Dissertation
submitted to the
Combined Faculties for the Natural
Sciences and for Mathematics
of the Ruperto-Carola University of
Heidelberg, Germany
for the degree of
Doctor of Natural Sciences

Put forward by

Diplom-Physiker: Caspar Jonas Goch
Born in: Essen

Oral examination: 02.07.2014

Expanding Graph Theoretical Indices to Include Medical Knowledge - An Assessment of Classification Accuracy in the Case of Autism Spectrum Disorders

Referees: Prof. Dr. Peter Bachert
Prof. Dr. Hans-Peter Meinzer

Zusammenfassung

Im letzten Jahrzehnt hat die graphtheoretische Analyse der Netzwerkarchitektur des menschlichen Gehirns, das sogenannte Konnektom, zunehmend an Interesse gewonnen. In dieser Arbeit erweitern wir Netzwerk Indizes, die für andere Anwendungen entwickelt wurden, um zusätzliche medizinische Informationen in Betracht zu ziehen. Anschließend evaluieren wir diese erweiterten Indizes auf einem Kollektiv von Autismuspatienten. Wir vergleichen die vorgeschlagenen Indizes mit anderen, traditionell in der Literatur benutzten, und zeigen, dass sie ein wirkungsvolles neues Werkzeug in der Analyse des Konnektoms sein können. Wir unterziehen sie dann ein weitergehenden Analyse um ihre Leistung in Abhängigkeit der Netzwerkdichte einschätzen zu können und den Dichtebereich zu ermitteln in dem sie die besten Informationen liefern.

Unser Ziel ist es durch so eine umfassende Analyse die Zahl der willkürlich gewählten Parameter in der Analyse des Konnektoms weiter zu reduzieren und zu helfen die Vergleichbarkeit verschiedener Studien zu erhöhen.

Abstract

Using graph theory to analyse the architecture of the human brain, the connectome, has gained increasing interest in the last decade. In this work we extend graph measures, which have previously been developed for other applications, to include prior medical information. These extended measures are then evaluated on a collective of Autism Spectrum Disorder patients. Examining their performance in comparison to other traditionally used measures we show, that they are a valuable new tool in the analysis of the human connectome. We then further evaluate their performance over a range of network densities in order to determine the range at which they supply the most valuable information.

By doing an in depth evaluation of these measures we aim to reduce the amount of guesswork in choosing variables in the analysis of the connectome and help to improve the comparability of different studies.

Contents

1 Introduction

1.1 Motivation

In the past decade the use of graph analysis to describe the architecture of the human brain has become more and more common. By using graph theoretical techniques known from other applications to examine these brain networks, called connectomes, a rich field has opened up which promises to yield new insights in the brain and the changes due to a variety of diseases.

This new technique offers up exciting possibilities, however there are many challenges still to be resolved. In contrast to many traditional applications of graph theory the connectomes that are currently used in research tend to be far smaller in scale, only having between less than a hundred and a few thousand nodes. At the same time the brain is an organ, which has long fascinated mankind and has a long history of medical research. Moving the connectome from a purely abstract description of the connectivity of different points in space and incorporating medical information can lead to a better understanding of the nature of the connectome and how to accurately represent and read it [55].

At the same time a lot of the research has been exploratory in nature as often happens when a new field opens up [87]. Using this groundwork to find applications which can potentially benefit the medical practice and provide new diagnostic tools is an important part of the continuing scientific progress.

By studying the research done on the changes in diseased connectome [45] the possibilities of adapting previous methodology to fit the situation in the human brain becomes apparent. By applying these new techniques to actual patient data it is possible to examine their performance and identify the possibilities and limitations in their use.

1.2 Objectives

The main objective of this work is to improve the understanding of the human connectome and examine the suitability of using a range of network indices for diagnostic purposes. A secondary objective is to supply researchers with a working open-source solution for a pipeline from the original data to the analysed connectome.

In order to achieve the main objective there are multiple avenues that are pursued in this work.

1. We try to establish a baseline of the diagnostic quality of existing measures, especially the commonly used ones. This is important in order to compare any

new measures developed in this work to pre-existing ones.

2. We want to improve upon the state of the art by extending existing graph theoretical measures to incorporate prior anatomical and disease symptomatology specific information.

3. Any new measure should be thoroughly investigated to understand its properties and judge its applicability.

The second objective is served by integrating the pipeline developed in this work with the Medical Imaging Interaction Toolkit [35], providing an easy-to-use open source implementation.

The evaluation of our work on a collective of Autism Spectrum Disorder (ASD) patients serves the idea of connecting symptoms with unknown causes with changes in the connectome. A new measure developed on the basis of these symptoms and showing the changes potentially responsible for this will advance our understanding of ASD.

1.3 Outline

Chapter 2 provides the technical and medical background necessary for this work. In it we introduce the concept of diffusion weighted Magnetic Resonance Imaging and detail the processes taking place between the acquisition of the data and the work detailed in this document. Additionally we provide a quick overview of the terminology and concepts of graph theory that are required for this work.

Chapter 3 gives an overview of the current state of the art regarding the construction of the connectome and its analysis. With a focus on the structural connectome we give an overview of current techniques and problems.

The methods developed as part of this work are described in chapter 4. We describe the global graph indices used for a baseline evaluation of current global techniques as well as the methods used to establish a diagnostic significance using classification. We then continue on to describing a new regional approach to network measures using prior anatomical knowledge. Lastly we examine the proposed regional approach in detail and examine some of the problems arising from the different methods for connectome construction and their impact on our proposed measures.

The results of these investigations are presented in chapter 5 and discussed in detail in chapter 6. We then continue to detail the impact of our work and suggest future courses of investigation based on our results.

2 Background

2.1 Medical Background

In this section we will give a short introduction into the human brain and Autism Spectrum Disorders (ASD) in order to set this work into a medical context.

2.1.1 The Brain

The human brain is the main information processing organ of the human body. It collects information from the various senses, integrates this information, sets it into context depending on previously made experiences, provides feedback and controls the movement of the various muscles of the body. In order to carry out its task its architecture provides it with significant processing power and data storage facilities. It is responsible for up to twenty percent of the bodies energy consumption.

The processing power is provided by billions of specialized cells, the neurons, which comprise the brain and are active concurrently. A neuron consists of three distinct parts. The branching dendrites receive information from other neurons. They can form a complex tree extending hundreds of micrometers and connecting to up to a thousand other neurons. The cell body or soma houses the organelles of the cell and continues into the axon, which can extend from a few millimetres to metres, depending on the role of the neuron. The axon propagates the signal to other neurons and can reach up to a thousand other neurons. Figure 2.1 shows a sketch of a neuron. This dense web of interconnecting neurons is responsible for the complex calculations the human brain is capable of.

Figure 2.2 shows a cross-section of the human head. An easily recognizable difference in the tissue of the human brain is the division into two distinctly coloured tissue types. The grey matter, making up the outer shell of the brain, is where the dendrites and cell bodies are located. The axons make up the inner white matter, they are densely packed in bundles connecting different areas of the grey matter. This dense packing is believed to be the main reason for the directional dependency of water diffusion in the white matter as water molecules can move along these tracts but not across the axon walls.

2.1.2 Autism Spectrum Disorder

Autism Spectrum Disorders (ASD) encompasses a range of conditions which are classified as neurodevelopmental disorders [37]. While the specific symptomatology of these conditions can be different, they have a range of symptoms in common.

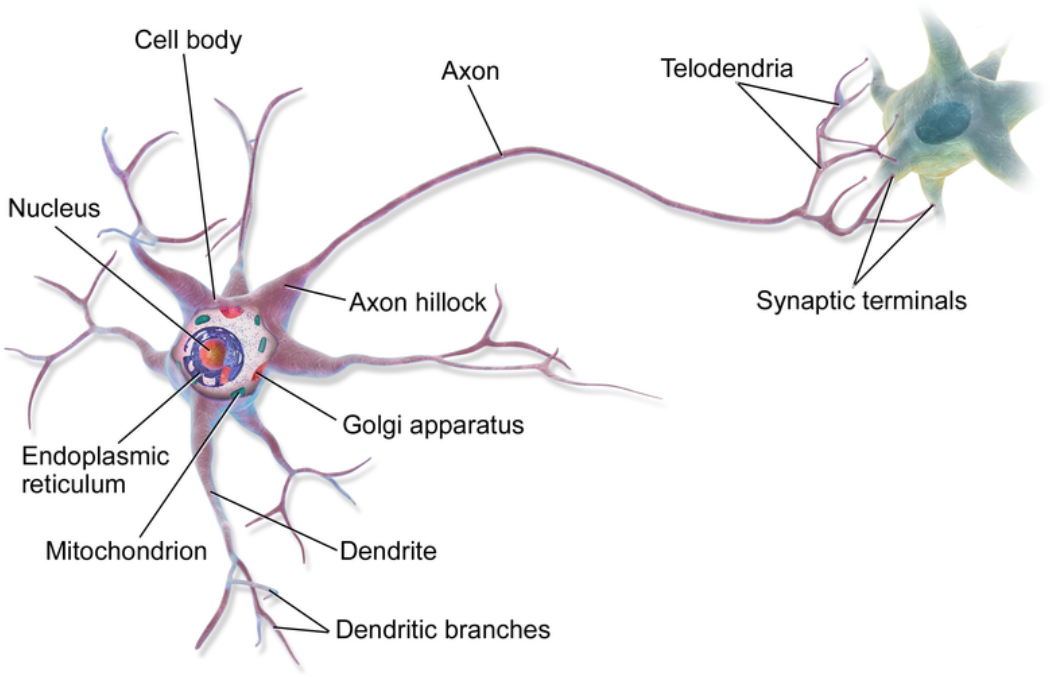

Figure 2.1: A schematic image of a neuron. Taken from Wikipedia (original author BruceBlaus)

Patients suffering from any of these disorders typically suffer from social deficits, reduced communicational capacity, repetitive behaviour and potentially cognitive delays [71, 44].

Specific causes for ASD are currently not known, resulting in the definition being based mainly on the symptoms. Using new techniques to gain a better understanding of these disorders will therefore be important as about 1 in 88 children in the US suffer from ASD and as many as 1.5 million Americans are believed to have some form of ASD [3].

2.2 MR Diffusion Imaging

The analysis of the human connectome is the last step in a long and complex pipeline. Our work is solely focussed on the last steps of this pipeline, namely the construction of the connectome and the analysis thereof. In this section we give an overview of the prior pipeline in order to place our work in context.

2.2.1 Acquisition

In contrast to Computer Tomography (CT), that is based on the absorption of gamma radiation by the tissue of the body, Magnetic Resonance Imaging (MRI) is a non invasive method based on the measurement of the response of non zero spin

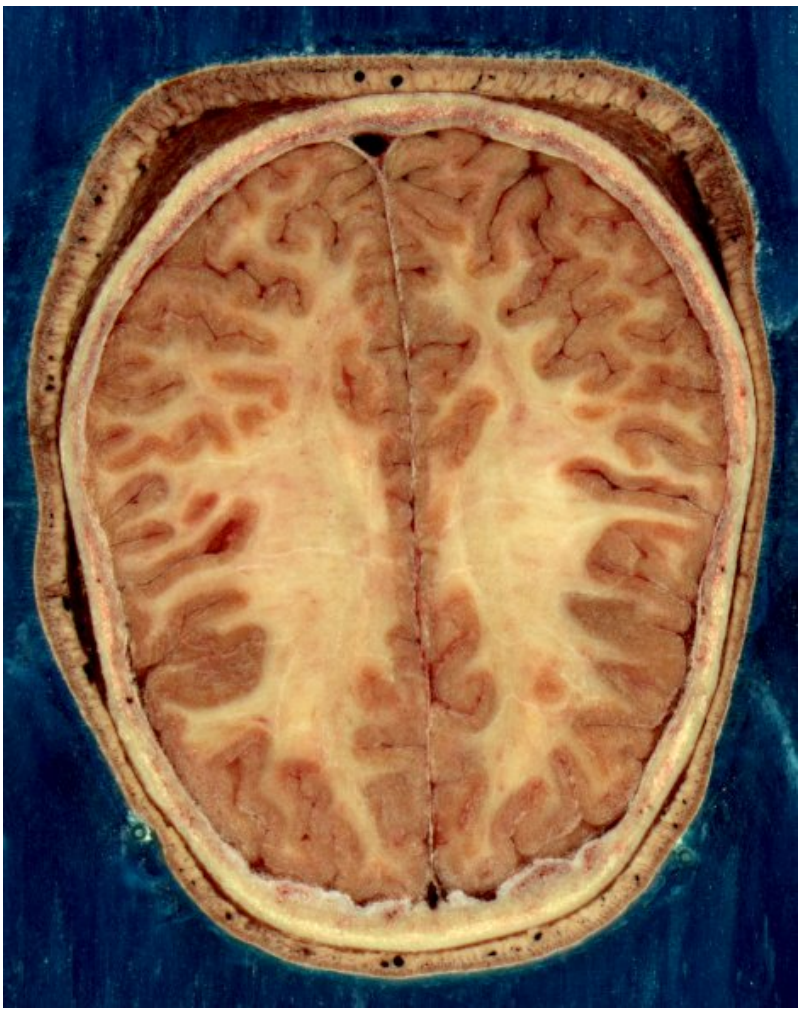

Figure 2.2: A cross section of the human head. The differences in the texture of the grey and white matter are easily visible. Image taken from Wikipedia(released to the public domain by the US government)

atoms to an external magnetic field. The most abundant such atom in the human body is the hydrogen atom. It is comparatively so much more abundant, that the main signal in the body will be due to ^1H atoms, unless specific contrast agents are introduced. We will henceforth consider only the ^1H atom, which can have a spin of either $m = +1/2$ or $m = -1/2$.

Spins in a Magnetic Field

Without an external magnetic field both spin states have the same energy and the spins will be equally distributed. By inserting a tissue sample or patient in a large electro-magnet, the Magnetic Resonance Tomograph, an external magnetic field is applied. This splits the spin states and introduces an energy difference between them.

The energy of a magnetic moment $\vec{\mu}$ in an external magnetic field $\vec{B}_0$ is given by:

$$E = -\vec{\mu} \cdot \vec{B}_0 \tag{2.1}$$

by choosing our coordinate system so the z-axis is aligned along $\vec{B}_0$ this simplifies to:

$$E = -\mu_z B_0 = -\gamma \hbar m B_0 \tag{2.2}$$

γ being the gyromagnetic ratio.
The energy difference between the two states is therefore equal to:

$$\Delta E = \gamma \hbar B_0 \tag{2.3}$$

Biasing the distribution of spin states towards the lower energy one.
The magnetic moment of the atoms will precess around the external field due to the torque exerted by it. This Larmor precession will happen at an angular frequency of

$$\omega = \gamma B_0 \tag{2.4}$$

Magnetic Resonance Imaging

If the external magnetic field is applied for a long enough time the net magnetization of all spins will be parallel to the magnetic field. Applying a radio frequency (RF) pulse of a frequency matching the resonance (Larmor) frequency will cause some of the spins to absorb the photons and flip to an excited state. This will reduce the longitudinal magnetization, while at the same time, due to phase coherence, establishing a magnetization perpendicular to the external magnetic field. The length of the RF pulse can be chosen to flip the magnetization by 90 ° in which case we speak of a 90 ° pulse. A pulse twice as long will flip the magnetization by 180 ° and is known as 180 ° pulse.

Given a long enough time the spins will return to the thermal equilibrium. The mean time for this to happen to an individual spin is designated T_1 and the process called T_1 relaxation. The length of T_1 is dependent on the strength of the magnetic field, but for our purposes can be seen as somewhere between 800 ms and 1 s. The second relaxation process, called T_2 relaxation is due to molecular interactions and leads to a decoherence of the spins and an exponential decay of the transverse magnetization with a time constant of T_2. The T_2 relaxation is in most circumstances the faster process with T_2 being in the order of 200 ms.

Inhomogeneities of the magnetic field lead to differences in B_0 depending on the location of the spin and lead to an even faster decoherence of the spins with a time

constant of T_2^*. By applying a 180 ° pulse some time after the 90 ° pulse the position of the spins will be switched, with faster processing ones now lagging behind in phase and after the same time interval the spins will be briefly in phase again leading to a spin-echo. This can be measured by receiver coils placed close to the patient.

In order to create a 3D image from 2D image slices an additional magnetic gradient can be applied so that the Larmor frequency will be dependent on the location of the spins along the gradient. This way only spins on one plane will be excited by a specific frequency. The thickness of this slice will depend on the bandwidth of the RF pulse and the magnitude of the field gradient.

Similar to the slice selection the readout of specific voxels uses additional gradients. The slices are in the x-y-plane and additional gradients G_x and G_y can be applied. Using this method each spin has a location dependent Larmor frequency and will acquire a phase Θ as follows:

$$\Theta(x, y) = 2\pi(k_x x + k_y y) \tag{2.5}$$

with

$$k_i = \gamma \int G_i(t)dt \tag{2.6}$$

for $i = x$ and $i = y$ respectively.

The resulting net magnetization across an excited slice can then be represented as a function $f(x, y)$. The signal $s(t)$ measured in the receiver coil is the integrated signal from all spins:

$$s(t) \propto F(k_x, k_y) = \int f(x, y)e^{i2\pi(k_x x + k_y y)}dx\,dy \tag{2.7}$$

Equation 2.7 is a fourier transformation and $f(x, y)$ can be reconstructed using an inverse fourier transformation, given enough samples of k_x and k_y.

Diffusion Weighted Magnetic Resonance Imaging

In 1994 a method was presented for using diffusion weighted Magnetic Resonance Imaging (dMRI) to estimate anisotropic diffusion of water along axons [6]. This opened a way to non-invasively examine the micro-structure of the human brain in-vivo. Using this information it became possible to quantify information about the fiber tracts in the brain, which previously had been done mainly in autopsies and using tracer experiments on monkeys.

The premise of diffusion-weighted imaging (DWI) is similar to MRI. It is an extension of the latter using information about the diffusion characteristics of water at a location to infer the micro-structure limiting the diffusion at the same location.

As has been noted previously applying a magnetic gradient during the MRI measurement will cause the gradients to loose coherence (see Eq. 2.5). If a reversed gradient is applied for the same amount of time the spins will then rephase. This however only works if the spins did not move between the gradients, otherwise the rephasing will not be perfect and the echo will be reduced. When this was first noted in an inhomogeneous magnetic field it was suggested that it could be used to measure the diffusion coefficient of a solution [48]. This idea was later taken to suggest a modification of the standard spin-echo sequence [15]. A refinement of this technique, called pulsed gradient spin-echo (PGSE), no longer applies the gradient during the entire acquisition, but instead uses two distinct gradient pulses [92]. The advantage is that this does allow a clear distinction between the encoding time δ and the time separating the two gradient pulses Δ. Varying the strength of the gradients allows for varying sensitivity to diffusion. The so called b-value reflects the gradient strength:

$$b = q^2\left(\Delta - \frac{\delta}{3}\right) = \gamma^2 G^2 \delta^2 \left(\Delta - \frac{\delta}{3}\right) \tag{2.8}$$

q is here the product of the strength and the length of the gradient pulses.

The first pulse takes place between the excitation pulse and the refocussing pulse, the second between the refocussing pulse and the spin echo. If we assume δ to be too short for diffusion to take place during the pulse the difference in phase will be a function of the location during the first pule x_1 and the second pulse x_2:

$$\Phi_2 - \Phi_1 = q(x_2 - x_1) \tag{2.9}$$

If the particles remained stationary between the pulses they will rephase and the difference will be zero. If they moved however the rephasing will not be perfect and the signal will be reduced.

The signal attenuation is calculated by dividing the observed signal using diffusion weighting $S(q)$ by the signal without diffusion weighting S_0:

$$E(q) = \frac{S(q)}{S_0} = \int \rho(x_1) \int P(x_1, x_2, \Delta) e^{-iq(x_2 - x_1)} dx_2 dx_1 \tag{2.10}$$

Where $\rho(x_1)$ is the spin density at x_1 during the application of the first gradient pulse. For most applications this will be set to 1. $P(x_1, x_2, \Delta)$ is the diffusion propagator and depends on the used diffusion model. It describes the likelihood of a particle at location x_1 during the first pulse to be at location x_2 during the second one. In the case of free diffusion this will be a Gaussian. $e^{-iq(x_2-x_1)}$ is the Fourier kernel used for the frequency response of a time-dependent signal.

If we assume free diffusion equation 2.10 simplifies to

$$E(q) = e^{-q^2 D(\Delta - \frac{\delta}{3})} = e^{-bD} \tag{2.11}$$

Which can be used to quickly calculate the diffusion coefficient from MRI measurements. As the diffusion is hindered in the brain the resulting diffusion coefficient will be lower than the free diffusion coefficient in water.

In the white matter the diffusion will usually depend on the direction of the diffusion encoding gradient. This is due to the packing of axons, which will hinder diffusion in certain directions, while allowing it in others.

2.2.2 Preprocessing and Modelling

Modelling the Diffusion

In the case of anisotropic diffusion the diffusion can no longer be expected to be a simple coefficient as it is direction dependent. Instead we can use a tensor to describe the diffusion.

$$\begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix} \tag{2.12}$$

The orthogonal elements of the tensor represent the diffusivity along the axes of the laboratory (scanner) frame of reference. The off-diagonal elements represent the correlation of the displacements along the axes. Figure 2.3 represents the diffusion tensor as ellipsoid in different cases.

The diffusion tensor is symmetric and as such has six independent variables. In order to calculate the tensor at least six different gradient directions have to be acquired in addition to an image without diffusion weighting.

Using the diffusion tensor we can describe regions of high anisotropy which are likely to contain densely packed axons. This model holds only true if the voxel contains a single fibre. It can however not distinguish between regions with lower anisotropy and those with a high anisotropy in multiple directions, such as crossings of densely packed fibres. Advanced techniques to more accurately describe such situations have been developed. One example is Q-ball imaging which uses the Funk-Radon Transform to represent the diffusivity in many directions [94]. The disadvantage is however the need to acquire more gradient directions, leading to longer acquisition times.

Artefact Correction

When working with dMRI data there are several pitfalls that should be considered in order to avoid artefacts due to the acquisition technique [57].

The most notable artefacts are induced eddy-currents due to the changing magnetic field during acquisition. These lead to a mis-registration as the image can be shifted or sheared, depending on the direction of the eddy-current. Given a long enough time constant eddy-current artefacts can be reduced by slice-by-slice global affine transformations of the image to compensate for the distortions.

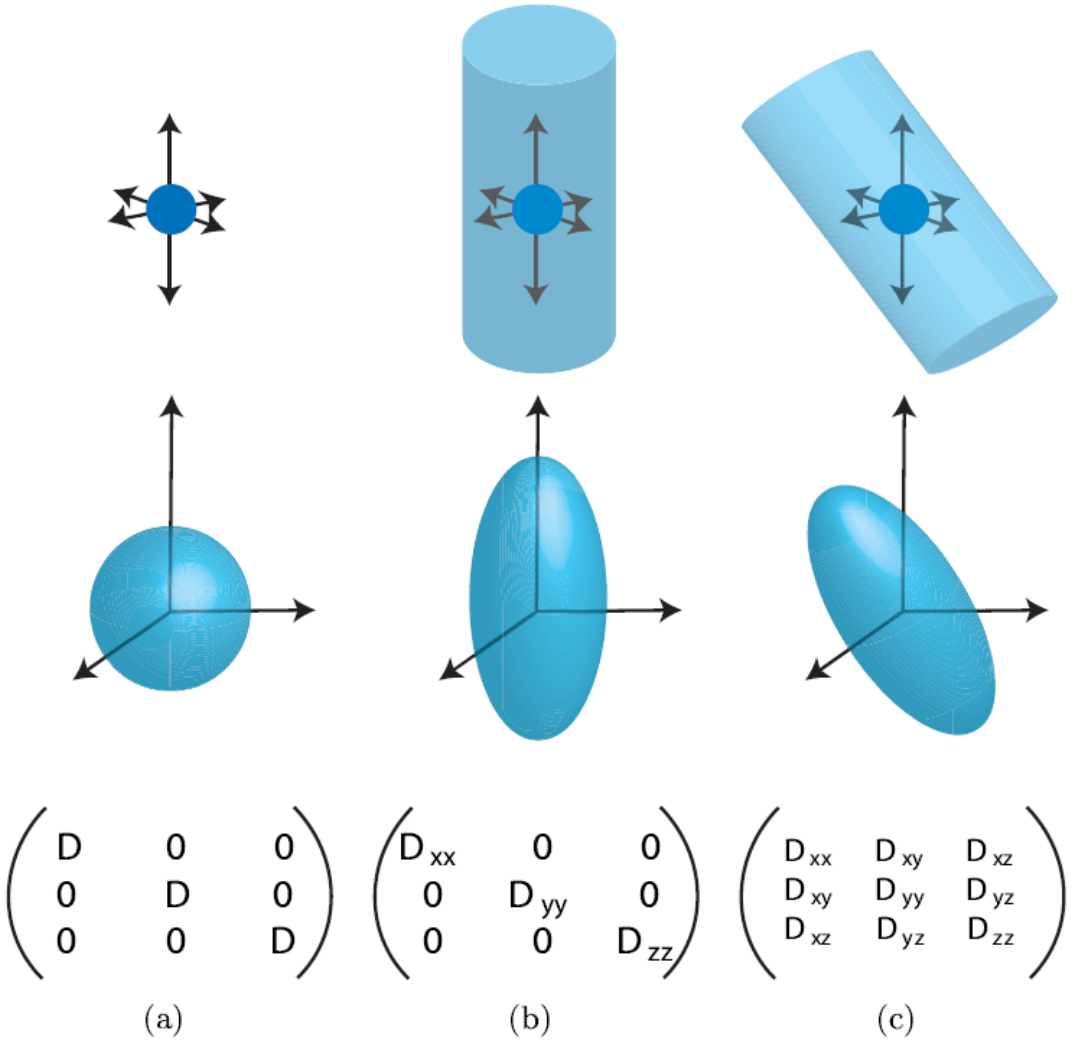

Figure 2.3: (a) Free diffusion: The diffusion ellipsoid is spherical. (b) and (c) Anisotropic diffusion: The diffusion ellipsoid is cigar shaped. In (b) the eigensystem of the tensor and the laboratory frame of reference align, in (c) they do not. Reused from [63]

Another common cause of artefacts is due to the movement of the subject during image acquisition. This becomes more pronounced the longer the acquisition takes and is especially notable in sequences which include a lot of gradients or repeat the same gradients in order to improve the signal to noise ratio. Possible solutions include registering the diffusion weighted images to the non-diffusion weighted image. Another option is to intersperse the diffusion-weighted images with multiple non-diffusion weighted images and interpolating a transformation based on the registrations of the non-diffusion weighted images. This second approach has the advantage of using a simpler metric for the registration.

2.2.3 Tractography

Diffusion-weighted images give an indication of the direction of axon walls in the white matter at a certain voxel. This information can be used to estimate the path of axon packages across longer distances in the brain. The methods used for this purpose are collectively known as tractography methods and the resulting image of fibres in the brain is known as tractogram.

The simplest method to trace a fibres path in the brain is by starting at a random seed point in the white matter and following the principal diffusion direction of the voxel until a voxel border is encountered. The direction is then changed to the principal direction of the new voxel and the process repeats until the anisotropy of the voxel drops below a threshold value. The result of this process is one continuous path from one low anisotropy voxel to another, which can be assumed to correspond to a densely packed bundle of axons in the brain. If this process is repeated for a large number of seed points we can expect to have a good representation of the principal fibre tracts in the human brain [21]. This process is very fast and delivers deterministic results, but due to the nature of the algorithm the resulting fibres are not able to cross regions with a lower anisotropy, such as crossing axon bundles. Also in regions where axon bundles split the smaller bundle has a high chance of being missed, as fibre will follow the stronger direction.

More complex approaches have been developed to deal with this problem [60, 81]. One of these algorithms is the Gibbs tracking. In contrast to the streamline approach it will distribute small fibre elements, usually cylinders, within the image. A fibre in this case consists of a link of many of these elements, which do not change their own location. By rotating these elements and linking them together an entire path is created which is then assigned an energy. This energy consists of an external part, which evaluates how well the proposed fibre conforms to the underlying image data, and an internal part which measures how well the proposal conforms to an internal model of what a fibre should look like. The internal model tries to ensure properties observed in anatomical fibre tracts, such as a smooth path, avoiding sharp turns which are not present in the brain. Using this method together with an optimisation approach, such as simulated annealing, will result in a fibre configuration with the lowest possible energy, which ideally conforms well to the image data and is biologically plausible. These fibres can bridge gaps in otherwise high anisotropy paths as the overall fibre will still represent a low energy solution.

2.3 Graph Theory

A graph in its simplest form is a number of points with lines connecting them. More formally correct it is a ordered triple (V(G), E(G), ψ_G) of a non-empty set V(G) of vertices or nodes, a set E(G) of edges and an incidence function ψ_G that associates an edge with a pair of vertices [10].

Historically graphs have been used to model a variety of objects and their relations.

From modelling the bridges in Königsberg [29] to the social interactions of people on on-line platforms [103]. The analysis of these graphs using graph theory has led to a large variety of tools that can be used to understand the architecture of a graph and the underlying principles of the data. As a review of the entire field is far beyond the scope of this work we limit ourselves to listing the measures relevant for this work.

Generally graph or network indices can be divided into two categories, global indices which provide a single value for an entire network and local indices which can be calculated for every single node or edge.

We introduce some terminology relevant for the understanding of these indices.

Neighbour
A neighbour to a node n is a node connected to n by an edge.

Neighbourhood
The neighbourhood of a node n is the set of all neighbours of n.

Shortest Path
The shortest path connecting two nodes n and m is the lowest possible number of edges which form a consecutive link from n to m. It is not necessarily unique as several combinations of the same number of edges can form a path.

Isolated Node
A node not connected to any other node by an edge.

Leaf Node
A leaf or end node is a node, that has only a single neighbour.

Connected Component
A set of nodes in which each node can be reached by any other node in the set by following a combination of edges.

Disconnected Graph
A graph with more than one connected components between which there is no connecting edge.

Central Point
A node with an eccentricity equal to the radius of the network.

Degree Matrix
A $n \times n$ matrix D where each diagonal element $d_{i,i}$ is the degree of node i and each non-diagonal element is 0.

Adjacency Matrix
A $n \times n$ matrix A

$$A = \begin{pmatrix} a_{1,1} & \dots & a_{1,n} \\ \dots & \dots & \dots \\ a_{n,1} & \dots & a_{n,n} \end{pmatrix} \tag{2.13}$$

where $a_{i,j}$ is equal to 1 if the nodes i and j are connected and 0 otherwise.

Laplacian Matrix

The laplacian matrix is the difference of the degree matrix and the adjacency matrix.

$$L = D - A \tag{2.14}$$

2.3.1 Local Measures

Degree

The degree of a node is the number of neighbours of the node k_n.

Clustering Coefficient

The clustering coefficient or clustering coefficient C of a node can be calculated by dividing the number of connections between neighbours of that node by the number of possible connections between them.

$$C = \frac{2E_n}{k_n(k_n - 1)} \tag{2.15}$$

Where E_n is the number of edges in neighbourhood of the node n.

Clustering Coefficient D

In contrast to the clustering coefficient C this includes connections between the nodes in the neighbourhood of n and n itself.

$$C = \frac{2(k_n + E_n)}{k_n(k_n + 1)} \tag{2.16}$$

Eccentricity

The eccentricity of a node is the longest shortest path connecting this node to any other node in the network.

Betweenness Centrality

The betweenness centrality of a node is the number of shortest paths between any two nodes in the network this node is part of.

2.3.2 Global Measures

For each local measure there is a global measure corresponding to the average of all local measures. We do not list these separately.

Number of Nodes

The number of nodes in the network.

Number of Edges

The number of edges in the network.

Clustering Coefficient E

The average of the clustering coefficient C disregarding isolated nodes.

Diameter

The maximum of all node eccentricities.

Radius

The minimum of all node eccentricities.

Average Path Length

Average shortest path between any pair of nodes in the network. In the case of disconnected graphs this value is infinite.

Efficiency

The inverse of the average path length. This value has the advantage of being defined for all graphs as it always remains between 0 and 1.

Small-Worldness

The small-worldness ω of a graph compares the average path length L_g and clustering coefficient C_g of a network to a random network [52]. It is defined as:

$$\omega = \frac{\frac{C_g}{C_{rand}}}{\frac{L_g}{L_{rand}}} \tag{2.17}$$

3 State of the Art

In the past two decades the way different areas in the brain are connected, the connectome, has been increasingly studied. Using graph theoretical analysis to understand the architecture of the brain and the changes in the connectome associated with diseases has yielded new insights into the human brain. The two main approaches to characterize this connection have been using structural, diffusion weighted imaging to identify connections based on the micro structural properties of the white matter, and functional, identifying connections due to correlated changes in activity in different areas [7, 12, 13, 87].

We focus mostly on the structural approach, using tractographies to determine the connections between areas. Our results are therefore not directly transferable to functional networks, as connections in the functional connectome do not necessarily require a connection in the structural connectome and vice versa. Recent research however points to a complex relationship between structural and functional connectivity [46, 87, 102].

Structural connectomics has since been used to examine a variety of changes in the human brain, ranging from the effect of ageing to diseases such as Alzheimer's disease, schizophrenia, epilepsy and Autism Spectrum Disorders [23, 45, 69, 95, 102]. An increase in the research activity of this field was spurred by the establishment of the Human Connectome Project [27] a research initiative designed to produce a large database of open data available for analysis [28, 72] until 2015.

The last years have also seen the development of a number of tools to facilitate the research on the human connectome. The lack of unified tools hinders the ability of others to reproduce a group's research. One of these tools is the Connectome Mapper [22] which provides a python based framework for the more commonly used diffusion imaging acquisition schemes (DTI, QBI and DSI). Other tools include the DTI focussed java based tool by Gray et al [43], the Nypipe [41] framework and the LONI pipeline [25], of which the latter two provide the ability to combine modules from commonly used diffusion imaging tools, such as FreeSurfer[1], FSL[2] and the diffusion toolkit[3]. While these tools have been recently presented, none of them has gained a real foothold in the research community yet and most researchers still use their own solutions.

[1] surfer.nmr.mgh.harvard.edu
[2] www.fmrib.ox.ac.uk/fsl
[3] www.trackvis.org/dtk

3.1 Graph Construction

The quality of a connectome and the information content for a specific hypothesis depend to a large part on the choice of methods for its construction. The two main decisions in this are the choice of parcellation scheme, represented by the nodes of the network, and the methods for determining the connections between different areas in the brain, represented by the edges in the network.

3.1.1 Parcellations

The choice of parcellation scheme depends on the subject under consideration. Different studies have opted for either atlas based parcellations which add anatomical information or randomly generated parcellations, which offer a finer parcellation scheme. Additional methods that have been used are functional definition of nodes and voxel-based approaches [32, 45, 59].

Subdividing the brain in anatomical areas provides several benefits. Due to the anatomical label associated with each node comparison across different networks is possible, as the same node can be identified in each of them. At the same time the atlases typically have between 60 and 200 nodes, ensuring a low computational burden. Furthermore this approach is easily reproducible as rerunning the same parcellation will generally produce the same results, making the observed differences in network architecture more easily repeatable [32]. Two of the more commonly used tools for the anatomical parcellation of the human brain are the Automated Anatomical Labelling package of the SPM package [77, 96] and FreeSurfer with its assorted atlases [24, 31].

The disadvantages of anatomical parcellations are mostly due to the origin of the anatomical atlases, which were created with a different focus, such as changes in the staining behaviour of tissue due to differences in the cytoarchitectural properties of neurons [11]. The nodes are generally not of the same size and very large compared to other methods. This can introduce artefacts where certain connections are selected for. If both connected nodes are large they can be expected to have larger than average numbers of connecting fibres, even using a completely random placement of fibres.

Random parcellations of the brain offer the option to choose a resolution most suited to the matter at hand. Depending on the parcellation method used this has the additional advantage of ensuring very uniformly sized nodes, eliminating most problems associated with due to their size inherently more prominent nodes. Depending on the chosen resolution the network properties of a network can vary considerably. A study of the variation of network properties for different resolutions found that common topological measures, such as the small-worldness, can differ by $\sim 95\%$ for the same individual when comparing a low resolution parcellation (90 nodes) to a high resolution one (4000 nodes) [32, 105]. Having to choose a resolution for the random parcellation adds another variable to the experiment, which can not easily be defended, as no number of nodes is inherently advantageous compared to

another. One possible answer to this problem is to repeat the experiments using different parcellation methods and at different resolutions [47]. This does however not guarantee that the observed changes remain stable across the range of tested resolutions. By only considering results who do appear at every resolution, possible real differences at a certain scale can be overlooked.

A further disadvantage of random methods lies in the fact that the nodes no longer correspond to each other. This raises several issues. Firstly the borders of the random nodes will not correspond to any borders due to anatomical or functional differences in the brain, if changes in the network are due to very defined changes in the function of certain areas the resulting changes in the network might be less than in an anatomical one. Secondly using a random parcellation allows only the comparison of global network properties across individuals. As individual nodes do not longer correspond to each other, any information about local changes in the network is lost.

Using a functional definition of nodes means to derive the nodes from the available data itself. This is a fairly recent technique, which has been used for the functional and structural connectome both. The idea is to define nodes based on DWI measures and node-specific connectivity fingerprints [2, 56, 76]. The ability to define the properties of nodes based on the hypothesis of the study allows for very specific node definitions which can be even more relevant for the issue in question than anatomical parcellations. At the same time the size of the nodes can be constrained very accurately. The parcellation method can also be chosen for reproducibility and nodes can potentially be identified between subjects based on their properties.

This method is however better suited for functional networks than for structural networks. Defining the nodes based on activation criteria lends itself very well to this approach. Diffusion data on the other hand needs more complex parameters. Additionally this approach can result in very specific node definitions for each study, potentially voiding any comparability between different studies. Depending on the chosen method there can be areas of the brain which do not belong to any node, excluding them from scrutiny.

Lastly using a voxel-based approach defines each voxel as its own node. This offers the highest possible resolution at a very high computing cost. Again this approach lends itself comparatively well to functional connectivity data [98], while for the tractogram based connectivity a very large percentage of voxels will not be connected at all, as the number of fibres is in most cases orders of magnitude lower than the number of voxels.

Similar to the random networks the nodes are not identifiable across patients and this approach allows mostly measurement of global properties.

3.1.2 Connections

Connectivity can be derived differently from the various modalities. The main definitions are structural connectivity, based on DWI and tractography, and functional, based on correlation between the activations of different areas in functional MRI

data [34, 32].

Additional methods of deriving structural connectivity information include correlation of differences in morphometric parameters of brain regions across patients, such as analysing the change in cortical thickness or grey matter volume [32, 49, 64], but these methods can not be used to construct a whole brain network for an individual subject.

We focus on structural connectivity using fibre tractography in this work. In contrast to functional connections, structural connections are generally undirected and homogeneous. They are undirected as the connectivity information using tractographies is based on the assumption, that water can more easily diffuse along densely packed neurons than across, this process is independent and indeed ignorant of the direction of travel of electrical impulses along the neuron. The information is homogeneous as the fibre tracts do not provide any additional properties beyond the number of fibres to distinguish between different types of connections, such as inhibitory or excitatory neurons.

The choice of tracking algorithm has a large influence on the resulting connectome [9]. The tractographies provided by global approaches tend to perform better than local alternatives. The global optimization approaches deal better with noisy image data and are able to cross intersections of fibres, where the local approaches terminate due to decreased anisotropy.

Not every edge in the connectome represents the same connection strength. Assigning a specific value to the connection strength allows a more detailed analysis of the architecture of the connectome. Currently there are two main methods for deriving the strength of a given connection using structural connectivity [32].

The first method is based on the number of fibres connecting different areas. While the exact number of fibres is an artefact of the tracking algorithm and does not represent single axons connecting the two areas, the assumption is that areas that are connected by more axons will be connected by more fibres as well. Additionally corrections can be made for the different grey matter volumes or distances of the regions in order not to favour large areas and short connections [53]. Beside the general assumption that more fibres translate to more axons there is currently no way to determine how much stronger one connection is than another based on this approach. Depending on the tractography algorithm noise, length of the fibres and crossings have an unknown, but potentially large influence on the exact number generated by the algorithm. The second approach tries to take the integrity of the overall fibre into account. Local measures of individual voxels along the fibre tract are integrated and taken as a measure for the connection capacity of the fibre. Possible measures for this are for example the fractional anisotropy or the mean, radial or axial diffusivity. A variant of this approach is to treat the measure as a bottleneck and assign a fibre the lowest value along the tract. Similar to the fibre count these measures are sensitive to image noise and the used diffusion model for the image [58].

Recent research has suggested using additional information to estimate the connection strength between different areas. One possible method is using a custom,

time consuming, DWI sequence and a four compartment model to determine axon diameter and density [1]. Other ideas include information from several modalities to calculate additional tissue properties of the white matter [32, 97]. These methods, while promising, however usually include longer acquisition times, which limit their usefulness.

Additional research in order to accurately determine the quantitative value of the connection strength of direct structural connections remains one of the great challenges of the field.

3.2 Global Network Assessment

When the idea of translating the graph theoretical treatment of social and biological systems to the connections in the human brain was first applied the general organisation of the brain architecture was one of the earliest properties to be examined. The early studies suggested that the brain was organized in small-world fashion [51, 86, 88, 89, 90]. It shares this property with several networks developed by biological or social processes, as these networks, combining a high clustering coefficient with a good efficiency, tend to perform well and be relatively robust for comparatively little cost. For the brain network this means that local processing can be as fast as possible due to high local interconnections, while keeping information flow between distant regions at an acceptable level with a limited number of available axons. Full connectivity allows for even faster information processing, but is not practical, as each axon requires space, consumes nutrients and produces waste, making the maintenance of a fully connected connectome extremely costly.

The small-world nature [101] of the connectome is very well accepted in the literature by now. The same can be said for the scale-free properties [19] of the network. However the exact nature of the graph model best used for representing the connectome is still a matter of research [67]. Determining the model most closely resembling the nature of the connectome will be important to understand changes in a variety of illnesses as well as judging their impact on the network as a whole.

Beyond understanding the general architecture of the human brain the early research focussed on observing changes of global network properties in case of disease. Many studies support the idea, that the efficiency [62] and dis-connectivity of the network can be used as an indication for pathological changes [16]. A recent study reviews the findings of the last decade as pertains to the change in network efficiency and clustering coefficient [45]. For high functioning autism spectrum disorders recent studies suggest increased characteristic path length, resulting in decreased efficiency and decreased clustering coefficient [4, 5, 54, 66].

To the best of our knowledge there has however not been an encompassing comparison of different global network indices for their capability to distinguish between the healthy and the diseased connectome, aside from the one presented in this work and previously published at a workshop [39]. Recent studies have suggested using some global and local characteristics to classify individuals based on sex and age

[26, 40].

3.3 Local and Regional Network Assessment

During recent years studies have increasingly examined the change of local network indices in addition to global indices. A recent review collected studies examining global and local indices for a variety of diseases [45].

One study showed that the local clustering coefficient and local characteristic path length of nodes correlate with the oral language skill in high functioning autism spectrum disorders [66]. This is in line with a prior study which found impaired micro-structure in the language pathways of tuberous sclerosis patients with autism spectrum disorders [65]. It has however been noted, that there are very few studies concerning themselves with the changes in specific areas in ASD in comparison to other diseases [82]. A reduction of regional efficiency has been found in multiple sclerosis patients [50].

Another avenue that has been increasingly pursued recently is the identification of hub nodes in the connectome [74] and the identification of sub-networks responsible for certain functions [18].

3.4 Influence of Density

There are two main techniques to examine the connectome by employing the different strength of connections to better understand it. Most studies either threshold the network, removing all edges below a certain strength, or employ weighted measures. Historically the early research was done exclusively using the thresholding approach, while in recent years the focus has shifted more to using weighted measures. The analysis of thresholded binary networks remains however an important part of current research [32]. Generally the use of weighted network metrics has been more commonly used in functional connectomics, as a reasonable definition of weight poses a greater problem in structural connectomics.

In the past there have been two main approaches to thresholding the connectome. The first method is to define an threshold t and delete all edges with a strength less than the threshold (e.g.[17, 30, 83, 84]). The exact value of t will depend on the strength distribution of the connections in the network and the measure used to define the strength of an edge. The assumption behind this approach is that the measured network will consist of the "true" sparse connectome and spurious edges due to noise and image or tractography artefacts. This approach has later lost in popularity, due to numerous flaws associated with it.

The problem with this approach is that many of the network measures depend on the number of nodes and edges of the network. Thresholding for a specific value will lead to different number of edges for different patients, resulting in possibly spurious results, which are due to thresholding artefacts rather than genuine difference in the

network architecture [99]. Advanced techniques for this approach include ensuring that the entire network remains connected by not deleting the strongest edge which ensures that a node or group of nodes remains connected to the rest of the network [8].

Rather than thresholding a specific threshold for all patients later studies have increasingly employed a target density κ, ensuring that all networks are comparable regarding node and edge count (e.g. [100]). The advantage of this is, that the network indices of different networks can be expected to have a similar value and differences in this are more likely due to changes in the actual network architecture.

On the other hand this approach will potentially disregard a great number of connections in a network with a very high connectivity before thresholding and include spurious connections in one with a very low connectivity. This could potentially disregard differences in the resulting tractography which are actually symptoms of the disease. Also, depending on the choice of κ removing strong edges can potentially change the topology of the network for networks with a high connectivity. Techniques have been developed to reduce the disadvantages of this approach while keeping the advantage of comparable network index values. The network can be thresholded at different density values (e.g. [61]) to ascertain that effects persists over a range of densities. This can even be extended to a very high resolution, where the network index is eventually integrated over the density [36]. This however yields its own share of problems, as there is no reason, that the effect persists at all densities, especially at high densities where more and more spurious connections are included in the connectome. This technique further requires the random removal of edges if several edges have the same weight, which has been shown to bias the network indices towards random networks [99].

Using weighted network indices has recently become more popular (e.g. [47, 83, 91]). These try to avoid the need for thresholding by using the full continuous strength information of the connections. The advantage of this approach lies mainly with the removal of arbitrary threshold values and it has been shown that the results are qualitatively similar to unweighted network analysis [68, 78]. Weighted network indices have shown some promise and additional research in this area will be a very important part of future connectomics.

However the problem of assigning a reasonable and meaningful weight to a connection and understanding what that weight represents in comparison to another connection of a different weight is one of the main obstacles, especially for structural connectomics [55]. Additionally the use of weighted network indices does not alleviate the problem of networks of inherently different connectivity. It is, at least in theory, possible to ascertain, that in a given study the networks will have comparable mean connection weights. When comparing different studies or data from different sources however, the mean connection weight can differ significantly and differences in the network architecture, that would appear when using binary networks, could be eclipsed by the difference in weighted statistics due to this weight difference.

One of the great challenges of this field remains to understand the meaning of different network measures in relation to the function of the human brain [55]. An

important step in this direction is to understand at which densities differences in architecture appear due to pathological changes associated with diseases and to determine which densities are therefore relevant for diagnostic purposes [36].

4 Methods

4.1 Graph Construction

This section gives an overview over the overall pipeline for generating networks representing the human connectome as used in this work. It can broadly be divided in developing the pipeline leading up to graph creation and graph construction and analysis. While this work relies on several third party tools and developments for the first part the most development work has taken place in the second part. Figure 4.1 gives a general overview of the graph construction pipeline.

4.1.1 Pre-Processing

Analysis of the human connectome is the last step in a long and complex pipeline starting with the raw data acquisition as described in section 2.2.1 and the subsequent artefact reduction and modelling as described in section 2.2.2.

In order to build a network of the human brain it is necessary to define the nodes and edges comprising the network.

Parcellation

There are several possible choices what nodes may represent. In this work we concentrate on using anatomical divisions of the human brain for defining nodes.

This has the advantage of dividing the entire brain in discrete areas. Every voxel and by extension every point in the image belongs to exactly one specific area.

One way of defining an anatomical parcellation of the brain is by registering to an atlas defined by medical professionals. We use a third party tool for this purpose. FreeSurfer and the atlas by Destrieux [24, 31, 85] are widely used in the literature. FreeSurfer provides an automatic pipeline which will perform motion correction, registration to a common atlas, skull stripping and application of several parcellations to the T1 image. For the definition of our network nodes we use the mapping of the cortical labels to the segmentation volume based on the atlas by Destrieux.

Registration

Due to long acquisition times in diffusion-weighted images the presence of head motion and other artefacts such as described in section 2.2.2 can noticeably reduce the quality of the images. In order to reduce these effects we perform eddy-current- and head-motion correction using FSL.

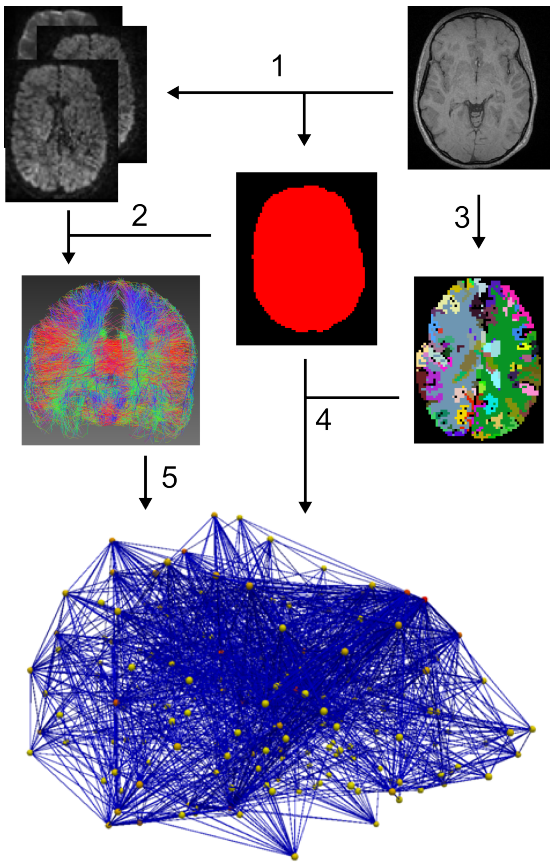

Figure 4.1: Preprocessing pipeline. 1. Registration of the b0 image and the T1 weighted image and brain masking 2. Motion correction and fibre tractography in the area defined by the brain mask 3. FreeSurfer parcellation of the brain 4. Registration of the parcellation to the diffusion weighted images by the transformation found in (1.). The parcels are used for the network nodes 5. Calculation of the connections between nodes based on the fibre image

As T1 images and the diffusion-weighted images do not necessarily occupy the same image coordinates they need to be registered in order to combine the information inherent in both images. The images used in this work have been acquired on the same day, allowing us to assume that they can be brought into alignment using only rigid transformations. We calculate the rigid transformation matrix for registering the extracted b0 image to the T1 image using ANTs and a mutual information metric. We then apply the inverse of this transformation matrix to the parcellation image in order to bring it into diffusion coordinate space.

We transform the parcellation image to diffusion image space in order to avoid transformation of all diffusion tensors, which would be necessary otherwise. This way we reduce the possible loss of information. All further processing steps happen in diffusion space.

Table 4.1: Tractography settings for feasibility study

Parameter	Value
Iterations	10^7
Particle Length	3.4 mm
Particle Width	1.2 mm
Particle Weight	0.0018
Start Temperature	0.1
End Temperature	0.001
Energy Balance	0
Minimum fibre Length	19 mm
Curvature Threshold	45 °

Table 4.2: Tractography settings for global classification and regional classification studies

Parameter	Value
Iterations	10^8
Particle Length	3.7 mm
Particle Width	0.1 mm
Particle Weight	0.0015
Start Temperature	0.1
End Temperature	0.001
Energy Balance	0
Minimum fibre Length	20 mm
Curvature Threshold	45 °

4.1.2 Tractography

After modelling the diffusion-weighted images as described in section 2.2.2 we use Gibbs tracking to create a tractography image of the brain. Gibbs tracking uses a global optimization approach to find possible fibre configuration solutions which fit the image data as well as an internal model of the brain. It seeds the brain with many fibre elements, small cylinders, which can be connected to form long chains of elements and finally fibres.

Depending on the experiment we choose different settings for the tractography algorithm.

Table 4.1 shows the settings used for the feasibility study, see section 5.3. The tractographies were generated on the entire brain mask, grey and white matter. The resulting tractographies were then directly used for the creation of the connectome.

Table 4.1 shows the settings used for the global classification, see section 5.2 and the study on the influence of density on classification, see section 5.4. By reducing the particle width and increasing the number of iterations when compared to the feasibility study (Table 4.1) we get tractograms with considerably more fibres (10x-20x) at the expense of considerably longer run time. On the used computer (Quad-Core i5) it took roughly 15 hours per tracking for a single core. To reduce noisy surface fibres in the grey matter due to low anisotropy we used a white matter mask generated from the parcellation image to mask the tracking volume.

In order to get more fibres we performed the tracking four times independently with the same parameters for each patient and combined the four separate trackings to get a combined tractogram. Figure 4.2 shows an example for the difference in the resulting tractogram depending on the settings.

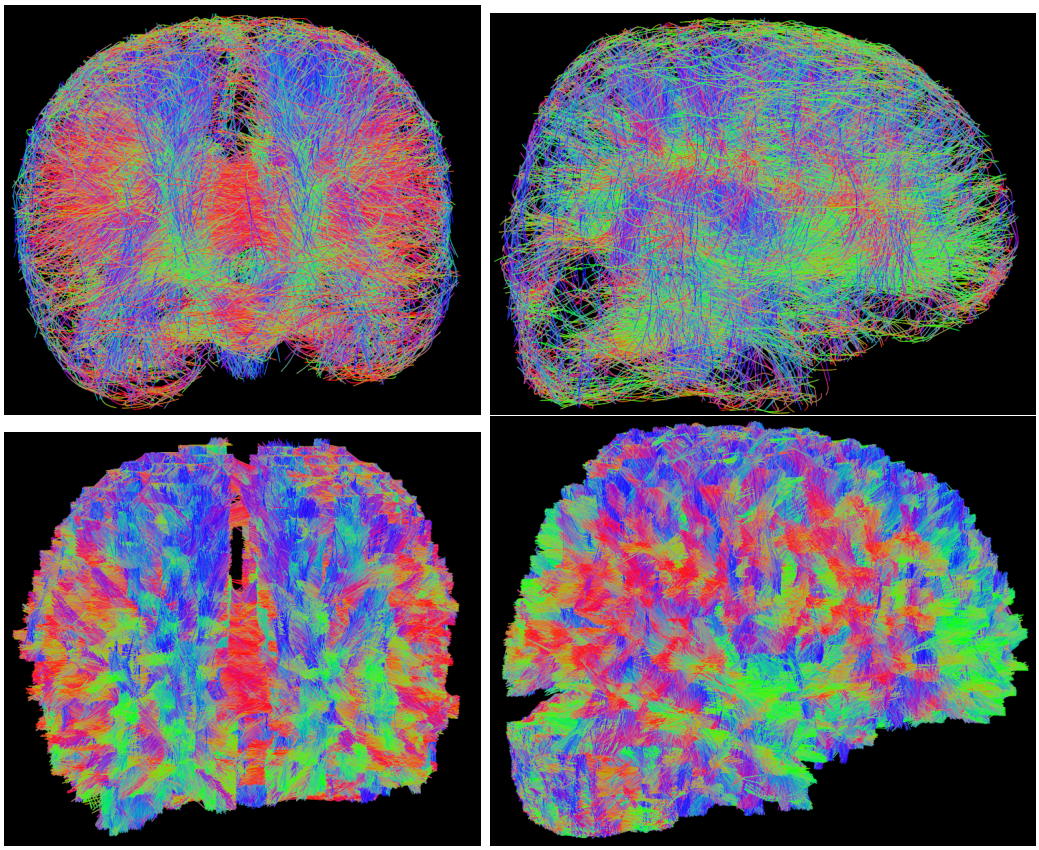

Figure 4.2: Example of the resulting tractogram for the different settings. Top left/right: Front and side view of a tractogram used for the feasibility study (10,054 fibres) Bottom left/right: Front and side view of the combined tractogram for the same patient (652,701 fibres)

Table 4.3: Information contained in the connectomics networks as implemented in the Connectomics module of MITK.

Section	Information per Element
Image geometry	
List of nodes	Node position
	Node id
	Node label
List of edges	Connected nodes
	Edge id
	Number of fibres

4.1.3 Network Implementation

We decided on basing our network implementation on the adjacency list implementation of the Boost Graph Library. The algorithms and data structures have been implemented as part of the Connectomics module of the open source toolkit MITK (Medical Imaging Interaction Toolkit) and are available in binary form in the MITK Diffusion release.

Information contained in the MITK connectomics networks is summarized in Table 4.3.

4.1.4 Network Creation

There are several strategies for creating a network from the anatomical information given by the tractogram and the parcellation image. This section describes the approaches used in this work. For the creation of the network we iterate over all fibres of the tractogram.

Nodes are defined by the parcellation image. We use two options for determining the locations of nodes. Nodes are either positioned at the centre of mass of their respective label segmentation or at the first point where a fibre is encountered crossing into the label segmentation. The second option is used primarily if the labelling scheme is not known.

As diffusion anisotropy tends to be lower closer to the grey matter there is a distinct likelihood for fibres in the tractogram ending before encountering grey matter. For the FreeSurfer parcellation used in this work we implemented a search algorithm in order to catch these prematurely ending fibres and extending them if possible. If a fibre end point is located in a volume whose label belongs to white matter as defined by the FreeSurfer atlas we try to extend it linearly for up to 20 mm. If during this extension no grey matter volume is encountered the fibre is discarded and will not be used.

Conversely a fibre might extend through the grey matter into the background. In

this case the fibre is retracted along its path until grey matter is encountered. If this retraction is unsuccessful the fibre is discarded.

If a fibre connects a grey matter label to itself it is discarded as well. We avoid self-connected nodes.

A labelled volume is added as a node if at least one fibre connects it to another node.

An edge is added to the network if a fibre connects two nodes which do not already have a connection. If a connection between the nodes already exists the fibre count for that connection is increased by one.

4.2 Global Network Assessment

Using connectomics for diagnostic purposes requires the understanding of how different network indices are affected by mental illnesses. In order to decide which network indices provide the best diagnostic potential we have to evaluate their individual performance in distinguishing between the healthy and the diseased connectome.

4.2.1 Extraction of Connectome Features:

We extracted 32 features for each patient's connectome. These features quantify the compactness, clustering, and spatial uniformity of the hypothesized connections within the brain. Graph features and their explanations are given in Table 4.4.

4.2.2 Classification and Validation:

Support vector machine (SVM) classification was employed for the feature selection and the resulting classification of the two groups. Though alternate supervised learning techniques may also be utilized, as we shall see in the next section, SVM classifier yielded the highest classification accuracy among the other well known candidates. We used radial basis function, also referred to as Gaussian kernel, in the form of $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\frac{|\mathbf{x}_i - \mathbf{x}_j|^2}{2\sigma^2})$ to transform the increase in the dimensionality of the data for better separability. We performed a parameter search to identify σ that achieves the highest classification accuracy. We sought σ in the set of candidate values that varied from 1.0 to 6.0 with 0.1 steps and determined that σ equalling 3.6 achieved the best performance in the identification of the patient's state.

The data is normalized so that the features have zero mean and unit variance to reduce the scale differences within different features. In order to obtain unbiased performance estimates, patient-based leave-one-out cross-validation was performed. The feature set was first divided into 32 disjoint partitions for each patient's data. For each patient, a classifier was trained with the remaining 31 patient's data and then tested on the retained data. The results for each patient were then combined to find the overall classification accuracy.

Table 4.4: Extracted graph features and their descriptions.

Feature Name	Description
Number of Nodes	Number of regions in brain
Number of Edges	Number of hypothesized communications
Average Degree	Number of edges per node
Clustering Coefficient C	Ratio of total number of edges among the neighbours of the node to the total number of edges that can exist among the neighbours of the node per node
Clustering Coefficient D	The average of the ratio of the links a node's neighbours have in between to the total number that can possibly exist
Clustering Coefficient E	Ratio of total number of edges among the neighbours of the node to the total number of edges that can exist among the neighbours of the node per node excluding the isolated nodes
Average Eccentricity	Average of node eccentricities, where the eccentricity of a node is the maximum shortest path length from the node to any other node in the graph
Diameter	Maximum of node eccentricities
Radius	Minimum of node eccentricities
Average Path Length	Average distance between the nodes of a graph, where the distance between two nodes is the number of edges in the shortest path that connects them
Average Betweenness	Average of node betweenness, where the betweenness of a node is the number of shortest paths from all nodes to all others that pass through that node
Giant Connected Component Ratio	Ratio between the number of nodes in the largest connected component in the graph and total the number of nodes
Number of Connected Components	Number of clusters in the graph excluding the isolated nodes
Average Connected Component Size	Number of nodes per connected component
Percentage of Isolated Points	Percentage of the isolated nodes in the graph, where an isolated node has a degree of 0
Percentage of End Points	Percentage of the end nodes in the graph, where an end node has a degree of 1
Number of Central Points	Number of nodes within the graph whose eccentricity is equal to the graph radius
Percentage of Central Points	Percentage of nodes within the graph whose eccentricity is equal to the graph radius
Spectral Radius	Largest valued eigenvalue of adjacency matrix
Second Largest	Second largest values eigenvalue of adjacency matrix
Adjacency Trace	Sum of the eigenvalues of adjacency matrix
Adjacency Energy	Sum of the squares of eigenvalues of adjacency matrix
Spectral Gap	Number of 0 valued eigenvalues of adjacency matrix
Laplacian Trace	Sum of the eigenvalues of laplacian matrix
Laplacian Energy	Sum of the squares of eigenvalues of laplacian matrix
Number of 0s	Number of eigenvalues that are equal to 0 in normalized laplacian matrix
Number of 1s	Number of eigenvalues that are equal to 1 in normalized laplacian matrix
Number of 2s	Number of eigenvalues that are equal to 2 in normalized laplacian matrix
Lower slope	The slope of the line fitted for the eigenvalues of the normalized laplacian matrix that are between 0 and 1 when sorted
Upper slope	The slope of the line fitted for the eigenvalues of the normalized laplacian matrix that are between 1 and 2 when sorted
Normalized Laplacian Trace	Sum of the eigenvalues of normalized laplacian matrix
Normalized Laplacian Energy	Sum of the squares of eigenvalues of normalized laplacian matrix

4.3 Local and Regional Network Assessment

While global network indices show overall changes in the network they can not show the changes in particular areas. By only using these global changes we do not incorporate information about the responsibility of different anatomical regions in the brain which is provided by the medical sciences.

Different areas have been identified as being mainly involved in certain tasks and we can reasonably expect pathological changes in these areas to affect the capability of the brain for these tasks. Conversely it is possible to identify possible areas which might be impaired by observing the changes in behaviour and capability of the patient.

Applying this knowledge to the connectome we expect certain illnesses to affect the role of certain nodes in the connectome more than the global network architecture.

4.3.1 Local Network Indices

There is a variety of network indices which can be calculated on an individual node basis. A lot of these indices however are not independent of each other. We examine a selection of those indices which are commonly used in the literature and additionally use the betweenness centrality, which has not been extensively used in examining the connectome.

Degree

The degree of a node is the number of other nodes it connects to. These nodes are called the neighbouring nodes, or the neighbourhood. A node with a high degree is likely a central node of the network and is commonly called a hub.

Clustering Coefficient

The clustering coefficient cc of a node measures the tendency of the neighbourhood to form a clique. It is defined as the percentage of nodes in the neighbourhood of the node which are also neighbours of each other.

$$cc(n) = \frac{2E_n}{k_n(k_n - 1)} \tag{4.1}$$

Where E_n is the number of edges in the neighbourhood of node n and k_n is the number of nodes in the neighbourhood of node n.

A high clustering coefficient indicates that the node is part of a tightly knit clique and implies a possible high segregation of the network.

Betweenness Centrality

The betweenness centrality bc is a measure for how important a given node is for the efficiency of the network. It quantifies how many shortest paths between any two nodes a and b pass through a given node n [33]:

$$bc(n) = \sum_{a \neq n \neq b} \frac{\sigma_{ab}(n)}{\sigma_{ab}} \tag{4.2}$$

where σ_{ab} is the number of shortest paths from node a to node b and $\sigma_{ab}(n)$ the number of those that pass through node n.

While a high degree likely corresponds to a high betweenness centrality the reverse is not true. A node with a high betweenness centrality and a low degree is likely an important link for the network.

We believe that the betweenness centrality has been under appreciated in the study of the human connectome in the literature. It is well suited to detect a changed role of an area in the overall connectome.

4.3.2 Regional Network Indices

Calculating the local network indices for each node can give important information about changes in a specific location of the brain, whose size is determined by the resolution of the used parcellation. It does not however wholly respect the way information is processed in the brain. Different anatomical locations might be involved in the processing of certain information. These areas can, but not necessarily do, neighbour each other. This is especially noticeable if the same area in the left and right hemisphere is involved in a certain task. In these cases merging the areas in the atlas is impractical, as they are not physically close and the same area might work with different areas for different tasks.

To take this into account we define these sub-networks involved in a certain task as a region and expand the local network indices to regional ones. We do this by averaging the local network indices of all nodes which are part of the region.

$$I_{reg} = \sum_{r_i \in R} \frac{I_{r_i}}{n} \tag{4.3}$$

Where I_{reg} is the regional network index, R is the set of nodes $r_1, ..., r_n$ contained in the region and n is the cardinality of R.

4.3.3 Lateralization of Indices

The connections in the human brain are typically not symmetrical. As the different hemispheres perform slightly different roles in the processing of information and abstract thought a perfect symmetry is not expected.

Changes in the asymmetry of the brain might result in some of the symptoms associated with different diseases. In the case of ASD several studies have reported marked changes in the lateralization of the brain when compared to typically developed controls [14, 42, 93]. As our previous definition of regional network indices spans areas in both hemispheres it is important to test whether changes in the regional network indices are only due to the changes in lateralization.

We define the lateralization index of a network index I as

$$\frac{I_{left}}{I_{right}} \tag{4.4}$$

and compare the differences in the lateralization index to the differences in the regional network index. I_{left} represents the averaged local indices of all nodes located in the left hemisphere and I_{right} of those in the right hemisphere.

If changes in the regional network indices are solely due to changes in the lateralization we expect the lateralization index to have a higher significance than the regional network indices.

4.3.4 Experimental Validation

In order to evaluate the proposed regional we examine the data described in section 4.5. Given the symptomatology of ASD, especially the reduced capability for communication we focus on areas involved with language processing. We examine three regions consisting of several nodes:

Wernicke's Area

Wernicke's Area's role in communication is to build meaningful sentences and extracting meaning from sentences. As the inability to easily understand speech is one possible symptom for ASD we expect this area to be less integrated into the connectome. In our set-up Wernicke's Area contains four nodes in the right hemisphere and four in the left hemisphere.

Broca's Area

Broca's Area's role in communication is the motor control of language, it is responsible for pronunciation and articulation. As these capabilities are not typically reduced in ASD we do not expect a large change in this area relevance for the network. In our set-up it consists of two nodes in the right hemisphere and two in the left hemisphere.

Motor Cortex

The Motor Cortex is responsible for general motor control. While some repetitive behaviour such as flapping can occur in ASD generally the motor control is not severely impaired. As such this area is used as an internal control. We expect no change to be visible. In our set-up it is comprised of one node in the right hemisphere and one in the left hemisphere.

It should be noted, that the above description of the areas' purpose is a simplification. We are focusing solely on the roles of these areas as are relevant for our hypothesis. We expect Wernicke's Area to be a good marker for ASD, whereas the other two areas are not expected to provide much information in that regard.

4.4 Influence of Density

In section 5.3 we found that regional network indices differ between ASD patients and typically developed controls. These experiments however were performed with a low number of fibres and showed mainly the feasibility of such an approach. In order to better understand the relevancy of these changes it is important to repeat the experiments with different parameters and study the effect at different densities.

As described in section 4.1.2 we performed four additional tractographies per subject and added all fibres together to get a tractogram with more than 500,000 fibres per patient. Using these larger tractograms it is possible to create connectomes with a significantly higher density. These connectomes in turn can be used to examine the observed the effects and their changes across different densities.

4.4.1 Thresholding

In order to create a binary network from a weighted network thresholding needs to be performed. There are several ways to threshold a network, for the purpose of this work we use the following:

Threshold Weight
: This method removes all edges from the network which have a fibre count lower than the specified one.

Smallest threshold below density
: This method selects the smallest threshold which will result in a connectome with a density below the specified one. All fibres below a certain fibre count are removed. The fibre count is incremented until the connectome density is below the target density.

Random removal
: In this method a target density is specified. It works similar to the smallest threshold below density. The only difference is in the last increment, instead of removing all fibres of the last fibre count, fibres of that fibre count are randomly removed until the target density is reached.

Many network indices, global, local and regional, are dependent on network properties such as node and edge count. The nodes for our networks are given by the used parcellation, the edge count however can vary. Due to this only networks of similar density should be compared to each other. We choose the three thresholding methods to ensure we compare compatible networks.

Network Index	0% density	100% density
Global Clustering Coefficient	*	1
Global Betweenness Centrality	0	0
Average Degree	0	n
Efficiency	0	1
Small Worldness	*	1
Regional Clustering Coefficient	*	1
Regional Betweenness Centrality	0	0
Degree	0	n

Table 4.5: Borderline values for different network indices for a network with n nodes. Values marked as * are not defined.

The "smallest threshold below density" (STB) and "random removal" (RR) methods ensure a similar density for the networks to be comparable.

On the other hand given that the data for all subjects is acquired the same way and tractography is performed using the same settings it can reasonably be argued that the different density of the networks can be viewed as an inherent property due to pathological changes in the case of illness which lead to a reduced fibre count in the tractogram.

To pursue this avenue the "threshold weight" (TW) method creates all binary networks with the same settings, disregarding possible differences in the resulting density.

4.4.2 Performance Across Density and Threshold

In order to better understand which network indices can be used to differentiate between the impaired and the healthy connectome it is important to understand how the regional network indices differ across a range of densities and thresholds. The border conditions for the different network indices are summarized in table 4.5. At these border conditions all networks will have the same values for these indices, whereas the value in between will change depending on the configuration of the network. The main purpose of our experiment is to evaluate how they change in case of ASD compared to healthy controls.

Analysing the density area where the connectome differs the most between ASD patients and controls will give us an indication at which densities we can expect our regional networks indices to be useful tools in diagnosing possible illnesses. At the same time the width of this area as well as how quickly diagnostic value drops off will give us an indication how sensitive these measures are to the inclusion of noise. As at lower threshold values we can expect more spurious fibre connections to be included.

4.5 Experiments

As mentioned in the previous sections we evaluated the the proposed methods on medical image data. This section will describe the study and used acquisition techniques. All evaluations share the same original image data.

4.5.1 Participants

The ASD group comprised 18 right handed children with a mean (SD) chronological age of 9.7 (2.1) years (range 6.1 - 12.8). Participants were recruited through outpatient clinics of the University departments of Mannheim, Heidelberg and Frankfurt. A diagnosis of ASD (Asperger Syndrome or High Functioning Autism) was established using the Autism Diagnostic Interview-Revised (ADI-R) [70], the Autism Diagnostic Observation Schedule (ADOS) [71] and supported by expert clinical diagnosis using ICD-10 criteria. The Social Responsiveness Scale (SRS [20]) was used to assess severity of social impairment and screening for social difficulties in control children. All participants had a Nonverbal IQ higher than 70 based on the Raven's Colored Progressive Matrices Test [80] and had fluent language skills, based on direct observation during the Autism Diagnostic Observation Schedule, Module 3 for participants with fluent language abilities. Control subjects were 18 normally developing children (mean age 9.7; SD 1.9; range 6.5 - 12.2). ADOS scores were not determined for the control subjects.

Controls and ASD children were matched rigorously and pairwise for age, sex and IQ (Tab. 4.6). Handedness was assessed by the Edinburgh Handedness Inventory [75]. Exclusion criteria for autistic participants and healthy comparison subject included: associated disorder, such as fragile-X syndrome or tuberous sclerosis, and presence of evidence of birth asphyxia, head injury, seizure disorders or other structural brain abnormalities. Potential control participants were also screened by experienced clinicians to exclude those with a family history of autism, developmental cognitive disorder, learning disability, affective disorder, anxiety disorder, schizophrenia, obsessive compulsive disorder, or other neurological or psychiatric disorder.

The study was approved by the Ethics committee of the University of Heidelberg. All parents gave written informed consent, children gave assent for participation in the study [79].

4.5.2 Data Acquisition

Data acquisition was done using a 1.5 T scanner (Siemens Avanto). T1 images for parcellation were taken with the following settings: MPRAGE TR/TE/TI/α = 1.9 s/4 ms/1.1 s/8 $^\circ$, FOV = 256 $\times$ 256 mm^2, matrix = 256 $\times$ 256, scan time 6 min). Diffusion weighted imaging was performed using single-shot EPI with a dual bipolar diffusion gradient and a double spin echo for reduction of eddy currents with the following parameters: TR/TE 4700/78, FOV 192 mm, data matrix of 96 $\times$ 96

Characteristics	ASD Group (n,18) Mean (SD)	TD Group (n,18) Mean (SD)
Sex	16 m/ 2 f	16 m/ 2 f
Age	9.7 (2.1)	9.7 (1.9)
IQ	111.0 (14.4)	112.8 (14.9)
SRS	107.7 (26.7)	16.1 (5.7)
ADOS Scores		
Interaction	7.1 (2.2)	
Communication and Language	3.4 (2.4)	
ADI-R Scores		
Interaction	18.3 (7.2)	
Communication and Language	11.6 (4.0)	
Repetitive behaviour	6.8 (2.7)	

Table 4.6: Subject characteristics, taken from [79]

yielding an in-plane resolution of 2.0 mm, 50 axial slices with a thickness of 2.0 mm and no gap, with 6 gradient directions (b=1000 s/mm^2) and a b=0 image. This scheme was repeated 15 times.

5 Results

5.1 Graph Construction

We implemented a fully functional pipeline for the construction and evaluation of the human connectome. An example for a connectome created the presented methods can be seen in figure 5.1. The code is freely available as part of the Medical Imaging Interaction Toolkit[1].

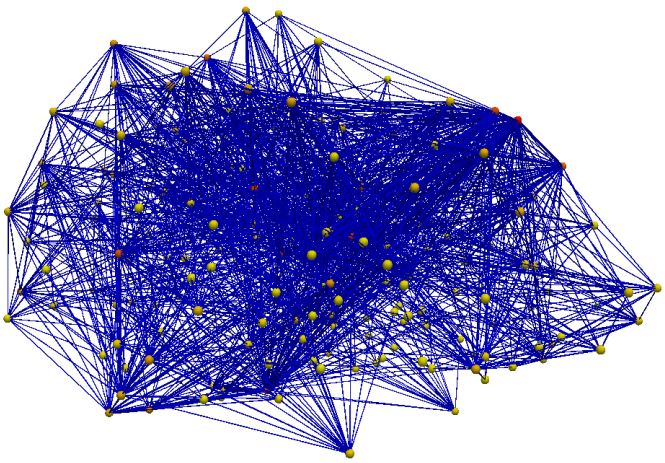

Figure 5.1: An example network created using our pipeline.

5.2 Global Network Assessment

We generated brain connectome networks as described previously for 32 patients each of which with four independent trackings. We then characterized the graphs using the 32 features described in Table 4.4 and using SVM classifier with RBF kernel we discriminated ASD patients from TD control with leave-one-patient-out cross-validation.

5.2.1 Classification Based on the Number of Features

Given the large number of features, we performed feature selection based on *t-statistic* to identify the most discriminative features. For a given feature i, the

[1] http://mitk.org

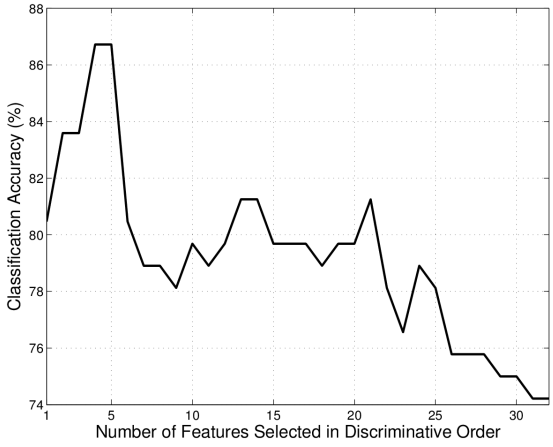

Figure 5.2: Influence of the number of discriminative features selected for classification on the classification accuracy. Highest grading accuracy achieved either the top four or five features selected.

t-statistic to test whether the population means are different is calculated as

$$t(i) = \frac{|\mu_1(i) - \mu_2(i)|}{\sqrt{\frac{\sigma_1^2(i)}{N_1} + \frac{\sigma_2^2(i)}{N_2}}} \tag{5.1}$$

where $\mu_k(i)$, $\sigma_k(i)$, and N_k are the sample mean, standard deviation, and size of the kth class ($k \in \{1, 2\}$) for ith feature, respectively. The features with high discriminative power get a higher score. We tested the grading accuracy of the feature sets constituted by the first M most discriminative features. We varied M from 1 to 32, and report the grading accuracy in Fig. 5.2. It is seen that a classification accuracy of 86.72% can be achieved using the top four or five features. When we investigated the results of this case, it is seen that 9 out of 64 ASD trackings were identified as TD control and eight out of 64 TD control trackings were classified as ASD, and the rest of the trackings were classified accurately.

5.2.2 Best Performing Learning Methods

In order to compare our result to our earlier study that only considered the betweenness centrality of speech related locations in the brain [38], we also performed classification using the average betweenness centrality alone. Our result showed 78.9% classification accuracy can be achieved using this feature alone. It is clear that considering additional features improved the classification accuracy significantly.

Independent of the learning method, we could achieve a consistent classification accuracy over 80%. Table 5.1 compares the classification accuracies of different classification methods. It is clear that SVM classifier achieves the highest overall accuracy in identifying the patient's neurological state. This is not unexpected as SVM classifiers are known to be highly successful in biomedical applications [104].

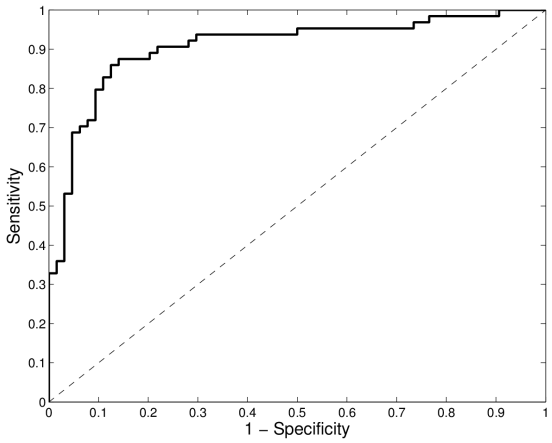

Figure 5.3: Receiver operating characteristics for the SVM classifier with RBF kernel. The area under the curve is 0.9067.

Table 5.1: Classification accuracy for different learning methods. SVM with RBF kernel yields the highest classification accuracy.

Learning Method	Classification Accuracy (%)
Support Vector Machines (RBF Kernel)	**86.72**
Support Vector Machines (Linear Kernel)	85.16
Linear Discriminant Analysis	84.38
Naïve Bayes Classifier	78.13
AdaBoost (Decision Stumps)	81.25

Table 5.2: Histogram of highest discriminative features where the frequency shows the number of times the feature was in the top five discriminative features according to t-statistic for a link threshold (N) ranging from 14 to 30.

Feature	Frequency
Giant Connected Component Ratio	17
Clustering Coefficient D	16
Normalized Laplacian Trace	15
Average Connected Component Size	11
Normalized Laplacian Energy	10
Second Largest Eigenvalue Adjacency	7
Clustering Coefficient C	4
Average Betweenness Centrality	1

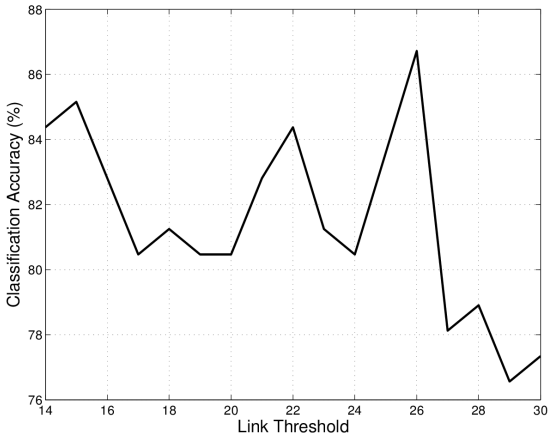

Figure 5.4: Classification performance across different thresholds using a well performing classifier.

5.2.3 Feature Performance Across Thresholds

We then investigated how often a feature was in the top five of features for classification for a range of thresholds where the discriminative influence of each feature was given by t-statistic. Table 5.2 shows the frequency of discriminative features that appear in the top five feature for different thresholds. The Giant Connected Component Ratio was consistently a discriminative feature for every threshold in the range. For threshold $N = 26$, with the highest classification accuracy the top five features with the highest t-statistics were Clustering Coefficient D, Giant Connected Component Ratio, Average Connected Component Size, Normalized Laplacian Trace, and Normalized Laplacian Energy.

Classification performance varied across different thresholds. Figure 5.4 shows the performance of a classifier.

Network Index	p-value
Global Clustering Coefficient	$<0.05^{\star}$
Global Efficiency	$<0.05^{\star}$
Clustering Coefficient (Wernicke)	$<0.05^{\star}$
Clustering Coefficient (Broca)	$<0.05^{\star}$
Clustering Coefficient (Motor Cortex)	0.15
Betweenness Centrality (Wernicke)	$<0.001^{\star\star}$
Betweenness Centrality (Broca)	0.27
Betweenness Centrality (Motor Cortex)	0.71

Table 5.3: P-values for the separation between ASD patients and TD for the global efficiency and the average clustering coefficient of the network as well as the clustering coefficient and the betweenness centrality of Wernicke's area, Broca's area and the primary motor cortex. Values with a * are statistically significant ($p < 0.05$) and those with ** have a p-value below 0.001.

5.2.4 Best Classification using Global Indices

Finally, we give the receiver operating characteristics (ROC) to evaluate the performance of the classification. ROC curve plots the *sensitivity* against the $1-$*specificity* at different threshold settings. For the SVM classifier, we used the distance from the maximum-margin hyperplane as the decision threshold. Figure 5.2 shows the ROC curve for our classifier. The area under the curve (AUC) is 0.9067, which is considered as a well-discriminating classifier.

5.3 Local and Regional Network Assessment

5.3.1 Feasibility and Internal Validation

To ascertain whether regional differences in general and the regional betweenness centrality in particular can be used to differentiate between ASD patients and typically developed controls (TD) we tested differences in regional areas for different regions and compared them to differences in global indices. As this was done to test for general possibility of using the indices we did not correct for changes in the found number of fibers, we did however test whether statistical difference existed even with the number of fibers as covariate, see section 5.3.2.

Tables 5.3 and 5.4 show the found statistical differences. It should be noted, that the motor cortex does not show significant changes in the regional parameters, this conforms to its role as internal validation area, as we did not expect changes in its embedding in the overall network.

The global efficiency as well as the global clustering coefficient is reduced in ASD patients compared to TD (Fig. 5.5). While the local clustering coefficient shows

Network Index	ASD	TD
Global CC	0.404 (0.051)	0.462 (0.064)
Global Eff	0.494 (0.026)	0.527 (0.039)
CC (Wernicke)	0.436 (0.093)	0.524 (0.071)
CC (Broca)	0.466 (0.058)	0.541 (0.041)
CC (Motor Cortex)	0.389 (0.033)	0.418 (0.064)
BC (Wernicke)	16.7 (6.7)	35.2 (6.8)
BC (Broca)	32 (17)	26 (11)
BC (Motor Cortex)	138 (56)	147 (57)

Table 5.4: Mean and standard deviation of each network index for ASD patients and typically developed subjects. Local and global clustering coefficient and global efficiency have a theoretical range between 0 and 1. The betweenness centrality is not normalized.

reduction in Broca's area as well as Wernicke's area (Fig. 5.6), the betweenness centrality is reduced only in Wernicke's area (Fig. 5.7).

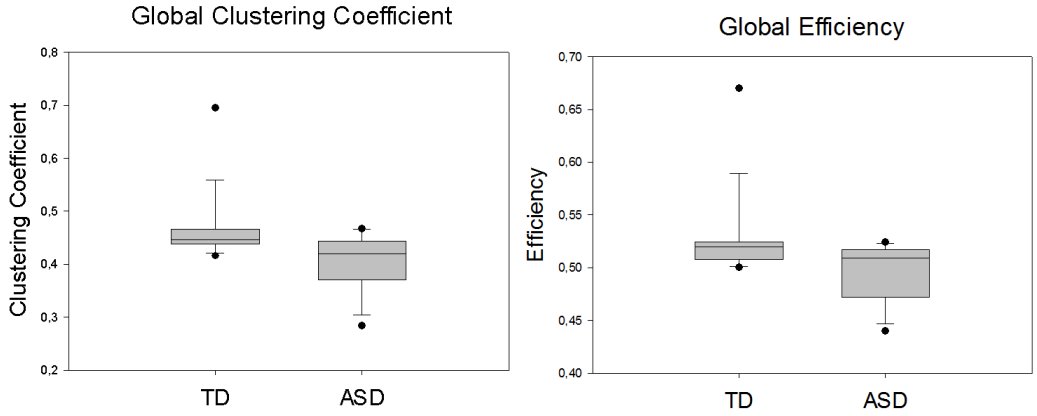

Figure 5.5: Global clustering coefficient and global efficiency in typically developed subjects (TD) and in patients suffering from ASD. The efficiency represents the inverse average shortest path length between nodes in the network. The global clustering coefficient describes to what extent the network forms clusters in which all nodes are interconnected. Global efficiency ($p = 0.01$) and the global clustering coefficient ($p = 0.013$) are significantly reduced for ASD patients.

5.3.2 Best Classification using Regional Indices

The discriminative power of the local indices was analyzed by means of ROC analysis (Fig. 5.8). With an area under curve of 0.97 the betweenness centrality of Wernicke's area outperforms the clustering coefficient of the same area with an area under curve of 0.77.

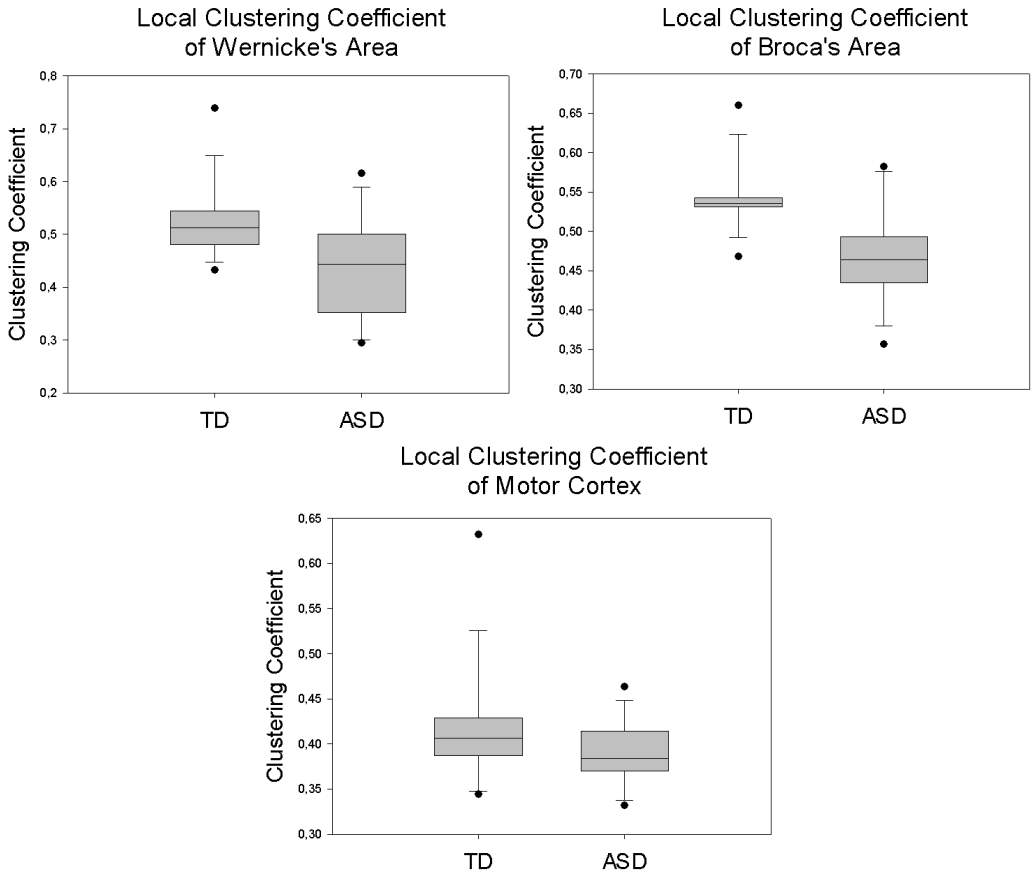

Figure 5.6: Local clustering coefficient of Wernicke's area (comprehension of speech), Broca's area (motor control of speech) and the motor cortex, comparing typically developed subjects (TD) and patients suffering from ASD. The clustering coefficient measures how many nodes in the neighborhood of a given node are connected to each other. It is significantly reduced for Wernicke's area ($p = 0.01$) and Broca's area ($p = 0.002$), but not for the primary motor cortex ($p = 0.15$) in ASD.

No significant difference between ASD patients and TD subjects can be found in the lateralization coefficients for the betweenness centrality of Wernicke's area ($p = 0.79$) and Broca's area ($p = 0.07$) or the clustering coefficient of Wernicke's area ($p = 0.68$) and Broca's area ($p = 0.48$).

Tractography results for the ASD group as described by the number of fibers (SD) are, at 6100 (1700), lower and more variable than for the control group at 7970 (770), though the difference is not significant ($p = 0.14$). The control average does not include a single outlier at 56089 fibers, which was included in the evaluation and produced centrality measures similar to the rest of the group. To ascertain that the observed differences in the betweenness centrality are not due to this difference in the number of fibers, a statistical analysis was performed with the number of fibers as covariate and showed that the observed significant differences persist ($t = 6.7$, p

< 0.001).

5.4 Influence of Density

5.4.1 Different Thresholding Methods

When thresholding different connectomes with a given weight, differences in classification accuracy can be due to the chosen method for thresholding. Two thresholding schemes used in this work are density based, whereas the last one is based on the number of fibres represented by a single connection. While density based thresholds can reasonably be assumed to not differ between ASD networks and TD networks the same is not necessarily true for thresholding based only on weight, as the number of fibres might be different in each case.

Figure 5.9 shows the change in density when thresholding is based on edge weight. As can be seen the density for a specific threshold seems to be consistently lower in case of ASD, suggesting a generally lower weight of connections in this case. To get a baseline for the performance of different network indices as base for classification we performed a classification based on this difference. The results for this can be seen in figure 5.10.

As can be seen in the figure it should be noted, that the classification accuracy based solely on the density for a given threshold remains relatively stable across a wide range of thresholds and below 0.8 for all tested thresholds. This gives us a minimum desired classification accuracy of 80 % for more complex network indices which should show a change in actual network architecture for weight threshold based approaches.

It should be noted again, that this only applies to classification based on a specific threshold weight, not classification based on a specific threshold density.

5.4.2 Changes in Classification Accuracy

Figures 5.11, 5.12, 5.13 and 5.14 show how well classification of ASD networks works for different thresholding strategies and different networks indices.

Figure 5.11 shows that global network indices can be reasonably used to classify between ASD and TD networks. It is especially noticeable, that the average betweenness centrality of all nodes proves to be a better classifier at low densities and perform significantly worse at high densities. This is to be expected as at high densities the expected betweenness centrality of all nodes will be very low as shortest paths tend to be very short, one or two steps as most. As such few nodes will actually lie on such paths. Interestingly the small worldness of the networks as well as the clustering coefficient perform relatively stable over a far larger range of densities suggesting differences at multiple scales.

For weight based thresholding the clustering coefficient is the only acceptable classifier, the betweenness centrality as well as the small worldness can not be rea-

sonably used as classifiers in this case. This is most likely due to the differences in density shown in figure 5.9.

Figure 5.12 shows the performance of the regional network indices for Broca's area. They do perform slightly better than chance ranging mostly from 0.5 to 0.8. It should be noted that all three indices perform similarly to each other and relatively stable over a wide range of densities.

In figure 5.13 we see the classification based on the regional indices for the motor cortex. The result here is similar to figure 5.12. Performance is relatively stable over a large range of densities and no classifier performs very well. It should be noted however, that not all three identifiers perform similarly well in this case, the clustering coefficient performs consistently worse than the other two.

Figure 5.14 in turn shows the performance for regional indices in Wernicke's area. It should be noted, that the classifier perfomance is significantly better than in any other presented cases. While the clustering coefficient results are comparable to the other areas the average degree and the betweenness centrality can reach classification accuracies above 0.9 . Also the dependency of classification accuracy on density is far more marked for these indices in Wernicke's area than in any other area. It should also be noted that the betweenness centrality performs consistently as good as or better than the average degree.

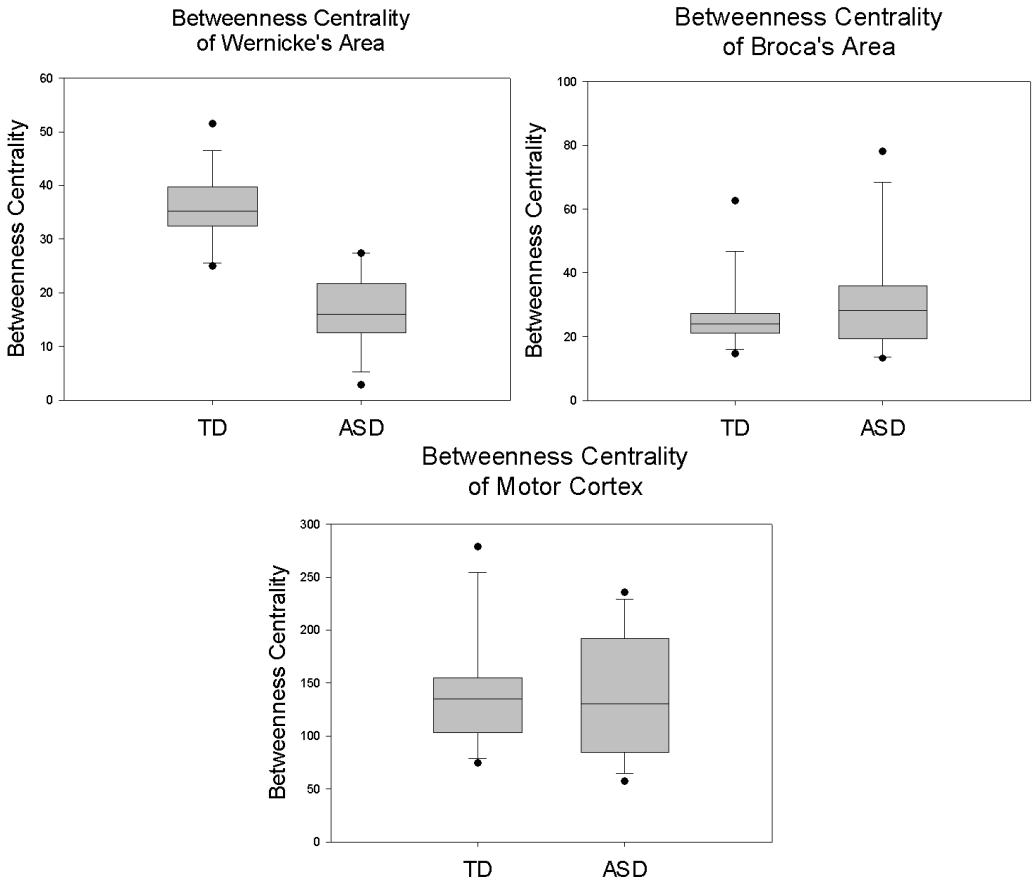

Figure 5.7: Betweenness centrality of Wernicke's area (comprehension of speech), Broca's area (motor control of speech) and the primary motor cortex (general motor control), comparing typically developed subjects and patients suffering from ASD. The betweenness centrality represents how integrated the corresponding area is in the connectome by quantifying the number of shortest paths traversing the area. The betweenness centrality for Wernicke's area is significantly reduced in ASD ($p < 0.001$), whereas it is not in Broca's area ($p = 0.27$) and the motor cortex ($p = 0.71$)

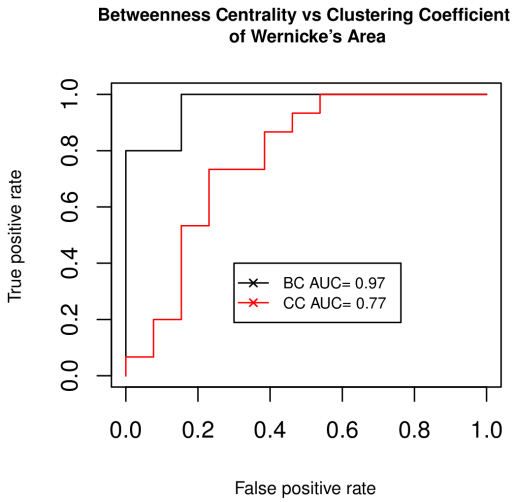

Figure 5.8: Receiver operating characteristic when classifying ASD patients and TD based on the betweenness centrality (BC) or based on the clustering coefficient (CC) of Wernicke's area. Classification based on BC performs considerably better than based on CC.

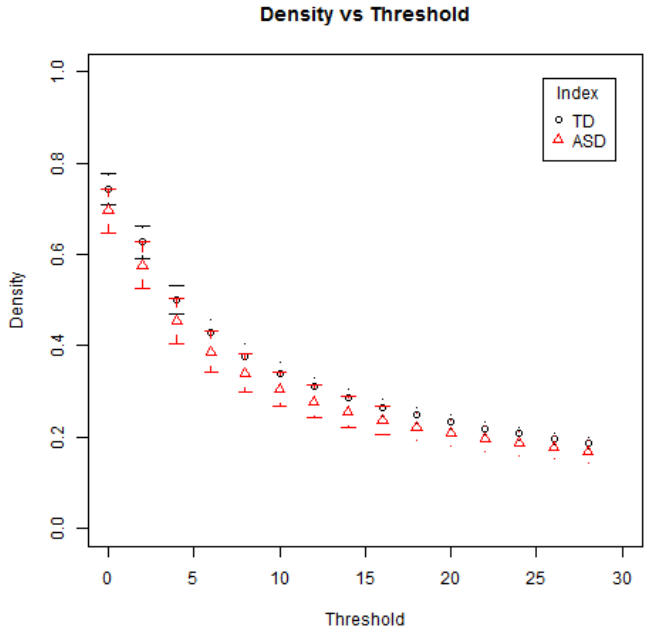

Figure 5.9: Resulting densities for different threshold values for autism spectrum disorder (ASD) patients and typically developed (TD) controls.

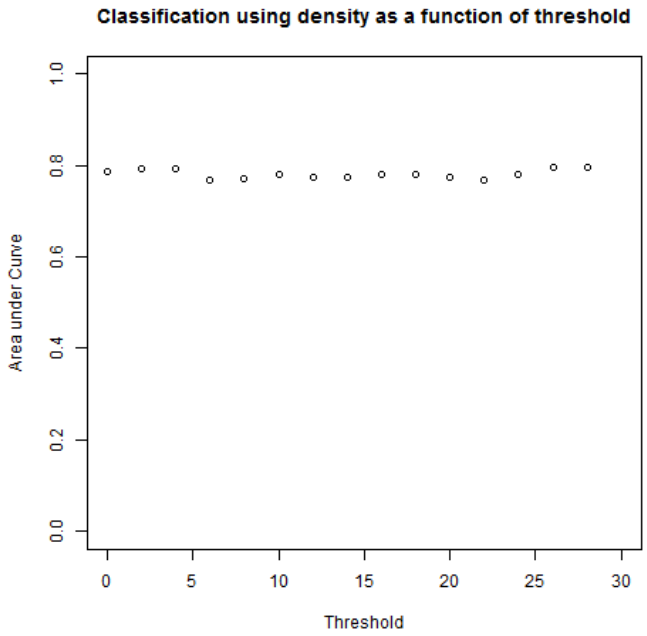

Figure 5.10: Classification based on the resulting density for a given threshold. It should be noted classification accuracy stays always below 80 % .

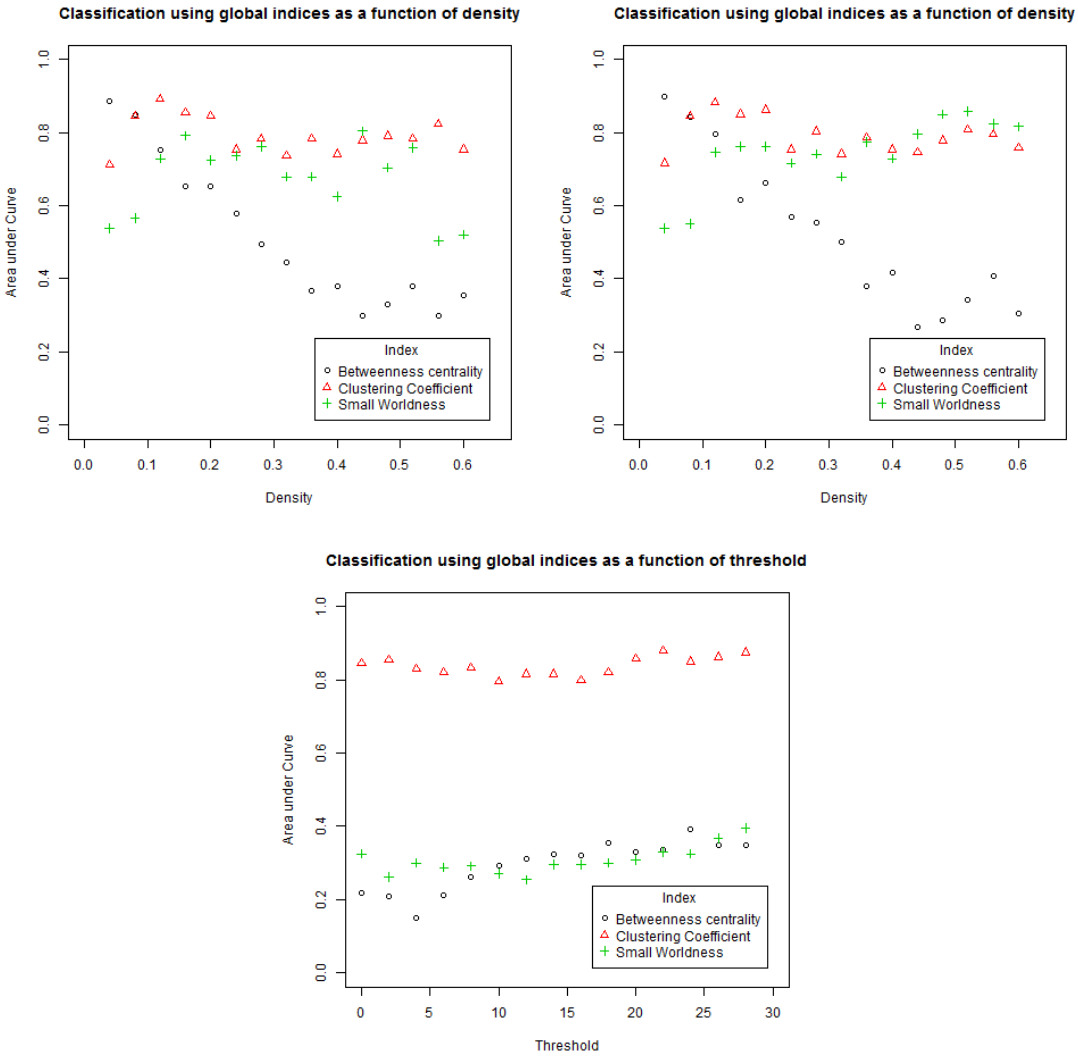

Figure 5.11: Classification based on global network indices. Top left: Smallest threshold resulting in a density below the target one. Top right: Random removal of weakest edges to reach target density. Bottom: Threshold based on edge weight.

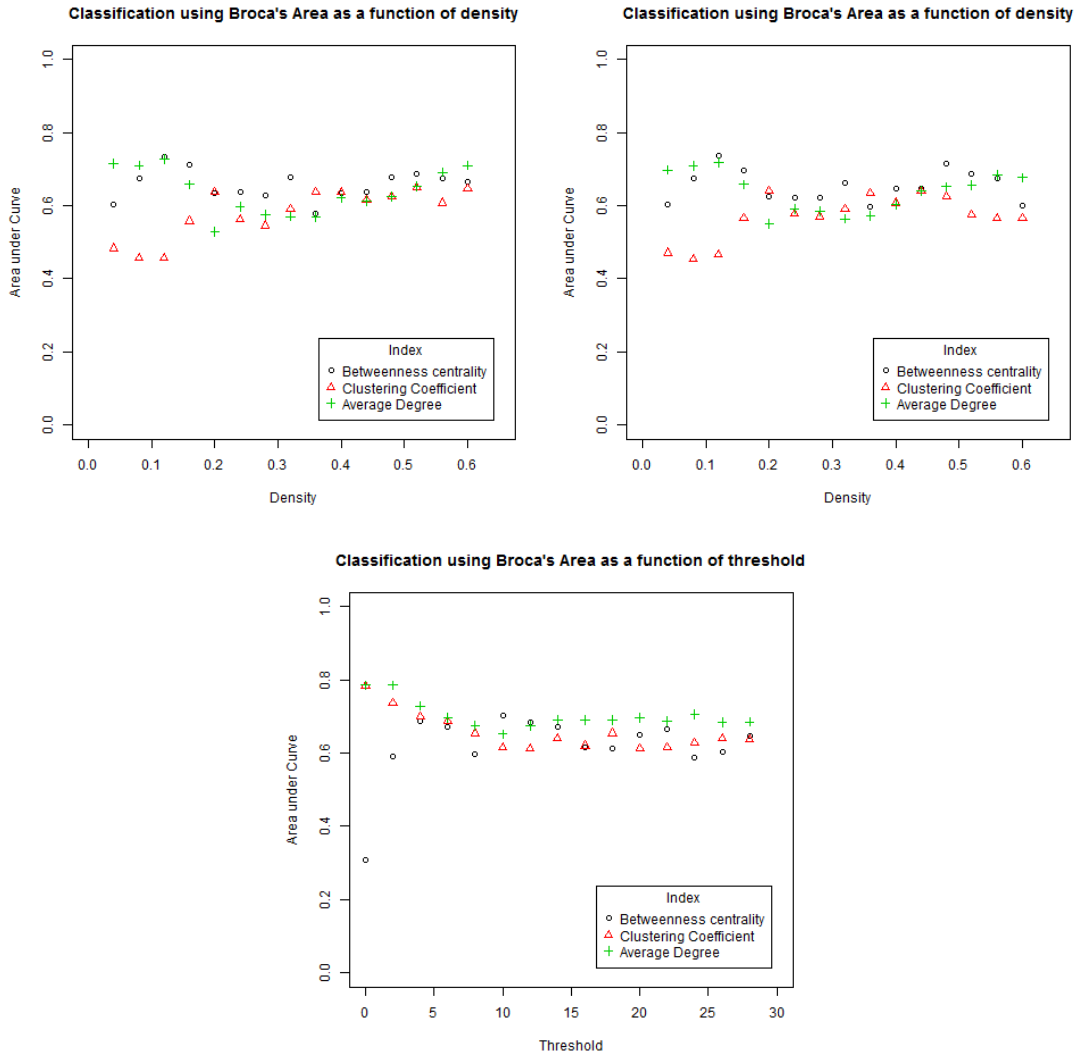

Figure 5.12: Classification based on local network indices for Broca's area. Top left: Smallest threshold resulting in a density below the target one. Top right: Random removal of weakest edges to reach target density. Bottom: Threshold based on edge weight.

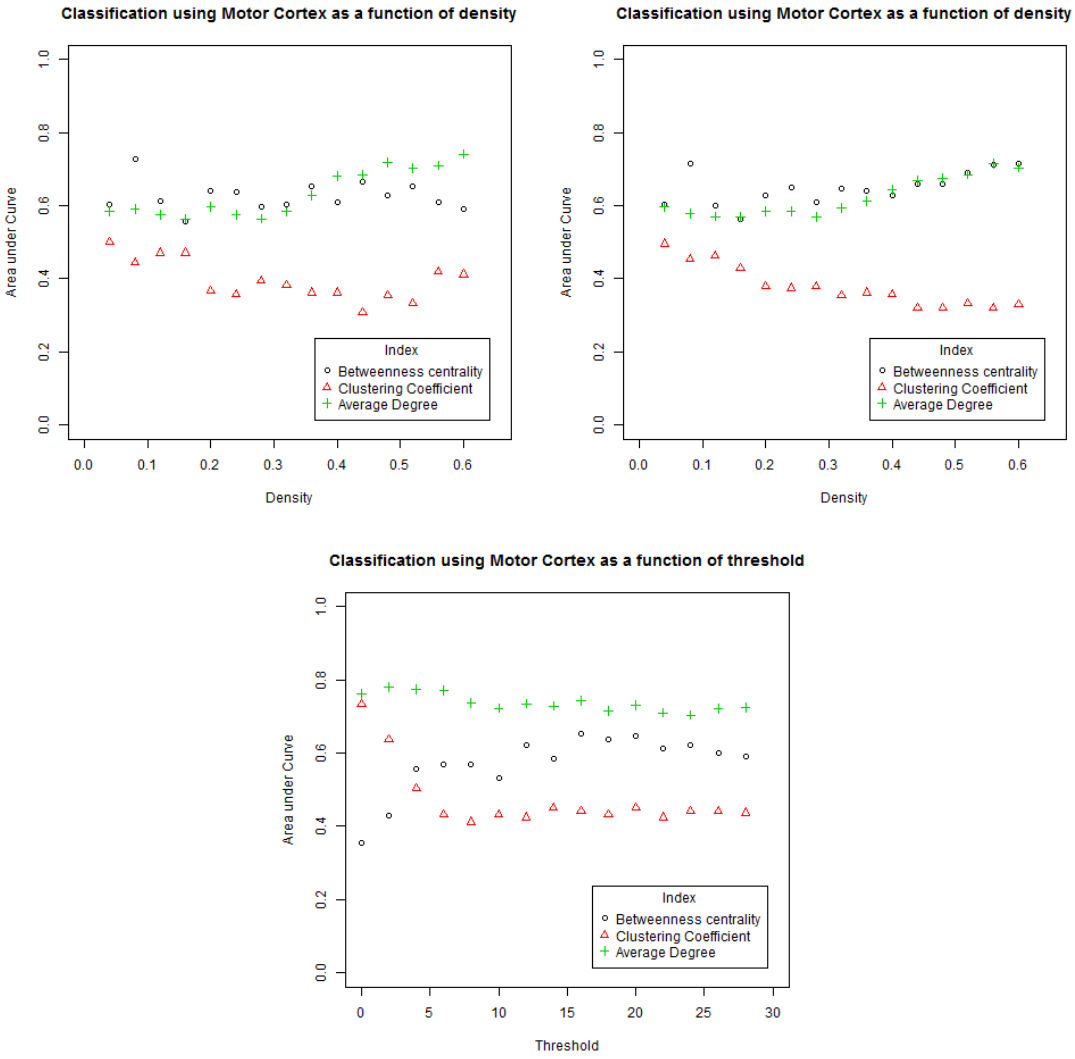

Figure 5.13: Classification based on local network indices for the motor cortex. Top left: Smallest threshold resulting in a density below the target one. Top right: Random removal of weakest edges to reach target density. Bottom: Threshold based on edge weight.

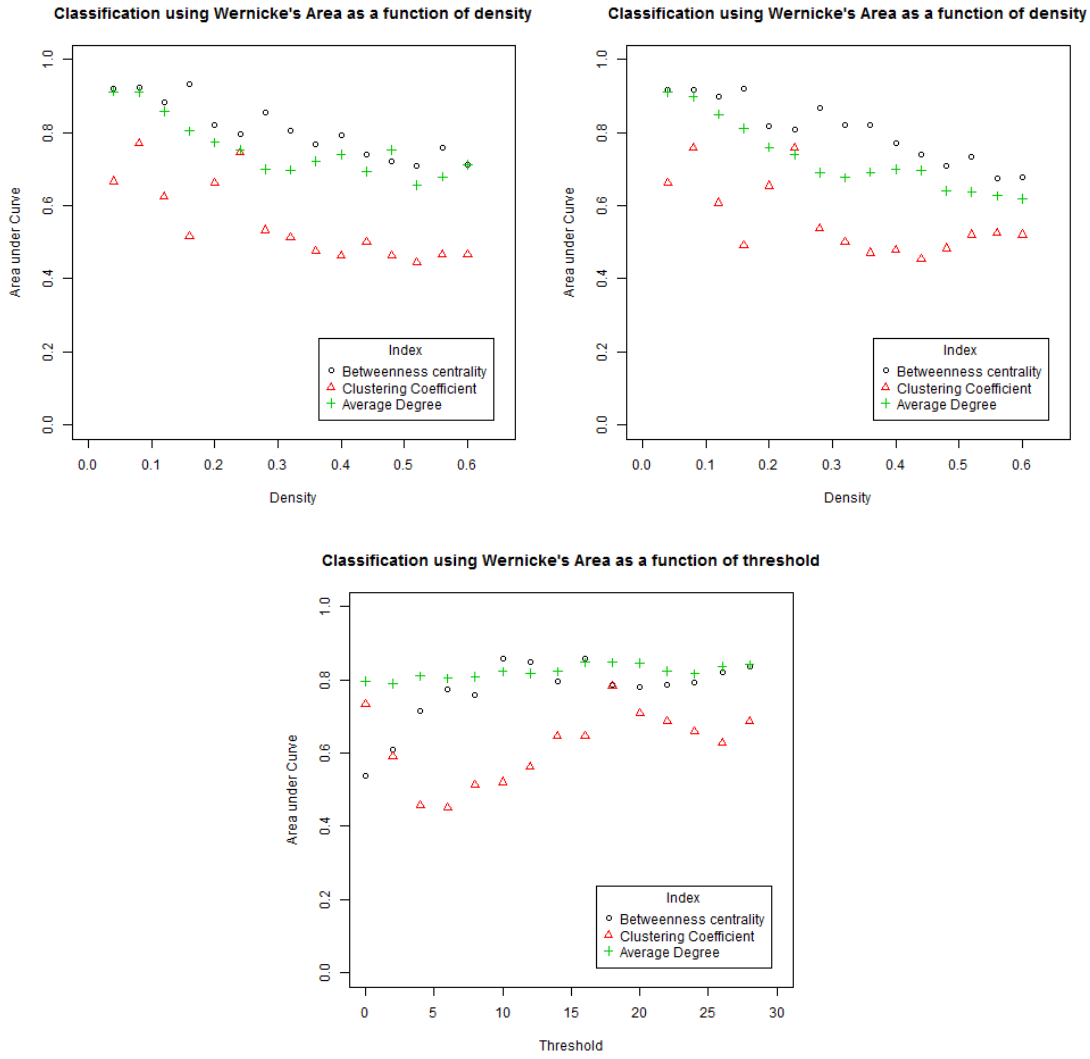

Figure 5.14: Classification based on local network indices for Wernicke's area. Top left: Smallest threshold resulting in a density below the target one. Top right: Random removal of weakest edges to reach target density. Bottom: Threshold based on edge weight

6 Discussion

6.1 Graph Construction

We have proposed a pipeline for the creation of connectome networks from structural DWI data and T1 images. While structural connectomes have been created before, we provide an easy to use open-source solution, that ties in with the Medical Imaging Interaction Toolkit. In addition to manual use we implemented an automated solution for batch processing.

Our solution has the further advantage of being able to use additional information about the used parcellation to reduce the problem of prematurely stopping fibres. This allows the user to concentrate on the fibre tractography in the high anisotropy region of the white matter and extrapolating from these results into the low anisotropy grey matter regions.

We have used the presented solution in our experiments and have found it to provide connectomes which can be potentially be used for diagnostic purposes with a reasonable reliability. At the same time our pipeline remains highly modular and allows the replacement of the different tools used for separate processing steps. Depending on the used parcellation, prediction of grey matter areas can be switched on or off. At the same time the used tractography method can easily be changed.

6.2 Global Network Assessment

We have analysed a wide range of network indices for their capability for classifying the connectome of children suffering from ASD. While several of these features perform reasonably well as classifiers no one feature stands out as singly important. Instead we note, that using a combined classifier taking into account more than one feature performs considerably better at a classification accuracy of about 90%. This exceeds the single feature classification based on the average betweenness centrality by 10%.

We also found that the classification accuracy does not continue to increase when including more and more features. Instead the optimal performance was reached when using 5 features for our classifier and dropped sharply after that. Together with our analysis of the number of times features were part of the top performing classifier for different thresholds this suggests that the best features for classification are Giant Connected Component Ratio, Clustering Coefficient D, Normalized Laplacian Trace, Average Connected Component Size and Normalized Laplacian Energy.

We used the same settings and automated pipeline for all networks including thresholding them at the same level to remove spurious connections. As can be

seen in figure 5.4 the thresholding does affect the performance of the classifier. At the same time we included the number of overall edges as possible feature and it did not perform well enough to be included in the top five features classifiers. This indicates that the difference in network architecture is due to a change in network configuration rather than a possible artefact due to a global change in anisotropy the resulting difference in fibre count.

The classification performance remains within four percent points of 82% for a large range of thresholds. While this is an acceptable accuracy it poses the problem of finding the optimal threshold for classification performance in a clinical setting. In such a setting it is not practicable to redo the analysis over a wide range of thresholds. Especially as the exact value of the threshold will depend on the acquisition and tractography settings, which are not necessarily available at the network analysis step. Additional noise and spurious fibres will reduce the classification accuracy as it can not necessarily be determined at which thresholding level they start to drop off.

A further limitation of this approach in a general diagnosis setting lies in the fact, that it detects a general change in network architecture rather than a more disease specific one. This might be used to support a possible prior diagnosis but will be of reduced help in a setting where a scan available due to other factors might be used to screen for potential other problems. More specific hypothesis driven approaches could ideally also result in a higher classification accuracy.

Our analysis also tries to alleviate the fact that we have a low number of unique subjects (32) compared to a large set of features (32) by creating four networks per patient from independent tractography runs. We believe that this technique helps reduce over fitting as during leave-one-out cross validation the four runs for each patient were left out together and not always assigned the same class, even though the likelihood of networks of one patient being classed together was higher than chance. Ideally the described analysis should be run across a significantly larger study to strengthen the indications found by us.

Despite its limitations we believe our work shows the need to consider additional networks measures, no only the most commonly used ones, when trying to analyse the changes in the human connectome due to disease and create diagnostic tools based on them.

6.3 Local and Regional Network Assessment

We have build on the concept of local network indices and expanded it to encompass anatomical regions which encompass several nodes. By introducing the concept of regions and focussing our statistical analysis on differences in certain regions we allow for a new prior knowledge based approach. In contrast to many classical networks, such as networks representing friendship of different persons, the nodes in the human connectome are not intrinsically similar. We know a priori that different nodes and groups of nodes are likely to be involved in certain tasks and we know at

the same time which tasks are affected in a disease.

Approaching the analysis of the human connectome from this angle opens two main ways to proceed. Firstly we can compare known symptoms of a disease to changes in the connectome and try to find their relevance. Secondly we can try to analyse the differences in the diseased connectome and try to profile a disease based on these differences.

While the second approach has a higher diagnostic value as it can give a completely new insight into an illness, it has its drawbacks as well. Due to the multitude of network indices and possible regions that can be looked at it needs a very large data sample to avoid over fitting the data and a high number of false positives. Furthermore this approach needs a solid base of studies of the first type to set the found differences in relation and compare them to known changes. Slowly more works of this kind become part of the body of literature, but many of them still suffer from the fact that the ground is not adequately prepared.

This work follows the first approach and starts from a clear hypothesis, to understand the known symptom of reduced communicational capability in ASD patients, and uses our new method to observe possible changes in language related regions of the brain. This experiment has been designed to show the basic feasibility of this method and contrast the changes in regions expected to show differences to those were we expect no such changes based on prior knowledge. At the same time we want to find network indices beside the most commonly used ones that can be used as an indicator of such differences.

The results are very promising and indicate that a diagnostic classification based on the regional betweenness centrality might be used to distinguish between the healthy connectome and the connectome of ASD patients. Of special interest is that we deliberately avoided fine tuning of tracking and network parameters. This is at the same time a strength and weakness of this experiment. On the one hand it is an approach that subjects potentially different data to an identical pipeline with the same settings, which is likely what will happen in a clinical setting, and tries to classify based on the results of the process. On the other hand this risks that any differences in the regional network indices are nothing more than artefacts of differences earlier in the pipeline, such as a generally lower anisotropy and a resulting lower fibre count. While this is fine from a purely diagnostic viewpoint, as the classification accuracy remains the same, it does not preclude the possibility that the network analysis step is superfluous and a classification of similar accuracy can be achieved at an earlier step in the pipeline.

This possibility was examined in a different part of this work and is discussed in section 6.4. There we examine how changes in the network parameters affect the classification accuracy.

By focusing on specific areas of the brain and not comparing all nodes across all patients, we try to avoid the multiple testing problems inherent with comparing such a large number of variables [73].

6.3.1 Medical Relevance

With the given unoptimized and unadjusted settings we achieved an area under curve of 0.97 for the classification based solely on the regional betweenness centrality of Wernicke's Area, which can be considered a very good classifier, albeit on a relatively small sample size (13 ASD, 15 TD) so it should be taken as an indication of the quality of the classifier only. It should be pointed out, that the subjects were rather old, 9.7 $\pm$ 2.1 years, an age, were ASD will usually have been detected by conventional means, such as questionnaires. This pipeline is however potentially applicable to the infant brain as well, although for very young children, less than a year, a different parcellation scheme should be used. Even though very promising results for older children do not guarantee similarly good results on younger children we believe that the presented method might be a useful technique.

However as ASD subjects and controls were matched for IQ our ASD subjects consisted of High Functioning Autism and Asperger patients only, which work comparably well in society. If used for diagnostic screening the differences could potentially be more pronounced for the average ASD patient and thus be even more easily distinguished.

6.4 Influence of Density

We have examined the limitations described in section 6.3 and tried to address the issues raised there to evaluate the performance of our suggested regional network indices across a range of settings. By using three different ways to get binary networks from weighted networks we try to cover both approaches to evaluating the performance of our method.

Firstly we remove all edges representing less than a given fibre count. This thresholding approach is methodically closest to the evaluation discussed in section 6.4. The entire pipeline is run with the same settings not correcting for possible differences in overall fibre count potentially introduced in a prior step.

Secondly we choose an individual threshold for each network to be as close as possible below the target density. This approach tries to correct for differences in fibre count while avoiding possible changes by random deletion.

Lastly we use an individual threshold to be as close as possible above the target density and then randomly remove the weakest links until a given density is reached. This way introduces randomness as a trade-off for having well comparable networks.

6.4.1 Classification Based on Threshold

In order to understand the differences prior to regional network analysis we calculated the density as a function of threshold and found that on average the ASD networks had a lower resulting density for the same threshold weight. This implies that the ASD connectomes have more weaker edges. This is not necessarily due to the corresponding tractograms having correspondingly less fibres, but instead can

also be due to a steeper fibre count distribution, where a larger number of fibres is accounted for by the "thickest" edges, or a larger number of fibres being discarded as they could not be correctly extrapolated.

As such a classification based on the density after thresholding still takes advantage of the presented framework, albeit not the regional network indices. It, however, gives us a reasonable classification accuracy that a classifier based on a network index should exceed in order to be considered to contain additional information. We can in turn consider any classifier below 80 % classification accuracy to not contain any additional information and be a bad classifier for our purposes.

6.4.2 Global Network Indices Across Densities

In addition to containing more information than the fibre count histogram, we expect a good regional network index to be better suited for classification than the global network indices. In order to compare the performance we examined three of the most commonly used global network indices.

We found that the classification accuracy using an universal threshold is relatively stable across a range of thresholds. This leads us to assume that the changes in network indices due to the different density dominate the classification as this difference is relatively stable. At the same time the area under curve is close to 0.8 or 0.2 respectively which we have identified as threshold for a good classifier based on the thresholding accuracy due to the density difference.

In contrast we can see a marked difference in the classification accuracy as a function of the target density. While classification accuracy for high densities does not exceed the 80 % mark we find that global classifiers based on the betweenness centrality and the clustering coefficient exceed this threshold for low densities. The best classification accuracy for the clustering coefficient is between a density of 0.1 and 0.3. The optimum for the betweenness centrality is below 0.1.

Both achieve a top performance of roughly 90 % which is a good classification accuracy. This in turn gives us a new threshold which we expect to exceed for our regional network index based classifiers.

6.4.3 Regional Network Indices Across Densities

We then tested our regional network indices to test whether our previous, promising results can be confirmed. We tested two commonly used network indices, the local degree and the clustering coefficient, as regional network indices, as well as the betweenness centrality, which we believe to be more useful in this context.

Our preliminary results suggested changes in the importance of Wernicke's Area and little to no change in the role of Broca's Area and the Motor Cortex. Our study seems to confirm these results.

Both Broca's Area and the Motor Cortex remain considerably below the 0.8 threshold for accuracy we set as minimum bar for additional information contained in the classifier. Especially the Motor Cortex based classifiers remain close to the

0.5 mark, which is an interesting result in itself. It suggests that the role of the Motor Cortex remains essentially the same and its embedding in the network does not change at all. This seems to confirm our original idea of using the Motor Cortex as internal validation for our method due to its unchanged importance. The only classifier deviating from this is the average degree in the weight based threshold case, which is most likely due to the difference in densitiy which is directly proportional to the difference in average degree as both are a direct function of the number of connections. The deviation for high densities could point to a higher percentage of "noisy" connections, if edges representing only few fibres are included.

Broca's Area based classifiers are not as close to random, but remain well below the 0.8 mark and can be discarded as adding additional information.

For Wernicke's Area based classifiers the situation is different. While here as well the weight threshold based classifiers remain relatively stable and close to the 0.8 mark, most likely again due to the difference in density, the density based classifiers show a different picture. The clustering coefficient remains below the mark and can be discounted as a useful classifier in this case, but the average degree and the betweenness centrality are markedly above the 0.8 mark for a range of densities and even exceed the 0.9 mark of the global classifiers for densities of about 0.05 to 0.15. The change in average degree at the same overall number of edges indicates that Wernicke's Area is relatively worse connected in ASD. The reduced betweenness centrality could be due to this reduction in connections, as fewer short shortest paths inherently run through a less well connected node, but the classifier based on the betweenness centrality consistently outperforms the one based on the average degree, indicating that it reflects a more global change in the role of Wernicke's Area than the purely local average degree.

6.4.4 Clinically Relevant Densities

Part of this experiment was to evaluate at which densities the analysis of the human connectome is clinically relevant. This question goes beyond the problem of edges introduced by noise and how to avoid them, although this is an important issue, and is relevant even for a completely noise-less connectome. At very high densities, close to 1, the binary connectome will trivially not contain useful information as in these boundary cases the network indices will have the same value for all networks. The same is true at very low densities, close to 0, where the network will be mostly disconnected.

However the question at which densities changes in the network due to illnesses are most pronounced can not trivially be answered. This question is linked to the question which kind of connections change the most in an illness, the very prominent ones, which can be easily found by fibre tractography and will be present with thousands of fibres, or the subtle ones, which represent only a few fibres and might be drowned by noise. We can not clearly answer this question, as it can be expected that our tractographies, and consequently connectomes, are not noise-free, but we at the very least give indications whether even prominent connections are altered in

Table 6.1: Density ranges of best classification performance for different classifiers. Top are global indices, bottom are regional indices of Wernicke's Area. Densities below 0.05 were not evaluated. Only network indices exceeding 85 % classification accuracy are listed.

Network Index	Density Range
Clustering Coefficient	0.1-0.3
Betweenness Centrality	< 0.1
Average Degree	< 0.15
Betweenness Centrality	< 0.15

diseased connectomes.

We have found, that the optimum classification window differs for different kinds of network indices. Table 6.1 shows the areas of best classification for different network indices.

Given that the acquisition protocol for ASD patients and TD subjects is identical we can assume that the acquisition process itself does not differentiate between the two groups. If the main difference between the two groups is due to an overall change in anisotropy we would assume that the fibre tractography results differ especially in the subtle connections and less in the main connections as these tracts can easily be found even if the anisotropy is slightly decreased. Instead we find that the main difference between the groups affects the top third of connections. While noise could potentially mask the differences in subtle connections, our findings suggests that the differences in the connectome are not due to some global change in anisotropy, but instead organizational changes.

The global clustering coefficient indicates that even at relatively high densities the network tends to form fewer tightly knit cliques in the case of ASD. Instead the connections are more evenly distributed. In conjunction with the reduced classification accuracy for lower densities we get an image, where both ASD patients and TD subjects form cliques of nodes that are strongly connected, but whereas these cliques are expanded by weaker connections in the healthy connectome to form larger cliques this is not the case in ASD, where these weaker connections tend to be more between different cliques.

In contrast to this the difference in the regional average degree as well as the regional and global betweenness centrality becomes more pronounced at lower densities. This suggests that the organization of the strong connections determines the main role of different areas in the network. While an area might be equally well connected with weaker links the reduction in strong links limits seems to limit the capability of these areas.

Our results suggest that for diagnostic purposes the involvement of a region at

low densities gives the best indication of a possibly reduced capability.

6.4.5 Changes in Autism Spectrum Disorders

In line with our earlier preliminary results these expanded experiments indicate that the reduced capability for oral communication in ASD reflects in the regional betweenness centrality of Wernicke's Area, whereas other regions are less affected. By focussing on specific regions of the brain we gain a specificity which the global measurements lack. We did only measure changes in certain regions of the brain in the scope of this work. The methods we implemented however can easily be adapted for other diseases and could be a useful tool for diagnosing diseases based on the affected regions.

6.5 Summary

We have presented, implemented and evaluated a new method to incorporate prior anatomical information into graph indices by adding an additional region layer between the local and global network indices. This work applies this technique to use the betweenness centrality of Wernicke's Area to classify ASD patients and typically developed controls.

Using this approach we have shown on a small set of patients that this regional betweenness centrality has a good diagnostic value and surpasses a purely global classification, even using more complex classifiers and learning approaches, as well as other indices traditionally more commonly used in current research.

Looking back to our original objectives we have fulfilled our goals. Incorporating anatomical and medical information into the region layer yields an improved index, which aligns well with the observed symptoms. At the same time the reduced centrality of Wernicke's Area points to a change in the network architecture which furthers our understanding of the changes in the role of different areas in ASD. Especially in combination with no apparent change in Broca's Area this points to a possible change in the development of the pathways in the brain.

Furthermore we incorporated our research into an openly accessible tool set, which has been designed to combine the power and modularity of C++ with the ease of use of binary releases for medical professionals.

6.6 Outlook and Further Work

The framework provided during the course of this work is easily applicable to cover other diseases. The primary result will be the answer to the narrow question of identifying regional indices which are well suited for describing a specific disease. But by extending our data pool to multiple diseases a secondary question that can be pursued is to identify regional indices which are unique to a certain disease. Contrary to the artificial lab environment a diagnostic tool in the clinic will not

have to decide whether a patient is suffering from a disease A or not, but instead will have to discriminate between a large number of possible diseases. Supplying such a tool with a fingerprint of expected changes in the brain associated with a disease is therefore of a high priority.

On the medical side extending our approach to younger children or infants suffering from ASD would provide a more useful diagnostic tool, as having an objective and quantifiable assessment is clinically more important at a younger age. By increasing the sample size and correcting for multiple tests a more complete analysis of different regions can discover associations to other symptoms.

Lastly using the established techniques to define a baseline accuracy for different techniques a more thorough investigation of using different weighting methods would provide a theoretically more grounded framework for estimating real edge weights from the image data and tractograms in the absence of ground truth.

Part I

Appendix

A Acknowledgement

This work was conducted at the Division of Medical and Biological Informatics (MBI) of Prof. Dr. H.-P. Meinzer at the German Cancer Research Center (DKFZ) in Heidelberg. It was financed by the Helmholtz International Graduate School.

I thank my supervisor Prof. Dr. H.-P. Meinzer for the chance to work on this project. Further thanks goes to my supervisor from the faculty of physics and astronomy of the university of Heidelberg Prof. Uwe Oelfke. I am also indebted to my advisor Dr. Klaus Maier-Hein who bore the brunt of the day to day supervision as well as being instrumental in the development of the research question itself as well as any answers to it. Special thanks also to Dr. Bram Stieltjes and Dr. Romy Henze who provided me with medical feedback and helped understanding the theoretical changes in context of the human brain and ASD. I want to express my gratitude to the colleagues at MBI for answering any question on any topic, a lot of support and great atmosphere at work. Especially the members of H834 and the Diffusion/CAD/MIC group have all helped the development of this work.

Finally I want to thank those people who supported me through my entire education and enabled me to complete it. Many thanks to my parents, and my family, for support and motivation. A special thank you to Simona Schmid, who was always there for me the past years.

To all those friends who have made the time in Heidelberg so enjoyable I want to say thank you.

B Lists

B.1 List of Figures

B.2 List of Tables

C Bibliography

[1] Daniel C. Alexander, Penny L. Hubbard, Matt G. Hall, Elizabeth A. Moore, Maurice Ptito, Geoff J.M. Parker, and Tim B. Dyrby. Orientationally invariant indices of axon diameter and density from diffusion MRI. *NeuroImage*, 52(4): 1374 – 1389, 2010. ISSN 1053-8119.

[2] A Anwander, M Tittgemeyer, DY von Cramon, AD Friederici, and TR Knösche. Connectivity-Based Parcellation of Broca's Area. *Cerebral Cortex*, 17(4):816–825, 2007.

[3] Jon Baio. Prevalence of Autism Spectrum Disorders - Autism and Developmental Disabilities Monitoring Network, 14 Sites, United States, 2008. *Department of Health and Human Services, Centers for Disease Control and Prevention.Morbitity and Mortality Weekly Report*, 2012.

[4] Pablo Barttfeld, Bruno Wicker, Sebastián Cukier, Silvana Navarta, Sergio Lew, and Mariano Sigman. A big-world network in ASD: Dynamical connectivity analysis reflects a deficit in long-range connections and an excess of short-range connections. *Neuropsychologia*, 49(2):254 – 263, 2011. ISSN 0028-3932.

[5] Pablo Barttfeld, Bruno Wicker, Sebastián Cukier, Silvana Navarta, Sergio Lew, Ramón Leiguarda, and Mariano Sigman. State-dependent changes of connectivity patterns and functional brain network topology in autism spectrum disorder. *Neuropsychologia*, 50(14):3653 – 3662, 2012. ISSN 0028-3932.

[6] P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance Series B*, 103:247–54, 1994.

[7] Danielle S Bassett and Edward T Bullmore. Human brain networks in health and disease. *Current Opinion in Neurology:*, 22:340–347, 2009.

[8] Danielle S. Bassett, Andreas Meyer-Lindenberg, Sophie Achard, Thomas Duke, and Edward Bullmore. Adaptive reconfiguration of fractal small-world human brain functional networks. *Proceedings of the National Academy of Sciences*, 103(51):19518–19523, 2006.

[9] Matteo Bastiani, Nadim Jon Shah, Rainer Goebel, and Alard Roebroeck. Human cortical connectome reconstruction from diffusion weighted MRI: The effect of tractography algorithm. *NeuroImage*, 62(3):1732 – 1749, 2012. ISSN 1053-8119.

[10] J.A. Bondy and U.S.R. Murty. *Graph Theory with Applications*. The Macmillan Press Ltd., 1976.

[11] Korbinian Brodmann. *Vergleichende Lokalisationslehre der Großhirnrinde : in ihren Prinzipien dargestellt auf Grund des Zellenbaues*. Barth, Leipzig, 1909.

[12] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*, 10:186–198, 03 2009.

[13] Edward T. Bullmore and Danielle S. Bassett. Brain Graphs: Graphical Models of the Human Brain Connectome. *Annual Review of Clinical Psychology*, 7 (1):113–140, 2011.

[14] Martha S. Burns and Jill Fahy. Broca's Area: Rethinking Classical Concepts From a Neuroscience Perspective. *Topics in Stroke Rehabilitation*, 17:401–410, 2010.

[15] H. Y. Carr and E. M. Purcell. Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments. *Physical Review*, 94:630–638, May 1954.

[16] Marco Catani and Dominic H. ffytche. The rises and falls of disconnection syndromes. *Brain*, 128(10):2224–2239, 2005.

[17] Zhang J. Chen, Yong He, Pedro Rosa-Neto, Jurgen Germann, and Alan C. Evans. Revealing Modular Architecture of Human Brain Structural Networks by Using Cortical Thickness from MRI. *Cerebral Cortex*, 18(10):2374–2381, 2008.

[18] Jonathan D. Clayden, Michael Dayan, and Chris A. Clark. Principal Networks. *PLoS ONE*, 8(4):e60997, 04 2013.

[19] Reuven Cohen and Shlomo Havlin. Scale-Free Networks Are Ultrasmall. *Phys. Rev. Lett.*, 90:058701, Feb 2003.

[20] John N. Constantino, Sandra A. Davis, Richard D. Todd, Matthew K. Schindler, Maggie M. Gross, Susan L. Brophy, Lisa M. Metzger, Christiana S. Shoushtari, Reagan Splinter, and Wendy Reich. Validation of a Brief Quantitative Measure of Autistic Traits: Comparison of the Social Responsiveness Scale with the Autism Diagnostic Interview-Revised. *Journal of Autism and Developmental Disorders*, 33(4):427–433, 2003. ISSN 0162-3257.

[21] Thomas E. Conturo, Nicolas F. Lori, Thomas S. Cull, Erbil Akbudak, Abraham Z. Snyder, Joshua S. Shimony, Robert C. McKinstry, Harold Burton, and Marcus E. Raichle. Tracking neuronal fiber pathways in the living human brain. *Proceedings of the National Academy of Sciences*, 96(18):10422–10427, 1999.

[22] Alessandro Daducci, Stephan Gerhard, Alessandra Griffa, Alia Lemkaddem, Leila Cammoun, Xavier Gigandet, Reto Meuli, Patric Hagmann, and Jean-Philippe Thiran. The Connectome Mapper: An Open-Source Processing Pipeline to Map Connectomes with MRI. *PLoS ONE*, 7(12):e48121, 12 2012.

[23] Madelaine Daianu, Neda Jahanshad, Talia M. Nir, Arthur W. Toga, Clifford R. Jack, Jr., Michael W. Weiner, and Paul M. Thompson for the Alzheimer's Disease Neuroimaging Initiative. Breakdown of Brain Connectivity Between Normal Aging and Alzheimer's Disease: A Structural k-Core Network Analysis. *Brain Connectivity*, 3(4):407–422, 2013.

[24] Rahul S. Desikan, Florent Ségonne, Bruce Fischl, Brian T. Quinn, Bradford C. Dickerson, Deborah Blacker, Randy L. Buckner, Anders M. Dale, R. Paul Maguire, Bradley T. Hyman, Marilyn S. Albert, and Ronald J. Killiany. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage*, 31(3):968 – 980, 2006. ISSN 1053-8119.

[25] Ivo Dinov, John Van Horn, Kamen Lozev, Rico Magsipoc, Petros Petrosyan, Zhizhong Liu, Allan MacKenzie-Graha, Paul Eggert, Douglass S Parker, and Arthur W Toga. Efficient, distributed and interactive neuroimaging data analysis using the LONI pipeline. *Frontiers in Neuroinformatics*, 3(22), 2009. ISSN 1662-5196.

[26] Julio M. Duarte-Carvajalino, Neda Jahanshad, Christophe Lenglet, Katie L. McMahon, Greig I. de Zubicaray, Nicholas G. Martin, Margaret J. Wright, Paul M. Thompson, and Guillermo Sapiro. Hierarchical topological network analysis of anatomical human brain connectivity and differences related to sex and kinship. *NeuroImage*, 59(4):3784 – 3804, 2012. ISSN 1053-8119.

[27] D.C. Van Essen and K. Ugurbil. The future of the human connectome. *NeuroImage*, 62(2):1299 – 1310, 2012. ISSN 1053-8119.

[28] D.C. Van Essen, K. Ugurbil, E. Auerbach, D. Barch, T.E.J. Behrens, R. Bucholz, A. Chang, L. Chen, M. Corbetta, S.W. Curtiss, S. Della Penna, D. Feinberg, M.F. Glasser, N. Harel, A.C. Heath, L. Larson-Prior, D. Marcus, G. Michalareas, S. Moeller, R. Oostenveld, S.E. Petersen, F. Prior, B.L. Schlaggar, S.M. Smith, A.Z. Snyder, J. Xu, and E. Yacoub. The Human Connectome Project: A data acquisition perspective. *NeuroImage*, 62(4):2222 – 2231, 2012. ISSN 1053-8119.

[29] Leonard Euler. Solutio problematis ad geometriam situs pertinentis. *Commentarii academiae scientiarum Petropolitanae*, 8:128–140, 1741.

[30] Luca Ferrarini, Ilya M. Veer, Evelinda Baerends, Marie-José van Tol, Remco J. Renken, Nic J.A. van der Wee, Dirk. J. Veltman, André Aleman, Frans G.

Zitman, Brenda W.J.H. Penninx, Mark A. van Buchem, Johan H.C. Reiber, Serge A.R.B. Rombouts, and Julien Milles. Hierarchical functional modularity in the resting-state human brain. *Human Brain Mapping*, 30(7):2220–2231, 2009. ISSN 1097-0193.

[31] Bruce Fischl, André van der Kouwe, Christophe Destrieux, Eric Halgren, Florent Ségonne, David H. Salat, Evelina Busa, Larry J. Seidman, Jill Goldstein, David Kennedy, Verne Caviness, Nikos Makris, Bruce Rosen, and Anders M. Dale. Automatically Parcellating the Human Cerebral Cortex. *Cerebral Cortex*, 14(1):11–22, 2004.

[32] Alex Fornito, Andrew Zalesky, and Michael Breakspear. Graph analysis of the human connectome: Promise, progress, and pitfalls. *NeuroImage*, 80(0):426 – 444, 2013. ISSN 1053-8119.

[33] Linton C. Freeman. A Set of Measures of Centrality Based on Betweenness. *Sociometry*, 40(1):35–41, March 1977.

[34] Karl J. Friston. Functional and Effective Connectivity: A Review. *Brain Connectivity*, 1(1):13–36, 2011.

[35] Klaus H. Fritzsche, Peter F. Neher, Ignaz Reicht, Thomas van Bruggen, Caspar Goch, Marco Reisert, Marco Nolden, Sascha Zelzer, Hans-Peter Meinzer, and Bram Stieltjes. MITK Diffusion Imaging. *Methods of Information in Medicine*, 51:441 – 448, 2012.

[36] Cedric E. Ginestet, Thomas E. Nichols, Ed T. Bullmore, and Andrew Simmons. Brain Network Analysis: Separating Cost from Topology Using Cost-Integration. *PLoS ONE*, 6(7):e21570, 07 2011.

[37] Sundar Gnanavel and Ruby Stella Robert. Diagnostic and statistical manual of mental disorders, fifth edition, and the impact of events scale-revised. *CHEST Journal*, 144(6):1974–1974, 2013.

[38] Caspar J. Goch, Bram Stieltjes, Romy Henze, Jan Hering, Hans-Peter Meinzer, and Klaus H. Fritzsche. Reduced centrality of Wernicke's area in autism. *Proc. SPIE*, 8672:867223–867223–6, 2013.

[39] Caspar J. Goch, Basak Oztan, Bram Stieltjes, Romy Henze, Jan Hering, Luise Poustka, Hans-Peter Meinzer, Bülent Yener, and Klaus H. Maier-Hein. Global Changes in the Connectome in Autism Spectrum Disorders. In Thomas Schultz, Gemma Nedjati-Gilani, Archana Venkataraman, Lauren O'Donnell, and Eleftheria Panagiotaki, editors, *Computational Diffusion MRI and Brain Connectivity*, Mathematics and Visualization, pages 239–247. Springer International Publishing, 2014. ISBN 978-3-319-02474-5.

[40] Gaolang Gong, Pedro Rosa-Neto, Felix Carbonell, Zhang J Chen, Yong He, and Alan C Evans. Age- and gender-related differences in the cortical anatomical network. *Journal of Neuroscience*, 29(50):15684–15693, December 2009.

[41] Krzysztof Gorgolewski, Christopher D. Burns, Cindee Madison, Dav Clark, Yaroslav O. Halchenko, Michael L. Waskom, and Satrajit S Ghosh. Nipype: A flexible, lightweight and extensible neuroimaging data processing framework. *Frontiers in Neuroinformatics*, 5(13), 2011. ISSN 1662-5196.

[42] Stephen J. Gotts, Hang Joon Jo, Gregory L. Wallace, Ziad S. Saad, Robert W. Cox, and Alex Martin. Two distinct forms of functional lateralization in the human brain. *Proceedings of the National Academy of Sciences*, 110(36): E3435–E3444, 2013.

[43] W.R. Gray, J.A. Bogovic, J.T. Vogelstein, B.A. Landman, J.L. Prince, and R.J. Vogelstein. Magnetic Resonance Connectome Automated Pipeline: An Overview. *Pulse, IEEE*, 3(2):42–48, March 2012. ISSN 2154-2287.

[44] Ellen Greimel, Barbara Nehrkorn, Martin Schulte-Rüther, GereonR. Fink, Thomas Nickl-Jockschat, Beate Herpertz-Dahlmann, Kerstin Konrad, and SimonB. Eickhoff. Changes in grey matter development in autism spectrum disorder. *Brain Structure and Function*, 218(4):929–942, 2013. ISSN 1863-2653.

[45] Alessandra Griffa, Philipp S. Baumann, Jean-Philippe Thiran, and Patric Hagmann. Structural connectomics in brain diseases. *NeuroImage*, 80(0):515 – 526, 2013. ISSN 1053-8119.

[46] P. Hagmann, O. Sporns, N. Madan, L. Cammoun, R. Pienaar, V. J. Wedeen, R. Meuli, J.-P. Thiran, and P.E. Grant. White matter maturation reshapes structural connectivity in the late developing human brain. *Proceedings of the National Academy of Sciences*, 2010.

[47] Patric Hagmann, Leila Cammoun, Xavier Gigandet, Reto Meuli, Christopher J Honey, Van J Wedeen, and Olaf Sporns. Mapping the Structural Core of Human Cerebral Cortex. *PLoS Biol*, 6(7):e159, 07 2008.

[48] E. L. Hahn. Spin Echoes. *Physical Review*, 80:580–594, November 1950.

[49] Yong He, Zhang J. Chen, and Alan C. Evans. Small-World Anatomical Networks in the Human Brain Revealed by Cortical Thickness from MRI. *Cerebral Cortex*, 17(10):2407–2419, 2007.

[50] Yong He, Alain Dagher, Zhang Chen, Arnaud Charil, Alex Zijdenbos, Keith Worsley, and Alan Evans. Impaired small-world efficiency in structural cortical networks in multiple sclerosis associated with white matter lesion load. *Brain*, 132(12):3366–3379, 2009.

[51] Claus C. Hilgetag, Gully A. P. C. Burns, Marc A. O'Neill, Jack W. Scannell, and Malcolm P. Young. Anatomical connectivity defines the organization of clusters of cortical areas in the macaque and the cat. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 355(1393): 91–110, 2000.

[52] Mark D. Humphries and Kevin Gurney. Network 'Small-World-Ness': A Quantitative Method for Determining Canonical Network Equivalence. *PLoS ONE*, 3(4):e0002051, 04 2008.

[53] Andrei Irimia, Micah C Chambers, Carinna M Torgerson, Maria Filippou, David A Hovda, Jeffry R Alger, Guido Gerig, Arthur W Toga, Paul M Vespa, Ron Kikinis, and John D Van Horn. Patient-tailored connectomics visualization for the assessment of white matter atrophy in traumatic brain injury. *Frontiers in Neurology*, 3(10), 2012. ISSN 1664-2295.

[54] Andras Jakab, Miklos Emri, Tamas Spisak, Anita Szeman-Nagy, Monika Beres, Sandor Attila Kis, Peter Molnar, and Ervin Berenyi. Autistic Traits in Neurotypical Adults: Correlates of Graph Theoretical Functional Network Topology and White Matter Anisotropy Patterns. *PLoS ONE*, 8(4):e60982, 04 2013.

[55] Heidi Johansen-Berg. Human connectomics - What will the future demand? *NeuroImage*, 80(0):541 – 544, 2013. ISSN 1053-8119.

[56] Heidi Johansen-Berg, Timothy E.J. Behrens, Emma Sillery, Olga Ciccarelli, Alan J. Thompson, Stephen M. Smith, and Paul M. Matthews. Functional-Anatomical Validation and Individual Variation of Diffusion Tractography-based Segmentation of the Human Thalamus. *Cerebral Cortex*, 15(1):31–39, 2005.

[57] Derek K. Jones and Mara Cercignani. Twenty-five pitfalls in the analysis of diffusion MRI data. *NMR in Biomedicine*, 23(7):803–820, 2010. ISSN 1099-1492.

[58] Derek K. Jones, Thomas R. Knösche, and Robert Turner. White matter integrity, fiber count, and other fallacies: The do's and don'ts of diffusion MRI. *NeuroImage*, 73(0):239 – 254, 2013. ISSN 1053-8119.

[59] Thomas R Knösche and Marc Tittgemeyer. The role of Long-Range Connectivity for the Characterization of the Functional-Anatomical Organization of the Cortex. *Frontiers in Systems Neuroscience*, 5(58), 2011. ISSN 1662-5137.

[60] B.W. Kreher, I. Mader, and V.G. Kiselev. Gibbs tracking: A novel approach for the reconstruction of neuronal pathways. *Magnetic Resonance in Medicine*, 60(4):953–963, 2008. ISSN 1522-2594.

[61] Nicolas Langer, Andreas Pedroni, and Lutz Jäncke. The Problem of Thresholding in Small-World Network Analysis. *PLoS ONE*, 8(1):e53199, 01 2013.

[62] Vito Latora and Massimo Marchiori. Efficient Behavior of Small-World Networks. *Phys. Rev. Lett.*, 87:198701, Oct 2001.

[63] F. B. Laun. *Diffusionstensor - Magnetresonanztomographie: Phantomentwicklung und Optimierung der Messtechnik für Anwendungen am Rückenmark.* PhD thesis, Ruprecht-Karls Universität Heidelberg, 2008.

[64] Jason P. Lerch, Keith Worsley, W. Philip Shaw, Deanna K. Greenstein, Rhoshel K. Lenroot, Jay Giedd, and Alan C. Evans. Mapping anatomical correlations across cerebral cortex (MACACC) using cortical thickness from MRI. *NeuroImage*, 31(3):993 – 1003, 2006. ISSN 1053-8119.

[65] William W. Lewis, Mustafa Sahin, Benoit Scherrer, Jurriaan M. Peters, Ralph O. Suarez, Vanessa K. Vogel-Farley, Shafali S. Jeste, Matthew C. Gregas, Sanjay P. Prabhu, Charles A. Nelson, and Simon K. Warfield. Impaired Language Pathways in Tuberous Sclerosis Complex Patients with Autism Spectrum Disorders. *Cerebral Cortex*, 23(7):1526–1532, 2013.

[66] Hai Li, Zhong Xue, Timothy M. Ellmore, Richard E. Frye, and Stephen T.C. Wong. Network-based analysis reveals stronger local diffusion-based connectivity and different correlations with oral language skills in brains of children with high functioning autism spectrum disorders. *Human Brain Mapping*, 35 (2):396–413, 2014. ISSN 1097-0193.

[67] Xiaojin Li, Xintao Hu, Changfeng Jin, Junwei Han, Tianming Liu, Lei Guo, Wei Hao, and Lingjiang Li. A Comparative Study of Theoretical Graph Models for Characterizing Structural Networks of Human Brain. *International Journal of Biomedical Imaging*, 2013:8, 2013.

[68] Yonghui Li, Yong Liu, Jun Li, Wen Qin, Kuncheng Li, Chunshui Yu, and Tianzi Jiang. Brain Anatomical Network and Intelligence. *PLoS Computational Biology*, 5(5), 2009.

[69] Sol Lim, Cheol E. Han, Peter J. Uhlhaas, and Marcus Kaiser. Preferential Detachment During Human Brain Development: Age- and Sex-Specific Structural Connectivity in Diffusion Tensor Imaging (DTI) Data. *Cerebral Cortex*, 2013.

[70] Catherine Lord, Michael Rutter, and Ann Couteur. Autism Diagnostic Interview-Revised: A revised version of a diagnostic interview for caregivers of individuals with possible pervasive developmental disorders. *Journal of Autism and Developmental Disorders*, 24(5):659–685, 1994. ISSN 0162-3257.

[71] Catherine Lord, Susan Risi, Linda Lambrecht, Jr. Cook, EdwinH., BennettL. Leventhal, PamelaC. DiLavore, Andrew Pickles, and Michael Rutter. The Autism Diagnostic Observation Schedule-Generic: A Standard Measure of Social and Communication Deficits Associated with the Spectrum of Autism. *Journal of Autism and Developmental Disorders*, 30(3):205–223, 2000. ISSN 0162-3257.

[72] Daniel Marcus, John Harwell, Timothy Olsen, Michael Hodge, Matthew Glasser, Fred Prior, Mark Jenkinson, Timothy Laumann, Sandra Curtiss, and David Van Essen. Informatics and Data Mining Tools and Strategies for the Human Connectome Project. *Frontiers in Neuroinformatics*, 5(4), 2011. ISSN 1662-5196.

[73] Djalel Eddine Meskaldji, Elda Fischi-Gomez, Alessandra Griffa, Patric Hagmann, Stephan Morgenthaler, and Jean-Philippe Thiran. Comparing connectomes across subjects and populations at different scales. *NeuroImage*, 80(0): 416 – 425, 2013. ISSN 1053-8119.

[74] Emil H. J. Nijhuis, Anne-Marie van Cappellen van Walsum, and David G. Norris. Topographic Hub Maps of the Human Structural Neocortical Network. *PLoS ONE*, 8(6):e65511, 06 2013.

[75] R.C. Oldfield. The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia*, 9(1):97 – 113, 1971. ISSN 0028-3932.

[76] Richard E. Passingham, Klaas E. Stephan, and Rolf Kotter. The anatomical basis of functional localization in the cortex. *Nat Rev Neurosci*, 3(8):606–616, 08 2002.

[77] William Penny, Karl Friston, John Ashburner, Stefan Kiebel, and Thomas Nichols. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Academic Press, 2006.

[78] S.C. Ponten, L. Douw, F. Bartolomei, J.C. Reijneveld, and C.J. Stam. Indications for network regularization during absence seizures: Weighted and unweighted graph theoretical analyses. *Experimental Neurology*, 217(1):197 – 204, 2009. ISSN 0014-4886.

[79] Luise Poustka, Christine Jennen-Steinmetz, Romy Henze, Kilian Vomstein, Johann Haffner, and Bram Stieltjes. Fronto-temporal disconnectivity and symptom severity in children with autism spectrum disorder. *The World Journal of Biological Psychiatry*, 13(4):269–280, 2012.

[80] J. C. Raven, J. H. Court, and J. Raven. *Coloured progressive matrices*. Oxford: Psychologist Press, 1995.

[81] Marco Reisert, Irina Mader, Constantin Anastasopoulos, Matthias Weigel, Susanne Schnell, and Valerij Kiselev. Global fiber reconstruction becomes practical. *NeuroImage*, 54(2):955 – 962, 2011. ISSN 1053-8119.

[82] Mikail Rubinov and Ed Bullmore. Fledgling pathoconnectomics of psychiatric disorders. *Trends in Cognitive Sciences*, 17(12):641 – 647, 2013. ISSN 1364-6613.

[83] Mikail Rubinov, Stuart A. Knock, Cornelis J. Stam, Sifis Micheloyannis, Anthony W.F. Harris, Leanne M. Williams, and Michael Breakspear. Small-world properties of nonlinear brain activity in schizophrenia. *Human Brain Mapping*, 30(2):403–416, 2009. ISSN 1097-0193.

[84] Raymond Salvador, John Suckling, Martin R. Coleman, John D. Pickard, David Menon, and Ed Bullmore. Neurophysiological Architecture of Functional Magnetic Resonance Images of Human Brain. *Cerebral Cortex*, 15(9): 1332–1342, 2005.

[85] F. Segonne, A. M. Dale, E. Busa, M. Glessner, D. Salat, H. K. Hahn, and B. Fischl. A hybrid approach to the skull stripping problem in MRI. *NeuroImage*, 22(3):1060 – 1075, 2004. ISSN 1053-8119.

[86] O. Sporns, G. Tononi, and G.M. Edelman. Theoretical Neuroanatomy: Relating Anatomical and Functional Connectivity in Graphs and Cortical Connection Matrices. *Cerebral Cortex*, 10(2):127–141, 2000.

[87] Olaf Sporns. The human connectome: Origins and challenges. *NeuroImage*, 80(0):53 – 61, 2013. ISSN 1053-8119.

[88] Olaf Sporns and Giulio Tononi. Classes of network connectivity and dynamics. *Complexity*, 7:2002, 2002.

[89] Olaf Sporns and JonathanD. Zwi. The small world of the cerebral cortex. *Neuroinformatics*, 2(2):145–162, 2004. ISSN 1539-2791.

[90] Olaf Sporns, Dante R. Chialvo, Marcus Kaiser, and Claus C. Hilgetag. Organization, development and function of complex brain networks. *Trends in Cognitive Sciences*, 8(9):418 – 425, 2004. ISSN 1364-6613.

[91] C. J. Stam, W. de Haan, A. Daffertshofer, B. F. Jones, I. Manshanden, A. M. van Cappellen van Walsum, T. Montez, J. P. A. Verbunt, J. C. de Munck, B. W. van Dijk, H. W. Berendse, and P. Scheltens. Graph theoretical analysis of magnetoencephalographic functional connectivity in Alzheimer's disease. *Brain*, 132(1):213–224, 2009.

[92] E. O. Stejskal and J. E. Tanner. Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient. *Journal of Chemical Physics*, 42:288–292, January 1965.

[93] Jerzy P. Szaflarski, Akila Rajagopal, Mekibib Altaye, Anna W. Byars, Lisa Jacola, Vincent J. Schmithorst, Mark B. Schapiro, Elena Plante, and Scott K. Holland. Left-handedness and language lateralization in children. *Brain Research*, 1433(0):85 – 97, 2012. ISSN 0006-8993.

[94] David S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004. ISSN 1522-2594.

[95] Olga Tymofiyeva, Christopher P. Hess, Etay Ziv, Patricia N. Lee, Hannah C. Glass, Donna M. Ferriero, A. James Barkovich, and Duan Xu. A DTI-Based Template-Free Cortical Connectome Study of Brain Maturation. *PLoS ONE*, 8(5):e63310, 05 2013.

[96] N. Tzourio-Mazoyer, B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot. Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain. *NeuroImage*, 15(1):273 – 289, 2002. ISSN 1053-8119.

[97] Martijn P. van den Heuvel, René C. W. Mandl, Cornelis J. Stam, René S. Kahn, and Hilleke E. Hulshoff Pol. Aberrant Frontal and Temporal Complex Network Structure in Schizophrenia: A Graph Theoretical Analysis. *The Journal of Neuroscience*, 30:15915–15926, 2010.

[98] M.P. van den Heuvel, C.J. Stam, M. Boersma, and H.E. Hulshoff Pol. Small-world and scale-free organization of voxel-based resting-state functional connectivity in the human brain. *NeuroImage*, 43(3):528 – 539, 2008. ISSN 1053-8119.

[99] Bernadette C. M. van Wijk, Cornelis J. Stam, and Andreas Daffertshofer. Comparing Brain Networks of Different Size and Connectivity Density Using Graph Theory. *PLoS ONE*, 5(10):e13701, 10 2010.

[100] Jinhui Wang, Liang Wang, Yufeng Zang, Hong Yang, Hehan Tang, Qiyong Gong, Zhang Chen, Chaozhe Zhu, and Yong He. Parcellation-dependent small-world brain functional networks: A resting-state fMRI study. *Human Brain Mapping*, 30(5):1511–1523, 2009. ISSN 1097-0193.

[101] Duncan J. Watts and Steven H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440 – 442, 06 1998.

[102] Chong-Yaw Wee, Pew-Thian Yap, Daoqiang Zhang, Kevin Denny, Jeffrey N. Browndyke, Guy G. Potter, Kathleen A. Welsh-Bohmer, Lihong Wang, and Dinggang Shen. Identification of MCI individuals using structural and functional connectivity networks. *NeuroImage*, 59(3):2045 – 2056, 2012. ISSN 1053-8119.

[103] Robert E. Wilson, Samuel D. Gosling, and Lindsay T. Graham. A Review of Facebook Research in the Social Sciences. *Perspectives on Psychological Science*, 7(3):203–220, 2012.

[104] Zheng Rong Yang. Biological applications of support vector machines. *Briefings in Bioinformatics*, 5(4):328–338, 2004.

[105] Andrew Zalesky, Alex Fornito, Ian H. Harding, Luca Cocchi, Murat Yücel, Christos Pantelis, and Edward T. Bullmore. Whole-brain anatomical networks: Does the choice of nodes matter? *NeuroImage*, 50(3):970 – 983, 2010. ISSN 1053-8119.