TECHNISCHE UNIVERSITÄT WIEN
Vienna University of Technology

Diplomarbeit

Identification and removal of signal components related to cardiac activity from high resolution resting-state fMRI data by means of blind source separation

Ausgeführt am

Zentrum für Medizinische Physik und Biomedizinische Technik
der Medizinischen Universität Wien
(Division: MR Physics, Research Focus: High Performance Computing)

unter der Anleitung von

Univ.Prof. Dipl.-Ing. Dr.techn. Gerald Badurek
Ao.Univ.Prof. Dipl.-Ing. Dr. techn. Ewald Moser
Dipl.-Ing. Roland Boubela, PhD

durch

Raphael A. Gunacker
Fillgradergasse 17/5
1060 Wien

Datum

Unterschrift

Abstract

BOLD-fMRI data analysis is challenging - in part due to the strong influences by non-white noise, some of which has physiological origins (e.g. cardiac activity or respiration). Noise becomes even more of a problem in resting state fMRI research, since BOLD fluctuations account for a smaller portion of signal than the BOLD response in specific task-related fMRI signals does.
Researchers have not always been aware of the great extent of this problem. Physiological noise is aliased into scans of a low temporal resolution (i.e. if the Nyquist criterion is not satisfied). But with the advent of multiband sequences, which facilitate higher sampling rates, it became possible to circumvent aliasing of high frequency noise into the signal of interest, if only the sampling frequency was set high enough. Since resting state networks have commonly been assumed to have a bandwidth lying in the range of 0.01-0.1 Hz, it seemed, as if scans of high temporal resolution could be freed from physiological noise by temporal filtering.
Recent studies have, however, revealed that resting state networks could also be identified in frequency ranges >0.1 Hz. Hence, the application of temporal filters turned out to be disadvantageous, since it eliminates useful information just as well as noise.
A novel approach to noise removal (based on blind source separation) is presented in this thesis. BOLD-fMRI data of high temporal resolution is decomposed into temporally independent components (time courses) and corresponding weights. The time courses are analysed in the frequency domain and components deemed to be related to physiological noise were excluded from subsequent data reconstruction.
fMRI data have been processed with this new algorithm and have (in another processing pipeline) been band passed (as a reference). The band passed data and the data processed by means of the proposed algorithm have both been subjected to seed-based correlation analysis in order to determine whether significant differences in functional connectivity could be observed. Results indicate that the different approaches did in fact lead to such significant differences. With the exception of the functional connectivity assessed for the auditory resting state network, the differences did however not appear to occur in specific patterns.
Other exploratory approaches (e.g. cluster analysis) could be used in future analyses in order to further validate this new noise removal approach and to compare results to those of similar approaches (e.g. CORSICA, PESTICA).

Zusammenfassung

Die BOLD-fMRI-Datenanalyse ist eine herausfordernde Aufgabe - mitunter bedingt durch den starken Einfluss von Störgrößen. Diese Störgrößen sind teilweise physiologischen Ursprungs (darunter fallen beispielsweise der Herzrhythmus oder aber auch die Atmung). In der sogenannten "resting state" fMRI-Forschung wird farbiges Rauschen zu einem noch größerem Problem, da BOLD-Fluktuationen dort einen kleineren Signalanteil ausmachen, als die sogenannte "BOLD response" in der "task-related" fMRI-Forschung.
Das Ausmaß dieses Problems ist in der Forschung noch nicht lange bekannt. Physiologisch bedingtes Rauschen wird durch Aliasing-Effekte in das Signal zeitlich niedrigauflösender Scans gespiegelt (wenn das Nyquist-Shannon-Abtasttheorem ignoriert wird). Mit dem Aufkommen von schnellen Multiband-Sequenzen ist es jedoch möglich geworden, solche Aliasing-Effekte zu vermeiden, sofern die Abtastfrequenz hoch genug eingestellt wird. Und nachdem lange Zeit angenommen wurde, dass "resting state" Netzwerke eine Bandbreite im Bereich von 0.01-0.1 Hz haben, schien es, als könnten zeitlich hochauflösende Scans durch die Anwendung einfacher Zeitfilter von physiologischem Rauschen befreit werden.
Aktuelle Studien haben jedoch gezeigt, dass "resting state" Netzwerke auch in Frequenzbereichen >0.1 Hz identifiziert werden können. Die Anwendung von Zeitfiltern erwies sich daher als nachteilig, da diese relevante Signalanteile ebenso eliminiert, wie Rauschen.
Ein neuer (auf einer sogenannten "Blind Source Separation"-Methode basierender) Ansatz zur Rauschentfernung wird im Rahmen dieser Arbeit vorgestellt. Zeitlich hochauflösende BOLD-fMRI-Daten werden in zeitlich unabhängige Komponenten zerlegt. Diese Zeitreihen werden im Frequenzbereich analysiert und Komponenten, die physiologischem Rauschen zugeordnet werden können, werden in die Datenrekonstruktion nicht miteinbezogen.
Dieser neue Algorithmus wurde auf reale fMRI Daten angewandt (und parallel dazu wurden die selben Daten mit einem Bandpass-Filter gefiltert). Die bandgepassten Daten und die mittels dem neuen Algorithmus analysierten Daten wurden einer sogenannten "seed-based correlation"-Analyse unterzogen um festzustellen, ob bezüglich der funktionellen Konnektivität signifikante Unterschiede festgestellt werden können. Die Ergebnisse haben gezeigt, dass dies sehr wohl der Fall ist. Mit Ausnahme der funktionellen Konnektivität, die speziell im Bezug auf das auditive resting state Netzwerk bestimmt wurde, scheinen diese signifikanten Unterschiede jedoch nicht in spezifischen Mustern aufzutreten.
Andere explorative Ansätze (z.B. eine Clusteranalyse) könnten Teil zukünftiger Analysen zur Validierung dieses neuen Ansatzes sein und zum Vergleich des Selbigen mit ähnlichen Ansätzen dienen.

Contents

1 Introduction

Magnetic resonance imaging (MRI) is an imaging modality used in neuroimaging, which – roughly stated – exploits the fact that the same nuclei (e.g. ^{1}H nuclei) have different magnetic properties in different molecules / tissues. Next to anatomical MRI scans, which fall into the category of structural imaging, there exists also functional magnetic resonance imaging (fMRI), which focuses on measuring brain activity based on changes in the blood oxygen level dependent (BOLD) contrast. This contrast mechanism is based on increased local oxygen consumption and therefore venous blood deoxygenation due to increased metabolic activity during a specific neural activity.
fMRI methods can be grouped into task-related- and resting state fMRI (rs-fMRI) methods (i.e. studies in which subjects have to perform specific tasks in opposition to studies in which subjects are resting). rs-fMRI research mainly focuses on functional connectivity of brain areas and has resulted in the discovery of multiple so called resting state networks such as the default mode network (see Raichle et al. [1]) and many others (see Damoiseaux et al. [2] and Robinson et al. [3]).
Besides model-driven fMRI data analysis approaches like the so called seed-based correlation analysis (which is, however, prone to biasing due to seed selection), exploratory approaches such as independent components analysis (ICA) are popular. Other exploratory approaches also fall within the category of blind source separation (BSS) methods – amongst which there is one approach referred to as second order blind identification (SOBI - [4]).
Analysis of fMRI data is, however, not trivial: on one hand, this is due to the huge amount of data obtained by even single scanning sessions (not to mention studies with large populations) as well as on the other hand the noisy and mixed nature of the BOLD signal.

This thesis focuses on a new approach based on blind source separation for the removal of physiological noise from resting state fMRI data with a high temporal resolution. The presence of non-white physiological noise in fMRI data and its characteristics have been repeatedly subjected to research (e.g. Liu et al. [5]). Damoiseaux et al. state that "[...] these resting fluctuations of interest are typically associated with strong power in the range of 0.01–0.1 Hz." [2] Due to this wide-spread assumption, physiological noise (e.g. noise due to cardiac activity and noise due to respiration), which typically lies in a higher frequency range, is commonly removed from the BOLD-fMRI signals with a high temporal resolution by simple band pass filtering. (Note, however, that this is not efficient in BOLD-fMRI data with a low temporal resolution, due to the fact that the noise will be aliased into the lower frequencies.) Boubela and Kalcher et al. have shown that

resting state networks are in fact *not* bound solely to a low frequency range but are much rather also identifiable in a frequency range >0.1 Hz. [6], [7] Since band pass filtering rigorously eliminates all portions of a signal that are recurring with a specified frequency, signal components, which are due to phenomena of interest (neural activity) are removed just as well as signal components related to undesired phenomena such as physiological noise. In short, band pass filtering is not an ideal means of physiological noise removal.

The new approach to physiological noise removal presented in this thesis can basically be outlined as follows: After preprocessing the data, blind source separation is performed in such a way that the results are temporally independent time courses (components) with corresponding weights for each voxel. Those time courses are then analyzed in the frequency domain. Components fulfilling certain criteria are identified as being potentially related to cardiac activity (as an example for physiological noise - the same could be done with different criteria for noise related to respiration) and are excluded from the subsequent data reconstruction. Thereby, noise can be eliminated from the data while retaining relevant signal components which lie in the same frequency range as said noise.

This novel approach is presented within this thesis as follows: Section 2 focuses on fundamentals related to NMR and MRI, as well as on the mathematical basics required to fully understand the underlying principles of the new approach. Section 3 gives specific details of how the new algorithm has been implemented and outlines the methods used to analyze whether data processed by this algorithm delivers more promising results in postprocessing analyses. Section 4 lists the actual results obtained by the application of said methods. Section 5 provides a brief review of the work and discusses the presented results. Future prospects are mentioned.

2 Fundamentals

This section aims at clarifying basic physical principles of nuclear magnetic resonance (NMR) and magnetic resonance imaging (MRI) as well as basic mathematical and statistical concepts required to understand the analytic methods which are applied in subsequent sections of the thesis. However, the degree of detail is obviously limited, so for more detailed insights into those topics consider consulting literature, some of which is referenced within this thesis.

2.1 Magnetic Resonance Imaging

2.1.1 Physical Principles of Nuclear Magnetic Resonance

2.1.1.1 Larmor precession

Nuclear Magnetic Resonance is a phenomenon referring to the interaction of the magnetic properties of atomic nuclei with three different electromagnetic fields. Nuclei have a non-zero *angular momentum* or *nuclear spin* $\vec{I}$ and a non-zero *magnetic moment* $\vec{\mu}$ if either their atomic number, their mass number or both are odd numbers. The relation between magnetic moment $\vec{\mu}$ and nuclear spin $\vec{I}$ is given by

$$\vec{\mu} = \gamma \vec{I}, \tag{2.1}$$

whereas the proportionality factor γ is called *gyromagnetic ratio* (note that $\vec{\mu}$ and $\vec{I}$ are vector quantities with identical directions, i.e. γ is a scalar quantity).
In the absence of an external magnetic field, nuclear spins and magnetic moments will have an indiscriminate orientation (due to Brownian motion), whereas in the vicinity of an external static magnetic field (often referred to as $\vec{B}_0$-field), they will tend to align themselves to the field in a quasi-parallel or quasi-antiparallel manner.
Note, however, that nuclear spins and magnetic moments can never fully align to the direction of the external field – they rather precess around the direction of the effective field. This is called *Larmor precession*. The angle between the direction of the $\vec{B}_0$-field and the direction of the magnetic moments / spins does *not* decrease when increasing the strength of the magnetic field (as one might be tempted to assume when having bar magnets in mind). Increasing the magnetic field strength does however increase the precession frequency (which is referred to as *Larmor frequency*). Also, the Larmor frequency is independent of said angle – it is given by

$$\omega = \gamma * B_0. \tag{2.2}$$

2.1.1.2 Zeeman Levels, Net Magnetization

According to the principle of induction the precessing spins generate a dynamic magnetic field which can be picked up by a coil generating a voltage. However, a single nucleus cannot be used to produce a detectable MRI signal: large amounts of nuclei are required for the generation of such a signal – thus, dealing with the whole topic on a more macroscopic level (i.e. considering *spin systems* and their *net magnetization* rather than individual nuclei and their spins and magnetic moments) is of importance. Consequently, classical physics models may suffice. The hydrogen nucleus (proton) has two possible quantum states (energy levels): Spin up (low energy state) and spin down (high energy state). Such energy levels are termed *Zeeman levels* and their energy difference is given by

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

with $\hbar$ being the reduced Planck constant ($\hbar = \frac{h}{2\pi}$).

As mentioned above, a single proton will not suffice to produce an MR image – instead, spin systems (consisting of large amounts of protons populating the energy states) are of interest. At thermal equilibrium, the ratio of populations of the two states is given by the so called *Boltzmann factor*:

$$\frac{n_{-\frac{1}{2}}}{n_{\frac{1}{2}}} = exp(-\frac{\gamma B_0 \hbar}{k_B T}), \tag{2.4}$$

with k_B denoting the *Boltzmann constant* and T denoting the absolute temperature (Kelvin) .
Considering that $\hbar$, k_B and γ are constants (γ being nucleus specific), it becomes obvious that the difference in populations between the energy states (and thus the net magnetization, which is the sum of the magnetic moments) can be influenced by changes in temperature and / or the strength of the magnetic field:

$$M \propto B_0 \text{ and } M \propto \frac{1}{T}, \tag{2.5}$$

with M denoting the net magnetization.

2.1.1.3 Nuclear Magnetic Resonance - Obtaining a Signal

The net magnetization vector precessing around the $\vec{B}_0$ field needs to be "tilted" out of equilibrium into the transverse plane in order to induce a voltage which can be picked up by a coil. The signal is emitted as soon as the magnetization vector returns to thermal equilibrium - quantum mechanically this corresponds to the excitation of spins from the low energy state to the high energy state and

then "falling" back to the low energy state (also termed *spin relaxation*) and thus inducing a voltage in a RF-coil tuned to the Larmor frequency.

Tilting the magnetization vector out of equilibrium (i.e. "exciting" the spins) is done with the help of a dynamic magnetic RF-field perpendicular to the $\vec{B}_0$-field. This perpendicular field (commonly referred to as $\vec{B}_1$-field) needs to be oscillating at the Larmor frequency of the nuclei which are to be excited (e.g. for ^{1}H in a 1 Tesla $\vec{B}_0$-field the Larmor frequency would be 42.576 MHz). Note that the nuclear spins actually precess around the sum $\vec{B}_0 + \vec{B}_1 = \vec{B}_{eff}$ rather than just around $\vec{B}_0$.

2.1.1.4 Spin Relaxation

As mentioned above, spin relaxation refers to the process of spins returning from an excited state to thermal equilibrium state. There are two basic relaxation mechanisms referred to as T1 relaxation (also *spin-lattice relaxation* or *longitudinal relaxation*) and T2 relaxation (also *spin-spin relaxation* or *transversal relaxation*). T1 relaxation usually takes longer than (or in some cases as long as) T2 relaxation.

T1 relaxation is a term used to describe the return of the net magnetization vector's z-component (longitudinal component) to it's equilibrium value (given that z is the direction of the $\vec{B}_0$-field). This mechanism is based on the interactions between a nucleus and its surroundings ("lattice"), i.e. energy dissipation. The associated time constant is called T1 and refers to the time it takes the net magnetization vector to return to $1 - \frac{1}{e}$ (i.e. approximately 63%) of its initial value. This corresponds to the so-called *Bloch equation*

$$M_z(t) = M_z(0) * (1 - e^{-\frac{t}{T1}}), \tag{2.6}$$

M_z being the net magnetization vector's longitudinal component.

T2 relaxation is a term used to describe the decay of the net magnetization vector's transversal component to it's equilibrium value due to spin-spin interactions. More specifically, after flipping the spins into the transversal plane with the help of a 90° pulse, the magnetic moments / spins precess around the $\vec{B}_{eff}$-field in phase at first, but will then dephase. This out-of-phase precession obviously results in a decrease of the net magnetization vector's transversal component and represents an increase in entropy. The associated time constant is called T2 and refers to the time it takes for the net magnetization vector's transversal component to decay to $\frac{1}{e}$ (i.e. approximately 37%) of its initial value. This corresponds to the Bloch equation

$$M_{xy}(t) = M_{xy}(0) * e^{-\frac{t}{T2}}, \tag{2.7}$$

M_{xy} being the net magnetization vector's transversal component.

Next to proton density (PD), the time constants T1 and T2 play a crucial role when it comes to the contrast of MRI scans. Further details on this topic will be treated in section 2.1.1.6 "Image Contrast".

2.1.1.5 Nuclear Magnetic Resonance Pulses and Pulse Sequences

The so called flip angle $\theta(t)$ (i.e. the angle formed by the net magnetization vector and the direction of $\vec{B}_{eff}$, when said net magnetization vector is tilted out of the equilibrium state) is influenced by the field strength of the $\vec{B}_1$-pulse and its on-time:

$$\theta(t) \propto B_1 * t, \tag{2.8}$$

Of special interest are 90° pulses, by which the magnetization vector can be tilted into the transversal plane from its initial position and 180° pulses, which are just as common in many pulse sequences – e.g. they are used as refocusing pulse in spin echo sequences. 180° pulses result from doubling either the amplitude or the on-time of a 90° pulse.

Apart from RF-pulses, gradient pulses are used in magnetic resonance imaging. Those temporarily active gradient fields encode spins in different locations precessing at different Larmor frequencies. Hence, gradient pulses are used for *spatial encoding*. Furthermore, in gradient echo sequences, they are not only used for spatial encoding but for the generation of echoes (i.e. signal generation) as well.

Specified trains of multiple RF-pulses and gradient pulses are called MRI sequences. Sequences can be differentiated according to their type, i.e. whether they are *spin echo (SE)* sequences or *gradient echo (GE)* sequences (those types create different image contrast). Sequences can however also be categorized according to the image types that they produce – e.g. T1 weighted images (T1WI), T2 weighted images (T2WI), proton density weighted images (PDWI). If multiple sequences are required during an examination, the succession of those sequences is called a *protocol*.

MRI sequences are determined by a multitude of parameters – e.g. the repetition time, echo time, inversion time (in the case of inversion recovery sequences),

flip angle and many more, all of which are affecting both image contrast and spatio-temporal resolution and may be easily changed on any MR scanner.

The duration between two successive pulse sequences applied to a slice is referred to as *repetition time (TR)*. The duration between the application of the initial 90° pulse and the maximum of the first signal echo is referred to as *echo time (TE)*. TR and TE are important parameters when it comes to image contrast (see section 2.1.1.6).

Echo Planar Imaging - commonly abbreviated as EPI - refers to a sequence type where the *readout gradient* (also known as *frequency encoding gradient*) is rapidly switched resulting in the generation of multiple echoes. Those facilitate the acquisition of more than just one (horizontal) line in *k-space*[1] at a time by adding phase-blips in phase encoding direction. Instead, greater sections of k-space are traversed (i.e. larger sections of the image / slice are acquired) within a single TR period. Capturing one slice (i.e. traversing the entire k-space) in one TR is referred to as *single shot EPI*, while capturing one slice over the period of multiple TRs is referred to as *multi shot EPI*.
The main advantage of EPI sequences are short acquisition times (hence the applications in e.g. cardiac imaging). Their main disadvantage is that the image quality is prone to suffer from susceptibility artifacts leading to blurring and geometrical distortion, particularly in the (slower) phase encoding direction.
EPI sequences can be implemented as spin echo- or gradient echo sequences.
Amongst common multi-slice imaging techniques are the so called multibanded EPI pulse sequences "[...] allowing simultaneous acquisition of multiple brain slices during a single EPI echo train [...]". [6]

2.1.1.6 Image Contrast, Weighting and MRI Applications

In MRI, image contrast depends basically on T1, T2 and proton density (PD) (for details on T1 and T2: see section 2.1.1.4). McRobbie et al. explain the influence of those parameters on image contrast as follows: "In PD images, high PDs give high signal intensities which in turn have bright pixels on the image. In T2-weighted images, tissues with long T2 give the highest signal intensities, producing a bright appearance. T1-weighted images are completely different; long T1 tissues give the weakest signal, i.e. bright pixels on T1 are associated with short T1s." [9]
McRobbie et al. also state that fluids (e.g. the cerebrospinal fluid, synovial fluid or edema) have the highest T1 and T2 time constants, followed by what the authors

[1] "The k-space represents the spatial frequency information in two or three dimensions of an object. The k-space is defined by the space covered by the phase and frequency encoding data." [8]

	short TE	long TE
short TR or large α	T1WI	poor contrast
long TR or small α	PDWI	T2WI

Table 1: Image contrast as a consequence of TE/TR or TE/α choice

refer to as "water-based tissues" (e.g. muscle, brain or cartilage), which have lower T1 and T2 time constants and finally, fat-based tissues (e.g. fat or bone marrow) have the lowest T1 and T2 time constants. For this reason, T1 weighted images are commonly used to represent anatomical structures in general (since boundaries between tissues are clearly visible). A more specific use would be the representation of fatty structures. T2 weighted images on the other hand may be used to pathologies such as edema or cysts.

Signal weighting is done by appropriate choice of TR and TE (in the case of spin echo sequences) and the flip angle and TE (in the case of gradient echo sequences), respectively (for details on TR and TE: see section 2.1.1.5).
Table 1 illustrates how different image contrasts depend on the choice of TE/TR (in the case of spin echo sequences) and TE/α (in the case of gradient echo sequences). For specific values concerning appropriate choices of TR, TE and the flip angle, consider reading [9].

2.1.2 Functional Magnetic Resonance Imaging

Functional magnetic resonance imaging (fMRI) uses the same hardware as standard clinical imaging (e.g. Moser et al. [10]), but refers to magnetic resonance imaging used to visualize brain *functions* (i.e. physiologic and pathologic processes) rather than visualizing brain *anatomy* (as is done in structural scans).

The term 'fMRI' refers to a variety of different methods used to obtain different kinds of image contrast (e.g. diffusion MRI, perfusion MRI - either with the help of contrast agents or by means of arterial spin labelling (ASL) - or so called BOLD-fMRI). Unless otherwise stated, within the scope of this thesis, the term fMRI will refer to BOLD-fMRI only.

2.1.2.1 BOLD Contrast

The BOLD contrast was first discovered by Ogawa et al. and reported in the year 1990. [11] 'BOLD contrast' is the abbreviation for 'blood oxygen level dependent contrast' and as the name already suggests, "[...] the BOLD signal is a measure of the ratio of oxygenated to deoxygenated hemoglobin" [12] - or, as Eloyan et al. put it - "Deoxyhemoglobin serves as an endogenous susceptibility contrast agent allowing MRI to report on local hemodynamic changes." [13] These hemodynamic changes are assumed to be related to neural activity via local oxygen consumption and blood flow. Ashby states: "The theory, which is not yet fully worked out, is that active brain areas consume more oxygen than inactive areas." [12]
In task-related fMRI, the so called BOLD *response* (which is a delayed reaction to the neural activity related to the specific task) is of interest, while in resting-state fMRI (rs-fMRI), one speaks of BOLD *fluctuations* (which are independent of any task).

One major problem with fMRI is, however, the noisy nature of the BOLD signal. As Davey et. al have stated: "For BOLD signals obtained at 3T, the signal of interest comprises less than 10% of the signal fluctuation, and is instead dominated by physiological noise (Kruger et al., 2001) and scanner drift (Bianciardi et al., 2009)." [14] Other sources (e.g. Damoiseaux et al. [2]) state that BOLD fluctuations in resting state networks lie within a range of 2-3% of the total signal.

The influence of the heart rate on the BOLD signal has been investigated by Chang et al. [15] They have found that modelling voxel time series by taking into consideration not only respiratory influences, but influences due to cardiac activity as well, will account for a higher percentage of total signal variance. Influences of the heart rate on the BOLD signal were modelled using a transfer function referred to as cardiac response function (CRF). Chang et al. state that the influence of cardiac activity on the BOLD signal affects primarily regions containing major vessels and cerebrospinal fluid (CSF).

2.2 Mathematical and Statistical Fundamentals of fMRI data analysis

So far, (mostly physics-related) fundamentals concerning (functional) magnetic resonance imaging have been addressed. Once image acquisition is completed, data processing begins. Since this thesis describes a novel approach to the removal of physiological noise from resting state fMRI data based on blind source separation (BSS), this subchapter will mainly focus on BSS methods - namely independent component analysis (ICA) and second order blind identification (SOBI).

Furthermore, the preliminary step of principal component analysis (PCA) and the underlying mathematical and statistical principles of all these methods will be discussed.

2.2.1 Statistical Relationships

In order to thoroughly understand how PCA and BSS algorithms work, one has to understand some basic principles of statistical relationships.
Correlation, covariance and dependence are measures of statistical relationships. While there are a variety of different correlation measures, the term is mostly used to refer to the so called Pearson correlation (this will also be the case for the remainder of this thesis).
Correlation and covariance are measures referring to *linear* relationships, which is exactly why uncorrelated variables are not necessarily independent, but can just as well be dependent (their dependence can be based on a non-linear relationship). While the covariance is dependent on the units of measurement of the respective variables, the correlation coefficient is not in any way affected by said units. This makes comparability of covariance values referring to data sets with different scales difficult. This is however not a problem with our fMRI data (all having the same units). Hence, the covariance may be used as a measure of linear relationship when performing fMRI data analysis.

2.2.1.1 Independence

BSS methods aim at separating mixtures of signals into their original source signals without (or with very little) a priori knowledge about the source signals and or the mixing process. ICA in particular works - as its name already suggests - based on the assumption that the source signals are *statistically independent*.

Statistical independence of two random variables X and Y is given if and only if their joint probability density function (pdf) is equal to the product of their marginal probability density functions:

$$p_{X,Y}(x, y) = p_X(x)p_Y(y) \ \forall \ x, y. \tag{2.9}$$

2.2.1.1.1 Estimating Independence

Estimating statistical independence is essential in the course of ICA since an exact solution of the problem could only be computed if the probability density functions of the sources were known (which in practice usually isn't the case).
There are basically three popular concepts of estimating statistical independence.

Their underlying principles will be briefly introduced in this subchapter, but for more detailed descriptions, consider reading [16]. Those concepts are somewhat related to each other – Hyvärinen even states: "In fact, almost all of these estimation principles can be considered as different versions of the same general criterion." [16]

- Minimization of mutual information
  This approach is based on information theory, in which one bit of information (or *entropy*) is defined as being a measure of uncertainty / unpredictability of an event.
  Joint entropy is the amount of information contained in a system of more than one variable.
  Mutual information is the information content that individual variables from a set of variables have on each other.
  These measures can easily be illustrated using a Venn diagram (as is common in set theory) – see figure 1. Members of a set of random variables are said to be statistically independent if and only if their mutual information equals 0.

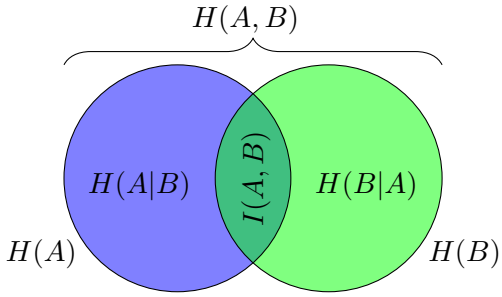

Figure 1: Venn diagram illustrating the relationship between different measures of information theory ($H(A)$ and $H(B)$ being the individual entropies of the variables A and B, respectively, $H(A,B)$ being the joint entropy, $I(A,B)$ being the mutual information)

- Maximization of non-gaussianity
  The central limit theorem basically states that (given certain conditions) the sum of non-gaussian random variables will be more gaussian than each individual constituent. Since ICA assumes that the individual components are non-gaussian (which will be explained in section 2.2.3.2.2), a linear combination of the original components will be maximally non-gaussian, if it represents one of the components. Hence, by maximizing non-gaussianity of the mixed signals matrix, one can estimate a model for the independent components (and thus also for the mixing matrix).
  When maximizing non-gaussianity, one could e.g. be interested in super-

gaussian distributions. Supergaussian distributions have a higher peak (at zero mean as well) and higher tails than gaussian distributions. They are referred to as *leptokurtic*, since their kurtosis is >3 (a gaussian distribution has a kurtosis of 3).
But since the kurtosis is a fourth-order statistic, using it as a measure for non-gaussianity is not ideal: Its sensitivity to outliers is greater than that of other measures.
Amongst all variables of identical mean and variance, a gaussian variable is the most unpredictable one – or, as one could also say – the gaussian variable has the highest entropy. Hence, entropy can be used as a measure for non-gaussianity. *Negentropy* has been introduced as a measure, which is zero for gaussian variables and otherwise always positive. Negentropy is mathematically defined as

$$J(x) = H(x_{gauss}) - H(x), \tag{2.10}$$

i.e. as the difference in entropy between the distribution of any given variable and the gaussian distribution with the same mean and the same variance as the variable of interest. The *International Organization of Standardization* uses a possibly more intuitive language, when stating that negentropy can be defined as "mean value of the information content of the events in a finite set of mutually exclusive and jointly exhaustive events [...]" [17]

- Maximum-likelihood approaches
  Maximum-likelihood estimation (MLE) aims at estimating the parameters of a model so that with those parameters, the observed data is most probable to be produced by the model. To be more specific, the parameters are chosen in such a way that they maximize a so called *likelihood function*.
  For computational reasons it is advantageous to calculate the logarithm of the likelihood function instead. This does not change the results in any way, since the *log*-function is strictly increasing and thus, its maximum coincides with the maximum of the likelihood function.

Those concepts of estimating independence are used in a variety of different ICA algorithms. Without going to much into details yet, ICA decomposes a matrix of observed (mixed) signals into a mixing matrix and a components matrix containing independent source signals (see equation 2.27 and section 2.2.3.2 for a thorough presentation of that topic). When it comes to minimizing mutual information, an algorithm would adapt the entries of the unmixing matrix until the mutual information of the rows in the components matrix becomes as small as possible. In fact, the fastICA algorithm (which will be presented in section 2.2.3.2.3) can be

regarded as an approach which is based on both: minimizing mutual information and maximizing negentropy. This is because mutual information can actually be computed from negentropy. Hyvärinen states: "[...] mutual information is minimized [...] when the sum of the negentropies of the components is maximized." [18] MLE-based ICA algorithms use the joint probability density function of the independent components as likelihood function. The elements of the mixing matrix then serve as parameters to maximize this likelihood function.

2.2.1.2 Whiteness and Uncorrelatedness

Uncorrelatedness and whiteness are weaker properties than independence in that variables can be uncorrelated or whitened and yet be dependent. Nevertheless, the concepts of uncorrelatedness and whiteness are important when it comes to blind source separation (for reasons which are discussed in section 2.2.1.2.3).

If the n components of a zero mean random vector X are uncorrelated, the covariance matrix of the random vector will be a diagonal matrix:

$$\begin{gathered} E[XX^T] = D \\ \text{with } d_{i,j} = 0 \text{ if } i \neq j \ \forall \ i, j \in \{1, 2, ..., n\} \end{gathered} \tag{2.11}$$

A zero mean random vector X' is said to be *whitened* if its components are decorrelated *and* its variances (i.e. the diagonal elements of the covariance matrix) all equal 1:

$$E[X'X'^T] = I, \tag{2.12}$$

I being the unity matrix.
Hence, as Hyvärinen simply states: "[...] whitening is essentially decorrelation followed by scaling [...]." [16]

Mathematically speaking, whitening is a linear transformation of a zero mean vector:

$$Y' = VY \text{ such that } E[Y'Y'^T] = I, \tag{2.13}$$

whereas V is referred to as *whitening matrix*

There are two common whitening transformations, which will be discussed in the following two paragraphs.

2.2.1.2.1 Whitening by means of Zero-Phase Component Analysis

Eigendecomposition of the covariance matrix of Y yields

$$E[YY^T] = Q\Lambda Q^T, \tag{2.14}$$

whereas Q is an orthogonal matrix containing the eigenvectors of $E[YY^T]$ and Λ is a diagonal matrix containing the eigenvalues of $E[YY^T]$ as diagonal elements. (Note that eigendecomposition will be explained in detail in section 2.2.2.2.)

Whitening by zero-phase composition analysis (ZCA) is then done by choosing

$$V = E[YY^T]^{-1/2} = Q\Lambda^{-1/2}Q^T \tag{2.15}$$

as a whitening matrix (whereas $\Lambda^{-1/2}$ can easily be determined, since it equals $diag(\lambda_1^{-1/2}, \lambda_2^{-1/2}, ..., \lambda_n^{-1/2})$).

The proof that this matrix whitens Y is simple:

$$\begin{aligned} E[Y'Y'^T] &= E[VY(VY)^T] = E[VYY^TV^T] = VE[YY^T]V^T \\ &= E[YY^T]^{-1/2}E[YY^T]E[YY^T]^{-1/2} = I. \end{aligned} \tag{2.16}$$

Note that since V is diagonal, $V = V^T$.

Geometrically, ZCA can be thought of as a transformation which rotates the data only minimally, thus preserving as much of its initial orientation as possible. This minimizes the Mahalanobis distance of the original data and the whitened data. The transformation is therefore also known as *Mahalanobis transformation.*

2.2.1.2.2 Whitening by means of Principal Component Analysis

Whitening by means of principal component analysis (which itself will be discussed in section 2.2.2) also starts off with the computation of the eigendecomposition of the covariance matrix.
The whitening matrix is then given by

$$V = \Lambda^{-1/2}Q^T. \tag{2.17}$$

The proof that using this matrix to orthogonally transform the original vector Y is as follows:

$$\begin{aligned} &E[Y'Y'^T] = E[VY(VY)^T] = E[VYY^TV^T] = VE[YY^T]V^T \\ &= \Lambda^{-1/2}Q^TE[YY^T]Q\Lambda^{-1/2} = \Lambda^{-1/2}Q^TQ\Lambda Q^TQ\Lambda^{-1/2} = I. \end{aligned} \tag{2.18}$$

Note that mere decorrelation (not whitening) can be obtained by

$$Y_d = Q^T Y, \tag{2.19}$$

Y_d denoting the vector containing the decorrelated elements of Y. This is essentially omitting the act of scaling which is mathematically represented by the multiplication with $\Lambda^{-1/2}$.

2.2.1.2.3 Whitening prior to Independent Component Analysis

One might wonder why so much attention has been paid to whiteness, although it is a weaker property than independence. The reason for this is that whitening data has two major advantages:

- One characteristic of BSS problems is that they are underdetermined. This leads to certain ambiguities, of which one is referred to as scaling ambiguity (see section 2.2.3.1). In this context, whitening normalizes the variances of the source signals, thus guaranteeing that their entire variance is adequately accounted for by the information contained in the mixing matrix.
- Whitening reduces the ICA problem to a simple orthogonal transformation. This drastically reduces the complexity of the problem, since it confines the solution space of the mixing matrix to orthogonal matrices (or unitary matrices, respectively, when referring to complex signal matrices) – i.e. only $n * (n-1)/2$ instead of n^2 parameters need to be estimated (n being the number of variables).

Considering that whitened data will be further rotated by ICA anyways, it is not necessary to pay attention to rotating the data as little as possible in the course of whitening (as is done in ZCA whitening). Hence, PCA whitening will work just as well in this context.

PCA of the data prior to ICA is also commonly used to reduce the dimensionality of the data. This is done as a means of noise reduction as well as it serves to reduce the computational effort during ICA. Further details are discussed in the next section.

2.2.2 Principal Component Analysis

Principal component analysis (PCA) is a multivariate statistical technique that decorrelates data. Hyvärinen et al. describe the purpose of PCA as follows: "Given a set of multivariate measurements, the purpose is to find a smaller set of

variables with less redundancy, that would give as good a representation as possible." [16] Mathematically speaking, PCA is an orthogonal transformation that projects the data into a new coordinate system. In this new coordinate system, the direction in which the data has the greatest variance is the direction of the first coordinate. The direction of second largest variance is then the direction of the second coordinate, which is also orthogonal to the first coordinate and so on. The result of PCA is low dimensional (depending on the number of chosen components) decorrelated data.
It is important to note that PCA is (in contrast to ICA) not a procedure that generates a statistical model: "[...] PCA is a straightforward orthogonal rotation of the data to a new coordinate system that requires no model fitting process. This is because PCA has no free parameters to estimate from the data. PCA simply allows the user to view his or her data from a different perspective." [12]

2.2.2.1 The PCA procedure

Figure 2 illustrates the steps comprising PCA. As shown, after having centered the data, PCA can be performed either by eigendecomposition or by singular value decomposition of the data. These two matrix factorizations will be explained in more detail in sections 2.2.2.2 and 2.2.2.3. Also, section 2.2.2.4 deals with how they are related to each other and finally, section 2.2.2.5 focuses on why it is often advantageous to perform PCA by means of SVD.
After having factorized the data matrix with either of the two decompositions, the dimensionality of the data may be reduced by using only a limited set of eigenvalue / eigenvector- or, respectively, singular value / singular vector pairs for further processing.
The data can now easily be projected into the new space (where they are decorrelated) and from there, back projected into the original space. Section 2.2.2.6 focuses on how to determine a useful degree of dimensionality reduction.

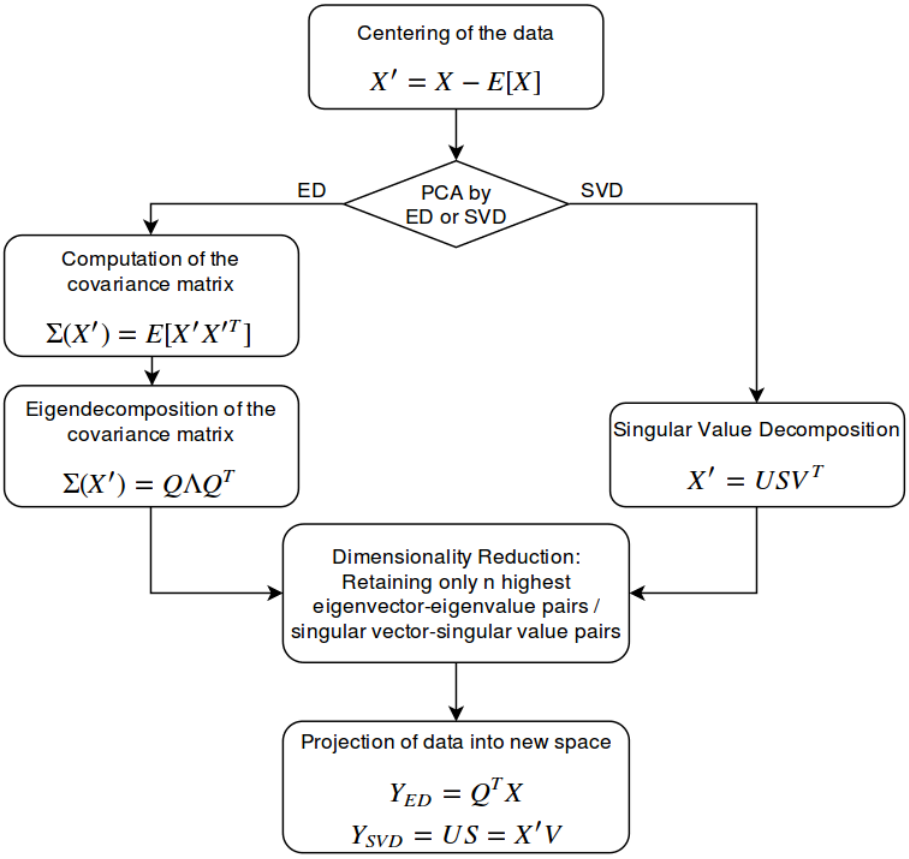

Figure 2: Basic steps of PCA.

2.2.2.2 Eigendecomposition

The eigendecomposition of a matrix is its factorization into three matrices containing its eigenvalues and eigenvectors. This representation is unique for each diagonalizeable matrix (except for the order of the eigenvalues and eigenvectors) and doesn't exist for non-diagonalizeable matrices.
Let A be a real or complex $n \times n$ matrix. The scalar $\lambda \in \mathbb{C}$ is called eigenvalue of A if and only if there exists a vector $\vec{v} \in \mathbb{C}^n, \vec{v} \neq \vec{0}$, which satisfies the linear equation

$$A\vec{v} = \lambda\vec{v}. \tag{2.20}$$

Every such vector is termed eigenvector of A with corresponding eigenvalue λ.
If the eigenvectors of a matrix A are linearly independent, A can be decomposed as

$$A = Q\Lambda Q^{-1}, \tag{2.21}$$

where the columns of Q are the eigenvectors of A and the non-zero entries of the diagonal matrix Λ are the corresponding eigenvalues.
Note, that the decomposition of a real valued symmetric matrix A' (covariance matrices have those properties by definition!) will yield an orthogonal eigenvector matrix. Hence, $A' = Q\Lambda Q^{-1} \Leftrightarrow A' = Q\Lambda Q^T$.

2.2.2.3 Singular Value Decomposition

The singular value decomposition of a matrix is its factorization into three matrices containing its singular values and its (left and right) singular vectors. This representation is unique for each (real or complex) matrix (except for the order of the singular values and singular vectors) and exists (in opposition to the eigendecomposition) for diagonalizeable as well as non-diagonalizeable matrices.
Let B be a real or complex $m \times n$ matrix. B can then be factorized as

$$B = USV^*, \tag{2.22}$$

whereas U is a real or complex $m \times m$ unitary matrix, V^* (the conjugate transpose of V) is a real or complex $n \times n$ unitary matrix and S is a $m \times n$ rectangular diagonal matrix with its diagonal elements being non-negative real numbers.[2] Those diagonal elements are the singular values of B; the columns of U are called the left singular values of B and the columns of V the right singular values of B.

2.2.2.4 Relation of ED and SVD

If and only if a matrix A is symmetric (or hermitan, if A is complex) and positive semi-definite, its singular value decomposition will be the same as its eigendecomposition.
Let's consider the matrices A^*A and AA^* and their decompositions:

$$\begin{aligned} A^*A &= (USV^*)^*USV^* = VS^*U^*USV^* = VS^*SV^* = VS^2V^*, \\ AA^* &= USV^*(USV^*)^* = USV^*VS^*U^* = USS^*U^* = US^2U^*. \end{aligned} \tag{2.23}$$

Since the matrices A^*A and AA^* are by definition symmetric and positive semi-definite, their singular value decompositions coincide with their eigendecompositions, whereas $\Lambda = S^2$ (i.e. the eigenvalues of A^*A and AA^* are the squared diagonal entries of S). Also, the columns of V and U are the eigenvectors of A^*A and AA^*, respectively.

[2] S has been used to denote the singular value matrix instead of using the equally common Σ in order to avoid confusion with covariance matrices, which are also commonly denoted by Σ.

2.2.2.5 Superiority of SVD over ED in PCA

This is of relevance, when it comes to principal component analysis. In PCA it is necessary to compute the eigenvalues and eigenvectors of the covariance matrix of the data (as explained in section 2.2.2.1). Since the covariance matrix of a data matrix X is given by

$$cov(X) = E[(X - E(X)(X - E(X))^T)] \tag{2.24}$$

and, thus, for a centered data matrix X', it is given by

$$cov(X') = E[X'X'^T] \tag{2.25}$$

and since covariance matrices are by definition always symmetric and positive semi-definite, it is obvious that instead of computing the eigendecomposition of $E[X'X'^T]$, one could just as well compute the singular value decomposition of X'.

This is highly advantageous considering the computational effort: E.g., if in a whole brain fMRI scan, 250000 voxels have been recorded (at t time points), the data matrix would have the dimensions $250000 \times t$ - resulting in the covariance matrix having the dimensions $(250000 \times t)(t \times 250000) = 250000 \times 250000$. Assuming that every voxel value is obtained at double precision (i.e. having 8 bytes), the covariance matrix would require $250000^2 * 8/1024^3 = 465$ Gigabytes of memory. Taking into consideration that for fMRI scans the temporal dimension is usually much lower than the spatial dimension, computing the singular value decomposition of the original data matrix in the course of PCA (instead of the eigendecomposition of the covariance matrix) makes much more sense regarding computational effort.

Apart from computational efforts, there is yet another reason to choose the singular value decomposition of a data matrix over the eigendecomposition of its covariance matrix. A measure used to determine how much a function's output changes with regard to its input (or, to put it differently: how sensitive the function is towards input errors) is the so called condition number κ. With respect to the spectral norm, the condition number is defined as being the ratio of the largest to the smallest singular value. And in the special case of normal matrices (i.e. $A^*A = AA^*$) it is given by the ratio of the largest to the smallest eigenvalue. Since the eigenvalues of a covariance matrix are (as mentioned above) equal to the squared singular values of the original data matrix, it is obvious that

$$\kappa(A^*A) = \kappa(AA^*) = \frac{\lambda_{max}}{\lambda_{min}} = (\frac{\sigma_{max}}{\sigma_{min}})^2 = \kappa(A)^2, \tag{2.26}$$

meaning that the PCA computed by means of singular value decomposition of the data matrix will lead to more precise results (i.e. a lower condition number) than when computing it by means of eigendecomposition of the covariance matrix (which would yield the higher condition number). And, as Narsky et al. state, "For this reason, PCA implementations often prefer SVD of X over EVD of $X^T X$." [19]

2.2.2.6 Dimensionality Reduction: Choosing the Number of PCs

After having decomposed the data matrix by means of ED or SVD, an important problem in PCA is how to choose the number of principal components. In his book "Principal Component Analysis" I.T. Jolliffe describes this problem as follows: "The major objective in many applications of PCA is to replace the p elements of x [x being a vector of p random variables] by a much smaller number m of PCs, which nevertheless discard very little information. It is crucial to know how small m can be taken without serious information loss." [20] There are a variety of different approaches, amongst which the most common are to base the decision on the cumulative percentage of total variation or to make the decision with the help of a so called scree plot. In the latter, the eigenvalues (which describe the variance in the data) are ordered according to their value and then plotted. Ideally, the scree plot would include a kink at some point, at which the shape of the curve changes. Assuming that "true" (desired) signal components are responsible for a higher variance of the data than noise components, it would make sense to keep only the principal components that are probably signal components rather than keeping all components – i.e. including those that are probably related to noise. However, as F. G. Ashby states, "[...] unfortunately, a kink is rarely obvious. [...] The key here is to be conservative. In most cases, it will be much better to leave some noise in the data than to throw away signal." [12]

2.2.3 Blind Source Separation Problems

The statistical concepts of independence, uncorrelatedness and whiteness have been thoroughly discussed and principal component analysis has been introduced as a crucial step prior to blind source separation. Based on all that, this section will now focus on BSS itself.

Blind source separation (BSS) problems deal with separating mixtures of signals into their original source signals without (or with very little) a priori knowledge about the source signals and / or the mixing process - or, as Ziehe puts it - "[r]estated in the matrix formulation, the BSS problem consists in factoring the observed signals data matrix X into the mixing matrix A and the source signals matrix S." [21]

This is represented by the equation

$$X = AS. \tag{2.27}$$

BSS problems are by nature underdetermined and thus, there is no unique solution to them. Even though some ambiguities will therefore always remain, useful results can be obtained from a multitude of approaches, if certain assumptions (depending on the BSS approach) about the nature of the source signals are made.

2.2.3.1 Ambiguities of BSS Problems

In BSS problems, there are basically two types of ambiguities:

- Scaling ambiguity
  Since BSS problems are described by equation 2.27, whereas X represents the observed data, A an unknown mixing matrix and S the latent (i.e. unobservable) source signals (or *components*), the variances of the components cannot be determined. Multiplying the source signals with a random scalar and simultaneously dividing the respective parameters (*weights*) of the mixing matrix by the same scalar leaves the mixed signals X unchanged.
  With respect to a commonly used normalization convention, Belouchrani et al. state: "Advantage can be taken of this ambiguity by assuming, *without any loss of generality*, that the source signals have unit variance so that the dynamic range of the sources is accounted for by the magnitude of the corresponding columns of A." [4] (This has already been briefly mentioned in section 2.2.1.2.3.)
  It is, however, important to "[n]ote that this still leaves the *ambiguity of the sign*." (Hyvärinen, [16])

- Permutation ambiguity
  The order of the computed components cannot be determined.
  Considering that the mixed signals are nothing else than the sums of the weighted source signals (as illustrated in equation 2.28) and considering that addition is always commutative, it becomes clear that the order of the individual terms can be arbitrarily changed.

$$x_i = a_{i,1}s_1 + a_{i,2}s_2 + ... + a_{i,n}s_n \tag{2.28}$$

It is important to note that whitening the observed data is *not* an indispensable prerequisite for BSS algorithms to yield useful results, but much rather a helpful preprocessing step. In the same way, centering of the data is not mandatory.

2.2.3.2 Independent Component Analysis

Independent component analysis is probably the most prominent BSS method. This section starts with a description of what ICA is meant to accomplish (based on the well known cocktail party problem). Furthermore, underlying assumptions are explained and finally, the fastICA algorithm is presented.

2.2.3.2.1 The Cocktail Party Problem

Most attempts to explain independent component analysis (ICA) start with the so called cocktail party problem. The term "cocktail party problem" was introduced in 1953 by Colin Cherry ([22]). In this context it was purely physiologically (rather than technically) referring to the ability of discriminating speech signals: "[H]ow do we recognize what one person is saying when others are speaking at the same time (the 'cocktail party problem')?"
With respect to ICA the cocktail party problem is mostly described in a more technical way: Several persons (signal sources) in a room talk at the same time, while their voices are being recorded by microphones (at least as many microphones are required as there are persons speaking simultaneously to make separation possible). The problem is then to separate the signal mixtures into the original source signals with little or no a priori knowledge about the source signals or the mixing process. The cocktail party problem is thus a classical BSS problem.
ICA can solve the cocktail party problem: This is due to the fact that speech signals are usually statistically independent - i.e. the fact that one person starts talking does not necessarily imply that others start talking at the same time. As its name already suggests, ICA attempts to separate the signal mixture into different components which are as independent as possible.
In addition to the independent components (source signals) the parameters of these components (weights), which are (in the case of the cocktail party problem) the amplitudes of the source signals at the individual microphones also need to be determined. As already mentioned, the cocktail party problem is a BSS problem and thus by nature an underdetermined problem - i.e. an infinite amount of solutions exist. Yet, apart from only a few ambiguities (see section 2.2.3.1), component estimates can be made, which come very close to the original signals.

2.2.3.2.2 Assumptions

There are two basic assumptions in ICA:

- Statistical independence of the source signals (components)
  As we recall from section 2.2.1.1, statistical independence of two random variables X and Y is given if and only if their joint probability density

function (pdf) is equal to the product of their marginal probability density functions:

$$p_{X,Y}(x,y) = p_X(x)p_Y(y) \text{ for all } x, y. \tag{2.29}$$

However, "[s]ince a closed form solution to the ICA problem would require exact determination of the pdfs, which are generally not available, the sources have to be estimated by approximating independence with an objective function." [23] Therefore, independence needs to be estimated. Means of doing this have been discussed in section 2.2.1.1.1.

- Gaussianity of at most one independent component
  A continuous random variable Z is said to have a *gaussian distribution* (or to be *normally distributed*) if its probability density function is given by

$$p(z) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(z-\mu)^2}{2\sigma^2}}, \tag{2.30}$$

whereas σ is the standard deviation of the distribution (i.e. σ^2 is its variance) and μ is its expected value (mean). A special case of the gaussian distribution is the *standard normal distribution*. It has $\mu = 0$ and $\sigma^2 = 1$ (note, that this is the case for whitened data!) and is thus given by

$$p(z') = \frac{1}{\sqrt{2\pi}} e^{-\frac{z'^2}{2}}. \tag{2.31}$$

Since statistical independence of a set of variables is given only if their joint probability density function (pdf) equals the product of their marginal pdfs, the joint pdf of two independent components (obtained from whitened data) is given by

$$p(s_1, s_2) = \frac{1}{\sqrt{2\pi}} e^{-\frac{s_1^2}{2}} * \frac{1}{\sqrt{2\pi}} e^{-\frac{s_2^2}{2}} = \frac{1}{2\pi} e^{-\frac{s_1^2+s_2^2}{2}} = \frac{1}{2\pi} e^{-\frac{||s||}{2}}. \tag{2.32}$$

The joint pdf of the mixed signals x_1 and x_2 can be computed from the joint pdf of the independent components as follows:

$$p(x_1, x_2) = \frac{p(s_1, s_2)}{|detA|}. \tag{2.33}$$

The division by $|detA|$ is a special case of the transformation theorem for nonsingular linear transformations (as is the case with A).

From $X = AS$ (equation 2.27), we obtain $S = A^{-1}X$. And since the data has been whitened, the mixing matrix A is orthogonal (see section 2.2.1.2.3), i.e. $A^{-1} = A^T$. Thus $S = A^T X$.
The joint pdf of the mixed signals therefore becomes

$$p(x_1, x_2) = \frac{1}{2\pi|detA|} e^{-\frac{||A^T x||^2}{2}}. \tag{2.34}$$

We will now consider two more properties of orthogonal matrices:

- $||Ax|| = ||x||$
  This is because isometry (preservation of lengths) and isogonality (preservation of angles) are characteristic properties of orthogonal matrices. Orthogonal matrices represent rotations, reflections or combinations of those.
- The determinant of an orthogonal matrix is always either $+1$ or -1.

Bearing those properties in mind, we obtain

$$p(x_1, x_2) = \frac{1}{2\pi} e^{-\frac{||x||^2}{2}} \tag{2.35}$$

for the pdf of the mixed signals x_1 and x_2.
From those considerations Hyvärinen concludes: "[W]e see that the orthogonal mixing matrix does not change the pdf, since it does not appear in this pdf at all. The original and mixed distributions are identical. Therefore, there is no way how we could infer the mixing matrix from the mixtures." [16]
If only one gaussian component is present in a mixture of otherwise non-gaussian components, "[...] we can estimate the model, because the single gaussian component does not have any other gaussian components that it could be mixed with." [16]

2.2.3.2.3 fastICA

fastICA is an ICA algorithm developed by Apo Hyvärinen. Different versions of the algorithm facilitate deflation-based and symmetric signal separation (i.e. sequential or simultaneous computation of the components). fastICA estimates independence by maximizing non-gaussianity (to be more specific: by maximizing the negentropy).

The basics steps of the fastICA algorithm are illustrated in figure 3. In this

example $X = AS$ with $A^T = W$, hence $S = WX$.

The approximation of the weight vectors is done using an approximation of negentropy:

$$w_n = E[xg(w^T x)] - E[g'(w^T x)]w, \tag{2.36}$$

whereas $g(.)$ is an objective function and $g'(.)$ is its derivative.

In [16] Hyvärinen states that good choices for the objective function are either

$$\begin{gathered} g_1(u) = \log\cosh(u) \text{ and therefore } {g_1}'(u) = \tanh(u) \text{ or} \\ g_2(u) = -e^{-\frac{u^2}{2}} \text{ and therefore } {g_2}'(u) = ue^{-\frac{u^2}{2}}. \end{gathered} \tag{2.37}$$

Apart from those two objective functions, one could also use

$$g_3(u) = \frac{1}{4}u^4 \text{ and therefore } {g_3}'(u) = u^3, \tag{2.38}$$

however, Hyvärinen states, that this last function is sensitive to outliers and therefore not an ideal choice.

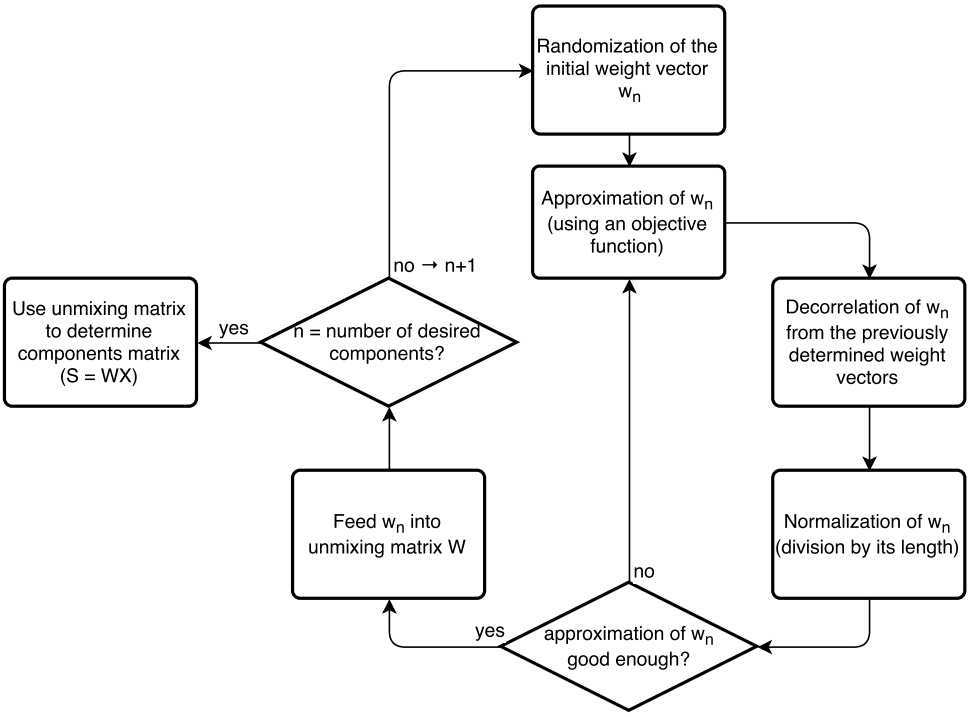

Figure 3: Basic steps of the fastICA algorithm.

2.2.3.3 Second Order Blind Identification

Apart from ICA, there is a diversity of BSS methods relying on joint diagonalization: Next to JADE (Joint Approximate Diagonalization of Eigen-matrices – [24]) and AMUSE (Algorithm for Multiple Unknown Signals Extraction – [25]) there is another approach called SOBI (second order blind identification), which will be the described in this section. Belouchrani et al. have introduced SOBI as an approach to solve BSS problems, which (as they state in said paper) "[...] relies only on stationary second-order statistics that are based on a joint diagonalization of a set of covariance matrices." [4] SOBI is thus superior to JADE and AMUSE in that it is a more robust algorithm (it is based on second order statistics only rather than on fourth order statistics, as is the case with JADE and it jointly diagonalizes a multitude of autocovariance matrices rather than just two, as is the case with AMUSE).
This section focuses on the underlying assumptions of SOBI and how they differ from those of ICA. Furthermore, approximate joint diagonalization is explained as a basis for understanding the actual algorithm, which is then given.

2.2.3.3.1 Assumptions

Compared to ICA, SOBI has significantly different prerequisites. In ICA, mutual independence of the components and gaussianity of at maximum one component is required (see section 2.2.3.2.2). SOBI on the other hand requires the components to be individually autocorrelated (to be more specific: weakly stationary) and mutually uncorrelated. It is worth mentioning again that uncorrelatedness is a much weaker criterion than independence. Also, multiple gaussian sources can be separated by SOBI.

2.2.3.3.2 Approximate Joint Diagonalization Problems

Eigendecomposition (as described in 2.2.2.2) is (at its core) a diagonalization problem - hence, when performing PCA by means of eigendecomposition of the covariance matrix of the data, one basically diagonalizes said covariance matrix.
We will now consider the autocovariance matrices of a data set (instead of its covariance matrix). Autocovariance simply means covariance of a signal with itself at different points in time.
Let x be zero mean mixed signals, s the source signals and A the mixing matrix, so that $x = As$ and let τ be a time lag. Then the autocovariance matrix of the mixed signals x with respect to the time lag τ will be

$$C_\tau(x) = E[x(t)x(t+\tau)^T]. \tag{2.39}$$

Substituting x with As on the right side of the equation leads to

$$C_\tau(x) \stackrel{\text{def}}{=} E[As(t)As(t+\tau)^T] = AE[s(t)s(t+\tau)^T]A^T = AC_\tau(s)A^T. \tag{2.40}$$

The matrix A is in fact already the solution to the diagonalization problem, since the source signals are assumed to be statistically independent (resulting in the off-diagonal elements of the covariance matrix $C_\tau(s)$ to be zero).

Note that, so far, we have only dealt with a simple diagonalization problem. We will now focus on *joint* diagonalization problems.
Joint diagonalization aims at simultaneously diagonalizing multiple matrices – i.e. minimizing the off-diagonal elements in every single one of those matrices. Using a set of autocovariance matrices as target matrices for joint diagonalization (rather than just a single autocovariance matrix) can solve BSS problems.
The squared Frobenius norm of the off-diagonal elements of autocovariance matrices is often used as a "joint diagonality criterion", which is to be minimized in order to solve the respective BSS problem. In the context of explaining SOBI Belouchrani et al. [4] define the "off" of a $n \times n$ matrix M as

$$\mathit{off}(M) \stackrel{\text{def}}{=} \sum_{1 \le i \ne j \le n} |M|^2. \tag{2.41}$$

They then define the joint diagonality criterion for any $n \times n$ matrix V as

$$C(\mathcal{M}, V) \stackrel{\text{def}}{=} \sum_{k=1,K} \mathit{off}(V^H M_k V) \tag{2.42}$$

with $\mathcal{M} = \{M_1, ..., M_k\}$ being a set of $n \times n$ matrices.

Belouchrani et al. state that "it is *not* required that the matrix set under consideration can be exactly simultaneously diagonalized by a single unitary matrix [...] it is not even required that the matrices in the set are *individually* unitarily diagonalizable." [4] Instead, the unitary matrix V serving as joint diagonalizer of the matrix set $\mathcal{M}$ is only meant to *minimize* the joint diagonality criterion. This is referred to as *approximate* joint diagonalization. Ziehe puts it this way: "The approximation is understood in the sense of minimizing a suitable diagonality criterion." [21]

To sum all this up: The SOBI algorithm computes an approximate joint diagonalization of a multitude of autocovariance matrices by minimizing the joint diagonality criterion given in equation 2.42. The exact steps of the algorithm are detailed in the following section.

2.2.3.3.3 The SOBI Algorithm

The SOBI algorithm has been presented by Belouchrani et al. in 1997. [4]
Since SOBI is based on joint diagonalization of autocovariance matrices and since autocovariance refers to the covariance of a signal with itself at different points in time, the classical BSS problem (as given in equation 2.27) is reformulated in such a way that it expresses the time dependency of the signals:

$$x(t) = As(t). \tag{2.43}$$

The steps of the actual algorithm are illustrated in figure 4.

The first three steps of SOBI comprise data whitening. Whitening has been included in the algorithm to confine the solution space for the mixing matrix A to unitary matrices. A can therefore be factored into the pseudoinverse of the whitening matrix and a unitary matrix (this can be seen in the last step of figure 4).

After having whitened the data, autocovariance matrices are formed for a set of time lags. As Miettinen et. al state, "[...] the choice of the number of lags and which lags to choose is still an open problem for SOBI." [26]

The next step is then to compute the unitary matrix which serves as a joint diagonalizer for the set of autocovariance matrices. For this purpose, Belouchrani et al. have presented a modified version of the Jacobi eigenvalue algorithm (which is an iterative diagonalization method). [4]

Considering that the product of the whitening matrix and the joint diagonalizer is the mixing matrix, one can now compute the source signals.

As with fastICA, there are deflation-based, as well as symmetric approaches to SOBI.

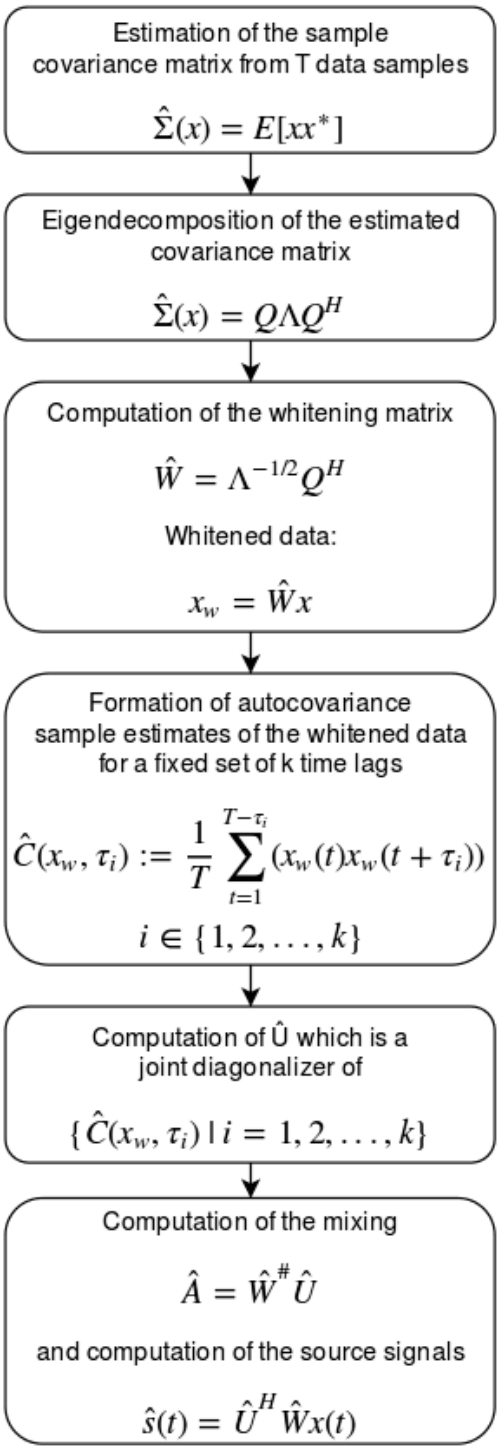

Figure 4: Basic steps of the SOBI algorithm.

2.3 Making use of all these fundamentals...

The first part of the fundamentals section basically dealt with the physical principles related to magnetic resonance imaging. Those principles help to understand what actually happens during (f)MRI scanning sessions and how images are formed.

However, before fMRI data analysis can be performed, a variety of processing steps are essential (e.g. motion correction, deobliqueing,...). Ideally, noise removal should also be part of this process. Since a novel approach to the removal of physiological noise from fMRI data based on blind source separation methods is the topic of this thesis, those methods and the mathematical and statistical background related to them have also been explained.

Before moving on to the next section, which goes into details on how this approach has actually been implemented, figure 5 roughly outlines the idea behind the approach: PCA (by means of SVD) is used for dimensionality reduction and decorrelation of the data. Subsequent blind source separation generates the independent components. Those are then analyzed (in the frequency domain) and only the ones which are assumed not to be related to cardiac activity are finally used for reconstruction.

Let X be a data matrix containing the (preprocessed) subject data. By means of SVD, it is decomposed such that

$$X_{t\times v} = U_{t\times c_1} S_{c_1\times c_1} (V_{v\times c_1})^T,$$

hence,

$$X_{t\times v} V_{v\times c_1} = U_{t\times c_1} S_{c_1\times c_1}.$$

Since the right singular vectors (columns of V) are the so-called *principal axes* or *principal directions*, projections of the data onto these axes (given by XV or, alternatively, US) are referred to as *principal components*.

Blind source separation of said principal components yields

$$(XV)_{t\times c_1} = C_{t\times c_2} (A_{c_1\times c_2})^T,$$

S denoting the component matrix (i.e. the columns of C are the independent time series) and A^T denoting the mixing matrix (i.e. the rows of A^T are the corresponding weights).
Note, that when using the current R implementation of the SOBI algorithm for blind source separation, c_1 will equal c_2.
Using spectral analysis, some of the independent components are identified as being related to cardiac activity. The respective columns of C and the corresponding weights (rows of A^T) are thus eliminated, such that

$$\widetilde{c_2} = c_2 \setminus c_{CA},$$

c_{CA} being the components related to cardiac activity. Reconstruction of the principal components is thus given by

$$(XV)_{t\times c_1} = C_{t\times \widetilde{c_2}} (A_{c_1\times \widetilde{c_2}})^T.$$

Projecting those reconstructed principal components into the original space produces a reconstructed data matrix (denoted by Y):

$$Y_{t\times v} = (XV)_{t\times c_1} (V_{v\times c_1})^T.$$

Figure 5: PCA by means of SVD, subsequent BSS and data reconstruction after removal of components related to cardiac activity. Matrix dimensions are denoted as subscripts, v being the number of voxels, t being the number of time points and c_1 and c_2 being the number of components computed by PCA and BSS, respectively.

3 Methods

3.1 Data Acquisition

The data analyzed in this thesis has been recorded within the course of a previous fMRI study of Boubela et al. [27] Within this study, a total of 18 subjects (9 male, 9 female, mean age: 34.6 years, SD = 12.2 years) have been scanned in the research facilities of the Medical University of Vienna using a Magnetom TIM Trio (3T) MRI scanner with a 32-channel head coil (Siemens Healthcare GmbH, Erlangen, Germany).

The resting-state BOLD fMRI scans were measured using a multiband EPI sequence (as described by Feinberg et al. in [28]) at a resolution of $1.7 \times 1.7 \times 1.7$ mm^3 with a slice gap set to 2 mm. 32 axial slices (matrix size 128×128, aligned with the AC-PC line) were recorded at TE/TR = 31/333 ms, a flip angle of 30°, with a multiband factor of 8 and a bandwidth of 1776 Hz/Pixel. 1200 volumes were recorded - corresponding to an overall scan time of approx. 7 minutes.

A high-resolution ($1 \times 1 \times 1.1$ mm^3) structural scan was also acquired (160 sagittal slices were recorded at TE/TR = 4.21/2300 ms, a flip angle of 90° and an inversion time (TI) of 900 ms).

The subjects heart rates have been recorded using the so called peripheral pulse unit (PPU - a Siemens product). Unfortunately, this has however only been done for 11 out of the 18 subjects.

3.2 Computing Environment

All computations have been performed on a server equipped with two Intel Xeon X5690 CPUs (3.47GHz, 6 cores) and 192 GB of main memory running on Ubuntu Linux (Version 14.04 'Trusty Tahr').

As for the software environment, R v3.1.1, AFNI version 'AFNI_2011_12_21_1014' and FSL v5.0 have been used.

3.3 Data Preprocessing

The data have been preprocessed using the neuroimaging tools / libraries FSL ([29]) and AFNI ([30]). The order in which the preprocessing steps have been performed, was based on the recommendation given in [31].

The resting-state data sets have been bias-corrected using FSL FAST and skull stripped using the FSL brain extraction tool (BET). They have then been motion corrected using FSL FLIRT. Subsequently, the data sets have been aligned to the anatomical images in MNI152 standard space (using AFNI's 3dAllineate). Deobliqueing was performed using AFNI's 3drefit. All data sets were then spatially blurred using a Gaussian kernel (FWHM = 8 mm) with AFNI's 3dmerge tool.
Furthermore, considering that each subject data set had a high spatial and temporal resolution ($1.7 \times 1.7 \times 1.7$ mm^3 with 1200 volumes being recorded), the group data matrix, which was constructed in order to perform group ICA (for ROI selection – see section 3.7.1) was a very large matrix (>75 GB for 18 subjects). This lead to computational limits, which were overcome by resampling the data to a spatial resolution of $2 \times 2 \times 2$ mm^3 using AFNI's 3dresample.

In order to assess the capabilities of the algorithm, functional connectivity analysis has been performed (*after* all the processing steps – see section 3.7). This analysis has not only been performed with the fully processed data, but as well with data of which physiological noise was removed by means of band pass filtering. This reference data was created by using the same original data, performing the exact same preprocessing steps on it and then band pass filtering it with cut-off frequencies of 0.01 and 0.1 Hz, respectively.

3.4 Principal Component Analysis

Principal component analysis has been performed by means of computing the singular value decomposition of the (centered) data matrix. This has been done using R's SVD package (since in comparison to R's IRLBA package, SVD allows for multi-threaded processing when using an OpenBLAS library instead of R's standard BLAS library).

The common considerations have been made with respect to the dimensionality determination for PCA (see section 2.2.2.6 for more details on that topic).
First, the scree plots for all 18 subjects have been computed. They have then all been scaled in such a way that their maximum values would all equal 100 (arbitrarily chosen value). In a next step, a mean scree plot has been computed from the scaled plots (this mean scree plot is depicted in section 4.1).
One common approach would then be to compute the cumulative variance for which the principal components account and to then take a random guess at the percentage of how much of this variance is probably noise related in order to finally choose the number of principal components. This, however, seemed unreasonably arbitrary. Also, projecting the original data of each subject (1200 dimensions) into 1000 dimensions would still retain only less than 90% of the total variance -

such an approach would thus probably tempt to retain too many components for further processing.
Another common approach would be to look for a kink in the scree plot (see section 2.2.2.6). Since the computed scree plot showed no kink, this approach was not helpful either.
Hence, the slope of the mean scree plot was computed at different points. Since scree plots are by definition strictly decreasing, their slopes can never be positive. Also, the absolute value of the slope of this scree plot is steadily decreasing (i.e. the curve doesn't get steeper). Therefore, it can be concluded that the changes in variance explained by the components constantly become smaller. At a certain point, they can be regarded as being negligible. Large numbers of components with negligibly different variances were assumed to be noise related. A threshold value for the slope (i.e. a value from which on the variances were defined as changing sufficiently little) has been defined (rather arbitrarily) in order to determine a reasonable number of principal components – results are given in section 4.2).

3.5 Blind Source Separation

Blind source separation was computed with the help of an R implementation of Belouchrani's SOBI algorithm ([4], [26]) and an R implementation of Aapo Hyvärinen's fastICA algorithm ([18]).

The R implementation of the SOBI algorithm automatically computes as many components as are given to it as input (i.e. in this case 100 components, since PCA was set to deliver 100 components) – it does not enable the user to define another number of components to be computed. So, in order to properly compare the results obtained from BSS by means of fastICA to the results obtained from BSS by means of SOBI, the number of independent components to be computed by the fastICA algorithm needs to be equal to the number of principal components computed by means of PCA.
Thus, 100 components (time series) and corresponding weights (spatial maps) were computed. However, so far, the choice of the number of components can hardly be justified – further research is needed in this area. As Hyvärinen et al. have stated: "[I]t is even more difficult to give any guidelines as to how many components should be estimated. Trial and error may be the only method applicable." [16]

Within the course of this work, symmetric SOBI and fastICA approaches have been utilized. It remains an open issue, though, whether different (and if so, better) results can be obtained by using deflation-based approaches.

The components obtained by BSS using fastICA were remarkably different to

those obtained by BSS using SOBI: While the Fourier transforms of some of the SOBI components were spectra with clearly distinguishable peaks (figure 6 is a good example of SOBI component with an obvious peak), the Fourier transforms of the fastICA components contained no such peaks (respective results are not shown, since displaying 100 components for 18 subjects - or even for a single subject - would go beyond the scope of this thesis).
Since the component identification algorithm (as described in section 3.6.1) fully relies on spectral analysis (to be more specific: on finding certain peaks within the spectra), the results obtained from the fastICA computations could not be used for further processing.

3.6 Component Identification and Extraction

3.6.1 Identification

In order to identify the components related to cardiac activity, the power spectral densities (PSDs) of all components were estimated for each subject by squaring the magnitude of their Fourier transform.
The local variability of the Fourier spectra has been estimated by means of computing the sliding window median absolute deviation (MAD). The MAD has been chosen as a measure of central tendency for this purpose, because unlike approaches using the mean or standard deviation, it is rather robust against outliers. The size of the sliding window has been set to 30 frequency bins (of which each has a size of $\sim 2.5 * 10^{-3}$ Hz – this size is of course determined by the sampling rate).
A magnitude threshold has then been used in order to determine the frequency range potentially related to cardiac activity: the threshold is basically dependent on the sum of the median and the median absolute deviation of the sliding window MAD. The latter can however be scaled by a factor f, which facilitates an adaption of the sensitivity of the algorithm:

$$threshold = median(sldg.mad) + f * mad(sldg.mad), \tag{3.1}$$

where a value of 5 for the factor f has been found empirically to deliver useful results and has thus been used in this work.

The highest frequency at which the sliding window MAD curve would rise above this threshold was defined as the upper limit of the above mentioned frequency range. The lower limit of this frequency range was defined as being the highest possible frequency at which the sliding window MAD curve would fall below the above mentioned threshold (provided that this frequency value is at least 50 frequency bins – i.e. ~ 0.125 Hz – lower than the upper limit).

Based on the sliding MAD curve within this frequency range, a component has been assumed to be primarily related to the heart rate,

- if the area under the sliding MAD curve within said frequency range (which corresponds to the signal power of this frequency range) is at least 50% of the area under the entire sliding window MAD curve (i.e. it comprises at least 50% of the total signal power) and
- if the lower limit of the frequency range is larger than 0.5 Hz.

If the first criterion is not fulfilled, the component might be related to cardiac activity and *additionally* to other physiological phenomena (there might be another relevant peak in the PSD curve) - e.g. neuronal activity, which has to be preserved in the data.
The second criterion simply ensures that low frequency peaks in the PSD curve related to other physiological phenomena are not mistaken for heart rate peaks (it is unlikely that a healthy subject has a resting heart rate of less than 30 bpm).

The measures used within the algorithm for identification of components related to cardiac activity are illustrated in figure 6 (the given component would be identified by the algorithm as being related to cardiac activity).
It is noteworthy that the behavior of the algorithm and thus also the results obtained by it (number and nature of extracted components) strongly depend on how some of the parameters are set (factor f in threshold computation, AUC criterion, sliding window size). So far, the values have been determined empirically.

3.6.2 Extraction

The components (and corresponding weights) related to cardiac activity have been saved separately. The spatial maps can be used to locate the brain areas in which cardiac activity predominantly contributes to the original data).
The remaining components were used for the reconstruction of the data matrix.

3.6.3 Heart rate determination as plausibility check

For every subject, a mean curve of all sliding window MAD curves of the components related to cardiac activity was computed. The maximum of this mean curve was assumed to be an estimator of the most prevalent heart rate. This value was compared to pulse data, which was recorded separately during the scanning sessions. (The pulse data itself was analyzed by computing its Fourier transform and then determining the maximum of the resulting spectrum.)

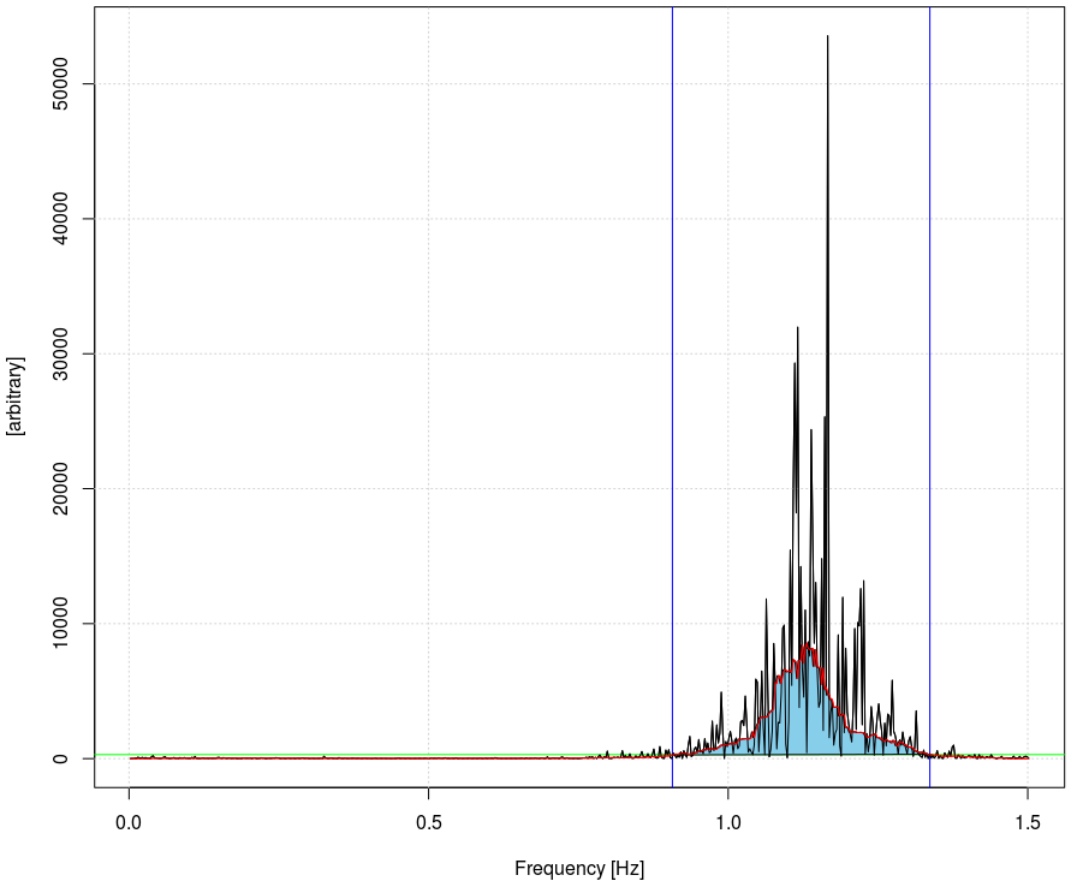

Figure 6: Measures used in the algorithm for identification of components related to cardiac activity: Component PSD (black curve), sliding window MAD (red curve), magnitude threshold (green horizontal line), frequency range limits (blue vertical lines), signal power within frequency range (blue area under the curve).

The similarities in the heart rates computed from the components related to cardiac activity and the heart rates computed from the pulse data were then assessed (high similarities were assumed to indicate the plausibility of the identified components to actually be related to cardiac activity).
However, it needs to be mentioned that specific heart rate values (like the ones used here, which are assumed to be the subjects most prevalent heart rates) computed from the components related to cardiac activity are strongly dependent on the behaviour of the component identification algorithm - which again depends on how its parameters are set (see section 3.6.1).
Results are listed in section 4.2.

Another approach of investigating the algorithms ability to correctly identify com-

ponents related to cardiac activity would be to compare the heart rate variability (HRV) computed from the photoplethysmography data to the HRV estimated from the components spectra. Heart rate variability is a physiological phenomenon referring to the variability in duration between heart beats (roughly stated: a high HRV is a sign of good health). There are different measures of HRV. One way of expressing it is by computing the standard deviation of NN intervals (i.e. intervals between heart beats). This measure is referred to as 'SDNN' and can be computed e.g. from photoplethysmography data. Estimating the HRV from the BOLD signal components Fourier spectra could be approached by computing the width of the frequency peak (e.g. between the frequency range limits described in section 3.6.1). The HRV measures obtained from the photoplethysmography data and those obtained from the BOLD signal components could then be compared and the ability of the algorithm to correctly identify BOLD signal components related to cardiac activity could be deducted. However, one drawback of using SDNN as a measure of HRV is that the standard deviation is sensitive to outliers – hence, an extremely accurate detection of peaks in the photoplethysmography data would be required.
This approach has not been followed through within the course of this thesis project, but should be subject to future research.

3.7 Analysis of Reconstructed Data

Seed-based correlation analysis (SCA) has been performed as a means of studying functional connectivity. In order to compare the functional connectivity of the band passed data with that of the reconstructed data (i.e. the data, from which the components related to cardiac activity have been removed), the same *regions of interest* (ROIs) have been used for said data sets. The correlations have been calculated using R and visualized using AFNI.

3.7.1 ROI computation for functional connectivity analysis

Although there seems to be no standardized way of ROI selection, it has become common practice to select ROIs based on the anatomy of standard atlases and to then use the selected ROIs for functional connectivity analysis of *all* subjects of a study. However, Sohn et al. have shown only recently that due to the variability of functional brain connectivity amongst subjects, this approach leaves room for improvement [32].
Hence, they describe a novel approach to ROI selection. This approach selects ROIs for each subject individually. To be more specific: "[...] the associated time-series for a network obtained in [group] ICA is used to reconstruct that network for a specific individual. To do this, we took the time component of each target

ICA network and identified the time-series for each individual. [...] Seed coordinates were determined by the peak z-values in the reconstructed networks for each individual, which resulted in a unique set of ROIs for each subject." [32]

Before performing the actual spatial ICA (with 30 components) on the group data, PCA has been performed (1000 components) - again, using R's SVD package. 1000 components have been used for PCA in order to retain a high degree of information (1000 components account for approx. 88% of the total variance). The number of PCA components is (in this context) however by far not as important as the number of ICA components (which was determined empirically). E.g. an attempt to use 100 ICA components (instead of 30) has proven to separate the data into too many subgroups, resulting in the fact that no known resting state networks could be identified.
The resulting ICA components (spatial maps) have been inspected visually (using AFNI) and categorized by comparison to literature data ([33], [2]).
Results are depicted in section 4.3.1.

The time series (weights) of the components of interest have then been used for correlation analysis with single subject voxel time series (using R's *cor()*-function, which computes Pearson's correlation coefficient): Every voxel time series of every subject has been correlated to said components' time series. Since exactly one r value is computed for each voxel time series of each subject, those r values can be easily visualized as spatial maps. In order to perform a plausibility check of the results obtained from the correlation analysis, a mean of those spatial maps (henceforth referred to as "mean correlation map") has been computed from all subjects for the previously selected components. Furthermore, the correlation between the individual correlation maps and the sICA components have been computed in order to investigate the consistency of the resting state networks among the subjects. The respective results are given in section 4.3.2.

A maximum correlation value has been determined for every correlation map and has been chosen as center voxel for the respective subject- and component-specific ROI. The Fisher transformation to z values ($z := arctanh(r)$) has been omitted, since the inverse hyperbolic tangent function is strictly monotonic increasing and hence, the maximum correlation value corresponds to the maximum z value.
To ensure that those center voxels really are part of the respective networks, voxel masks have been computed from the spatial ICA maps of the group data. The threshold to determine whether a voxel would be part of the mask or not has been set to ± 2, respectively, for every component of interest. However, due to the scaling ambiguity of BSS problems, which is discussed in section 2.2.3, the

threshold needs to be determined anew for every new analysis. Using those masks, the computation of the center voxels was confined to the selected networks. Results are given in section 4.3.3.

3.7.2 Seed-Voxel Correlation Maps

The ROIs around the center voxels have been set to be spherical with a radius of 4 mm. With a voxel size of $2 \times 2 \times 2$ mm^3, this corresponds to 33 voxels within each ROI. As usually done in seed-based correlation analysis, the mean time series has been computed from all the time series of the voxels within the ROI. This mean time series has then been correlated to all voxel time series of the respective scan. This has been done for the reconstructed and the band passed data of every subject.

In order to examine the differences between the seed-voxel correlation maps of the reconstructed and the band passed data, their absolute values have been subtracted from each other as follows:

$$\Delta_r = |r_{reconstructed}| - |r_{bandpassed}|, \tag{3.2}$$

with $r_{reconstructed}$ denoting a seed-voxel correlation map of the reconstructed data, $r_{bandpassed}$ denoting a seed-voxel correlation map of the band passed data.
For every component of interest, the mean of the difference maps of all subjects has been computed. Showing that this mean is significantly different from 0 would imply that the seed-voxel correlation maps of the reconstructed and the band passed data are significantly different. This test, which is essentially a paired t-test, has been carried out voxelwise using AFNIs *3dttest++*. AFNI uses False Detection Rate (FDR) correction in order to control Type-I errors. The fact that such a correction is necessary (since the test procedure would otherwise represent a classical multiple comparisons problem) has been described in [34].

Fisher's z-transformation has been computed (before performing AFNIs 3dttest++) in order to meet the normality assumption of the t test.

4 Results

4.1 Dimensionality determination for PCA

The steps leading to the computation of the mean scree plot (depicted in figure 7) are described in section 3.3.

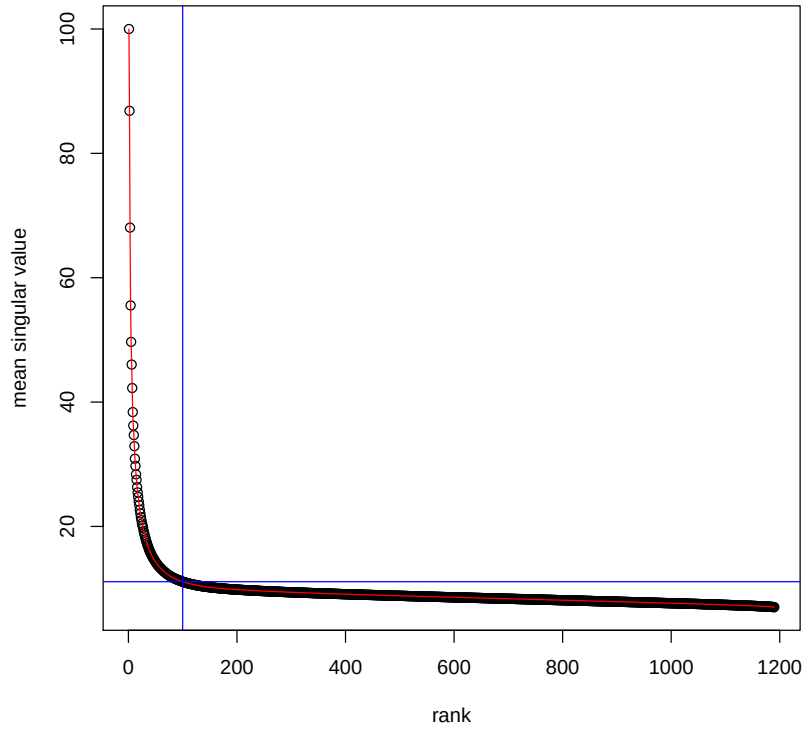

Figure 7: Mean scree plot computed from the scree plots of all subjects (the vertical blue line marks the singular value ranked 100 and the horizontal blue line marks the actual 100th singular value).

The first one hundred singular values account for 17.4% of the total variance. The fact that the slope of the curve at the 100th singular value lies at -0.02 and that it (and its absolute value) steadily decreases from there on indicates that the variance accounted for by the remaining 1100 components changes only negligibly (the curve's slope at the 1st singular value lies at 13.16). The more or less similar variance of these last 1100 components leads to the assumption that those components might represent noise.
Hence, the first 100 principal components have been computed for the data sets of all subjects. This may seem arbitrary (and to a certain extent it actually is), however, as Ashby states, "[...] PCA, by itself, offers no statistical basis for deciding how many components to eliminate." [12]

4.2 Heart rates computed from components related to cardiac activity

The methods used to determine the most prevalent heart rates of the subjects from the components related to cardiac activity are described in section 3.6.3. Said section also describes the method used to determine the subjects most prevalent heart rates from the data recorded with the Siemens peripheral pulse unit (PPU).

Those heart rates are listed in table 2.

Subject ID	HR computed from components	HR computed from pulse data	Percent Error [%]
1	78	78	0
2	72	66	9.1
3	42	42	0
4	66	66	0
5	60	60	0
6	54	54	0
7	66	42	57.1
8	72	66	9.1
9	54	54	0
10	66	72	8.3
11	54	54	0
$\mu \pm \sigma$	62±10	59±12	7.6±16.9

Table 2: Heart rates computed from components which are related to cardiac activity versus heart rates as obtained from the Siemens PPU. The last row contains the means (μ) and standard deviations (σ). Heart rates are given in beats per minute.

The percent error has been calculated assuming that the heart rates computed from the PPU data correspond to the "true" values.
In seven out of eleven subjects the heart rate computed from the BOLD signal was identical to the heart rate computed from the PPU signal. In another three subjects, the heart rate computed from the BOLD signal deviated by ±0.1 Hz (i.e. ±6 bpm) from that of the PPU data. This corresponds to percent errors of 8.3% and 9.1%, respectively. In one subject, the heart rate computed from the BOLD signal was 0.4 Hz (i.e. 24 bpm) off - that corresponds to a percent error of 57.1%.

The given results show that the described method can not be used to determine a

subject's exact heart rate. However, the results do suggest that using the method for identifying components related to cardiac activity is legitimate. As the the data are taken from brain scans, other factors like vessel compliance (decreasing with increasing age) may also play a role.

Changing the parameters of the identification algorithm in a way which makes it 'less strict' (i.e. more components are identified as being related to cardiac activity) is possible and increases the accuracy of the determined heart rates. On the other hand, this would also lead to an 'extraction' of more data and thus might result in the loss of valuable information contained in the signal.

4.3 ROI computation

4.3.1 Components of interest

Spatial ICA on the group data matrix yielded 30 spatial components (with corresponding weights), of which only a few have been chosen (upon comparison with literature – see [2] and [35]) for ROI determination.
The chosen components are depicted in figures 8 through 12.[3] Said figures also depict the maps referred to as mean correlation maps (further details on those are to be found in section 4.3.2).
The MNI152 T1 template has been chosen as an underlay.

The sub-figures (a)-(c) of figures 8 through 12 have been thresholded at $\pm$ 2.
A more detailed look at the atlas regions covered by the component depicted in figure 8 reveals that it primarily covers parts of the superior temporal gyri (Brodman Area 22), the insular cortex and - to some extent - sections of the inferior frontal gyri - especially the pars orbitalis (BA 47). This is in accordance to literature (Damoiseaux et. al [2]), where said network is described as covering parts of the superior temporal gyri, the insular cortex and the postcentral gyri. The given source describes those cortical regions as being "[...] acknowledged to represent the auditory cortex."
This is consistent with the fact that some activation was also found in the transversal temporal gyri (also known as Heschl's gyri – BA 41 and BA 21), which comprise the primary auditory cortex.

The component depicted in figure 9 represents the primary visual network. It basically covers the cuneus. Using functional rather than structural terminology,

[3] All 2d maps within this thesis are given together with coordinates according to neurological convention (i.e. LPI axes orientation). The coordinates refer to the distance (in mm) to the location of the anterior commissure in the MNI152 T1 template.

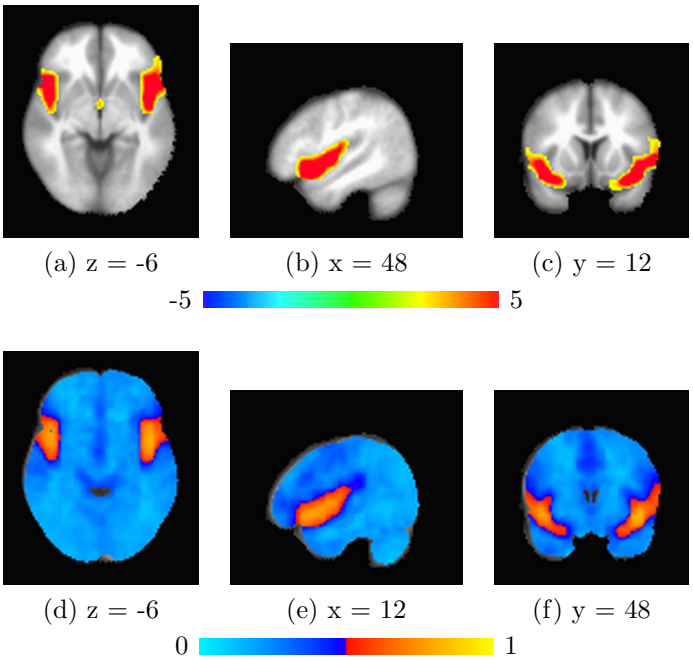

Figure 8: Auditory network computed by group sICA (a-c) and corresponding mean correlation map (d-f)

the covered section is referred to as primary visual cortex (or striate or visual area V1). The corresponding Brodmann area is BA 17. Parts of the primary visual network also cover the parastriate (BA 18).
These anatomical classifications were not only determined with the help of the underlying atlas, but have also been reported about in the before mentioned publication by Damoiseaux et al.

Due to the strong resemblance of the component depicted in figure 10 (a-c) to the network referred to as "auditory" in [35], this component was first assumed to actually represent the auditory network. However, an in-depth analysis showed that the areas covered by the component are primarily the precentral gyrus and - from a cytoarchitectural point of view - BA 4 and BA 6. BA 4 refers to the primary motor cortex (it coincides with the precentral gyrus) while BA 6 is located anterior to BA 4 and comprises the premotor- and the supplementary motor cortex.
Out of all five selected components, this one shows the greatest variability amongst subjects (as indicated by the strongly smeared mean beta map depicted in figure 10 (d-f)). Beckmann et al. ([33]) have reported about a network to which they refer as sensory-motor network. The smeared beta map (especially the coronal view which illustrates strong correlation of the voxel time series of even the superior

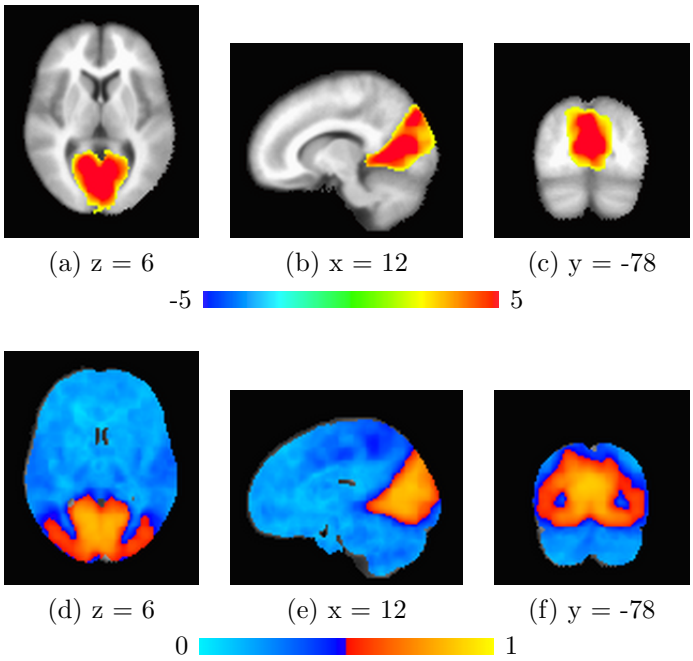

Figure 9: Primary visual network computed by group sICA (a-c) and corresponding mean correlation map (d-f)

sections of the pre- and postcentral gyri with the sICA component's time series) shows a strong resemblance to Beckmann's sensory-motor component.
Hence, it seems plausible that this network is related to sensory-motoric neural activity.

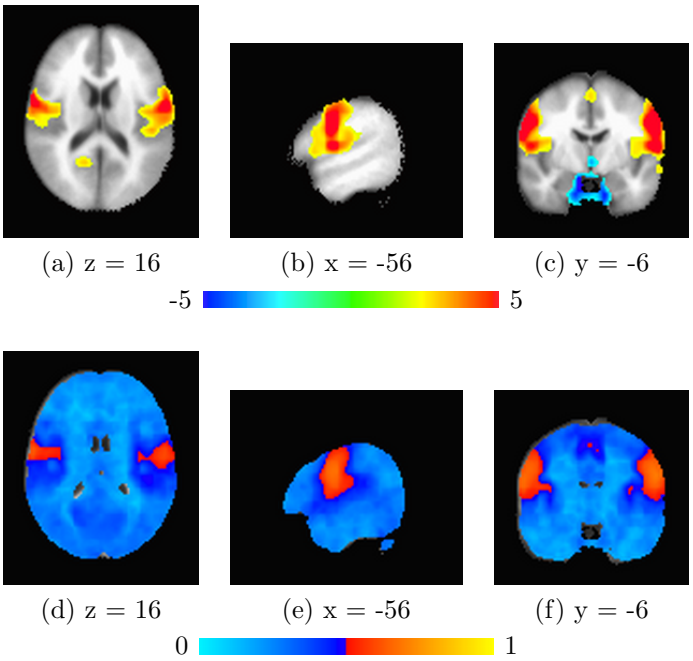

Figure 10: Sensory-motor network computed by group sICA (a-c) and corresponding mean correlation map (d-f)

The component depicted in figure 11 covers a variety of areas belonging to the occipital lobe: primarily the superior and middle occipital gyri as well as the lingual gyri. From a cytoarchitectural point of view, BA 18 and BA 19 are comprised in this component - which together form the extrastriate (or peristriate) cortex (i.e. visual areas V1-V5).
Damoiseaux et al. have found this same network (covering the exact same areas) and have also stated, that those areas "[...] are recognized as part of the visual cortex." [2].
The network is commonly referred to as lateral visual network.

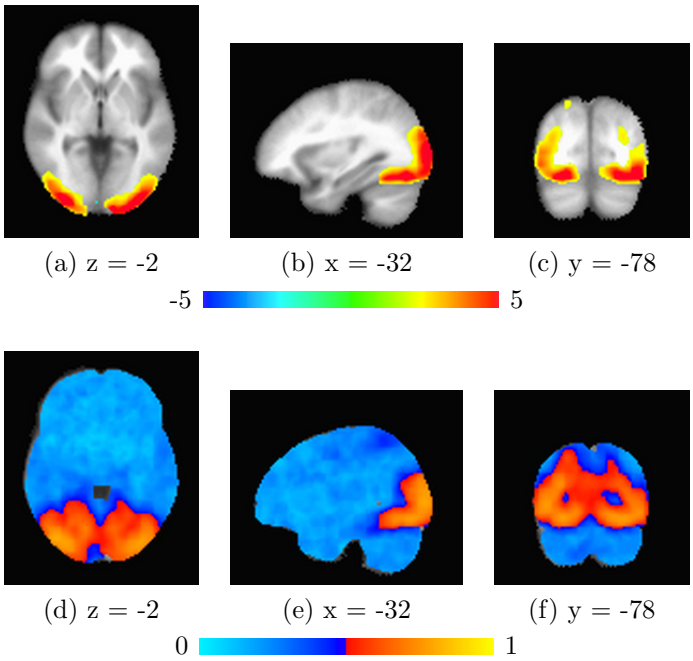

Figure 11: Lateral visual network computed by group sICA (a-c) and corresponding mean correlation map (d-f)

The component depicted in figure 12 covers parts of the medial frontal gyrus, the superior frontal gyri, as well as parts of the anterior cingulate cortex. In cytoarchitectural terms, BA 9 and BA 10 lie within the covered area.
Damoiseaux et. al have found this same network and have analogously stated that prefrontal cortical areas as well as the cingulate cortex are involved ([2]).
This network represents the anterior part of a network commonly referred to as default mode network. Buckner et al. describe the functional properties of this network as follows: "In particular, the default network is the most active brain system when individuals are left to think to themselves undisturbed. The default network also increases activity during mental explorations referenced to oneself including remembering, considering hypothetical social interactions, and thinking about ones own future." [36]

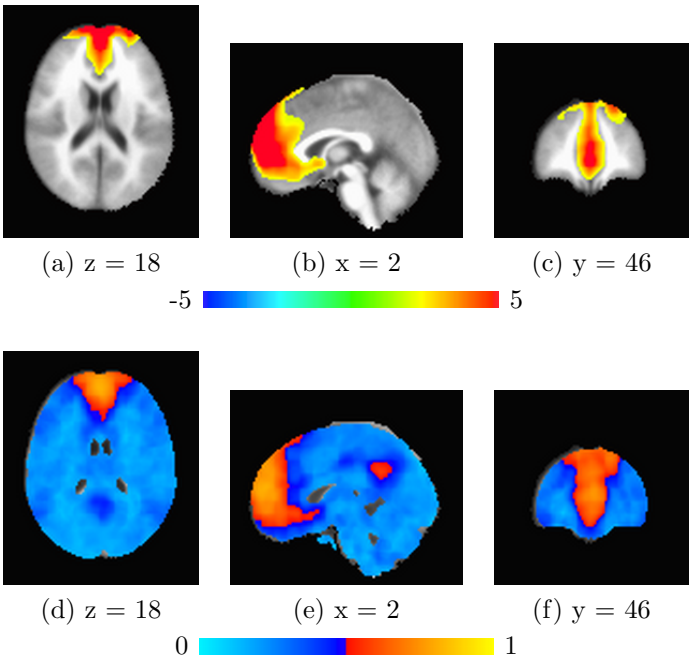

Figure 12: Default mode network (anterior part) computed by group sICA (a-c) and corresponding mean correlation map (d-f)

4.3.2 Mean correlation maps and network consistency among subjects

As described in section 3.7.1, correlation analysis yielded correlation maps, of which a mean map has been computed from all subjects for each of the selected components. Those mean maps are also depicted in figures 8 through 12 (d-f) – together with the corresponding maps obtained from group data sICA (see section 4.3.1).

The similarity between those mean correlation maps and the maps computed by group data sICA is obvious. This indicates that the networks found by means of sICA are not merely patterns artificially produced by sICA, but that the correlations of the respective time series are actually present in the subjects.

In order to investigate the consistency of the networks amongst subjects (i.e. to ensure that the patterns do not arise due a single occurrence in one of the subjects), the individual correlation maps have been correlated with the corresponding sICA components. This has also been described in section 3.7.1. The results of this analysis are listed in table 3 and illustrated in form of a box plot in figure 13. Initially, the r values might seem to indicate a low correlation, but the reason for

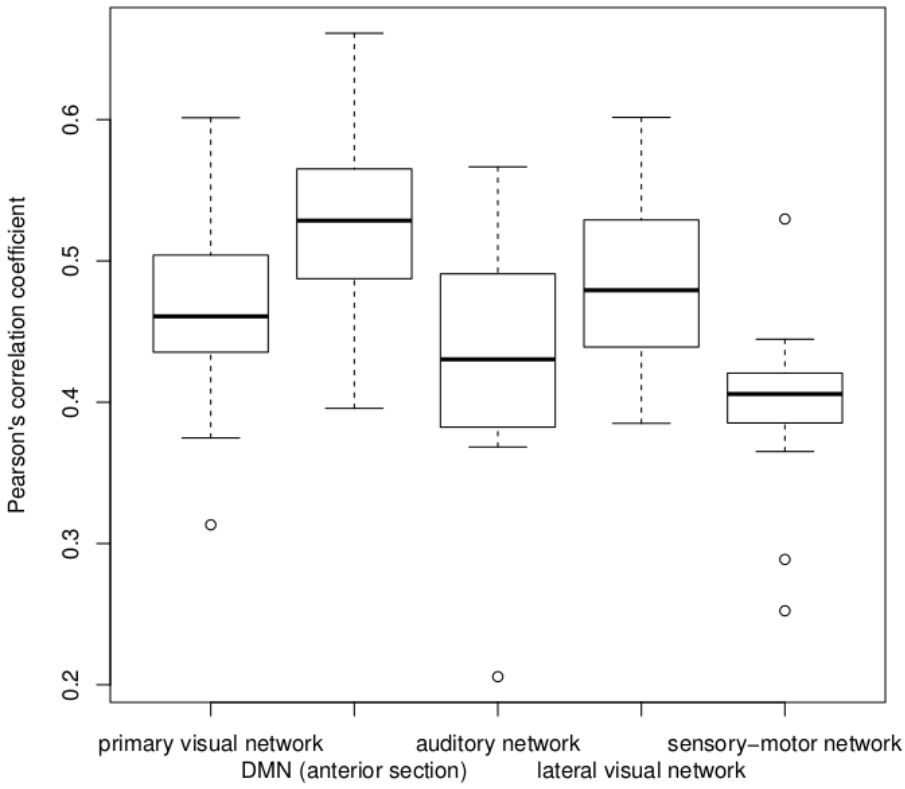

Figure 13: Pearson correlation between the individual subjects' correlation maps and the sICA components

this is simply that only a portion of all voxels of a brain scan actually contribute to each of those networks. When investigating the consistency of resting state networks among 1000 subjects (data obtained from the human connectome project), Kalcher et al. have obtained comparable results. [37]

However, the magnitude of the individual correlation coefficients only accounts for the degree to which the resting state networks are present in the individual subjects. The *consistency* of the presence of these networks among the individual subjects has to be derived from the spread of the correlation coefficients of all subjects. As depicted in figure 13, the interquartile range (IQR) of the data with respect to the the auditory network reaches a value of 1.0 (i.e. higher than that of any other network). This indicates that the auditory network seems to be the least consistent amongst subjects. The sensory-motor network seems to be the most consistent among all subjects, having an IQR of only 0.03 as well as the lowest minimum-to-maximum range.

	Min	1^{st} Quart.	Median	3^{rd} Quart.	Max
primary visual network	0.31	0.44	0.46	0.50	0.60
DMN (anterior section)	0.40	0.49	0.53	0.56	0.66
auditory netowork	0.21	0.38	0.43	0.49	0.57
lateral visual network	0.39	0.44	0.48	0.52	0.60
sensory-motor network	0.25	0.39	0.41	0.42	0.53

Table 3: Five-number summary of the correlation coefficients between the subject specific correlation maps and the sICA components

A more thorough investigation of the consistency of resting state networks would go beyond the scope of this thesis.

4.3.3 Subject specific ROIs

The resulting subject specific ROIs for each of the selected components have been computed as described in section 3.7.1 and are depicted in figures 14 through 18. For the sake of brevity, not all 90 ROIs (one for each subject-component combination) can be visualized individually within this thesis. However, a combination of an axial view, a sagittal (or coronal) view (depending on which was deemed to be better suited) and a three dimensional view for each of the resting state networks should suffice to offer insights into the locations of the subject specific ROIs. More specifically, the atlases depicted contain the subject specific ROIs of all subjects (although not every single one will be visible due to their different locations) for every one of the chosen resting state networks.

The ROIs seem to appear in more or less dense clusters in case of some networks (this is for example clearly visible in the case of the ROIs determined for the default mode network – see figure 18) and are more spread across the brain areas in the case of other networks (e.g. the sensory-motor network – see figure 16).

The MNI152 T1 template has been chosen as an underlay for the 2d views. Analogously, an MNI152 template has also been chosen as underlay for the 3d views.

Figure 14 shows that the majority of subject specific ROIs computed for the auditory network are located in the anterior section of the superior temporal gyri and – to a certain extent – in the insular cortex as well.

Figure 15 reveals that the subject specific ROIs computed for the primary visual network are not as densely clustered as e.g. in case of the default mode network.

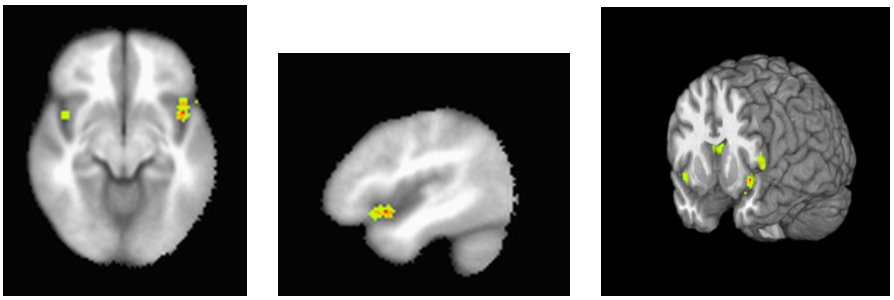

(a) z = -10 (b) x = -44 (c)

Figure 14: Subject specific ROIs computed for the auditory network

In the case of the subjects of this study, almost all of the ROIs appear to be located in the right cuneus.

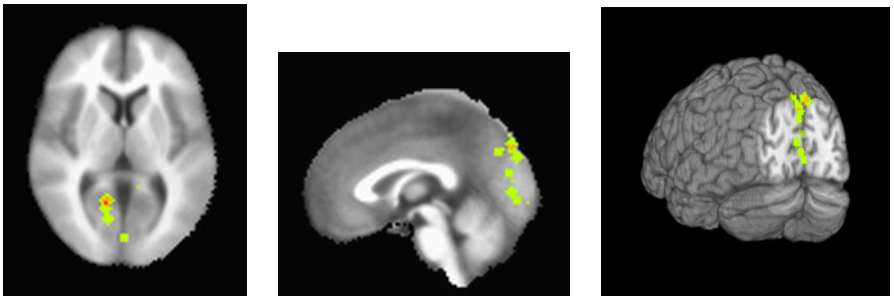

(a) z = 6 (b) x = 0 (c)

Figure 15: Subject specific ROIs computed for the primary visual network

Figure 16 shows that the subject-specific ROIs computed for the sensory-motor network are spread over the entire precentral gyri (this is especially visible in the coronal view (b) and the 3d view (c)).

The subject specific ROIs computed for the lateral visual network are mostly located in the middle occipital lobe (mostly in the area of the lingual gyrus). Some clustering appears in the left BA 18. This is illustrated in figure 17.

Figure 18 clearly illustrates that the subject specific ROIs computed for the anterior default mode network are all located in the anterior section of the medial frontal gyri - or, from a cytoarchitectural point of view - within BA 10.

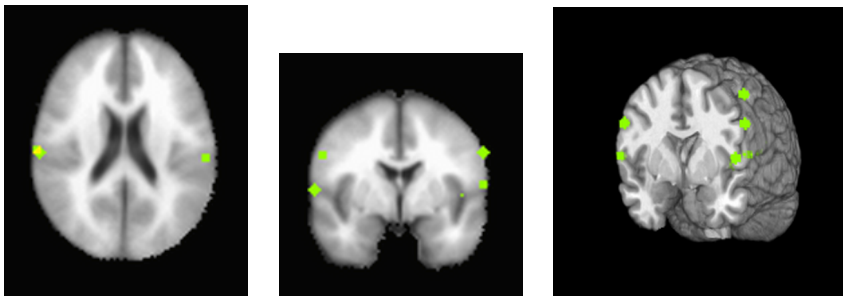

(a) z = 20 (b) y = -2 (c)

Figure 16: Subject specific ROIs computed for the sensory-motor network

(a) z = -18 (b) x = 26 (c)

Figure 17: Subject specific ROIs computed for the lateral visual network

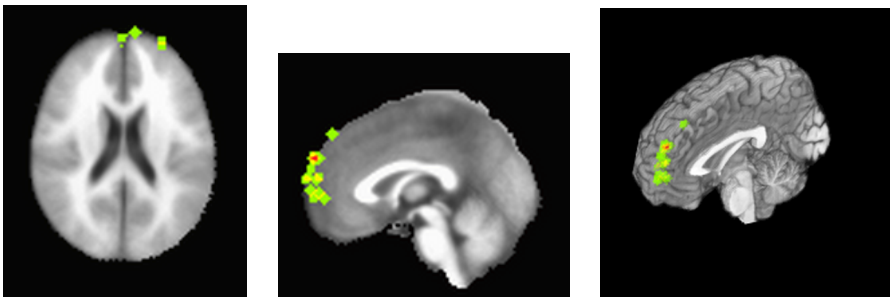

(a) z = 18 (b) x = 0 (c)

Figure 18: Subject specific ROIs computed for the default mode network (anterior part)

4.4 Analysis of Functional Connectivity

Analysis of functional connectivity has been carried out by means of seed-based correlation analysis using the previously computed subject specific ROIs. The specific steps are outlined in section 3.7.2.

Displaying the seed-voxel correlation maps of the band passed data and the reconstructed data (containing an r value for each voxel) for every subject with respect to every one of the five chosen resting state networks would go beyond the scope of this thesis. This, however, is not necessary, since the five maps (one for each network) of the FDR corrected results of *3dttest++* (containing a t statistic for each voxel) portray the differences between the band passed data sets and the reconstructed data sets for each of the networks.

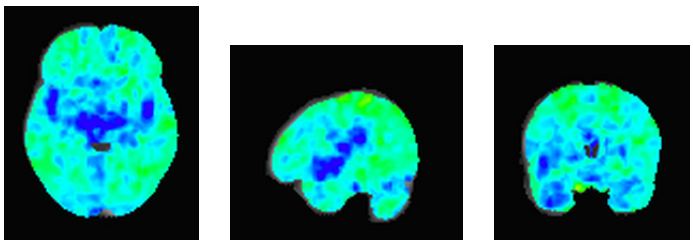

(a) z = -6 (b) x = 48 (c) y = 12

Figure 19: Differences in seed-based correlation analysis of band passed and reconstructed data (using the auditory network ROIs)
t <0 → higher correlation in band passed data
t >0 → higher correlation in reconstructed data

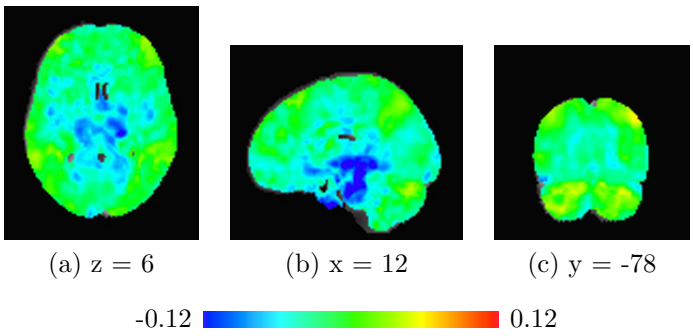

(a) z = 6 (b) x = 12 (c) y = -78

Figure 20: Differences in seed-based correlation analysis of band passed and reconstructed data (using the primary visual network ROIs)
t <0 → higher correlation in band passed data
t >0 → higher correlation in reconstructed data

Since the difference maps are computed as being the voxelwise differences between the absolute correlation values of the band passed data and the absolute correlation values of the reconstructed data (see equation 3.2), the fact that the results

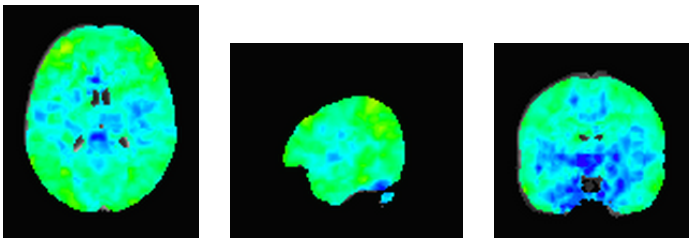

(a) z = 16 (b) x = -56 (c) y = -6

-0.12 0.12

Figure 21: Differences in seed-based correlation analysis of band passed and reconstructed data (using the sensory-motor network ROIs)
t <0 → higher correlation in band passed data
t >0 → higher correlation in reconstructed data

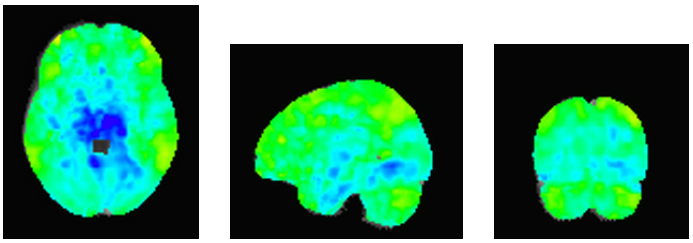

(a) z = -2 (b) x = -32 (c) y = -78

-0.12 0.12

Figure 22: Differences in seed-based correlation analysis of band passed and reconstructed data (using the lateral visual network ROIs)
t <0 → higher correlation in band passed data
t >0 → higher correlation in reconstructed data

depicted in figures 19 through 23 mainly consist of negative values means that the correlation (or anti-correlation, respectively) in the band passed data tends to be higher than in the reconstructed data. This was originally unexpected, since the higher information content of the reconstructed data was assumed to lead to higher (anti-)correlations. However, on second thought, this makes sense, considering that Chang et al. have already reported about the removal of heart rate components from fMRI data to result in reductions of functional connectivity (a photoplethysmograph was used to obtain the subjects heart rates). [15]
Yet, another important factor could play a role: As Davey et al. have put it - "[...] the fMRI preprocessing step of temporal filtering increases the sample variance of correlation, thereby inflating the false positive rate and artificially inducing connectivity." [14]. Higher (anti-) correlations could thus be induced by the process

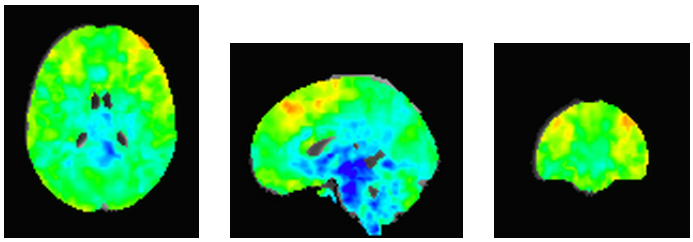

(a) z = 18 (b) x = 2 (c) y = 46

-0.12 0.12

Figure 23: Differences in seed-based correlation analysis of band passed and reconstructed data (using the anterior default mode network ROIs)
t <0 → higher correlation in band passed data
t >0 → higher correlation in reconstructed data

of band pass filtering and hence reflect spurious connectivity rather than actual functional connectivity. Further research on this matter and potentially corrections for filtering-induced correlations are necessary.

The differences in the seed-voxel correlation maps computed for the auditory network from the reconstructed data and the band passed data seem to coincide with the structure of the network itself (compare figures 8 and 19 - especially the axial and sagittal views).

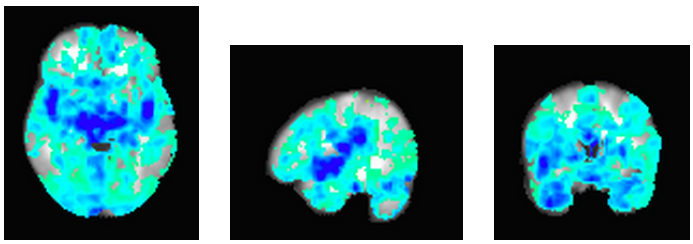

(a) z = -6 (b) x = 48 (c) y = 12

-0.12 0.12

Figure 24: Differences in seed-based correlation analysis of band passed and reconstructed data (using the auditory network ROIs) – threshold set to a q value of 0.05

Figure 24 illustrates the results of *3dttest++* for that network once more - this time with the threshold being set to a q value of 0.05 (the MNI152 T1 template has once again been chosen as underlay). The fact that the (anti-)correlations within the regions of the network are *significantly* different in the band passed

data and the reconstructed data leads to the assumption that the different means of removing noise related to cardiac activity could potentially have an impact on functional connectivity analysis. However, this assumption should be reflected upon with caution, since on the one hand seed-voxel correlations may be induced by temporal filtering (as already mentioned above) and on the other hand similar findings were not made for the other resting state networks.

5 Discussion

The goal of this work was to remove physiological noise related to cardiac activity from resting state fMRI data with high temporal resolution by means of an algorithm based on blind source separation rather than by simple band pass filtering. Band pass filtering removes not only physiological noise from the data but rather *all* signal components at some predefined frequencies. The new approach, however, separates the data into temporally independent components, analyses their frequency spectra, removes only the components related to cardiac activity and then uses the remaining components for data reconstruction. It thus retains more useful signal components.

The proposed noise identification approach is based solely on the assumption that time series components fulfilling certain criteria (see section 3.6.1) are primarily related to physiological noise (in this case: cardiac activity). This seems to be a legitimate assumption, since, as Beall et al. state, "[...] ICA can identify components in BOLD-weighted MRI data whose principal temporal behavior is attributable to physiologic noise-either cardiac, respiratory, or an unknown combination of these (Beckmann et al., 2005; De Luca et al., 2006)." [38]
Other approaches based on blind source separation algorithms are often based on stricter assumptions.
CORSICA (CORrection of Structured noise using spatial Independent Component Analysis), which has been developed and reported by Perlbarg et al. in 2007, e.g. does (in opposition to many other approaches) work for short TR (333 ms) and long TR (>1000 ms) scans, however, it requires a priori knowledge of where physiological noise is to be found in the brain: "[...] a set of noise-characteristic signals is defined by extracting time courses from regions that are known to exhibit major physiological fluctuations [...]". [39] This leads - as they also state in said source - to the fact that "[...] the design of the masks of interest seems to be a sensitive part of the procedure because such masks must allow one to select only signals related to noise but not to other processes of interest on the functional point of view."
Other approaches make even stricter assumptions. PESTICA for example (Physiologic EStimation by Temporal Independent Component Analysis) makes the assumption that "[...] linear additive physiologic noise, in the form of respiratory and cardiac-rate related fluctuations, will be sufficiently represented in *one* independent temporal component [...] [emphasis added]". [38]
Therefore, with respect to it's fewer and more general assumptions (not requiring a priori knowledge about the location of noise sources and not limiting the sources to single components only), the proposed approach seems to be a promising alter-

native. Nevertheless, comparing those different approaches with respect to their performance (ability to correctly identify noise, computational efforts, etc.) would be of interest and should addressed in future work.

Analysis of the proposed algorithm's performance (with respect to the signal changes related to noise removal) was approached by comparing the functional connectivity of its output data with a band passed version of the same original data. This was done by means of seed-based correlation analysis.

Results indicated that significantly higher (anti-)correlations are to be found in the band passed data. This is in accordance with what Chang et al. have found: They reported that the removal of signal components related to physiological noise from fMRI data reduces functional connectivity. [15] It is worth mentioning that for the case of the auditory network, seed-based correlation analysis was able to show that the significantly higher (anti-)correlations in the band passed data seem to coincide (location-wise) with the areas covered by the actual resting-state network.
This does, however, not necessarily imply that band passing is superior to the proposed approach, since temporal filtering is known to induce spurious (anti-) correlations. [14] Hence, corrections will need to be implemented to compensate for those artificially induced correlations.

Yet another unsolved problem, which crucially influences the performance of the proposed algorithm, is the choice of the number of components / weights that are to be computed. Beall et al. sum this up as follows: "Reliable identification of the physiologic source can obviously depend on the appropriate choice of the number of sources. Overestimation of the number of sources can lead to a subdivision of the physiologic sources into uninteresting or trivial temporal effects. Likewise, underestimation could lead to incomplete separation of the sources." [38]
This also should be subject to future research.

As already stated in section 3.6.3 the HRV could be computed from photoplethysmography data and compared to the HRV estimated from BOLD signal components, which are identified as being related to cardiac activity. This comparison could potentially improve the proposed algorithms ability to correctly identify BOLD signal components related to cardiac activity.

Finally, the developed algorithm has only been used for a training data set yet and future study data could be used as test data and validation data in order to further evaluate the algorithms performance.

6 References

[1] M.E. Raichle and A.Z. Snyder. A default mode of brain function. *NeuroImage*, 37:1083–1090, 2007.

[2] J.S. Damoiseaux, S. A. R. B. Rombouts, F. Barkhof, P. Scheltens, C.J. Stam, S.M. Smith, and C.F. Beckmann. Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences*, 103(37):13848–13853, 2006.

[3] S. Robinson, G. Basso, N. Soldati, U. Sailer, J. Jovicich, L. Bruzzone, I. Kryspin-Exner, H. Bauer, and E. Moser. A resting state network in the motor control circuit of the basal ganglia. *BMC Neuroscience*, 10(137), 2009.

[4] A. Belouchrani, K. Abed-Meraim, Cardoso J.-F., and E. Moulines. A blind source separation technique using second-order statistics. *IEEE Transactions on Signal Processing*, 45(2):434 – 444, 1997.

[5] C.S. Liu, A. Miki, J. Hulvershorn, L. Bloy, E.E. Gualtieri, G.T. Liu, J.S. Leigh, J.C. Haselgrove, and M.A. Elliott. Spatial and temporal characteristics of physiological noise in fmri at 3t. *Academic Radiology*, 13(3):313–323, 2006.

[6] R.N. Boubela, K. Kalcher, W. Huf, C. Kronnerwetter, P. Filzmoser, and E. Moser. Beyond noise: using temporal ICA to extract meaningful information from high-frequency fMRI signal fluctuations during rest. *Frontiers in Human Neuroscience*, 7(168), 2013.

[7] K. Kalcher, R.N. Boubela, W. Huf, L. Bartova, C. Kronnerwetter, B. Derntl, L. Pezawas, P. Filzmoser, C. Nasel, and E. Moser. The spectral diversity of resting-state fluctuations in the human brain. *PLoS ONE*, 9(4), 2014.

[8] D. Moratal, A. Valls-Luch, L. Mart-Bonmat, and M.E. Brummer. k-Space tutorial: an MRI educational tool for better understanding of k-space. *Biomedical Imaging and Intervention Journal*, 4(1):e15, 2008.

[9] W.D. McRobbie, E.A. Moore, M.J. Graves, and M.R. Prince. *MRI – From Picture to Proton (Second Edition)*. Cambridge University Press, The Edinburgh Building, Cambridge CB2 8RU, UK, 2006.

[10] E. Moser, A. Stadlbauer, C. Windischberger, H.H. Quick, and M.E. Ladd. Magnetic resonance imaging methodology. *European Journal of Nuclear Medicine and Molecular Imaging*, 36(Suppl 1):S30–S41, 2009.

[11] S. Ogawa, T.M. Lee, A.R. Kay, and D.W. Tank. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings of the National Academy of Sciences of the United States of America*, 87:9868–9872, 1990.

[12] F.G. Ashby. *Statistical Analysis of fMRI Data*. The MIT Press, Cambridge, Massachusetts, 02142, 2011.

[13] A. Eloyan, S. Li, J. Muschelli, J.J. Pekar, S.H. Mostofsky, and B.S. Caffo. Analytic programming with fMRI data: A quick-start guide for statisticians using R. *PLoS ONE*, 9(2):e89470, 2014.

[14] C.E. Davey, D.B. Grayden, G.F. Egan, and L.A. Johnston. Filtering induces correlation in fMRI resting state data. *NeuroImage*, 64:728740, 2013.

[15] C. Chang, J.P. Cunningham, and G.H. Glover. Influence of heart rate on the bold signal: the cardiac response function. *Neuroimage*, 44(3):857–869, 2009.

[16] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. John Wiley and Sons, Inc., 605 Third Avenue, New York, NY 10158-0012, 2001.

[17] Information technology – vocabulary. ISO Standard 2382:2015(E), International Organization for Standardization, Geneva, CH, 2015.

[18] A. Hyvärinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, 1999.

[19] I. Narsky and F.C. Porter. *Statistical Analysis Techniques in Particle Physics: Fits, Density Estimation and Supervised Learning*. Wiley-VCH Verlag GmbH & Co. KGaA, 2013.

[20] I.T. Jolliffe. *Principal Component Analysis, Second Edition*. Springer Series in Statistics. Springer-Verlag New-York, Inc., 175 Fifth Avenue, New York, NY 10010, USA, 2002.

[21] A. Ziehe. *Blind Source Separation based on Joint Diagonalization of Matrices with Applications in Biomedical Signal Processing*. PhD thesis, University of Potsdam, 2005.

[22] E.C. Cherry. Some experiments on the recognition of speech, with one and two ears. *The Journal of the Acousitcal Society of America*, 25:975–979, 1953.

[23] J. Ylipaavalniemi. Variability of independent components in functional magnetic resonance imaging. Master's thesis, Helsinki University of Technology, 2005.

[24] J.F. Cardoso and A. Souloumiac. Blind beamforming for non-gaussian signals. *IEE Proceedings F (Radar and Signal Processing)*, 140(6):362–370, 1993.

[25] L. Tong, V.C. Soon, Y.F. Huang, and R. Liu. AMUSE: a new blind identification algorithm. *IEEE International Symposium on Circuits and Systems, 1990*, 3:1784 – 1787, 1990.

[26] J. Miettinen, K. Nordhausen, H. Oja, and S. Taskinen. Deflation-based separation of uncorrelated stationary time series. *Journal of Multivariate Analysis*, 123:214–227, 2014.

[27] R.N. Boubela, K. Kalcher, W. Huf, E.-M. Seidel, B. Derntl, L. Pezawas, C. Nael, and E. Moser. fMRI measurements of amygdala activation are confounded by stimulus correlated signal fluctuation in nearby veins draining distant brain regions. *Scientific Reports*, 5(10499), 2015.

[28] D.A. Feinberg, S. Moeller, S.M. Smith, E. Auerbach, Ramanna S., M.F. Glasser, K.L. Miller, K. Ugurbil, and E. Yacoub. Multiplexed echo planar imaging for sub-second whole brain fMRI and fast diffusion imaging. *PLoS ONE*, 5(12):e15710, 2010.

[29] S.M. Smith, M. Jenkinson, M.W. Woolrich, C.F. Beckmann, T.E.J. Behrens, H. Johansen-Berg, P.R. Bannister, M. De Luca, I. Drobnjak, D.E. Flitney, R.K. Niazy, J. Saunders, J. Vickers, Y. Zhang, N. De Stefano, J.M. Brady, and P.M. Matthews. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage*, 23(1):208–219, 2004.

[30] R.W. Cox. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29(3):162–173, 1996.

[31] A. Weissenbacher, C. Kasess, Lanzenberger R. Gerstl, F., E. Moser, and C. Windischberger. Correlations and anticorrelations in resting-state functional connectivity mri: A quantitative comparison of preprocessing strategies. *NeuroImage*, 47(4):1408–1416, 2009.

[32] W.S. Sohn, K. Yoo, Lee Y.-B., S. W. Seo, D.L. Na, and Y. Jeong. Influence of ROI selection on resting state functional connectivity: an individualized approach for resting state fMRI analysis. *Frontiers in Neuroscience*, 9, 2015.

[33] C.F. Beckmann, M. DeLuca, J.T. Devlin, and S.M. Smith. Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360:10011013, 2005.

[34] C.M. Bennett, A.A. Baird, Miller M.B., and Wolford G.L. Neural correlates of interspecies perspective taking in the post-mortem atlantic salmon: An argument for proper multiple comparisons correction. *Journal of Serendipitous and Unexpected Results*, 1(1):1–5, 2009.

[35] I.M. Veer, C.F. Beckmann, M.J. van Tol, L. Ferrarini, J. Milles, D.J. Veltman, A. Aleman, M.A. van Buchem, N.J. van der Wee, and S.A.R.B. Rombouts. Whole brain resting-state analysis reveals decreased functional connectivity in major depression. *Frontiers in Systems Neuroscience*, 4, 2010.

[36] R.L. Buckner, J.R. Andrews-Hanna, and D.L. Schacter. The brains default network - anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, 1124:1–38, 2008.

[37] K. Kalcher, W. Huf, R.N. Boubela, P. Filzmoser, L. Pezawas, B. Biswal, S. Kasper, E. Moser, and C. Windischberger. Fully exploratory network independent component analysis of the 1000 functional connectomes database. *Frontiers in Human Neuroscience*, 6(301), 2012.

[38] E.B. Beall and M.J. Lowe. Isolating physiologic noise sources with independently determined spatial measures. *NeuroImage*, 37:1286–1300, 2007.

[39] V. Perlbarg, P. Bellec, J.-L. Antond, M. Plgrini-Issac, J. Doyon, and Benali H. Corsica: correction of structured noise in fmri by automatic identification of ica components. *Magnetic Resonance Imaging*, 25(1):35–46, 2007.

Index

Acknowledgments

First and foremost I would like to thank Gerald Badurek for his willingness to supervise my work.
I would also like to thank Ewald Moser for giving me the opportunity to work on my master thesis project at the High Field MR Center, Vienna, and thereby enabling me to gain valuable experiences in such an interesting field of research.

Special thanks also go to Roland Boubela and Klaudius Kalcher who have both spent much of their precious time introducing me to high performance (statistical) computing and fMRI research, answering my countless questions and finally reviewing my work. They have contributed a lot to this work with their ideas and their support. I also attribute it to them that my stay here has been a very pleasant experience.

I am just as well grateful to my brother Patrik Gunacker and my fellow student Lukas Traxler, who have both invested a lot of time explaining to me just about everything I had difficulties with understanding during my years of study.
Thanks also go to my entire family and friends for encouraging me in times where I wasn't as motivated concerning my studies as I should have been.

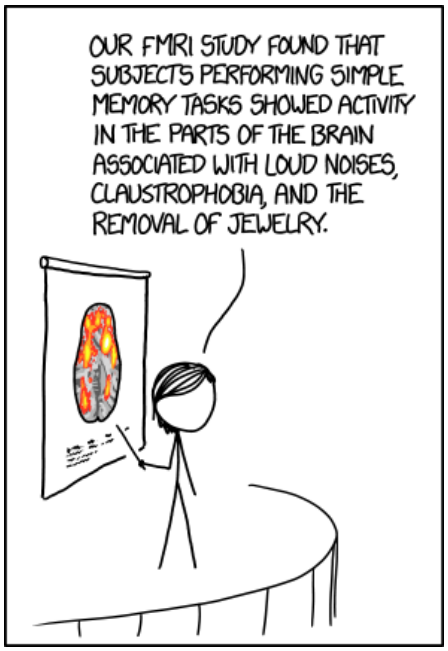

Just for the sake of amusement: An xkcd-comic on fMRI.
(http://xkcd.com/1453/)