hauberg, S., 2014. Probabilistic shortest path tractography in DTI using Gaussian Process ODE solvers. *MICCAI 17 (Pt 3)*, 265–272.
- Sotiropoulos, S. N., Bai, L., Morgan, P. S., Constantinescu, C. S., Tench, C. R., Feb 2010. Brain tractography using Q-ball imaging and graph theory: Improved connectivities through fibre crossings via a model-based approach. *NeuroImage* 49 (3), 2444–2456.
- Sotiropoulos, S. N., Moeller, S., Jbabdi, S., Xu, J., Andersson, J. L., Auerbach, E. J., Yacoub, E., Feinberg, D., Setsompop, K., Wald, L. L., Behrens, T. E. J., Ugurbil, K., Lenglet, C., Dec 2013. Effects of image reconstruction on fiber orientation mapping from multichannel diffusion MRI: reducing the noise floor using SENSE. *Magn Reson Med* 70 (6), 1682–1689.
- Steinbach, R., Loewe, K., Kaufmann, J., Machts, J., Kollewe, K., Petri, S., Dengler, R., Heinze, H.-J., Vielhaber, S., Schoenfeld, M. A., Stoppel, C. M., 2015. Structural hallmarks of amyotrophic lateral sclerosis progression revealed by probabilistic fiber tractography. *J Neurol*, 1–14.
- Tang, P.-F., Ko, Y.-H., Luo, Z.-A., Yeh, F.-C., Chen, S.-H. A., Tseng, W.-Y. I., Jun 2010. Tract-specific and region of interest analysis of corticospinal tract integrity in subcortical ischemic stroke: reliability and correlation with motor function of affected lower extremity. *Am J Neuroradiol* 31 (6), 1023–1030.
- Tournier, J.-D., Calamante, F., Connelly, A., May 2007. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35 (4), 1459–1472.
- Varentsova, A., Zhang, S., Arfanakis, K., May 2014. Development of a high angular resolution diffusion imaging human brain template. *NeuroImage* 91, 177–186.
- Vorburger, R. S., Reischauer, C., Boesiger, P., Feb 2013. BootGraph: Probabilistic fiber tractography using bootstrap algorithms and graph theory. *NeuroImage* 66, 426–435.
- Voronoi, G., 1908. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Deuxième mémoire. Recherches sur les paralléloèdres primitifs. *Journal für die reine und angewandte Mathematik* 134, 198–287.
- Wakana, S., Caprihan, A., Panzenboeck, M. M., Fallon, J. H., Perry, M., Gollub, R. L., Hua, K., Zhang, J., Jiang, H., Dubey, P., Blitz, A., van Zijl, P., Mori, S., Jul 2007. Reproducibility of quantitative tractography methods applied to cerebral white matter. *NeuroImage* 36 (3), 630–644.
- Whitford, T. J., Kubicki, M., Pelavin, P. E., Lucia, D., Schneiderman, J. S., Pantelis, C., McCarley, R. W., Shenton, M. E., Jan 2015. Cingulum bundle integrity associated with delusions of control in schizophrenia: Preliminary evidence from diffusion-tensor tractography. *Schizophr Res* 161 (1), 36–41.

- Wozniak, J. R., Mueller, B. A., Bell, C. J., Muetzel, R. L., Lim, K. O., Day, J. W., Apr 2013. Diffusion tensor imaging reveals widespread white matter abnormalities in children and adolescents with myotonic dystrophy type 1. *J Neurol* 260 (4), 1122–1131.
- Xia, S., Li, X., Kimball, A. E., Kelly, M. S., Lesser, I., Branch, C., Nov 2012. Thalamic shape and connectivity abnormalities in children with attention-deficit/hyperactivity disorder. *Psychiatry Res* 204 (2-3), 161–167.
- Yendiki, A., Panneck, P., Srinivasan, P., Stevens, A., Zlei, L., Augustinack, J., Wang, R., Salat, D., Ehrlich, S., Behrens, T., Jbabdi, S., Gollub, R., Fischl, B., 2011. Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy. *Front Neuroinform* 5, 23.
- Zalesky, A., Oct 2008. DT-MRI fiber tracking: a shortest paths approach. *IEEE Trans Med Imaging* 27 (10), 1458–1471.
- Zhang, Y., Brady, M., Smith, S., Jan 2001. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging* 20 (1), 45–57.

5.2

Finding significantly connected voxels based on histograms of connection strengths

Niklas Kasenburg^a, Morten Vester Pedersen^a and Sune Darkner^a

^aDepartment of Computer Science, University of Copenhagen, Denmark

ABSTRACT

We explore a new approach for structural connectivity based segmentations of subcortical brain regions. Connectivity based segmentations are usually based on fibre connections from a seed region to predefined target regions. We present a method for finding significantly connected voxels based on the distribution of connection strengths. Paths from seed voxels to all voxels in a target region are obtained from a shortest-path tractography. For each seed voxel we approximate the distribution with a histogram of path scores. We hypothesise that the majority of estimated connections are *false-positives* and that their connection strength is distributed differently from *true-positive* connections. Therefore, an empirical null-distribution is defined for each target region as the average normalized histogram over all voxels in the seed region. Single histograms are then tested against the corresponding null-distribution and significance is determined using the false discovery rate (FDR). Segmentations are based on significantly connected voxels and their FDR. In this work we focus on the thalamus and the target regions were chosen by dividing the cortex into a prefrontal/temporal zone, motor zone, somatosensory zone and a parieto-occipital zone. The obtained segmentations consistently show a sparse number of significantly connected voxels that are located near the surface of the anterior thalamus over a population of 38 subjects.

1. INTRODUCTION

Subcortical brain regions, like the thalamus, are highly connected with the cortex of the human brain. It is possible to segment these subcortical brain regions based on connection strengths to different cortical areas.¹⁻³ Analysis of volume, fractional anisotropy (FA) or other changes in these subregions obtained by such segmentations can give insight into pathological effects of certain diseases, e.g. Huntington’s disease.³

Typically, such segmentations are computed based on the number of “fibres” connecting a seed region to a specific target region of interest (ROI)^{1,2} or on the strength of a single connection. A voxel in the region to be segmented is assigned to the target ROI to which it has the most connections or the highest connection strength. However, it is known that noise in the data can increase the amount of false-positive connections.^{4,5} Existing segmentation methods do not account for these connections and do not consider whether seed voxels are *significantly* connected. Here, we present a new approach for segmenting subcortical brain regions based on significantly connected voxels, where significance is determined from a voxel’s distribution of connection strengths.

2. METHODS

The framework for testing whether a seed voxel is significantly connected to a specific target ROI is based on the distribution of connection strengths of tracts or paths from the seed voxel to that target ROI. Paths between a seed voxel and the target ROI are computed using a shortest-path tractography method (see Section 2.2). Since this method returns a path to every voxel in the target ROI, we assume that most of these connections do not reflect an existing connection and are due to noise and poor resolution of the data, which is reflected by a low connection strength. Furthermore, we assume that for significantly connected voxels we observe a higher amount of strong connections than for non-significantly connected voxels.

Further author information: Niklas Kasenburg and Morten Vester Pedersen contributed equally to this work (contact N.K. or M.V.P. for correspondence)

N.K.: E-mail: niklas.kasenburg@di.ku.dk

M.V.P.: E-mail: mvp288@nyu.edu

Here, we therefore determine significance of a seed voxel by measuring whether the amount of high-scoring paths in its distribution of connection strengths is significantly different from a distribution that only consists of noise. In the following, we call the noise-distribution also the null-distribution. We further think of a voxel's distribution as a mixture of the null-distribution and a distribution of true, high-scoring connections, and therefore call it a mixture distribution. The goal is now to determine how similar or different a voxel's distribution is from the null-distribution, that is how strong the contribution of the null-distribution is in a voxel's mixture distribution.

Before explaining how we achieve this, we briefly revisit the large-scale simultaneous hypothesis testing proposed by Efron,⁶ which inspired our framework.

2.1 Large-scale simultaneous hypothesis testing

Efron presented a method to simultaneously test a large number of hypotheses.⁶ Given the a likelihood value, denoted z_i , for each of the N hypotheses $i = 1, \dots, N$, a histogram of these z -values is created. The goal is now to calculate the probability that a hypothesis stems from a null-distribution based on the z -value. Efron shows that this probability is equal to the FDR defined as follows⁶

$$FDR(z) = f_0(z)/f(z), \quad z \in \mathbb{R} \quad , \quad (1)$$

where f_0 denotes the null-distribution, either empirically estimated or theoretically chosen, and f denotes the mixture distribution. In his paper, f is a curve fitted to the histogram and different ways of computing f_0 are presented. To define whether a hypothesis i is *significant*, Efron computes $FDR(z_i)$ and checks whether it is below a certain threshold.⁶ In this way he seeks those z for which $f(z)$ is significantly different from the null-distribution $f_0(z)$. This implies that the null hypothesis can be rejected.

In the following we explain how this method can be used to find significantly connected voxels based on histograms of connection strengths. Note, however that we perform the simultaneous inference separately for every voxel, since we are only interested in whether *any* of the bins for a specific voxel is significantly different from the null-distribution. Furthermore, we do not estimate or fit f_0 by a function, but instead use an empirical null-distribution computed directly from the histograms as the average normalized histogram over all voxels.

2.2 Tractography

Histograms are constructed from the connections between each seed voxel and the target ROI and their corresponding connection strengths. Here, we perform shortest-path tractography (SPT) described by Kasenburg et al.⁷ to find these connections. First, the diffusion weighted image (DWI) is modelled as a weighted, undirected graph G consisting of all grey and white matter voxels within the brain, where each voxel is connected via an edge to adjacent white matter voxels in its $3 \times 3 \times 3$ neighbourhood. Edge weights quantify the diffusion strength along the corresponding direction and reflect the probability of a fibre bundle tangential to that direction. Tractography between two voxels \mathbf{x} and \mathbf{x}' in graph G is then performed by finding the most probable path $\pi_{\mathbf{x},\mathbf{x}'}$ between the pair of voxels. The connection strength s is given by the product of edge weights along the path normalized by the number of edges n contained in the path:⁷

$$s(\pi_{\mathbf{x},\mathbf{x}'}) = \sqrt[n]{\prod_{i=1}^n w(e_i)} \quad , \quad (2)$$

where e_i is the i -th edge along the path and $w(e_i)$ its corresponding weight. Note that finding the most likely path in G is equivalent to finding the shortest path in a graph with negative log-transformed edge-weights.

2.3 Histogram of connection strengths

For a pair consisting of a seed voxel and a target ROI, the shortest paths in G from the seed voxel to all voxels in the target ROI are computed with SPT as described in Section 2.2. Since SPT is guaranteed to find a path between any pair of voxels, it returns as many paths as there are voxels in the target ROI for each seed voxel. This results in thousands of paths and corresponding connection strengths for each seed voxel. To estimate the distribution of connection strengths for voxel \mathbf{x} to a target ROI R , we bin the strengths of all found paths in a histogram $H_{R,\mathbf{x}}(i)$ with 1000 bins, where i is the the bin index (see Figure 1 for examples).

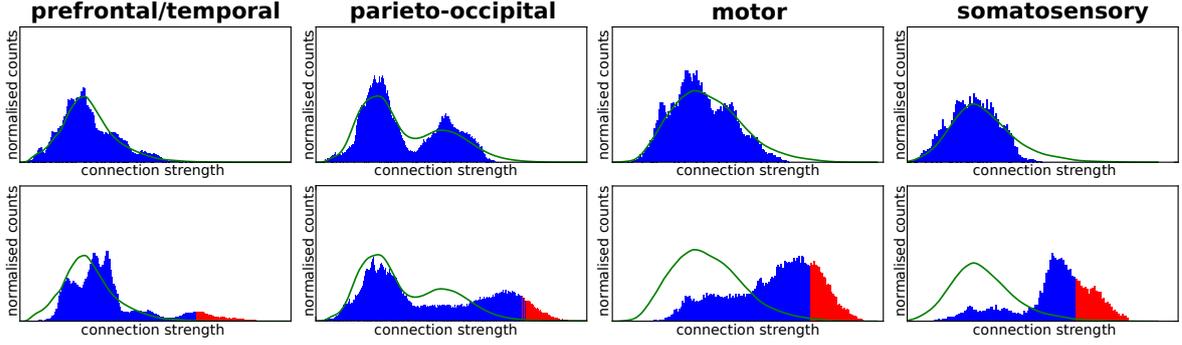


Figure 1. Histograms for all target regions of interest for single seed region voxels for one subject (ID: 100307) together with the average histogram (green line). The top row shows non-significantly connected voxels and the bottom row shows significantly connected voxels using a threshold of 0.05. Significant bins (see Equation 4) are coloured in red.

2.4 Finding significantly connected voxels

We now describe how the concept of large-scale simultaneous hypothesis testing (see Section 2.1) can be used for analysing histograms of connection strengths. In our framework the null-distribution f_0 is independently defined as the normalized average histogram over all voxels in the seed region for each target ROI (see Figure 1). Under the assumption that most connections are *false-positives*, the average histogram is a good approximation of the null-distribution of connection strengths of non-significantly connected voxels. In the following, the average histogram is denoted as $H_R(i)$, where i denotes the bin of the histogram and R is a numerical label for the ROI.

The mixture distribution f is dependent on the location of the seed voxel \mathbf{x} and is defined as the histogram $H_{R,\mathbf{x}}(i)$ for every target ROI R . If the histogram contains the scores of many *true-positive* connections, the shape is expected to be different compared to the null-distribution $H_R(i)$, whereas it would be similar in shape if most connections are *false-positive*. Similar to (1), we define an FDR for every bin i , or connection strength, by dividing the null-distribution (average histogram $H_R(i)$) by the mixture distribution (voxel-specific histogram $H_{R,\mathbf{x}}(i)$)

$$FDR_{R,\mathbf{x}}(i) = \frac{H_R(i)}{H_{R,\mathbf{x}}(i)} . \quad (3)$$

We can now determine whether a voxel \mathbf{x} is significantly connected to R by testing whether the FDR for any of the bins i is lower than a predefined significance threshold thr , that is whether the corresponding strength is significant, with

$$\exists i \geq i_{max,R} : FDR_{R,\mathbf{x}}(i) < thr \Rightarrow \mathbf{x} \text{ is significantly connected} . \quad (4)$$

Since we assume that a low connection strength reflects noise or false-positive connection, we only include those bins larger than or equal to the maximum bin in the null-distribution $i_{max,R} = \operatorname{argmax}_i \{H_R(i)\}$. In our experiments we used thresholds of 0.05, 0.01 and 0.005 equivalent to 95, 99 and 99.5% significance levels to study the effect of a decreasing threshold.

2.5 Creating segmentation maps

To test the level of consistency in location and number of significantly connected voxels between different subjects, probability overlap maps for each target ROI are constructed in *standard MNI space*. These maps show the frequency of a single voxel \mathbf{x} of the seed region being found to be significantly connected to the corresponding ROI R (see Equation 4) over a set of subjects \mathcal{S} . The probability overlap map for region R is defined as the number of subjects in which \mathbf{x} is significantly connected to R divided by the number of subjects $|\mathcal{S}|$.

In addition, an FDR-value was assigned to each significantly connected voxel in order to compare it between different target ROIs. It is defined as the averaged FDR over all bins satisfying (4):

$$B_{R,\mathbf{x}} = \{i : FDR_{R,\mathbf{x}}(i) < thr \wedge i \geq i_{max,R}\} \quad FDR_R(\mathbf{x}) = \frac{1}{|B_{R,\mathbf{x}}|} \sum_{i \in B_{R,\mathbf{x}}} FDR_{R,\mathbf{x}}(i) . \quad (5)$$

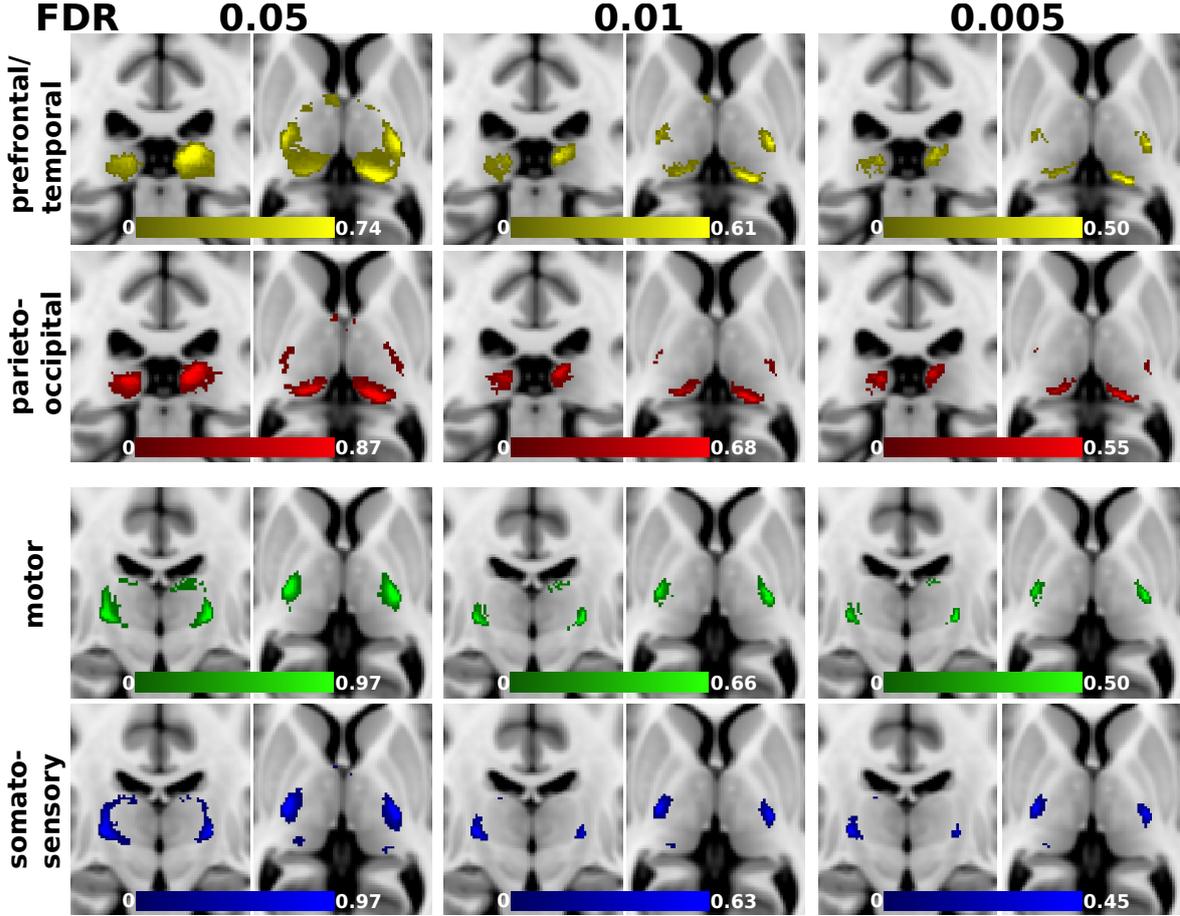


Figure 2. Probability overlap maps for the prefrontal/temporal zone (top row, yellow), parieto-occipital zone (second row, red), motor zone (third row, green) and the somatosensory zone (bottom row, blue) in standard MNI space (top and second row: $x = 90, y = 96, z = 77$, third and bottom row: $x = 90, y = 110, z = 75$), zoomed in on the thalamus. Each column shows results for different FDR thresholds (from left to right: 0.05, 0.01 and 0.005).

From these voxel-wise FDR-values a segmentation map Seg_S of the thalamus can be constructed for every subject $S \in \mathcal{S}$ by assigning the region label with the lowest FDR to each voxel \mathbf{x} in the image

$$Seg_S(\mathbf{x}) = \begin{cases} 0, & \text{if } |\{FDR_R(\mathbf{x}) \neq 0\}| = 0, \\ \operatorname{argmin}_R \{FDR_R(\mathbf{x}) \neq 0\}, & \text{otherwise.} \end{cases} \quad (6)$$

2.6 Data

We used the preprocessed data⁸ of 38 subjects from the Q3 release of the Human Connectome Project^{9,10} (HCP). Tissue segmentations were performed on the provided structural images with FAST¹¹ after brain extraction with BET.¹² Voxel-wise fibre orientation distribution functions were computed with the DiPy library¹³ from the DWIs using constrained spherical deconvolution¹⁴ with order 8.

The target ROIs were extracted from the MNI atlas¹⁵ and the Juelich atlas¹⁶ similar to those described in Behrens et al.:¹ prefrontal/temporal zone, motor zone, somatosensory zone, parieto-occipital zone (see Figure 3 A). Tractography was performed for all subjects as described in Section 2 and segmentation maps (see Equation 6) of the thalamus were created and saved in subject space. Probability overlap maps (see Section 2.5) were created from the segmentations of all subjects warped into *standard MNI space* using FNIRT¹⁷ provided in FSL.¹⁸

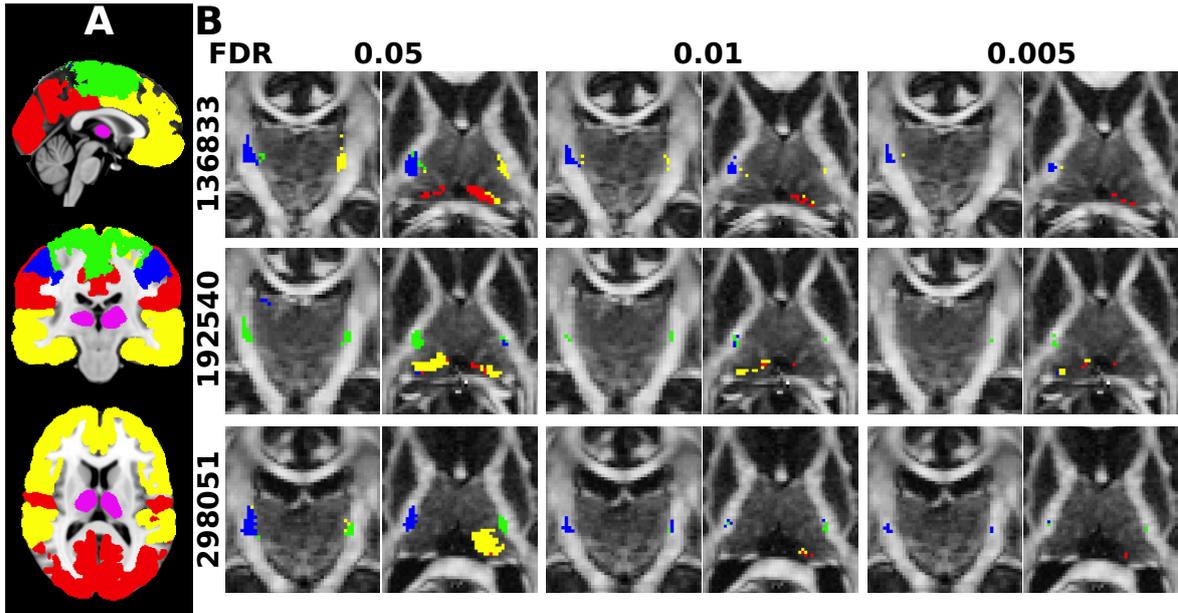


Figure 3. **A** Target regions of interest (yellow: prefrontal/temporal zone, green: motor zone, blue: somatosensory zone, red: parieto-occipital zone) shown together with the thalamus (magenta) in standard MNI space ($x = 89, y = 103, z = 84$). **B** Zoom on the thalamus segmentations for three selected subjects in subject space ($x = 72, y = 87, z = 61$) and three different FDR thresholds (from left to right: 0.05, 0.01 and 0.005). The same colours as in Figure 3A are used for the four target regions.

3. THALAMUS SEGMENTATIONS AND SIGNIFICANT VOXELS

Both the probability overlap maps (see Figure 2) and the subject-specific segmentation maps (see Figure 3 B) show that significantly connected voxels are mostly found on the surface of the thalamus and in similar locations as in Behrens et al.¹ Moreover, it appears that the posterior part of the thalamus is rarely significantly connected to any target ROI (see Figures 2 and 3 B).

On one hand, the subject-specific segmentations show a high variability between the labels chosen for similar locations in different subjects (see Figures 3 B). On the other hand, the probability overlap maps show that the number of significantly connected voxels is consistent across subjects and their location shows reasonably little variance as can be seen by the high overlap probabilities (see Figure 2). This difference can be explained by the fact that most voxels are significantly connected to more than one region. This is especially true for the prefrontal/temporal zone and the parieto-occipital zone or for the motor zone and the somatosensory zone (see Figure 2). So when choosing a label for a voxel for the subject-specific segmentations, it is very often a close call between two target ROIs as the voxel is significantly connected to both. This can in part be explained by the resolution of DWIs, at which it is currently impossible to make distinctions between neuronal fibres that are located close to each other. Especially if the target ROIs are in close proximity to one another, it is likely that the majority of connecting fibres follow, in part, a common path from and to the thalamus.

We also investigated the effect of different thresholds on location and number of significantly connected voxels. Figures 2 and 3 B visualises that when decreasing the threshold of the FDR to 0.01 or 0.005, less voxels are found to be significantly connected. Furthermore, fewer significant voxels are found that do not lie on the surface. Note that the agreement across subjects decreases with the threshold, since subject-specific segmentations maps get extremely sparse (see Figure 3 B).

4. CONCLUSION

We presented a method to find significantly connected voxels and applied it to segmenting the thalamus. Results are consistent across all subjects in the sparse number of significant voxels and their location on the surface of

the anterior part of the thalamus. The proposed framework provides a simple foundation from which probability maps of the connectivity can be created in a standard space from a large sample population. Furthermore the framework is in no way limited to the thalamus or the four target ROIs, but is applicable for any choice of regions for which one wishes to investigate the connectivity.

ACKNOWLEDGMENTS

Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

REFERENCES

- [1] Behrens, T. E. J., Johansen-Berg, H., Woolrich, M. W., Smith, S. M., Wheeler-Kingshott, C. A. M., Boulby, P. A., Barker, G. J., Sillery, E. L., Sheehan, K., Ciccarelli, O., Thompson, A. J., Brady, J. M., and Matthews, P. M., “Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging,” *Nat Neurosci* **6**, 750–757 (Jul 2003).
- [2] Johansen-Berg, H., Behrens, T. E. J., Sillery, E., Ciccarelli, O., Thompson, A. J., Smith, S. M., and Matthews, P. M., “Functional-anatomical validation and individual variation of diffusion tractography-based segmentation of the human thalamus,” *Cereb Cortex* **15**, 31–39 (Jan 2005).
- [3] Bohanna, I., Georgiou-Karistianis, N., and Egan, G. F., “Connectivity-based segmentation of the striatum in huntington’s disease: vulnerability of motor pathways,” *Neurobiol Dis* **42**, 475–481 (Jun 2011).
- [4] Fillard, P., Descoteaux, M., Goh, A., Gouttard, S., Jeurissen, B., Malcolm, J., Ramirez-Manzanares, A., Reisert, M., Sakaie, K., Tensaouti, F., Yo, T., Mangin, J.-F., and Poupon, C., “Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom,” *NeuroImage* **56**, 220–234 (May 2011).
- [5] Neher, P. F., Laun, F. B., Stieltjes, B., and Maier-Hein, K. H., “Fiberfox: facilitating the creation of realistic white matter software phantoms,” *Magn Reson Med* **72**, 1460–1470 (Nov 2014).
- [6] Efron, B., “Large-scale simultaneous hypothesis testing,” *Journal of the American Statistical Association* **99**(465), 96–104 (2004).
- [7] Kasenburg, N., Liptrot, M., Reisle, N. L., Ørting, S. N., Nielsen, M., Garde, E., and Feragen, A., “Training shortest-path tractography: Automatic learning of spatial priors,” *NeuroImage* (2016).
- [8] Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Essen, D. C. V., Jenkinson, M., and WU-Minn HCP Consortium, “The minimal preprocessing pipelines for the Human Connectome Project,” *NeuroImage* **80**, 105–124 (Oct 2013).
- [9] Essen, D. C. V., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., Ugurbil, K., and WU-Minn HCP Consortium, “The WU-Minn Human Connectome Project: An overview,” *NeuroImage* **80**, 62–79 (Oct 2013).
- [10] Sotiropoulos, S. N., Moeller, S., Jbabdi, S., Xu, J., Andersson, J. L., Auerbach, E. J., Yacoub, E., Feinberg, D., Setsompop, K., Wald, L. L., Behrens, T. E. J., Ugurbil, K., and Lenglet, C., “Effects of image reconstruction on fiber orientation mapping from multichannel diffusion MRI: reducing the noise floor using SENSE,” *Magn Reson Med* **70**, 1682–1689 (Dec 2013).
- [11] Zhang, Y., Brady, M., and Smith, S., “Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm,” *IEEE Trans Med Imaging* **20**, 45–57 (Jan 2001).
- [12] Jenkinson, M., Pechaud, M., and Smith, S., “BET2: MR-based estimation of brain, skull and scalp surfaces,” in *Eleventh annual meeting of the organization for human brain mapping*, **17** (2005).
- [13] Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., van der Walt, S., Descoteaux, M., Nimmo-Smith, I., and Dipy Contributors, “Dipy, a library for the analysis of diffusion MRI data,” *Front Neuroinform* **8**, 8 (2014).
- [14] Tournier, J.-D., Calamante, F., and Connelly, A., “Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution,” *NeuroImage* **35**, 1459–1472 (May 2007).

- [15] Mazziotta, J., Toga, A., Evans, A., Fox, P., Lancaster, J., Zilles, K., Woods, R., Paus, T., Simpson, G., Pike, B., Holmes, C., Collins, L., Thompson, P., MacDonald, D., Iacoboni, M., Schormann, T., Amunts, K., Palomero-Gallagher, N., Geyer, S., Parsons, L., Narr, K., Kabani, N., Goualher, G. L., Boomsma, D., Cannon, T., Kawashima, R., and Mazoyer, B., “A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM),” *Philos Trans R Soc Lond B Biol Sci* **356**, 1293–1322 (Aug 2001).
- [16] Eickhoff, S. B., Paus, T., Caspers, S., Grosbras, M.-H., Evans, A. C., Zilles, K., and Amunts, K., “Assignment of functional activations to probabilistic cytoarchitectonic areas revisited,” *NeuroImage* **36**, 511–521 (Jul 2007).
- [17] Andersson, J. L. R., Jenkinson, M., and Smith, S., “Non-linear registration, aka spatial normalisation,” tech. rep., FMRIB (2010).
- [18] Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., and Smith, S. M., “FSL,” *NeuroImage* **62**, 782–790 (Aug 2012).

5.3

STRUCTURAL PARCELLATION OF THE THALAMUS USING SHORTEST-PATH TRACTOGRAPHY

Niklas Kasenburg¹, Sune Darkner¹, Ute Hahn², Matthew Liptrot^{1,3}, Aasa Feragen¹

¹DIKU, University of Copenhagen; ²Department of Mathematics, Aarhus University; ³DTU Compute

ABSTRACT

We demonstrate how structural parcellation can be implemented using shortest-path tractography, thereby addressing some of the shortcomings of the conventional approaches. In particular, our algorithm quantifies, via p -values, the confidence that a voxel in the parcellated region is connected to each cortical target region. Calculation of these statistical measures is derived from a rank-based test on the histogram of tract-based scores from all the shortest paths found between the source voxel and each voxel within the target region. Using data from the Human Connectome Project, we show that parcellation of the thalamus results in p -value maps that are spatially coherent across subjects. Comparing to the state-of-the-art parcellation of Behrens et al. [1], we observe some agreement, but the soft segmentation exhibits better stability for voxels connected to multiple target regions.

1. INTRODUCTION

Structural parcellation [2, 3, 4, 5, 6, 7] is the data-driven segmentation of a source region defined by structural connectivity to a set of pre-defined target regions of interest (ROIs). We propose a parcellation algorithm which statistically quantifies the connectivity to each target ROI for each voxel in the region to be parcellated. In contrast to classical majority voting methods [1], our method outputs a soft segmentation based on these confidence measures. We demonstrate the approach on 5 subjects from the Human Connectome Project [8].

Our approach, based on shortest-path tractography (SPT), overcomes all of the following problems, which remain unsolved by traditional approaches to structural parcellation:

- a) Traditional fibre tracking methods start at a source point and randomly walk in the most likely direction as defined by a fibre orientation function (fODF). These methods may have a hard time retrieving connections from a voxel v to a specific cortical region C if the

voxel v is *also* connected to other regions, whose connections are easier to track. In particular, fibre tracking has a bias towards finding connections between nearby regions and will typically find more connections to large target regions than to small ones, even if this is not supported by the data [9, 10].

- b) Traditional structural parcellation assumes that every voxel in the region to be segmented (in our case, the thalamus) is connected to one of the target ROIs. However, this is not always the case. Moreover, there may not be enough signal in the data to support such an assignment even if the connections exist. This is especially true for inside the thalamus, where fibre directions are usually hard to resolve.
- c) Traditional structural parcellation obtains a hard segmentation of the source region through voting over connections to the cortical ROIs. However, this is suboptimal for several reasons [3]: First, if a source voxel is not physically connected to either of the target regions, the found connections are based only on noise and may therefore be very unstable. Second, both anatomically and because of partial volume effects, many voxels are with high probability connected to multiple cortical regions. Since a hard segmentation is not able to model this, it may again, lead to unstable results.

Shortest-path tractography finds the most likely trajectory for a fibre connecting any two locations in the brain [11]. This formulation lends itself well to structural parcellation, which precisely seeks connections between two pre-specified ROIs: The source region to be segmented (in this paper, the thalamus) and the cortical target region, one at a time. Since SPT will *always* find a most likely tract connecting any two voxels, it avoids problem a): Connections will be found both to small regions and between regions that are any distance apart.

However, this property of SPT also introduces a new problem: A connection will always be found, even if it is not there. Many SPT methods assign a score to the found paths which can be used to threshold unlikely paths, but such a threshold will reintroduce the biases from problem a). In this paper, we therefore propose an alternative approach which, assuming *some* thalamic voxels *are* physically connected to each target region, aims to define these as statistically more sig-

A.F. is funded in part by the Danish Council for Independent Research (DFR), Technology and Production Sciences. Data were provided by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

nificant than those which are not. This is obtained through a statistical test over source region voxels.

Inspired by Kasenburg et al. [12] we quantify the confidence with which a voxel is significantly connected to each of the target regions, by assigning a p -value to each source voxel. This also solves problem b), as we allow voxels within the source region to have a p -value close to 1 for every single target region, i.e. not being strongly connected to any target region. Using the p -values as confidence maps, we obtain a soft structural parcellation of the thalamus, which allows for overlap between the segments corresponding to different target ROIs. In this way, we also solve problem c).

2. METHODOLOGY

Shortest-path tractography (SPT). In SPT [11], the diffusion weighted image (DWI) is turned into a graph where all grey- and white matter voxels from the brain are nodes. Edges are formed between pairs of voxels in a $3 \times 3 \times 3$ neighbourhood, whenever at least one of the voxels is white matter. Each edge \vec{e} pointing out of a voxel v is given a weight reflecting the probability of a fibre bundle tangential to \vec{e} . This probability is defined by integrating the fODF estimated for the voxel v over the set of directions out of v which are closer to \vec{e} than to any other edge pointing out of v . The integral is estimated by sampling. To obtain an undirected graph, edge weights are averaged over their start and end nodes, obtaining probabilistic edge weights $p(e)$.

From this weighted graph, tractography between two points v and w is phrased as finding the most probable path from v to w in the graph, that is, the path maximising $p(\pi) = \prod_{i=1}^n p([v_{i-1}, v_i])$, where $\pi = [v_0, v_1 \dots, v_n]$ is a path in the brain graph from $v_0 = v$ to $v_n = w$, and $[v_{i-1}, v_i]$ is the edge connecting v_{i-1} and v_i . Most probable paths are computed using Dijkstra's shortest path algorithm after log-transforming the weights into edge lengths $l(e) = -\log(p(e))$.

Significance of SPT tracts. The main disadvantage of SPT is that for any two brain voxels v and w , SPT *will* find a path connecting them, whether the data really supports this or not. This can be alleviated by thresholding the path-length corrected path probability score $s(\pi) = [p(\pi)]^{\frac{1}{|\pi|}}$, where $|\pi|$ is the number of nodes on π . However, this requires a choice of a threshold. A test for the statistical significance of SPT tracts would therefore be desirable. In this paper, we present a statistical test for the significance of connectivity from a source voxel $v \in R_1$ to a target region R_2 .

Assumptions. We seek a segmentation of a given source region R_1 (the thalamus) defined by its connections to a set of cortical target regions. Focusing on one target region R_2 , consider all paths found by SPT from a voxel v to any voxel w in the target region R_2 . We assume that the majority of voxels in R_1 are *not* physically connected to the target region R_2 and that the scores of SPT paths found from such $v_{unc} \in R_1$ to

voxels $w \in R_2$ describe noise. This noise may be specific to the regions R_1 and R_2 , but we assume that it is independent of the source voxel v_{unc} . We also assume that voxels $v_{con} \in R_1$ that actually are physically connected to region R_2 show *more* high scoring SPT paths than expected for v_{unc} . Thus, their score distribution should be skewed towards the right, as the distribution is a mixture of noise from the target voxels w in R_2 that are not connected to v_{con} , and high scores from those target voxels w which *are* connected to v_{con} .

Histograms and cumulative histograms. For each source voxel $v \in R_1$ we extract a histogram H_v of scores corresponding to SPT paths from v to any $w \in R_2$ as follows: Scores are divided into N bins and the number $H_v(i)$ of scores falling into bin i is counted for $i = 1, \dots, N$. Examples of these histograms are shown in the top row of Fig. 1 together with the average histogram over all source voxels. Some histograms (top right in Fig. 1) indeed represent a distribution of scores that is more skewed to the right compared to the average histogram, suggesting that the corresponding voxel is physically connected to the target region.

Such raw histograms are, however, very jagged and noisy, therefore they are rarely used for statistical testing. Instead, one usually transforms them into cumulative normalized histograms (also known as empirical cumulative distribution functions) prior to testing without losing information. The cumulative normalized histogram of voxel v is given by a vector $(C_v(1), \dots, C_v(N))$, where

$$C_v(i) = \sum_{j=1}^i \frac{H_v(j)}{\text{norm}(H)}, \quad \text{norm}(H) = \sum_{i=1}^N H_v(i). \quad (1)$$

Assigning a p -value to histograms. To quantify if and how much the histogram of voxel v deviates from the null distribution, we compare it to a simulated sample of s histograms representing noise. When a histogram is skewed to the right, the corresponding cumulative histogram lies below a typical cumulative histograms representing the null distribution (see the right of Fig. 1).

Inspired by the envelope rank tests of Myllymäki et al. [13], we present a one-sided rank-based test for the cumulative histograms. Our test computes a p -value for each source voxel v given by

$$p = \frac{1}{s+1} \sum_{k=1}^{s+1} \mathbf{1}(\text{rank}_k \leq \text{rank}_1), \quad (2)$$

where $\mathbf{1}(x)$ returns 1 if x is true and 0 otherwise, and the rank_k is the average rank from below over all bins of the k^{th} sample from the null distribution. That is,

$$\text{rank}_k = \frac{1}{N} \sum_{j=1}^N \#\{k' = 1, \dots, s+1 | C_{k'}(j) < C_k(j)\}, \quad (3)$$

where $k = 1$ is the index of the observed histogram. The rank thus measures how skewed the corresponding original

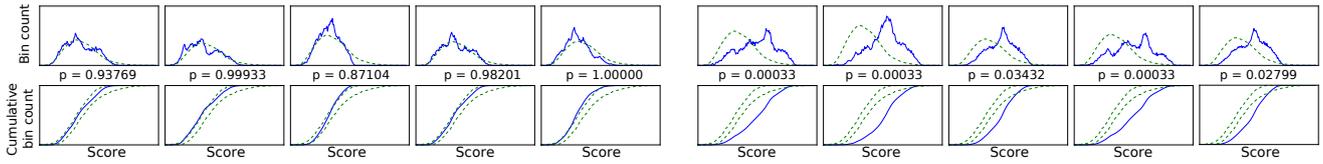


Fig. 1. Top: Normalized histograms of SPT tract scores for single source region voxels. The dotted green line is the average histogram over all source voxels. **Bottom:** The corresponding cumulative histograms. The dotted green lines are the minimum and maximum envelopes of the cumulative histograms resulting from the normalized, bootstrapped samples. The first five voxels were found to have a p -value very close to 1, whereas the latter five have lower p -values.

histogram is to the right, compared to the other histograms in the sample of $s + 1$ histograms. The original test by Myllymäki et al. [13] was based on minimum ranks rather than the average rank, which led to an envelope interpretation of the test. For our application, the mean rank provides additional stability, but lacks the envelope interpretation.

Drawing from the null distribution of cumulative histograms. In order to define a statistical test for significantly connected source voxels, we need to be able to draw samples from the null distribution. Under the assumption that the majority of source voxels are *not* physically connected to the target region, this could be solved by drawing entire histograms from the population. However this could lead to an unnecessarily conservative test, since we risk to draw histograms from voxels that actually *are* physically connected.

Instead, we propose bootstrapping the histograms as follows: To draw a cumulative histogram C from the null distribution, we first draw a histogram H from the null distribution of score histograms, giving a cumulative sample C from H as in (1). The sample H is drawn by randomly sampling a bin value $H(i)$ from the i^{th} bin values in the entire population for every bin i in H . This gives a histogram H whose bin values $H(i)$ are i^{th} bin values drawn from different source voxels v .

3. EXPERIMENTS

We used the preprocessed DWIs [14, 15, 16] from 5 subjects (Q3 release) of the Human Connectome Project (HCP) data [8]. Fibre orientation distribution functions (fODFs) were computed for every voxel using constrained spherical deconvolution [17] with order 8 using the DiPy [18] package. The graph required for SPT was constructed from the given fODFs. The voxel-wise diffusion parameters necessary for probabilistic tractography were generated using FSL’s BedpostX [19] with the zeppelin model and 3 fibres per voxel.

The thalamic source region was extracted from the MNI atlas [20] provided in FSL [21]. Target regions were extracted both from the MNI atlas and the Juelich atlas [22] similar to Behrens et al. [1]: prefrontal/temporal zone (frontal and temporal lobe from the MNI atlas), motor zone (primary motor cortex and premotor cortex from the Juelich atlas), somatosensory zone (primary somatosensory cortex from the

Juelich atlas) and parieto-occipital zone (occipital and parietal lobe from the MNI atlas). The atlas in MNI space was warped into the respective subject-specific spaces using the warps provided in the HCP data.

For every source voxel, SPT was performed to obtain the most likely paths to all voxels in the given target region. Path scores were binned into a histogram with 1000 bins. Histograms for all source voxels were analysed for each target region separately as described in Section 2. The corresponding p -value for each voxel is shown on the left of Fig. 2.

We compared to the state-of-the-art hard segmentation by Behrens et al. [1]. Here, each thalamus voxel was labelled with the target region to which it had the highest number of streamlines based on probabilistic tractography (FSL’s probtrackx [23], 5000 samples, 0.5 mm step length, maximum inter-step curvature 80°). The resulting segmentations are shown in the right side of Fig. 2.

4. DISCUSSION AND CONCLUSION

We present a new method for structural parcellation based on structural connectivity as an alternative to hard segmentation. Our method addresses the problems mentioned in Sec. 1:

a) We choose shortest-path tractography (SPT) to avoid problems like path-length dependency common in fibre tracking methods [10]. We overcome the problem that SPT always finds a path by assigning a confidence p -value to every voxel in the parcellated region.

b) Most voxels in the thalamus are assigned with a low confidence (high p -value) for all target regions by our method (see left side of Fig. 2). This occurs mostly inside the thalamus and implies that the diffusion signal for those voxels is not strong or clear enough to assign them to a region with high confidence. This is supported by the low fractional anisotropy (FA) inside the thalamus, suggesting that the resolution is simply not high enough to resolve the direction of fibres inside the thalamus.

c) Standard hard segmentation of the thalamus [1], as shown in Fig. 2 (right) assumes that each voxel is connected to only one target region. However, our results indicate that high confidence regions overlap strongly between the motor and somatosensory zone, as well as between the pre-

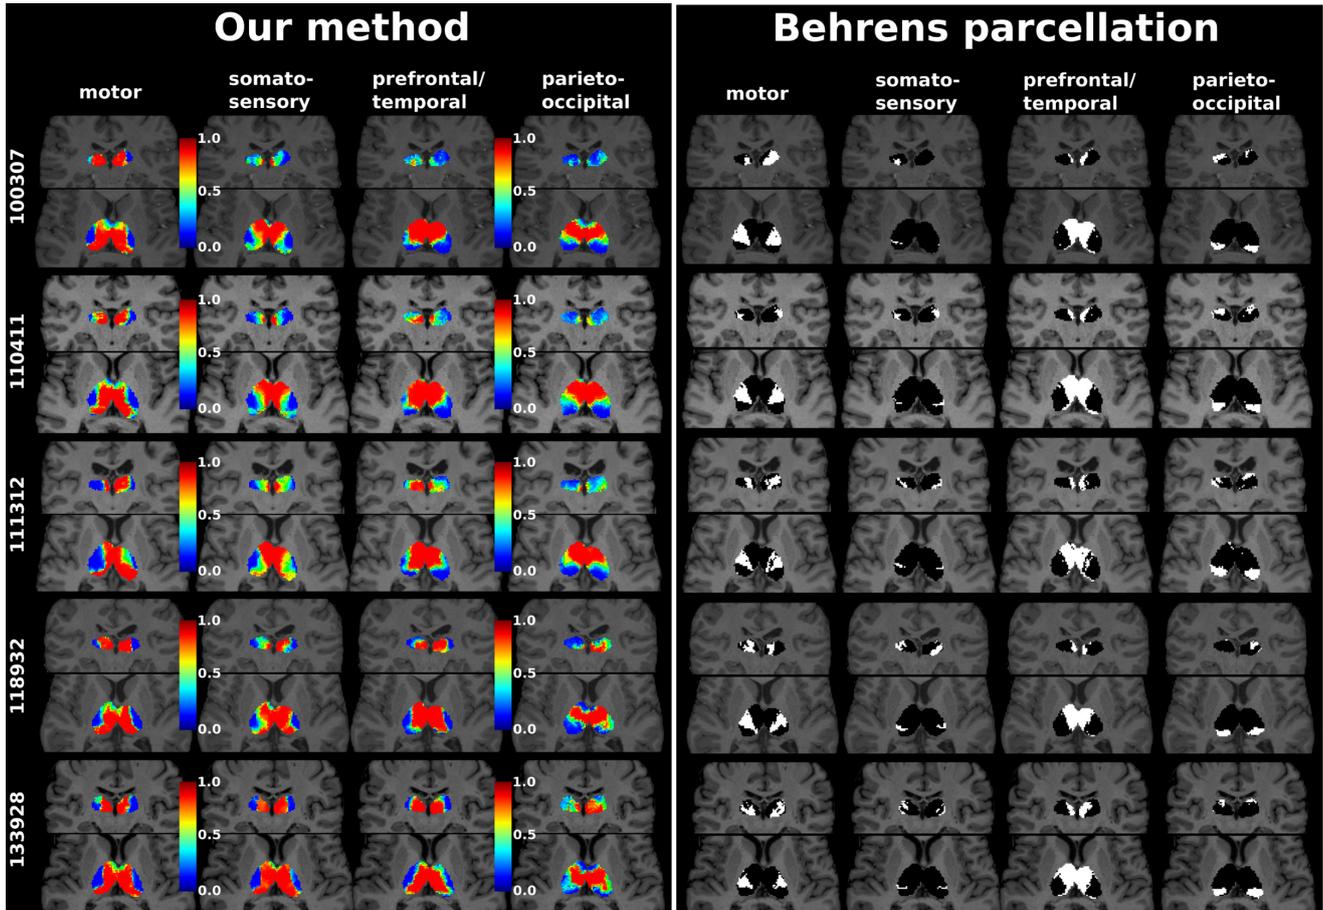


Fig. 2. Soft (left) and hard [1] (right) segmentations of the thalamus for all target regions for all 5 subjects. **Left:** Low confidence p -values (blue) reflect regions that are more likely connected to the respective target region than high p -values (red). **Right:** We compare to the state-of-the-art parcellation by Behrens et al. [1], where the thalamus is segmented by which target region is connected with the most streamlines. These voxels are highlighted in white, while the rest of the thalamus is coloured black.

frontal/temporal and parieto-occipital zone. This makes it especially difficult for the hard segmentation to pick up any voxels connected to the somatosensory zone (see Fig. 2). In contrast, our soft segmentation allows for a voxel to be connected to multiple target regions.

While the soft segmentations overlap strongly with the hard segmentations (see Fig. 2) for the motor, somatosensory and parieto-occipital zone, there is only a small overlap for the prefrontal/temporal zone. One possible reason is that a large part of the hard segmentation lies in the region of high uncertainty, where the resolution of the data is not good enough to resolve the fibre directions. Additionally, the distribution of histograms can be a mixture of multiple underlying distributions corresponding to several major tracts that connect the thalamus to the cortex. Here, this may affect the p -values of the projection of the prefrontal and temporal cortex onto the thalamus, indicating that the target regions should be suffi-

ciently spatial coherent and reasonably small. A division of the prefrontal/temporal zone into two separate regions as in Behrens et al. [1] will potentially resolve this problem.

Here we are not testing which source voxels are *significantly* connected to the target region based on the confidence p -values, as this would require to address multiple hypothesis testing. Although this could be done via Bonferroni correction, structural connectivity of neighbouring source voxels is strongly correlated. Bonferroni correction would therefore be far too conservative and would affect the segmentation in an unknown way. Future work therefore includes multiple hypothesis testing that respects correlation between source voxels to find those significantly connected to the target region. Furthermore, it would be interesting to apply the soft segmentation to other brain structures like the corpus callosum and striatum or to extend the method to be able to parcellate the whole cortex based on whole brain tractography.

5. REFERENCES

- [1] T. E. J. Behrens and et al., “Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging,” *Nat Neurosci*, vol. 6, no. 7, pp. 750–757, July 2003.
- [2] I. Bohanna and et al., “Connectivity-based segmentation of the striatum in huntington’s disease: vulnerability of motor pathways,” *Neurobiol Dis*, vol. 42, no. 3, pp. 475–481, Jun 2011.
- [3] N.S. Gorbach and et al., “Hierarchical information-based clustering for connectivity-based cortex parcellation,” *Frontiers in Neuroinformatics*, vol. 5, no. 18, 2011.
- [4] L.L. Cloutman and et al., “Connectivity-based structural and functional parcellation of the human cortex using diffusion imaging and tractography,” *Frontiers in Neuroanatomy*, vol. 6, no. 34, 2012.
- [5] J. V. Stough and et al., “Automatic method for thalamus parcellation using multi-modal feature classification,” *MICCAI*, vol. 17, no. Pt 3, pp. 169–176, 2014.
- [6] D. Moreno-Dominguez and et al., “A hierarchical method for whole-brain connectivity-based parcellation,” *Human Brain Mapping*, vol. 35, no. 10, pp. 5000–5025, 2014.
- [7] S. Eickhoff and et al., “Connectivity-based parcellation: Critique and implications,” *Human Brain Mapping*, p. 22, 2016.
- [8] D.C. Van Essen and et al., “The WU-Minn Human Connectome Project: An overview,” *NeuroImage*, vol. 80, pp. 62–79, Oct 2013.
- [9] D.K. Jones, “Challenges and limitations of quantifying brain connectivity in vivo with diffusion MRI,” *Imaging in Medicine*, vol. 2, no. 3, pp. 341–355, 2010.
- [10] M.G. Liptrot and et al., “Addressing the path-length-dependency confound in white matter tract segmentation,” *PLoS ONE*, vol. 9, no. 5, pp. e96247, 2014.
- [11] N. Kasenburg and et al., “Training shortest-path tractography: Automatic learning of spatial priors,” *NeuroImage*, 2016.
- [12] N. Kasenburg and et al., “Finding significantly connected voxels based on histograms of connection strengths,” *Proceedings SPIE Medical Imaging: Image Processing*, 2016.
- [13] M. Myllymäki and et al., “Global envelope tests for spatial processes,” *eprint arXiv:1307.0239*, June 2013.
- [14] M.F. Glasser and et al., “The minimal preprocessing pipelines for the Human Connectome Project,” *NeuroImage*, vol. 80, pp. 105–124, 2013.
- [15] J.L.R. Andersson and et al., “How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging,” *NeuroImage*, vol. 20, no. 2, pp. 870–888, Oct 2003.
- [16] S. N. Sotiropoulos and et al., “Effects of image reconstruction on fiber orientation mapping from multi-channel diffusion MRI: reducing the noise floor using SENSE,” *Magn Reson Med*, vol. 70, no. 6, pp. 1682–1689, Dec 2013.
- [17] J.-D. Tournier and et al., “Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution,” *NeuroImage*, vol. 35, no. 4, pp. 1459–1472, May 2007.
- [18] E. Garyfallidis and et al., “Dipy, a library for the analysis of diffusion MRI data,” *Frontiers in Neuroinformatics*, vol. 8, 2014.
- [19] T. E. J. Behrens and et al., “Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?,” *NeuroImage*, vol. 34, no. 1, pp. 144–155, Jan 2007.
- [20] J. Mazziotta and et al., “A probabilistic atlas and reference system for the human brain: International Consortium for Brain Mapping (ICBM),” *Philos Trans R Soc Lond B Biol Sci*, vol. 356, no. 1412, pp. 1293–1322, Aug 2001.
- [21] M. Jenkinson and et al., “FSL,” *NeuroImage*, vol. 62, no. 2, pp. 782–790, Aug 2012.
- [22] S.B. Eickhoff and et al., “Testing anatomically specified hypotheses in functional imaging using cytoarchitectonic maps,” *NeuroImage*, vol. 32, no. 2, pp. 570–582, 2006.
- [23] T. E. J. Behrens and et al., “Characterization and propagation of uncertainty in diffusion-weighted MR imaging,” *Magn Reson Med*, vol. 50, no. 5, pp. 1077–1088, Nov 2003.

5.4

Supervised hub-detection for brain connectivity

Niklas Kasenburg^a, Matthew Liptrot^{a,b}, Nina Linde Reislev^c, Ellen Garde^c, Mads Nielsen^a and Aasa Feragen^a

^aDepartment of Computer Science, University of Copenhagen, Denmark

^bDTU Compute, Copenhagen, Denmark

^cDanish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital Hvidovre, Denmark

ABSTRACT

A structural brain network consists of physical connections between brain regions. Brain network analysis aims to find features associated with a parameter of interest through supervised prediction models such as regression. Unsupervised preprocessing steps like clustering are often applied, but can smooth discriminative signals in the population, degrading predictive performance. We present a novel hub-detection optimized for supervised learning that both clusters network nodes based on population level variation in connectivity and also takes the learning problem into account. The found hubs are a low-dimensional representation of the network and are chosen based on predictive performance as features for a linear regression. We apply our method to the problem of finding age-related changes in structural connectivity. We compare our supervised hub-detection (SHD) to an unsupervised hub-detection and a linear regression using the original network connections as features. The results show that the SHD is able to retain regression performance, while still finding hubs that represent the underlying variation in the population. Although here we applied the SHD to brain networks, it can be applied to any network regression problem. Further development of the presented algorithm will be the extension to other predictive models such as classification or non-linear regression.

1. INTRODUCTION

The nature of how structural brain connectivity is associated with parameters such as pathology, gender or age, is a cornerstone of modern neuroscience.¹⁻³ Structural connectivity networks are typically derived from diffusion weighted imaging (DWI) using tractography, and analysed using supervised learning methods like regression or classification. Supervised learning is commonly preceded by network dimensionality reduction methods that cluster nodes into an atlas of *hubs* common to the population.^{1,2,4} However, these algorithms are generally unrelated to the learning problem and tend to smooth out the discriminative signal, leading to lower-dimensional brain networks with reduced predictive power,^{4,5} and therefore reduced applicability to e.g. automated diagnosis or prognosis. We provide a novel approach to hub-detection where the supervised learning problem is incorporated in the hub-detection algorithm.

1.1 Previous work

Clayden et al.⁶ perform eigendecomposition of the adjacency matrix to generate so called 'principal networks', and Robinson et al.¹ extract hubs from the main principal components, followed by a maximum uncertainty linear discriminant analysis. In other related works, hierarchical clustering on spatial locations selects groups for structured sparsity in regression and classification⁷ or matrix factorization generates hubs used for permutation testing on hub-hub connections.² In all these approaches, hubs are selected *without taking the desired learning problem into account*. While hub-detection eases interpretability, hubs and their connections are optimized to describe within-population variation. This often leads to a smoothing of latent discriminative features, thereby reducing the power of predictive models.^{4,5}

We address this problem by including the learning problem in the hub-detection process. This forces a small number of hubs to describe population variation well *and* be associated with the parameter of interest. We apply this method to detecting hub-network features that are associated with increasing age in structural connectivity networks.

Corresponding author: Niklas Kasenburg (niklas.kasenburg@di.ku.dk)

2. METHODS

Let $\mathbf{S}^{(m)} \in \mathbb{R}^{n \times n}$ be a symmetric, non-negative matrix specifying the strength of connectivity between n brain regions, for the m -th out of M subjects. This matrix specifies a graph with n nodes corresponding to the brain regions, and n^2 edges connecting pairs of nodes. The (i, j) -th entry of $\mathbf{S}^{(m)}$ is the edge *weight* describing the strength of connectivity between regions (or network nodes) i and j .

We extend the unsupervised hub-detection by Ghanbari et al.,² where each subject matrix $\mathbf{S}^{(m)}$ is approximated through matrix factorization as $\mathbf{S}^{(m)} \approx \mathbf{U}\mathbf{\Lambda}^{(m)}\mathbf{U}^T$, where $\mathbf{U} \in \mathbb{R}^{n \times k}$ is a non-negative matrix assigning hub-membership scores to each of the n regions for each of the k hubs, and $\mathbf{\Lambda}^{(m)} \in \mathbb{R}^{k \times k}$ is a non-negative, symmetric matrix capturing subject specific hub-hub connectivity. The reconstruction of $\mathbf{S}^{(m)}$ is obtained through the objective function²

$$J_{\text{recon}}(\mathbf{U}, \mathbf{\Lambda}^{(m)}) = \sum_{m=1}^M \|\mathbf{S}^{(m)} - \mathbf{U}\mathbf{\Lambda}^{(m)}\mathbf{U}^T\|_F^2 + \beta \left(\|\mathbf{U}\|_F^2 + \sum_{m=1}^M \|\mathbf{\Lambda}^{(m)}\|_F^2 \right), \quad (1)$$

where constraints are applied such that the properties of \mathbf{U} and $\mathbf{\Lambda}^{(m)}$ are ensured. The first term of (1) is the reconstruction error, while the second is a regularizing term.

Given a real-valued label $y^{(m)} \in \mathbb{R}$ for each subject (here age), we add a linear regression term

$$J_{\text{reg}}(\mathbf{U}, \mathbf{\Lambda}^{(m)}, \mathbf{W}, w) = \gamma \sum_{m=1}^M \left(\text{Tr}(\mathbf{X}\mathbf{W}) + w - y^{(m)} \right)^2 + \alpha \|\mathbf{W}\|_F^2 + \alpha w^2, \quad (2)$$

where $\mathbf{X}_+ \in \mathbb{R}^{k \times k}$ is a symmetric matrix representing the features for regression, $\mathbf{W} \in \mathbb{R}^{k \times k}$ are the regression weights, w is the regression offset, α is a regularization factor for \mathbf{W} and w , and γ is a weighting factor for the regression term. Adding (1) and (2), we get (due to symmetry of $\mathbf{S}^{(m)}$, $\mathbf{\Lambda}^{(m)}$ and \mathbf{W})

$$J(\mathbf{U}, \mathbf{\Lambda}^{(m)}, \mathbf{W}, w) := J_{\text{recon}} + J_{\text{reg}} = \sum_{m=1}^M \text{Tr} \left(\left(\mathbf{S}^{(m)} - \mathbf{U}\mathbf{\Lambda}^{(m)}\mathbf{U}^T \right)^2 \right) + \gamma \sum_{m=1}^M \left(\text{Tr}(\mathbf{X}\mathbf{W}) + w - y^{(m)} \right)^2 + \beta \left[\text{Tr}(\mathbf{U}\mathbf{U}^T) + \sum_{m=1}^M \text{Tr}(\mathbf{\Lambda}^{(m)2}) \right] + \alpha \text{Tr}(\mathbf{W}^2) + \alpha w^2. \quad (3)$$

In the following we choose to use the original connections, weighted by the hub-membership scores of the nodes as regression features

$$\mathbf{X} = \mathbf{U}^T \mathbf{S}^{(m)} \mathbf{U}. \quad (4)$$

In this way inter-hub connections summarize the original connections between the nodes of any pair of hubs, while intra-hub connections are reflected by connections between nodes within the same hub. Note that \mathbf{X} is a dimensionality reduction of the original connections (features) and regression is therefore performed on a lower dimensional network.

2.1 Optimization of the objective function J

We minimize (3) through an iterative update scheme for \mathbf{U} , $\mathbf{\Lambda}^{(m)}$, \mathbf{W} and w using a constrained version of the quasi-Newton method of Broyden et al. (L-BFGS-B)⁸ implemented in SciPy*. In contrast to gradient descent, the BFGS converges superlinearly and uses an estimated Hessian during the optimization. The gradients^{2†} required for each update step are the following for \mathbf{U} and $\mathbf{\Lambda}^{(m)}$

$$\frac{\partial J}{\partial \mathbf{U}} = -4 \sum_{m=1}^M \left[\left(\mathbf{S}^{(m)} - \mathbf{U}\mathbf{\Lambda}^{(m)}\mathbf{U}^T \right) \mathbf{U}\mathbf{\Lambda}^{(m)} \right] + 2\beta\mathbf{U} + \frac{\partial J_{\text{reg}}}{\partial \mathbf{U}}, \quad (5)$$

*<http://www.scipy.org/>

†Note that [2] has a typographical error in (4), where it states $4\beta\mathbf{U}$.

$$\frac{\partial J}{\partial \mathbf{\Lambda}^{(m)}} = -2\mathbf{U}^T \left(\mathbf{S}^{(m)} - \mathbf{U}\mathbf{\Lambda}^{(m)}\mathbf{U}^T \right) \mathbf{U} + 2\beta\mathbf{\Lambda}^{(m)} . \quad (6)$$

The gradients $\frac{\partial J}{\partial \mathbf{W}}$ and $\frac{\partial J}{\partial w}$ are straightforward

$$\frac{\partial J}{\partial \mathbf{W}} = 2\gamma \sum_{m=1}^M \left[\mathbf{X} \left(\text{Tr}(\mathbf{X}\mathbf{W}) + w - y^{(m)} \right) \right] + 2\alpha\mathbf{W} , \quad (7)$$

$$\frac{\partial J}{\partial w} = 2\gamma \sum_{m=1}^M \left[\left(\text{Tr}(\mathbf{X}\mathbf{W}) + w - y^{(m)} \right) \right] + 2\alpha w . \quad (8)$$

We now need to compute the gradient of J_{reg} with respect to \mathbf{U}

$$\frac{\partial J_{\text{reg}}}{\partial \mathbf{U}} = 2\gamma \sum_{m=1}^M \frac{\partial \text{Tr}(\mathbf{X}\mathbf{W})}{\partial \mathbf{U}} \left(\text{Tr}(\mathbf{X}\mathbf{W}) + w - y^{(m)} \right) . \quad (9)$$

With \mathbf{X} as defined in (4) we get

$$\frac{\partial \text{Tr}(\mathbf{X}\mathbf{W})}{\partial \mathbf{U}} = \frac{\partial \text{Tr}(\mathbf{U}^T \mathbf{S}^{(m)} \mathbf{U} \mathbf{W})}{\partial \mathbf{U}} = \mathbf{S}^{(m)} \mathbf{U} \mathbf{W} + \mathbf{S}^{(m)T} \mathbf{U} \mathbf{W}^T = 2\mathbf{S}^{(m)} \mathbf{U} \mathbf{W} . \quad (10)$$

Using (5), (9) and (10) the resulting gradient with respect to \mathbf{U} is the following

$$\frac{\partial J}{\partial \mathbf{U}} = -4 \sum_{m=1}^M \left[\left(\mathbf{S}^{(m)} - \mathbf{U}\mathbf{\Lambda}^{(m)}\mathbf{U}^T \right) \mathbf{U}\mathbf{\Lambda}^{(m)} \right] + 2\beta\mathbf{U} + 4\gamma \sum_{m=1}^M \mathbf{S}^{(m)} \mathbf{U} \mathbf{W} \left(\text{Tr}(\mathbf{X}\mathbf{W}) + w - y^{(m)} \right) . \quad (11)$$

After their initialisation, \mathbf{U} , $\mathbf{\Lambda}^{(m)}$, \mathbf{W} and w are iteratively updated using the L-BFGS-B. All parameters, except for \mathbf{W} , are initialized by drawing from a normal distribution with zero mean and standard deviation of one. The regression weights \mathbf{W} are initialized with unity. Non-negative entries are guaranteed for \mathbf{U} and all $\mathbf{\Lambda}^{(m)}$, as well as symmetry of each $\mathbf{\Lambda}^{(m)}$. The method stops when either the solution or the gradient does not change (up to machine precision).

2.2 Construction of the structural connectivity matrix

The structural connectivity matrix $\mathbf{S}^{(m)}$ of each subject contains a connection strength for each pair of ROIs, as defined above. It can be seen as a structural connectivity graph where the nodes represent the ROIs and the connection strengths represent the edge weights.

Connection strengths between pairs of brain regions are generated using shortest-path tractography (SPT), modelling the DWI as a weighted brain-graph $G = (V, E, w_E)$. The nodes V consist of all DWI voxels within the brain, edges E link adjacent voxels, and edge weights $w_E: E \rightarrow \mathbb{R}_+$ quantify the probability of connection along edges. Note that in contrast to the graph represented by $\mathbf{S}^{(m)}$, which models global connections between brain regions, the graph G models local connections between voxels based on their underlying fibre orientation distribution function (fODF) on a finer level.

In SPT, brain fibres are estimated as *most probable paths* connecting voxels in the brain. This becomes a shortest-path problem by negative log-transforming the edge connectivity probabilities. SPT computes pathways between all pairs of voxels for two given regions and provides a single score for each pathway. This measure is averaged over all possible pathways found between two ROIs i and j to obtain the corresponding (i, j) -th connection matrix entry for $\mathbf{S}^{(m)}$. We use the SPT implementation previously presented by Kasenburg et al.⁹ which estimates connection probability as a sample-based integral of the fODF associated to any given diffusion model.

Table 1. **Left:** Mean absolute error over all LOOCV folds in years for the original network and the hub networks from the SHD and UHD, where k is the number of hubs. The hubs are visualised in Figures 1, 2 and 3. **Right:** Mean reconstruction error (J_{recon}) without the regularizing terms over all LOOCV folds.

k	Baseline	SHD	UHD	k	SHD	UHD
10	10.0±7.5	9.6±7.7	37.0±27.2	10	0.29±0.36	0.19±0.00
15	10.0±7.5	10.9±9.3	16.2±11.7	15	0.22±0.10	0.21±0.34
20	10.0±7.5	10.5±7.4	11.8± 8.7	20	0.18±0.06	0.15±0.24

3. EXPERIMENTS AND RESULTS

We evaluate supervised hub-detection (SHD) by searching for structural connectivity hubs associated with increasing age in a healthy population, compared with unsupervised hub detection (UHD).²

3.1 Data and preprocessing

We analysed 54 healthy subjects (23 females and 31 males) of ages between 18 to 81. Diffusion weighted images¹⁰ (10 at $b = 0$ s/mm², 61 directions at $b = 1200$ s/mm²) were acquired using a twice-refocused balanced spin echo sequence to minimise eddy current distortion (TR = 8200 ms; TE = 100 ms; field of view (FOV)= 220 × 220 mm; matrix = 96 × 96; GRAPPA: factor 2, 48 lines; 61 transverse slices; no gap; 2.3 mm isotropic voxels). Whole brain T1-weighted magnetization prepared rapid acquisition gradient echo (MPRAGE) scans (voxel dimension of 1 × 1 × 1 mm, FOV= 256 × 256 × 192 mm, matrix = 256 × 256 × 192, TR = 1540 ms; TE = 3.93 ms ; inversion time (TI) = 800 ms, and a flip-angle of 9 degrees) were also acquired.¹⁰

Linear and non-linear registration between DWI and structural T1-weighted scans, and to MNI space, was performed with FLIRT¹¹ and FNIRT¹² supplied in FSL.¹³ Grey- and white matter tissue classifications of the T1-weighted images were obtained with FAST¹⁴ after brain extraction using BET.¹⁵ ExploreDTI¹⁶ was used to reduce distortions due to eddy currents, motion and susceptibility differences. Fibre orientation distribution functions based on constrained spherical deconvolution¹⁷ with order 8 were computed with the DiPy package.¹⁸ Brain regions were selected from the lateralized and sub-cortical Harvard Oxford atlas¹⁹ supplied in FSL.¹³ A boundary was computed using only those voxels lying on the white matter boundary (cortical) or the boundary voxels of a region (sub-cortical). The resulting 111 ROIs were used as target and seed regions for tractography to construct the structural connectivity matrices for every subject (see Section 2.2).

3.2 Evaluation scheme

Regression performance for SHD and UHD was compared using leave-one-out cross-validation (LOOCV), where for each subject the regression was trained on the remaining subjects. Hyper-parameter selection for all methods was done with 5-fold cross-validation on the training data. The model was then retrained with the optimal hyper-parameters on the whole training data.

Hyper-parameters were selected from the following ranges: $\gamma \in [0.001, 0.005]$ with 0.0005 steps and $\alpha, \beta \in \{0.001, 0.01, 0.1\}$, where $\alpha = \beta$. For the UHD, the value for $\alpha = \beta$ was similarly trained. Here \mathbf{W} and w were obtained by performing a ridge regression with regularization parameter α on the features $\mathbf{X} = \mathbf{U}^T \mathbf{S}^{(m)} \mathbf{U}$ learned by the hub-detection. Different values for the number of hubs k were tested independently to show different levels of dimensionality reduction. The regression performance of the UHD and SHD lower dimensional networks (see Equation 4) were compared with that of the original networks $\mathbf{S}^{(m)}$, training regression with the same values for α .

Table 1 shows the regression results for three different choices of k and the corresponding SHD and UHD hubs are visualized individually in Figures 1, 2 and 3 showing the hub membership score for every node emphasizing both the anatomical location and the relative importance of each hub.

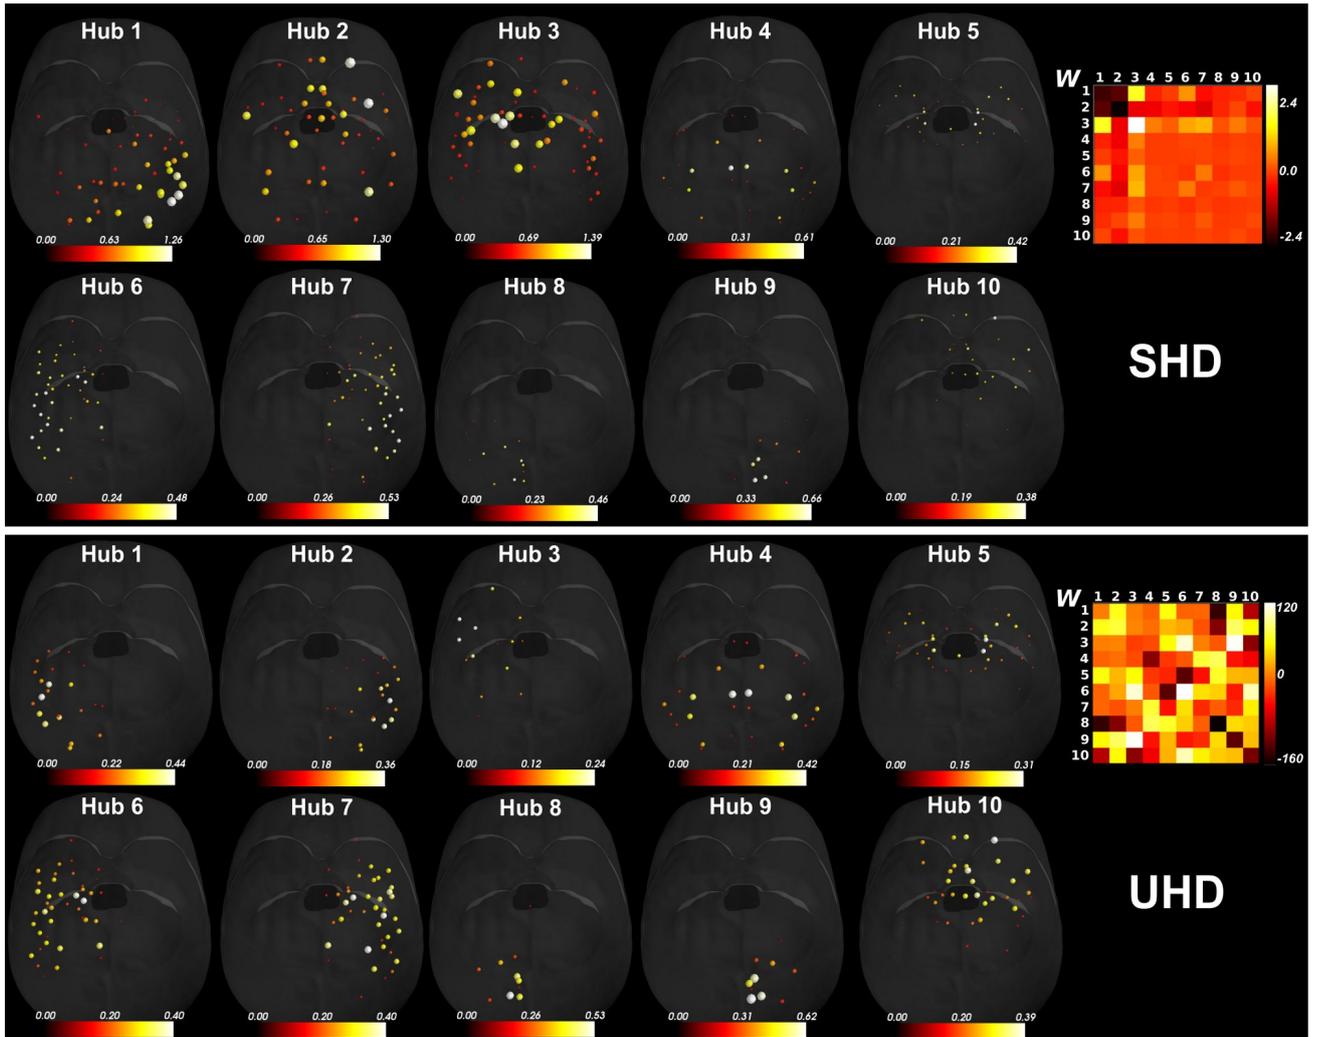


Figure 1. Hubs detected by SHD (top) and UHD (bottom) together with the hub-connections' regression weights \mathbf{W} (right) trained on the whole data for $k = 10$ with the best parameters (SHD: $\alpha = 0.1$, $\beta = 0.1$, $\gamma = 0.0025$; UHD: $\alpha = \beta = 0.001$) over all folds. The hub-membership scores for each node are represented by a colour scale (within hub) and by size (across all hubs). Note that the orientation of the coordinate system in which the brains are shown is from right to left.

4. DISCUSSION

It is clear from Table 1 that the lower-dimensional networks found by SHD retain the same regression performance as the original networks, while those found by UHD get a much higher RMSE with a higher standard deviation. It is expected that the mean reconstruction errors, computed as the squared norm of the differences divided by the number of subjects, is lower for the SHD than the UHD, since the SHD is optimizing for both reconstruction and regression performance. Nevertheless, the mean reconstruction error for SHD is relatively similar to that of the UHD, showing that the SHD hubs still represent the population well.

Figure 1 shows that both SHD and UHD extract symmetric hubs (hubs 4 and 5 in both) or hubs with a symmetric counterpart (SHD: hubs 6 & 7, 8 & 9; UHD: hubs 1 & 2, 6 & 7 and 8 & 9) for $k = 10$. The SHD hubs are more distributed throughout the brain (hubs 2 and 3) and less focal than UHD hubs (e.g. hubs 8 and 9). The same observations can be made for other choices of k (see Figures 2 and 3), but due to the larger number of hubs, the symmetric hubs consist of fewer more localized nodes. Nevertheless, additional hubs found by the

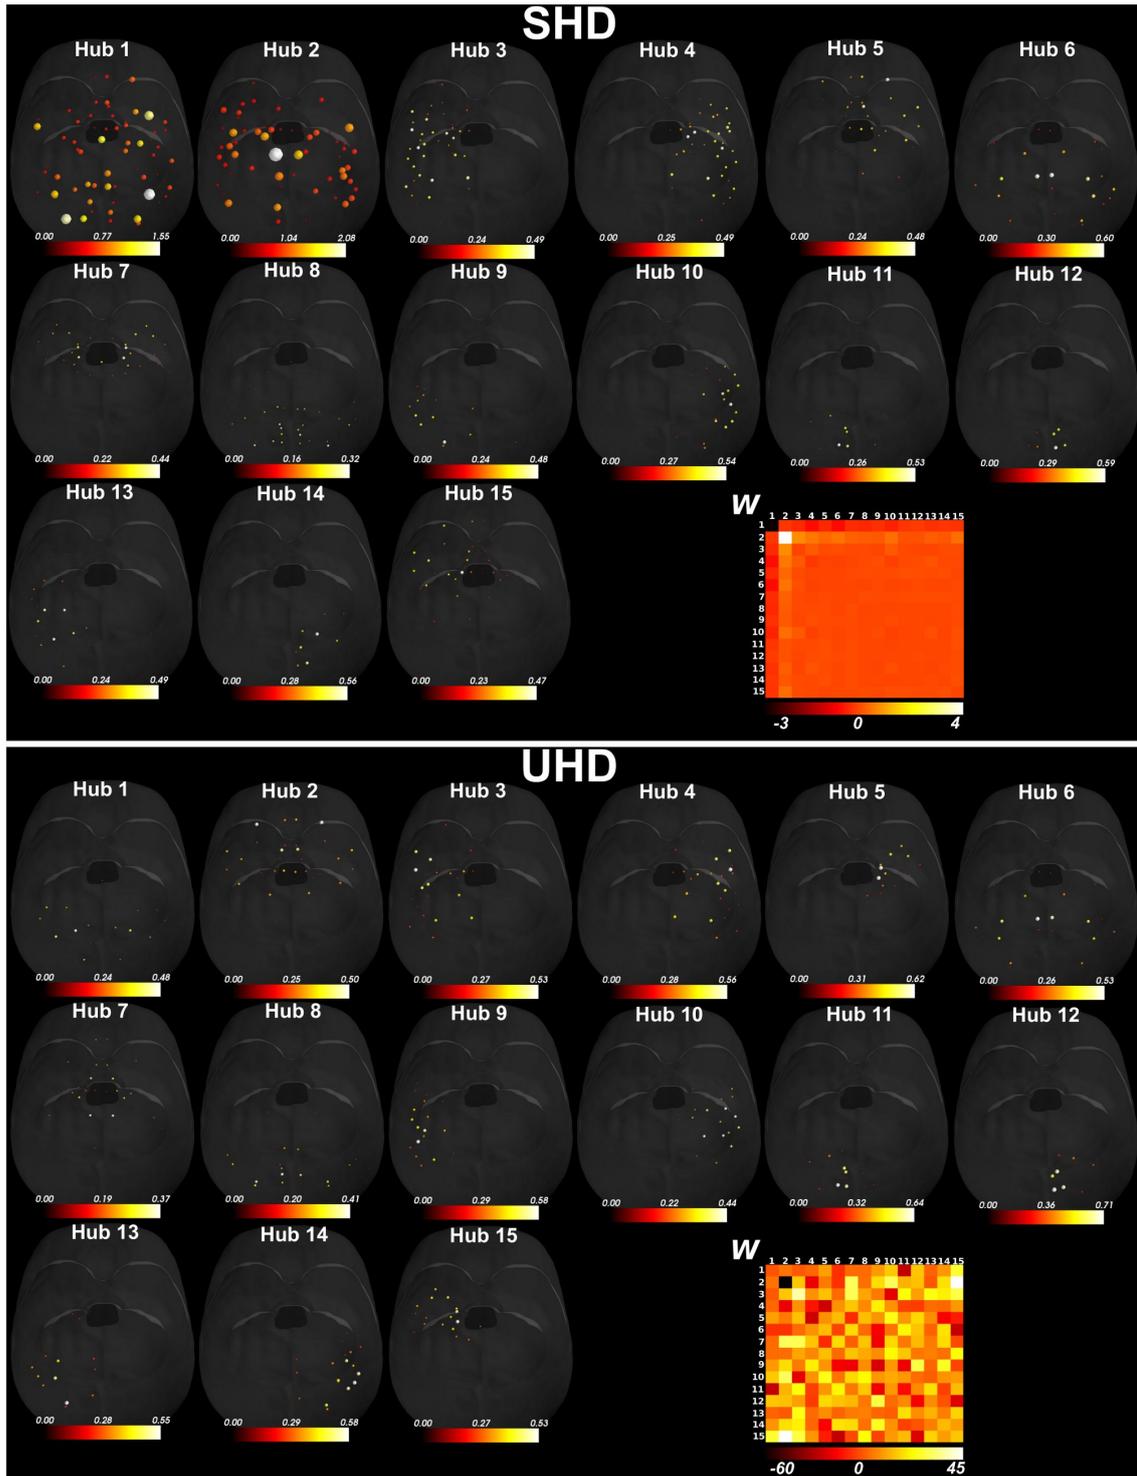


Figure 2. Hubs detected by SHD (top) and UHD (bottom) together with the hub-connections' regression weights \mathbf{W} (bottom right) trained on the whole data for $k = 15$ with the best parameters (SHD: $\alpha = 0.1$, $\beta = 0.1$, $\gamma = 0.0025$; UHD: $\alpha = \beta = 0.01$) over all folds. The hub-membership scores for each node are represented by a colour scale (within hub) and by size (across all hubs). Note that the orientation of the coordinate system in which the brains are shown is from right to left.

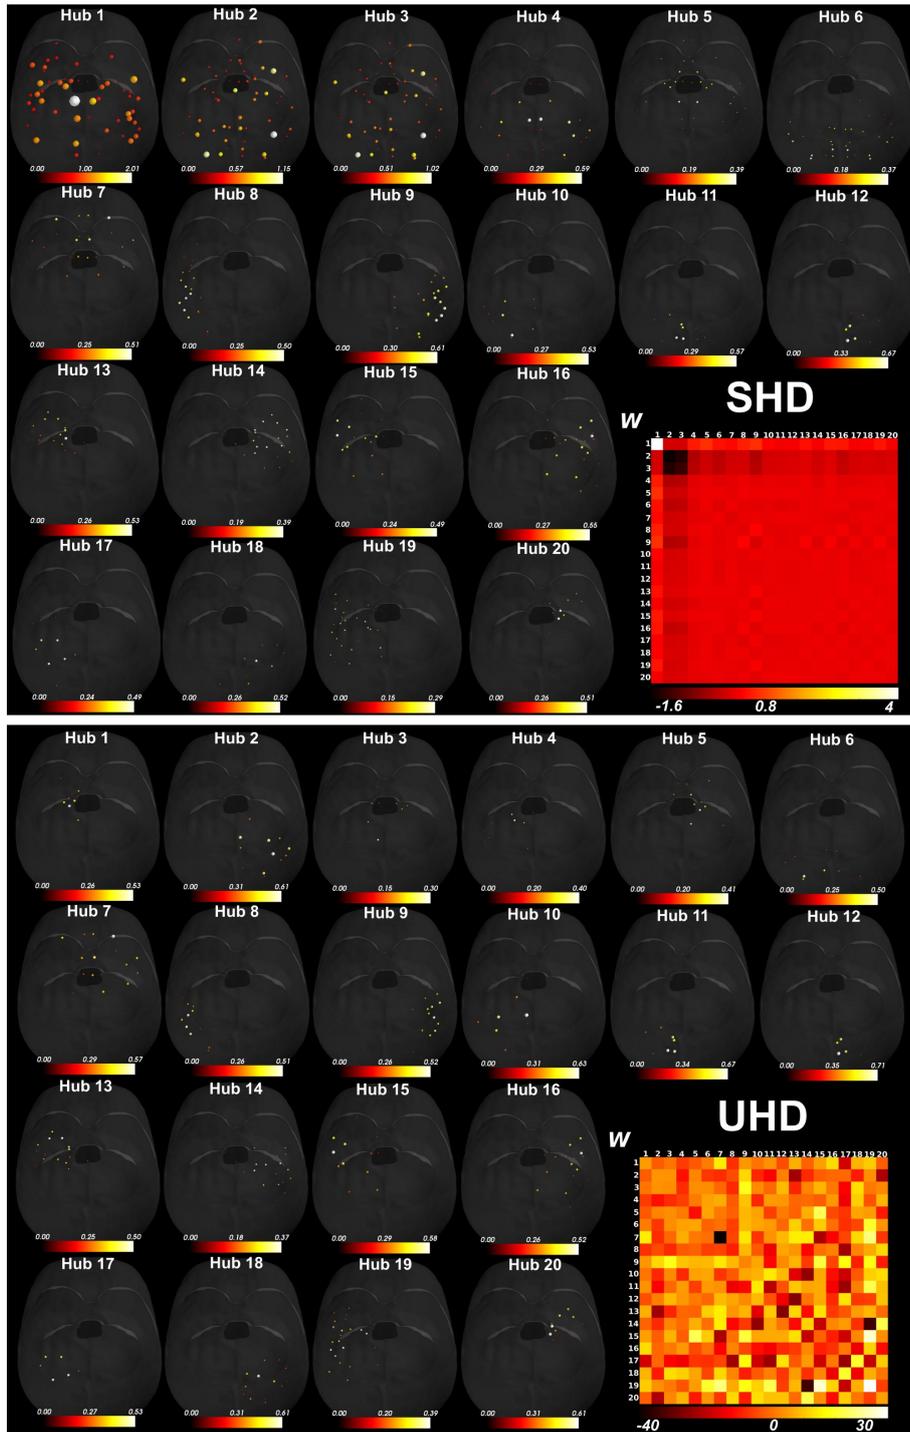


Figure 3. Hubs detected by SHD (top) and UHD (bottom) together with the hub-connections' regression weights \mathbf{W} (bottom right) trained on the whole data for $k = 20$ with the best parameters (SHD: $\alpha = 0.1$, $\beta = 0.1$, $\gamma = 0.0025$; UHD: $\alpha = \beta = 0.01$) over all folds. The hub-membership scores for each node are represented by a colour scale (within hub) and by size (across all hubs). Note that the orientation of the coordinate system in which the brains are shown is from right to left.

SHD are still globally distributed over the whole brain for larger choices of k .

In addition, global SHD hubs have more impact on the regression than focal hubs, reflected by their high absolute regression weights ($k = 10$: intra-hub connection of SHD hubs 1, 2 and 3, and inter-hub connection between SHD hubs 1 & 2 in Figure 1; $k = 15$: intra-hub connection of SHD hubs 1 and 2 in Figure 2; $k = 20$: intra-hub connection of SHD hubs 1, 2 and 3, and inter-hub connection between SHD hubs 2 & 3 in Figure 3). This may indicate that age related changes are not focal, but affect tracts throughout the brain (in agreement with Lebel et al.²⁰). In fact, human brain development is a complex process, continuing into early adulthood with a relative plateau during middle-age, followed by variable decline.²¹ This also explains the large standard deviation in prediction performance (see Table 1).

Although found hubs are consistent between different choices of k , the regression performance is only stable for the SHD while it improves for the UHD (see Table 1). It is expected that for large k the regression performance for both hub-detection methods gets closer to the baseline as fewer nodes are within a hub up to a single node per hub when the number of hubs is equal to the number of nodes ($k = n$). While the SHD achieves the best regression performance for $k = 10$, a larger number of hubs ($k = 15$) allows the symmetric hubs to be smaller and more localized (see Figure 2), which reflects meaningful anatomical units. In summary, an optimal choice for k lies between 10 and 15.

5. CONCLUSION

We have presented a supervised hub-detection algorithm for dimension reduction of brain connectivity networks, which clusters network nodes into *hubs*, while retaining predictive performance in a given regression task. The presented algorithm was applied to structural connectivity networks, but it applies to any network regression problem. Future work includes development of SHD algorithms for other predictive models such as classification and non-linear regression to further extend the applicability.

REFERENCES

- [1] Robinson, E. C., Hammers, A., Ericsson, A., Edwards, A. D., and Rueckert, D., “Identifying population differences in whole-brain structural networks: A machine learning approach,” *NeuroImage* **50**(3), 910–919 (2010).
- [2] Ghanbari, Y., Bloy, L., Shankar, V., Edgar, J., Roberts, T., Schultz, R., and Verma, R., “Functionally driven brain networks using multi-layer graph clustering,” in [*Medical Image Computing and Computer-Assisted Intervention – MICCAI 2014*], *Lecture Notes in Computer Science* **8675**, 113–120, Springer International Publishing (2014).
- [3] Fornito, A., Zalesky, A., and Breakspear, M., “The connectomics of brain disorders,” *Nat Rev Neurosci* **16**(3), 159–172 (2015).
- [4] Cheplygina, V., Tax, D. M. J., Loog, M., and Feragen, A., “Network-guided group feature selection for classification of autism spectrum disorder,” in [*Machine Learning in Medical Imaging*], *Lecture Notes in Computer Science* **8679**, 190–197, Springer International Publishing (2014).
- [5] Rubinov, M. and Bullmore, E., “Fledgling pathoconnectomics of psychiatric disorders,” *Trends in Cognitive Sciences* **17**(12), 641–647 (2013). Special Issue: The Connectome.
- [6] Clayden, J. D., Dayan, M., and Clark, C. A., “Principal networks,” *PLoS One* **8**(4), e60997 (2013).
- [7] Jenatton, R., Gramfort, A., Michel, V., Obozinski, G., Eger, E., Bach, F., and Thirion, B., “Multiscale mining of fMRI data with hierarchical structured sparsity,” *SIAM Journal on Imaging Sciences* **5**(3), 835–856 (2012).
- [8] Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C., “A limited memory algorithm for bound constrained optimization,” *SIAM Journal on Scientific Computing* **16**(5), 1190–1208 (1995).
- [9] Kasenburg, N., Liptrot, M., Reisle, N. L., Ørting, S. N., Nielsen, M., Garde, E., and Feragen, A., “Training shortest-path tractography: Automatic learning of spatial priors,” *NeuroImage* (2016).
- [10] Ramsøy, T. Z., Liptrot, M. G., Skimminge, A., Lund, T. E., Sidaros, K., Christensen, M. S., Baar, W., Paulson, O. B., Jernigan, T. L., and Siebner, H. R., “Healthy aging attenuates task-related specialization in the human medial temporal lobe,” *Neurobiol Aging* **33**, 1874–1889 (Sep 2012).

- [11] Jenkinson, M., Bannister, P., Brady, M., and Smith, S., “Improved optimization for the robust and accurate linear registration and motion correction of brain images,” *NeuroImage* **17**, 825–841 (Oct 2002).
- [12] Andersson, J. L. R., Jenkinson, M., and Smith, S., “Non-linear registration, aka spatial normalisation,” tech. rep., FMRIB (2010).
- [13] Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., and Smith, S. M., “FSL,” *NeuroImage* **62**, 782–790 (Aug 2012).
- [14] Zhang, Y., Brady, M., and Smith, S., “Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm,” *IEEE Trans Med Imaging* **20**, 45–57 (Jan 2001).
- [15] Jenkinson, M., Pechaud, M., and Smith, S., “BET2: MR-based estimation of brain, skull and scalp surfaces,” in [*Eleventh annual meeting of the organization for human brain mapping*], **17** (2005).
- [16] Leemans, A., Jeurissen, B., Sijbers, J., and Jones, D. K., “ExploreDTI: a graphical toolbox for processing, analyzing, and visualizing diffusion MR data,” in [*17th Annual Meeting of Intl Soc Mag Reson Med*], 3537 (2009).
- [17] Tournier, J.-D., Calamante, F., and Connelly, A., “Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution,” *NeuroImage* **35**, 1459–1472 (May 2007).
- [18] Garyfallidis, E., Brett, M., Amirbekian, B., Rokem, A., van der Walt, S., Descoteaux, M., Nimmo-Smith, I., and Dipy Contributors, “Dipy, a library for the analysis of diffusion MRI data,” *Front Neuroinform* **8**, 8 (2014).
- [19] Desikan, R. S., Sgonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., Buckner, R. L., Dale, A. M., Maguire, R. P., Hyman, B. T., Albert, M. S., and Killiany, R. J., “An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest,” *NeuroImage* **31**, 968–980 (Jul 2006).
- [20] Lebel, C., Gee, M., Camicioli, R., Wieler, M., Martin, W., and Beaulieu, C., “Diffusion tensor imaging of white matter tract evolution over the lifespan,” *NeuroImage* **60**(1), 340–352 (2012).
- [21] Westlye, L. T., Walhovd, K. B., Dale, A. M., Bjørnerud, A., Due-Tønnessen, P., Engvig, A., Grydeland, H., Tamnes, C. K., Ostby, Y., and Fjell, A. M., “Life-span changes of the human brain white matter: diffusion tensor imaging (DTI) and volumetry,” *Cereb Cortex* **20**(9), 2055–2068 (2010).