

Automatic Venous Vessel Segmentation in High Field, Multi-echo Susceptibility Weighted Imaging

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Biomedical Engineering

eingereicht von

Albert Rechberger, BSc

Matrikelnummer 0626957

an der Fakultät für Physik

der Technischen Universität Wien

Betreuung: **Univ.Prof. Dipl.-Ing. Dr.techn. Gerald Badurek**
Assoc.Prof. Dr. Simon Robinson

Wien, 20. Oktober 2016

Albert Rechberger

Gerald Badurek

Automatic Venous Vessel Segmentation in High Field, Multi-echo Susceptibility Weighted Imaging

DIPLOMA THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Biomedical Engineering

by

Albert Rechberger, BSc

Registration Number 0626957

to the Faculty of Physics

at the Vienna University of Technology

Supervised by: **Univ.Prof. Dipl.-Ing. Dr.techn. Gerald Badurek**
Assoc.Prof. Dr. Simon Robinson

Vienna, 20th October, 2016

Albert Rechberger

Gerald Badurek

Erklärung zur Verfassung der Arbeit

Albert Rechberger, BSc
Wehrgasse 18/10, 1050 Wien

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 20. Oktober 2016

Albert Rechberger

Acknowledgements

At this point I would like to thank everybody who supported me during my work on this thesis. First and foremost I want to thank my supervisors from the Medical University of Vienna, Simon Robinson (HFMRG) and Georg Langs (CIR Lab) for the opportunity to work on this interesting intersection between the fields of machine-learning, medical image processing and high-field MRI research. They continuously provided me with expert advice in their respective fields and supported me in all aspects of this work. I also want to thank Karin Poljanc and Gerald Badurek for the straight-forward and uncomplicated supervision on part of the Vienna University of Technology.

Furthermore I want to thank all my colleagues from the CIR Lab for the great cooperation and support throughout the development and writing of the thesis. Thanks Markus Krenn, Johannes Hofmanninger, Markus Holzer, and Marcus Eger for helping me solve many problems and for making it a comfortable working atmosphere. I'm also grateful for the technical and sometimes mental support I received from my colleagues at the MR center, thank you Korbinian Eckstein, Paul Schennach, Barbara Dymerska, and Pedro Cardoso.

A special thank you goes to my wonderful partner Johanna. Not only for your endless patience and support throughout, but also for your positive nature that helped keeping me motivated.

The biggest thank you goes to my parents Hermann and Johanna and my brother Thomas. Thanks for all your unconditional support and encouragement during the course of my years of studying. I do consider myself indescribably lucky to be part of this family and I am grateful for everything you have done for me.

Abstract

Susceptibility-Weighted Imaging (SWI) is a Magnetic Resonance Imaging (MRI) technique that exploits both the magnitude and phase of the complex MRI signal to increase contrast for tissues of different susceptibilities. Deoxygenated blood in in venous vessels is more paramagnetic than the surrounding tissues, therefore veins can be depicted in SWI without the need for external contrast agents. Identifying and segmenting the venous vessels in whole-brain SWI scans facilitates the creation of three-dimensional models of the cerebral venous vasculature. However, manual segmentation of veins in whole brain SWI datasets is unfeasible due to the amount of manual labor required.

To date automatic segmentation approaches of veins from ultra-high field SWI datasets have predominately been performed using only the magnitude images because of the non-local and orientation dependent properties of the phase. However, in recent years, dedicated algorithms have been established which aim to turn the complex phase information into maps of the local susceptibility, a process that is known as Quantitative Susceptibility Mapping (QSM).

In this project, a new approach to automatic venous vessel segmentation was developed that uses information from magnitude images, phase images and the derived QSM images of a multi-echo T_2^* -weighted gradient echo scan. A Random Forest (RF) classifier was used to segment veins based on a combination of appearance and shape features that are computed separately from magnitude images, phase images and QSM images. This supervised machine-learning approach also allowed us to investigate the importance of each feature for the segmentation task. This not only gives insight to the importance of magnitude, phase and QSM images for venous vessel segmentation, but because the features were computed from multiple echoes, the feature importance findings can also be used to suggest echo time settings for future data acquisition.

The segmentation approach was tested on datasets of five different healthy subjects, two of which were partially annotated to serve as ground truth for training the RF and for quantitatively evaluating the segmentation performance. In all of the performance metrics used within our experiments, the RF approach yielded higher scores than either of those features used individually. Specifically, the RF approach outperformed the common vesselness filtering approach in all similarity measures that were computed against the manual annotations. Visual assessment of 3D renderings of the surface veins confirmed that the segmentations obtained by the RF approach did look very similar to renderings

of the manual annotations. The results of the feature importance measurements indicate that most of the information that is needed for surface vein segmentation is already contained in the first echo, which potentially enables quicker data acquisition.

Overall the developed RF segmentation approach enables the generation high-quality, patient-specific 3D models of the cerebral venous vasculature, which have the potential to aid neurosurgeons in presurgical planning by helping them to localize brain regions that need to be spared in order to minimize the risk of post-operative neurological deficits.

Contents

Abstract	ix
Contents	xi
1 Introduction	1
1.1 Motivation	1
1.2 Related Work	2
1.3 Aims and Hypothesis	4
2 Background: Susceptibility-Based MR Imaging	5
2.1 Principles of Nuclear Magnetic Resonance	5
2.1.1 Magnetic Spins	5
2.1.2 Macroscopic Magnetization	7
2.1.3 Radio-Frequency (RF) Excitation	7
2.1.4 Relaxation	9
Longitudinal Relaxation	9
Transversal Relaxation	10
2.2 Magnetic Resonance Imaging	11
2.2.1 MRI hardware	11
2.2.2 MRI - Signal Localisation	12
Slice Selection	12
Phase and Frequency encoding	12
3D sequences	13
2.2.3 Gradient Echo Imaging	13
2.3 Magnetic Susceptibility and it's effect in MRI	14
2.3.1 Magnetic Field Perturbation of a Sphere and a Cylinder	16
2.3.2 Effects of Magnetic Susceptibility in MRI	18
2.4 Susceptibility Weighted Imaging	20
2.4.1 Phase Unwrapping	20
2.4.2 Combination of Magnitude and Phase	21
2.4.3 Minimum Intensity Projection	21
2.5 (Quantitative) Susceptibility Mapping	22
2.5.1 Basic Concept	22
	xi

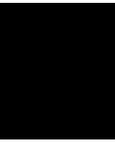
2.5.2	Venous Vessel Appearance in QSM	23
3	Background: Vessel Extraction	25
3.1	Hessian-Based Vessel Enhancement	25
3.1.1	Vesselness Filter (Frangi)	26
	Hessian matrix and Scale Space	26
	Hessian Eigenvalue Analysis	26
	Dissimilarity measures	27
	Frangi Vesselness Function	28
3.1.2	Vessel Enhancing Diffusion	29
	Smoothed Vesselness Filter (Manniesing Vesselness)	29
	Nonlinear, Anisotropic Vessel Enhancing Diffusion	30
3.2	Segmentation	30
3.2.1	Thresholding	30
3.2.2	Classifier	31
3.3	Random Forest	32
3.3.1	The Random Forest Algorithm	32
	Node Splitting and Impurity Measure	33
	Segmentation / Prediction	34
3.3.2	Feature Importance	34
4	Materials and Methods	37
4.1	Acquired Data	37
4.2	Performance Metrics	38
5	Analysis	41
5.1	Susceptibility Based Image Processing	42
5.1.1	Brainmask generation	42
5.1.2	Biasfield Correction	43
5.1.3	Phase Processing and QSM computation	43
5.2	Vessel Extraction	43
5.2.1	Manual Annotation	44
5.2.2	Data normalization	47
5.2.3	Frangi Vesselness	47
5.2.4	Random Forest Segmentation	48
	Features	48
	Training a RF Model	49
	Segmentation with RF model	49
	Validation	50
5.2.5	Feature Importance	51
6	Results	53
6.1	Susceptibility Based Image Processing	53
6.2	Frangi Vesselness Parameter Iteration	55

6.3	Segmentation Performance of RF using all Features (multi-echo)	57
6.3.1	Random Forest Parameter Optimization	57
6.3.2	Quantitative Segmentation performance compared to Manual An- notation	58
6.3.3	Qualitative Segmentation Performance	60
6.4	Feature Ranking	65
6.4.1	Gini Importance	65
6.4.2	Segmentation Performance of Feature Subsets	65
7	Discussion and Conclusion	69
7.1	Susceptibility-based image processing	69
7.2	Vesselness Parameter Iteration	70
7.3	Segmentation Performance of the Multi-Echo RF-Approach	71
7.4	Feature Importance	72
7.5	Conclusion	72
A	Vesselness Parameter Iteration	75
A.1	Dice	76
A.2	Accuracy	78
A.3	Matthews Correlation Coefficient (MMC)	80
A.4	Area under the Curve: Receiver Operating Characteristics	82
	List of Figures	85
	List of Tables	86
	Bibliography	87

Acronyms

AUC	Area under Curve
BET	Brain Extraction Tool
BOLD	Blood Oxygen Level Dependent
CRF	Conditional Random Field
DICE	Dice Similarity Coefficient
DTI	Diffusion Tensor Imaging
GE	Gradient-echo
fMRI	functional Magnetic Resonance Imaging
MMC	Matthews Correlation Coefficient
MIP	Maximum Intensity Projection
mIP	minimum Intensity Projection
MR	Magnetic Resonance
MRF	Markov Random Field
MRI	Magnetic Resonance Imaging
NMR	Nuclear Magnetic Resonance
OOB	Out-Of-Bag
QSM	Quantitative Susceptibility Mapping
RF	Random Forest
ROC	Receiver Operating Characteristic
ROI	Region Of Interest
SNR	Signal to Noise Ratio
SPM	Statistical Parametric Mapping suite
SSS	Superior Sagittal Sinus
SWI	Susceptibility Weighted Imaging
TA	Acquisition Time
TE	Echo Time
TGV	Total Generalized Variation
TR	Repetition Time
VED	Vessel Enhancing Diffusion

This page intentionally left blank.



Introduction

1.1 Motivation

Around 3 or 4 persons per 100,000 are diagnosed with brain tumors annually. For many of those patients, surgery is the best therapeutic option. The main goal of the resulting neurosurgical procedures is to maximize tumor removal while minimizing post-operative neurological deficits. Imaging techniques such as functional Magnetic Resonance Imaging (fMRI) and Diffusion Tensor Imaging (DTI), alongside anatomical MRIs are used to localize essential brain functions prior to surgery and provide insight about the complex organization of the cerebral white matter[1]. Provided with this imaging informations, neurosurgeons are able to plan which tissue needs to be resected and which needs to be spared.

In order to correspond to the intraoperative camera view, images are usually registered to the patient's skull. During surgery, however, the brain moves and changes shape in response to surgical manipulation and anesthesia [2], making it increasingly difficult to apply imaging information as the operation proceeds. Venous vessels on the brain's surface are clearly visible to the neurosurgeon during operation and constitute a mesh which moves and deforms with the underlying tissue. Therefore, using cerebral veins as landmarks to register and relate other imaging information to has the potential to aid neurosurgeons in intraoperative navigation.

With Susceptibility Weighted Imaging (SWI), veins can be depicted in high (sub-millimeter) resolution without the need for external contrast agents [3]. The contrast in SWI originates from susceptibility differences between deoxygenated (venous) blood and the surrounding tissues and the effect of this susceptibility differences on the magnitude and phase of the complex T2*-weighted MRI signal.

In order to create 3D models that visualize the venous vasculature which than can be used for presurgical planning, veins within the image need to be segmented, i.e. voxels

that belong to venous vessels have to be labeled. Due to the density and complexity of the cerebral venous vasculature it is impossible to manually label all veins of a whole-brain 3D SWI datasets within a reasonable time. Accurate automatic segmentation of veins in conventional SWI, however, is challenging for the following reasons:

- SWI is also sensitive to other iron-rich brain structures such as the basal ganglia, which therefore have similar gray level values.
- Surface veins and the surrounding skull both appear hypo-intense in SWI, which makes it challenging to distinguish between the two.
- The changes in phase due to susceptibility differences are dependent on the geometry and orientation to the main magnetic field and phase changes extend beyond areas of altered susceptibility [4].

Recently developed methods of Quantitative Susceptibility Mapping (QSM) aim to overcome orientation-dependence and non-locality of the phase image contrast and create an image that depicts the underlying source of these phase changes: the local magnetic susceptibility [4, 5]. Using QSM images, either instead of conventional magnitude images, or as additional information, might increase the accuracy of automatic venous vessel segmentation in susceptibility-based imaging and enable the generation of reliable maps of surface veins to which neurosurgeons can then reference other imaging information.

1.2 Related Work

Table 1.1 shows a list of publications, which have dealt with automatic venous vessel detection in susceptibility-based imaging. With the exception of Koopmans et. al. [6], all those approaches were performed on images acquired at field strength of 3 T or lower. This study was performed at 7 T, at which field strength susceptibility-differences produce stronger field changes. This leads to more contrast in the phase images, which, due to the orientation-dependent and non-local properties of the phase, can lead to artifacts in fully processed SWIs. [6] For that reason Koopmans et. al. [6] and Jin et. al. [7] use only the magnitude images for vessel extraction instead of fully processed SWIs. The other publications listed in Table 1.1 all use fully processed SWIs.

Deistung et. al. [8], Koopmans et. al. [6], and Jin et. al. [7] use scale-space methods such as Frangi vesselness filtering [12] and Vessel Enhancing Diffusion (VED) [13]. These shape-based methods extract vessels based on a measure of local tubularity¹. While these approaches are particularly robust and produce acceptable results inside the brain, they tend to undersegment veins on the cerebral surface. This is because in SWI, veins and skull both appear hypo-intense and thus the surface veins do not show a fully developed 3D “tube-like” contrast.

¹Frangi vesselness filtering and VED are described in detail in Chapter 3.

Table 1.1: List of publications that have performed automatic vein segmentation in susceptibility-based MR imaging. In addition, information about the imaging method, the segmentation approach, and (if available) the reported evaluation scores is given.

Author / Year	MRI Method	Approach	Field Strength	Evaluation	Additional Information
Deistung 2006 [8]	SWI	Frangi vesselness filtering	1.5 T	only visual	Only Abstract
Koopmans 2008 [6]	SWI (Mag. only)	Vessel Enhancing Diffusion (VED)	3 T, 7 T	Visual Comparison with Manual Annotation	
Jin 2010 [7]	SWI (Mag. only)	Frangi vesselness filtering	3 T	only visual	Very similar to Koopmans 2008
Bériault 2014 [9]	SWI	Expectation Maximization + Markov Random Field (MRF)	3 T	DICE on different ROIs: subcortical & deep veins=0.85, surface=0.81	Quant. comparison with Frangi: subcort.& deep=0.79, surf.=0.48
Ward 2015 [10]	SWI + QSM	Gaussian mixture model + novel MRF that uses Gabor filters	3 T	DICE =0.52 compared to manual annot. of SWI	Only Abstract; reported DICE for Frangi on QSM=0.27, Frangi on SWI=0.46
Bériault 2015 [11]	SWI	Conditional Random Fields, using appearance, shape and location information	3 T	DICE: subcortical & deep = 0.83, mid sag.& surface= 0.82	Quant. comparison with Frangi& spherical flux

In [9], Bériault et. al. performed venous vessel segmentation from SWI with a statistical approach that is based on voxel appearance. An initial segmentation is formed by estimating the intensity distributions of veins and background voxels using an expectation maximization algorithm. This initial segmentation is then smoothed using an anisotropic Markov Random Field (MRF) model that integrates spatial dependencies between neighboring voxels. As final step a skull stripping procedure is computed to distinguish between surface veins and skull. Bériault et. al. show that this approach is able to segment veins on the cortical surface [9]. However, they also mention that this method is not able to distinguish veins from other structures that appear hypo-intense in SWI, such as basal ganglia.

In a recently (Dec. 2015) published paper [11], Bériault et. al. describe an enhanced method for venous vessel segmentation from SWI. The approach uses a Conditional Random Field (CRF) model that combines appearance, shape, and location information. This sophisticated segmentation approach reportedly performs well both inside the brain and on surface veins. One limitation that Bériault et. al mention in [11] is that they work with standard SWI reconstruction that is implemented scanner. Thereby the phase and magnitude information get merged together. Also the individual echoes (5) get averaged

to increase the signal to noise ratio. They mention that processing magnitude and phase separately and treating the intensities specific to each echo as a vector of observations rather than a signal average, does provide extra information which could potentially be useful for the generation of segmentations that better describe the venous vasculature. [11]

1.3 Aims and Hypothesis

In this thesis, a new approach to automatic venous vessel segmentation from ultra-high field susceptibility-based MR imaging was developed and tested. Instead of segmenting veins on the basis of conventionally processed SWI images, or using only the magnitude image, a Random Forest [14] classifier is used to segment veins from a combination of appearance and shape features that are computed separately from magnitude images, phase images, and the derived QSM images of a multi-echo T_2^* -weighted Gradient Echo scan.

The performance of segmentations obtained with the multi-echo Random Forest model is compared with the performance of segmentations obtained by using only a single image (magnitude, phase, or QSM) of a particular echo as basis for segmentation. Special emphasis is put on veins of the cerebral surface, which are especially challenging to segment in conventional SWI, but, as mentioned earlier, have great potential to serve as landmarks in neurosurgical procedures.

The Random Forest algorithm also has a built-in measure of feature importance, which is used to investigate the importance of the magnitude, phase, and QSM features for venous vessel segmentation. By using features that are computed from each echo of a multi-echo sequence, we are further able to analyze which echo provides the most amount of information for vessel segmentation and whether a multi-echo approach is beneficial.

Background: Susceptibility-Based MR Imaging

Magnetic susceptibility is an intrinsic material property which can vary between different types of tissues. With specialized Magnetic Resonance Imaging (MRI) acquisition and post-processing techniques those susceptibility differences can be visualized. Since deoxygenated blood has higher susceptibility values than its surrounding tissues, such susceptibility-based imaging methods allow the depiction of venous vessels without the need for external contrast agents.

In this chapter, the fundamentals of susceptibility-based magnetic resonance imaging are discussed. First, the physical principles behind Nuclear Magnetic Resonance (NMR) are explained in Section 2.1 and Section 2.2 introduces the hardware and sequence design in Magnetic Resonance Imaging (MRI). Section 2.3 then focuses on magnetic susceptibility and its effect in MRI. Finally, Sections 2.4 and 2.5 introduce two susceptibility-based MR techniques: Susceptibility Weighted Imaging (SWI) and Quantitative Susceptibility Mapping (QSM).

2.1 Principles of Nuclear Magnetic Resonance

2.1.1 Magnetic Spins

MRI makes use of the magnetic dipole moment of nucleons, which arises from their intrinsic angular momentum, the so called spin. The relation between spin \mathbf{L} and magnetic moment $\boldsymbol{\mu}$ for elemental particles is described by:

$$\boldsymbol{\mu} = \gamma \mathbf{L} \tag{2.1}$$

with γ being the gyromagnetic ratio ($\gamma \approx 267.513 \times 10^6 \text{ rad s}^{-1} \text{ T}^{-1}$ for protons¹).

When placed inside an external magnetic field \mathbf{B}_0 , the nuclear spins precess around the field lines with an angular frequency that is proportional to $|\mathbf{B}_0|$ and the gyromagnetic ratio γ :

$$\omega_0 = -\gamma |\mathbf{B}_0|. \quad (2.2)$$

This precession frequency ω_0 is also referred to as Larmor frequency. The interaction of the magnetic moment with the external field also yields a magnetic energy, which is determined by the orientation of $\boldsymbol{\mu}$ to the magnetic field:

$$E = -\boldsymbol{\mu} \cdot \mathbf{B}_0 \quad (2.3)$$

In classical physics $\boldsymbol{\mu}$ and \mathbf{B}_0 could align along any orientation. Eq. (2.3) could therefore be written as

$$E = -|\boldsymbol{\mu}| |\mathbf{B}_0| \cos \theta \quad (2.4)$$

where θ is the angle between $\boldsymbol{\mu}$ and \mathbf{B}_0 . The magnetic energy would thus be at a minimum when the spins are aligned with \mathbf{B}_0 ($\theta = 0$), and at a maximum when $\boldsymbol{\mu}$ and \mathbf{B}_0 are anti-parallel ($\theta = 180^\circ$).

However, according to quantum mechanics, both, the magnitude and the orientation of \mathbf{L} to an external field are quantized. That means only certain values are allowed:

$$|\mathbf{L}| = \hbar \sqrt{I(I+1)} \quad I = 0, \frac{1}{2}, 1, \frac{3}{2}, \dots \quad (2.5)$$

$$L_z = \hbar m \quad m = -I, -I+1, \dots, I-1, I \quad (2.6)$$

where \hbar is equal to $h/2\pi$, with h being the Planck's constant ($\approx 6.626 \times 10^{-34} \text{ J s}$) and I and m being quantum numbers. The spin quantum number I is particle-specific, and is equal to $\frac{1}{2}$ for protons, neutrons, and electrons.

With no external magnetic field applied, the spatial orientation of the spin is indiscriminate, whereas in an external magnetic field the magnetic moment (and the spin with it) can only be at a number of $2I + 1$ discrete angles with the field direction. Assuming an external magnetic field in z-direction, the z-component of the magnetic moment therefore is (with Eqs. (2.1) and (2.6)):

$$\mu_z = \gamma \hbar m. \quad (2.7)$$

The magnetic energy of a nucleon in an external magnetic field \mathbf{B}_0 (Eq. (2.3)) then becomes

$$E = -\gamma \hbar m |\mathbf{B}_0| \quad m = -I, -I+1, \dots, I-1, I \quad (2.8)$$

leaving only discrete energy levels, which depend on the quantum number m . For particles with $I = \frac{1}{2}$ (i.e. protons, electrons, neutrons), this allows only two possible energy levels

¹Due to the great abundance of water molecules in biological tissues, the proton in the hydrogen atom has become the essential nucleon in MRI.

$m = \pm\frac{1}{2}$, with $m = +\frac{1}{2}$ being the low energy state (also called *parallel* or α state), and $m = -\frac{1}{2}$ the high energy (*anti-parallel* or β) state. The energy difference ΔE between the two states is increasing with the field strength (Zeeman effect):

$$\Delta E = \gamma\hbar |\mathbf{B}_0|. \quad (2.9)$$

2.1.2 Macroscopic Magnetization

Within the \mathbf{B}_0 magnetic field, the whole spin ensemble will tend towards a state of minimal total energy. Spin states with lower energy will therefore be favored compared to those of higher energy. The spin population distribution is described by Boltzmann statistics. For a nucleus with I equal $\frac{1}{2}$ (e.g. a proton) the ratio of the two energy states is described by

$$\frac{n_{1/2}}{n_{-1/2}} = \exp\left(\frac{\Delta E}{k_B T}\right) = \exp\left(\frac{\gamma\hbar |\mathbf{B}_0|}{k_B T}\right) = \exp\left(\frac{\hbar\omega_0}{k_B T}\right) \quad (2.10)$$

where $n_{\pm 1/2}$ is the number of spins in the $+\frac{1}{2}$, and $-\frac{1}{2}$ state respectively, ΔE is the energy difference between the two states (2.9), k_B the Boltzmann constant ($\approx 1.381 \times 10^{-23} \text{ J K}^{-1}$), T the absolute temperature, and ω_0 the Larmor (or resonance) frequency (2.2). With a temperature of 310.15 K (human body temperature) and a magnetic field of 7 T, the relative population difference between the two states amounts only 0.0023%. This means in an ensemble of 1 million protons, only 23 more will be in the low energy state than in the high energy state.

The net magnetic moment \mathbf{M} of a macroscopic sample (e.g. a voxel) consists of the vector sum of the individual magnetic moments $\boldsymbol{\mu}$ within the sample. \mathbf{M} can be characterized by a magnitude and a phase, as illustrated in Figure 2.1. The population difference between the spin states leads to a net magnetization the z-direction, the direction of \mathbf{B}_0 (Fig. 2.2).

At thermal equilibrium, the spins have no phase coherence in the transverse plane, therefore the net component in the x-y plane is 0. The signal can, as such, not be detected by regular coils. In order to observe nuclear magnetization, the precessional motion needs to be detected. Excitation pulses are used to generate a component in the transverse plane.

2.1.3 Radio-Frequency (RF) Excitation

In a stationary external magnetic field \mathbf{B}_0 the macroscopic magnetization is static and points in the direction of \mathbf{B}_0 (see Fig. 2.2). For imaging it necessary to flip the magnetization towards the transverse plane (xy-direction). This is done by temporarily applying an additional electromagnetic field (RF-pulse) in the transverse plane, which is oscillating in the radio frequency (MHz) range. During the pulse, protons absorb energy from the pulse at a particular frequency, the resonance frequency. A protons resonance

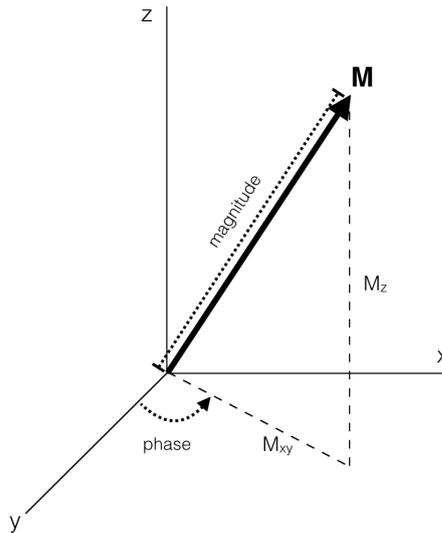


Figure 2.1: Magnitude and phase of the magnetization vector. M_z is the longitudinal magnetization (parallel to \mathbf{B}_0) and M_{xy} is the transverse magnetization (perpendicular to \mathbf{B}_0)

frequency ω is defined by the Larmor equation

$$\omega = \gamma |\mathbf{B}| \quad (2.11)$$

where γ is the protons gyromagnetic ratio and \mathbf{B} is the magnetic field experienced by the proton. Only protons that meet the resonance criteria, i.e. spin with the same frequency as the RF-pulse, respond to the RF-pulse.

Application of the RF-pulse has two effects on the spins. First, by absorbing energy, protons are excited from the low energy state to the high energy state. The longitudinal part M_z of the macroscopic net magnetization therefore decreases as the two energy states become more evenly populated and eventually points in the opposite direction. Second, the external field forces the spins into a state of phase coherence, which leads to a transverse magnetization M_{xy} . Both effects together lead to a spiral movement of the macroscopic net magnetization towards the xy-plane. In a rotating frame of reference, the net magnetization vector tilts down during excitation. The amount of tilt is determined by the amplitude and duration (t) of the RF-pulse.

$$\alpha(t) = \gamma B_1 t. \quad (2.12)$$

A RF-pulse with a duration and amplitude such that the equilibrium magnetization is tilted entirely into the transverse plane is known as 90° pulse ($\alpha = 90^\circ$). A 180° pulse inverts the initial longitudinal magnetization. Figure 2.3 illustrates the effect of a 90° pulse on the spins and the macroscopic magnetization.

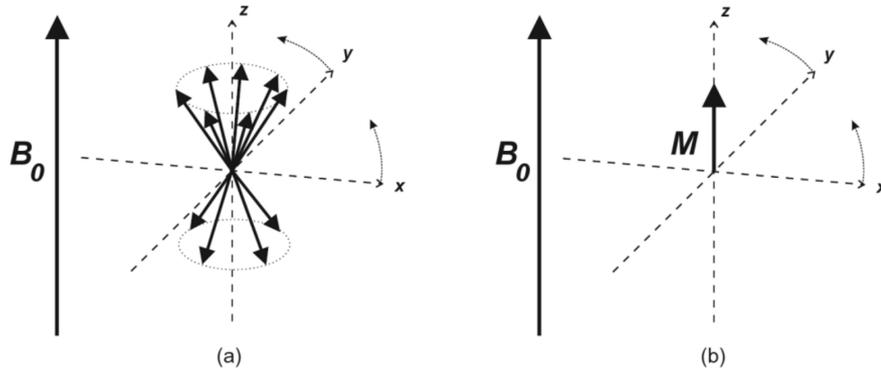


Figure 2.2: Microscopic (a) and macroscopic (b) pictures of a collection of spins in the presence of a stationary external magnetic field \mathbf{B}_0 . Each individual spin precesses around the magnetic field. In a frame of reference that is rotating with ω_0 around the z-axis, the magnetic moments appear stationary. The x-y-components of the spins are randomly distributed, whereas the z component is only one of two values. There are more spins in the +z direction than in the -z direction, which results in a nonzero net magnetic vector \mathbf{M} in the direction of \mathbf{B}_0 . (Image reproduced from Ref .[15])

To conclude, the excitation pulse tilts the net magnetic vector away from the equilibrium state towards the transverse plane. The resulting nonzero transverse component rotates around the z-axis at Larmor frequency. This time-dependent magnetic motion induces an electrical current in the receiver coils, and can thereby be detected.

2.1.4 Relaxation

After excitation by the RF-pulse, the spin system gradually returns to a state of thermal equilibrium. This process is known as relaxation and there are two independent mechanisms that occur simultaneously: longitudinal relaxation and transversal relaxation.

Longitudinal Relaxation

During excitation, the spin system absorbed energy from the RF-pulse, which led to an increased population of the high-energy configuration and thereby a decreased longitudinal magnetization M_z . Longitudinal relaxation refers to the release of that added energy to the protons surroundings, the *lattice*. The return of M_z to the equilibrium value M_0 follows an exponential growth process

$$M_z(t) = M_0 \left(1 - \exp\left(\frac{-t}{T_1}\right) \right) \quad (2.13)$$

where t is the time following the RF pulse and T_1 is the characteristic time constant for longitudinal relaxation. T_1 values in biological tissues range from a few hundred milliseconds to a few seconds and are specific for a tissue or fluid. [17]

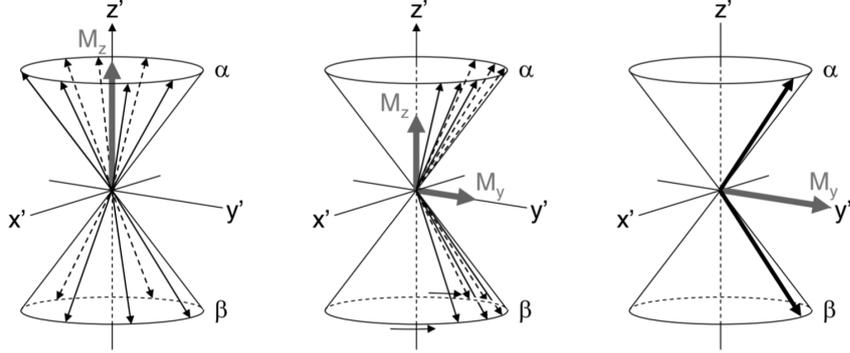


Figure 2.3: Effect of a 90° pulse on a spin ensemble within a static field \mathbf{B}_0 in z direction; shown in a rotating frame of reference. *left:* before the pulse: Spins are at thermal equilibrium, with no phase coherence; net magnetization in $+z$ direction. *middle:* additional field B_1 along x' rotates generates phase coherence between individual spins and thereby rotates the net magnetization towards y' . *right:* endpoint of the 90° pulse: equal distribution in the spin states and complete phase coherence resulting in a net magnetization along y' . (Image reproduced from [16])

Transversal Relaxation

Following the RF pulse, the contributing spins precess in phase, thereby generating a net transverse magnetization. However, stochastic fluctuations of the magnetic field at a microscopic level due to *spin-spin* interaction leads to slightly altered precession rates for individual spins. This causes the spin isochromats to fan out in time, reducing the transverse component of the net magnetization vector M_{xy} . This irreversible decay of M_{xy} by an exponential function with a characteristic time T_2 , the spin-spin relaxation time [17]:

$$M_{xy}(t) = M_{xy}(t=0) \exp\left(\frac{-t}{T_2}\right). \quad (2.14)$$

Additionally to spin-spin relaxation, there are other effects that contribute to the decay of M_{xy} . Inhomogeneities of the main magnetic field (e.g. due to imperfect shimming) and varying magnetic properties² within the tissue cause variations in the precession speeds, which again fans out the spins. The transverse magnetization therefore decays faster than predicted by pure T_2 relaxation which results in an effective transverse relaxation time T_2^* :

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \quad (2.15)$$

²The field inhomogeneities induced by magnetic susceptibility differences are the root cause of the image contrast in susceptibility-based imaging techniques and will be discussed in more detail in the following sections.

where T_2' denotes the time constant for inhomogeneity-induced relaxation.[17, 18]

While a signal loss due to intrinsic T_2 relaxation cannot be avoided, signal loss due to field inhomogeneities can be reversed by applying a second RF pulse known as refocusing pulse (a 180° pulse) [15]. However, as we will discuss in the following sections, the accelerated decay of the signal due to local field inhomogeneities is actually desirable for some imaging purposes such as venous vessel detection. Susceptibility-based MR imaging techniques therefore use a gradient echo sequence, which does not have a refocusing pulse.

2.2 Magnetic Resonance Imaging

Up to now, we have discussed the physical principles of nuclear magnetic resonance; how spins are excited to form a measurable net magnetic vector and the main relaxation processes. In order to generate Magnetic Resonance Images (MRIs) however, we also need to have spatial information, i.e. where in the body which signal comes from. This necessitates spatial encoding of signal. In the following, the main constituents of an MRI scanner are presented in Section 2.2.1 and the basic principles behind spatial encoding are discussed in Section 2.2.2. Then the Gradient Echo sequence, which is used for susceptibility-based image acquisition, is presented in Section 2.2.3.

2.2.1 MRI hardware

The main constituents of an MRI scanner are the main magnet that provides a large static magnetic field, B_0 , gradient coils, and the radiofrequency coils that transmit the RF pulse and receive the NMR signal.

Most MRI systems use superconducting electromagnets to generate the main field. These consist of a coil that has been made superconductive by cooling it to 4 K using liquid helium. The resulting field strength of B_0 ranges up to 3 T in clinical practice and up to 15 T in research. A higher field strength leads to a higher Signal to Noise Ratio (SNR), which can be used for higher spatial resolution or decreased scan time. [18]

The gradient coils enable the generation of linear gradient fields in all three spatial dimensions: x,y and z. These gradients are used for the spatial encoding and localization of the MRI signal, which is discussed in Section 2.2.2.

Transmit coils produce the radio-frequency electromagnetic field B_1 that is perpendicular to the B_0 field. Transmit coils are usually volume coils that are able to deliver uniform excitation throughout the scanned volume. They are either located inside the bore of the magnet (for whole body scans up to 3 T) or closer to the body part that is to be measured in the form of specialized head, knee or ankle coils. The transmit coils can also be used to receive the MR signal. Dedicated arrays of receiver coils, however, are able to produce better signal to noise ratios.

2.2.2 MRI - Signal Localisation

To localize the signal, spatial information needs to be encoded into the NMR signal. This is done with the help of three sets of gradient coils, which are able to provide linear magnetic field gradients along the x, y and z axis. In the following a brief description of the way these gradients are used to encode spatial information in a prototypical 2D imaging sequence is provided. More in depth explanations can be found in one of the many excellent textbooks such as [19, 15, 20, 21].

Slice Selection As a first step (in 2D imaging), the excitation of the spins is limited to a thin slice. This is done by applying a gradient simultaneously with the Radiofrequency Pulse (RF-Pulse). The gradient G_z (assuming an axial slice selection) alters the precession frequency in z-direction:

$$\omega(z) = \gamma B_0 + \gamma G_z \quad (2.16)$$

The RF-pulse only excites spins which match the resonance criteria, i.e. only spins that are located in defined positions along the z-axis are excited. The slice thickness depends on the amplitude of G_z and the bandwidth BW of the RF-pulse [20]:

$$\Delta z = \frac{BW}{\gamma G_z}. \quad (2.17)$$

Phase and Frequency encoding Subsequent to exciting a specific slice, the spatial information along the x and y direction needs to be encoded. Similar to (2.16), the linear gradients field gradients G_x and G_y can alter the precession frequency of the spins along the x and y axis respectively, thus they can establish a relationship between the position of the spins and their precession frequency:

$$\omega(\mathbf{r}) = \gamma (B_0 + \mathbf{G} \cdot \mathbf{r}). \quad (2.18)$$

This process is called frequency encoding. The MR signal induced in a coil $S(t)$ is the sum of all the excited spins in the volume, each resonating at a frequency corresponding to it's position along the gradient directions [21]:

$$S(t) = \int M_{xy}(\mathbf{r}) \exp(-i\Phi(\mathbf{r}, t)) d\mathbf{r} \quad (2.19)$$

with the phase angle

$$\Phi(\mathbf{r}, t) = \int \omega(\mathbf{r}) dt \quad (2.20)$$

representing the accumulated phase changes. For the sake of simplicity (2.19) does not consider relaxation phenomena. If we demodulate the signal (i.e. remove the rapid oscillation caused by B_0), the accumulated phase angle due to gradient imposition can be written as

$$\Phi(\mathbf{r}, t) = \gamma \int \mathbf{G} \cdot \mathbf{r} dt \quad (2.21)$$

With the definition

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int \mathbf{G}(t) dt \quad (2.22)$$

(2.19) writes as [21]:

$$S(t) = \int M_{xy}(\mathbf{r}) \exp(-i2\pi\mathbf{k}(t) \cdot \mathbf{r}) d\mathbf{r} \quad (2.23)$$

This equation shows that the position variable \mathbf{r} and the “spatial frequency” $\mathbf{k}(t)$ are a Fourier pair and the signal we measure is the Fourier transform of the transverse magnetization $M_{xy}(\mathbf{r})$.

Thus after enough data points in the two dimensional k-space are acquired by manipulating the gradients G_x and G_y , the image can be obtained via 2D Fourier transformation as illustrated in Fig. 2.4. Various different strategies for traversing the k-space have been developed; a good overview can be found in [21]. In Section 2.2.3 the gradient echo sequence, which is used in susceptibility-based imaging is described.

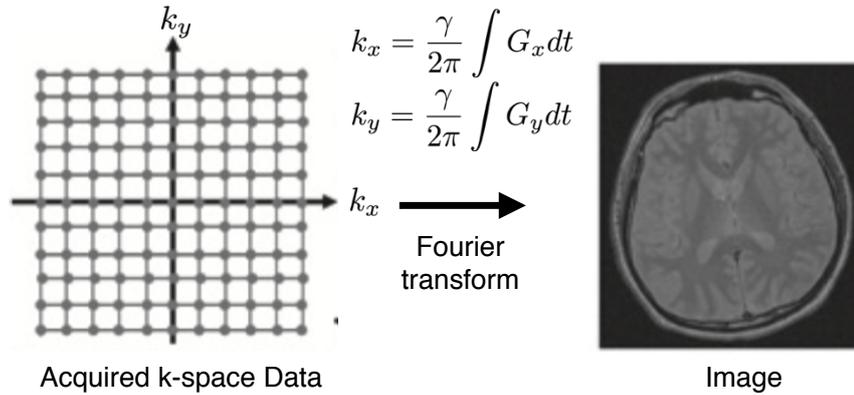


Figure 2.4: Image Acquisition process in 2D. Datapoints in k-space are acquired by altering G_x and G_y . A Fourier transform yields the desired image (magnitude and phase). (Modified from [20])

3D sequences To obtain information in the third dimension, one can either repeat the 2D acquisition at shifted slice positions or use a 3D sequence. In 3D sequences a larger volume or slab is excited and the gradient in the third dimension is used to fill a three dimensional k-space (similar to the 2D). A 3D Fourier transform then yields the images.

2.2.3 Gradient Echo Imaging

In susceptibility-based imaging, high resolution 3D Gradient-echo (GE) sequences are typically used for data acquisition[22]. GE sequences utilize time-varying gradient fields to dephase and rephase the MR signal in such a way that one or multiple echo signals can be created[22]. Figure 2.5 shows a generic 2D single-echo GE sequence. Within each

cycle ($TR =$ repetition time), the following steps occur: First, the slice select gradient G_S together with the radiofrequency pulse select a particular slice. Then the phase encoding gradient G_P selects a particular k_y in k-space (a different one for each repetition) as described in Section 2.2.2. Simultaneously, the negative readout gradient G_R dephases the spins along the x direction, which puts us to the left side of the k-space. Then the data readout starts, which is indicated in Fig. 2.5 by the Analog Digital Converter (ADC) being active. During the readout, G_R is active in the opposite direction. This traverses the k-space in positive k_x direction by first rephasing the spins along x, eventually creating an echo (at time interval TE after the RF-pulse), and then dephasing them again. Thus within each cycle one line in a 2D k-space gets filled.

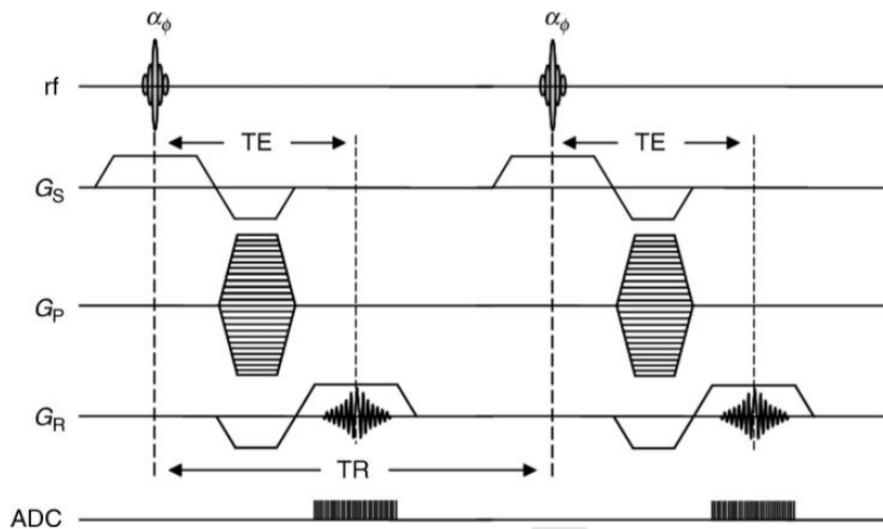


Figure 2.5: Pulse timing of a generic 2D Gradient Echo Sequence. (Reproduced from [22])

In practice, the generic GE-sequence described above is altered to form more sophisticated sequences that fit the particular imaging needs. By repeatedly reversing the readout gradient, for example, the echo creation process can be repeated multiple times during one TR . Thereby multiple images, which differ in TE , can be formed. For 3D data acquisition, the RF-pulse and G_S are adjusted to excite a thicker slab and G_S additionally performs phase encoding steps similar to G_P . For more details on GE-sequence design, including spoiling and flow-compensation, the interested reader is referred to textbook sources such as [18, 21].

2.3 Magnetic Susceptibility and its effect in MRI

As mentioned in Section 2.1.2, materials become magnetized in the presence of an external magnetic fields. The amount (and direction) of this magnetization depends on

the magnetic properties of the material. Magnetic Susceptibility χ is the measure which relates the induced magnetization to the applied field. In most materials (including most tissues in the body) the magnetization \mathbf{M} depends linearly on the applied field \mathbf{H} [15]:

$$\mathbf{M} = \chi \mathbf{H}. \quad (2.24)$$

Based on their susceptibility values, materials may be broadly classified into three large groups [23]:

Diamagnetic materials are materials with negative susceptibility values $-1.0 < \chi < 0$.

The diamagnetic response is present in all materials and is usually very small. It is the only magnetic response for most biological materials including water ($\chi = -9.05 \times 10^{-6}$) and almost all human tissue.

Paramagnetic materials have components with unpaired electrons which results in a magnetic moment that overrides the diamagnetic response. Prominent examples for paramagnetic components are iron, magnesium and gadolinium. Most important for vessel detection in susceptibility-based imaging is the fact that deoxygenated haemoglobin is paramagnetic while oxygenated haemoglobin is diamagnetic. The magnetic susceptibility of blood therefore is a function of the oxygenation level. This effect is called **BOLD effect** (Blood Oxygen Level Dependent).

Ferromagnetic materials have large susceptibility values $\chi \gg 1$ and may exhibit magnetization even in the absence of an applied field. Eq. (2.24) does not apply for those materials.

Due to the magnetization \mathbf{M} , the actual magnetic field \mathbf{B} inside a material within an external magnetic field $\mathbf{B}_0 = \mu_0 \mathbf{H}_0$ is given by

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}) \quad (2.25)$$

where \mathbf{B} is measured in Tesla (T), \mathbf{H} is measured in Ampere/meter (A/m) and μ_0 is the absolute permeability of free space ($\mu_0 = 4\pi \times 10^{-7}$ T m/A) [18].

With Eq. (2.24), equation (2.25) can be written as

$$\mathbf{B} = \mu_0 (1 + \chi) \mathbf{H} = \left(\frac{1 + \chi}{\chi} \right) \mu_0 \mathbf{M}. \quad (2.26)$$

The actual magnetic field inside a substance therefore is a function of the substance's susceptibility.

Furthermore, the field distribution outside the substance will also be perturbed by the induced magnetization. The magnetic field \mathbf{B}_{out} at a position \mathbf{r} outside the object is a function of the induced magnetization \mathbf{M} , and the object's shape, volume, and orientation to B_0 [18]:

$$\mathbf{B}_{\text{out}}(\mathbf{r}) = \mathbf{B}_0 + f(\mathbf{M}, \text{object shape, object size, } \mathbf{r}). \quad (2.27)$$

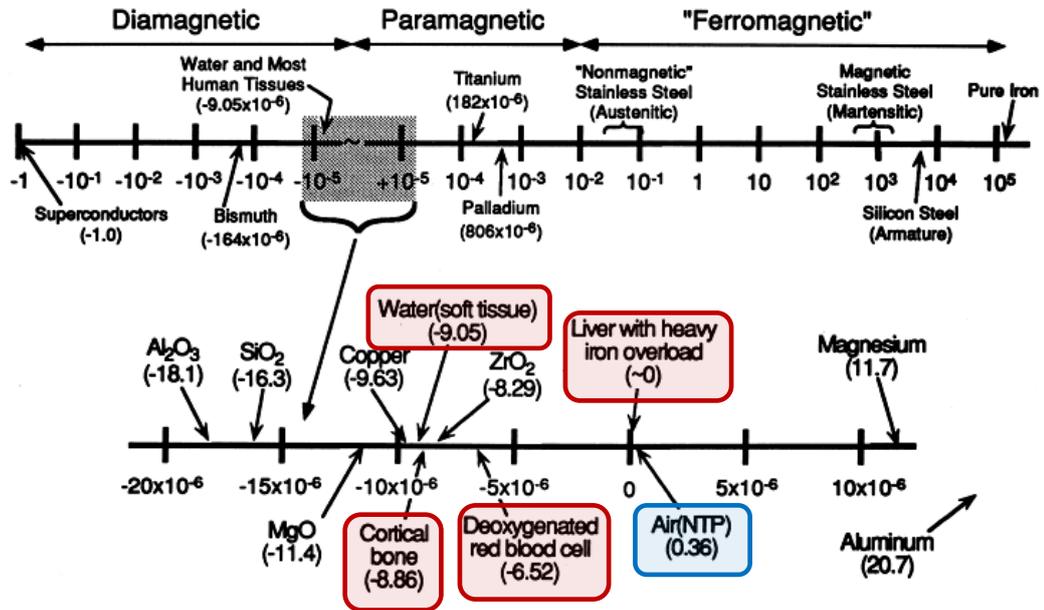


Figure 2.6: Magnetic susceptibility values of different materials and tissues. The upper diagram uses a logarithmic scale to visualize the broad range of observed susceptibility values. The bottom diagram shows the 10^{-6} range at linear scale. Susceptibility values of human tissue (marked red) are mostly in the range from -7.0 to -11.0 ppm. Notice that deoxygenated blood is more paramagnetic compared to water and other soft tissue. (Figure modified from [23])

Now, why are the local magnetic field changes that are caused by susceptibility differences of interest? Because those field changes alter the precession frequency of the protons in MRI (see Larmor equation: (2.11)). In order to understand the effects of susceptibility differences on MRI, Section 2.3.1 describes susceptibility induced field perturbations for some simple geometries and Section 2.3.2 describes the effects on the phase and magnitude of the MR signal.

2.3.1 Magnetic Field Perturbation of a Sphere and a Cylinder

To understand the nature of the field perturbation that we can expect in MRI, it is instructive to look at field perturbation solutions for some simple geometries. Table 2.1 shows expressions for the internal and external magnetic field of a sphere and an infinitely long cylinder in vacuum. In those equations, \mathbf{r} denotes the position vector of the point of observation and a the radius of the cylinder or sphere. For the spherical case θ is the angle that the position vector makes with the magnetic field \mathbf{B}_0 . In the cylindrical case θ represents the angle between the long axis of the cylinder and \mathbf{B}_0 , and φ and ρ denote the position vector in cylindrical coordinates. [18]

Table 2.1: Equations for magnetic field inside and outside a sphere and an infinitely long cylinder in vacuum. These equations assume susceptibility values of $\chi \ll 1$ which is the case for most human tissue. [18]

	Internal Field	External Field
Sphere	\mathbf{B}_0	$\mathbf{B}_0 + \frac{\chi \mathbf{B}_0}{3} \cdot \frac{a^3}{ \mathbf{r} ^3} \cdot (3 \cos^2 \theta - 1)$
Cylinder	$\mathbf{B}_0 + \frac{\chi \mathbf{B}_0}{6} \cdot (3 \cos^2 \theta - 1)$	$\mathbf{B}_0 + \frac{\chi \mathbf{B}_0}{2} \cdot \frac{a^2}{\rho^2} \cdot \sin^2 \theta \cos 2\varphi$

When embedded in some external medium, the susceptibility term χ in the equations in Table 2.1 is replaced by a term representing the susceptibility differences of the materials $\Delta\chi = \chi_{\text{sphere or cylinder}} - \chi_{\text{outside material}}$ and an additional field is added which is dependent on the shape of the outside compartment [18].

Blood vessels can be modeled as cylindrical structures, and, in case of venous vessels, their susceptibility differs from the surrounding tissue (see Fig 2.6). It is therefore particularly important for vein detection purposes to understand the field perturbations for cylindrical structures. Figure 2.7 illustrates the field perturbations for a cylindrical structure oriented parallel and perpendicular to the magnetic field \mathbf{B}_0 . As we will see and discuss in the next sections, these images look very similar to the phase images close to venous vessels.

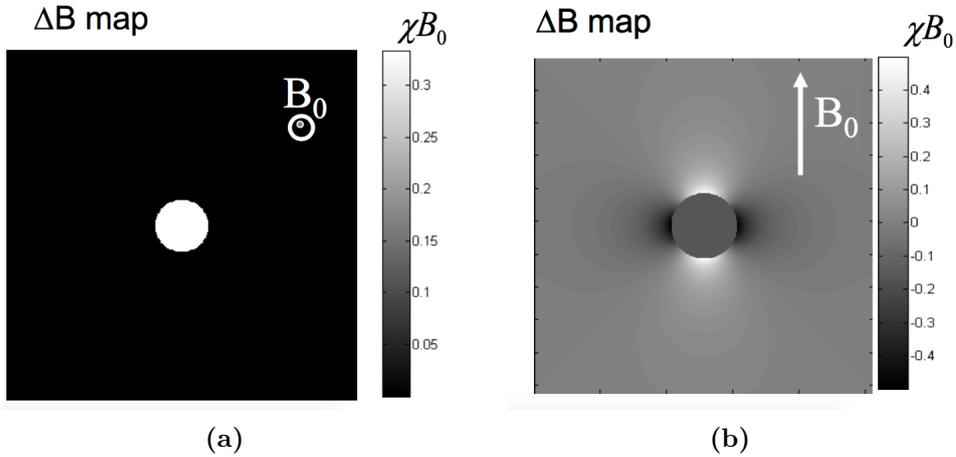


Figure 2.7: Simulated magnetic field perturbation of an infinite cylinder with finite susceptibility oriented (a) parallel, and (b) perpendicular to the magnetic field. It is clearly visible that the orientation to \mathbf{B}_0 has a strong effect on the field. For the non-parallel case the susceptibility effects also alter field around the object (see equations in Table 2.1). (Image reproduced from Ref. [24])

2.3.2 Effects of Magnetic Susceptibility in MRI

The signal received by the MR scanner is a complex vector that represents the transverse component of the magnetic moment. The signal of each voxel can be written as

$$S(\mathbf{r}, t) = M(\mathbf{r}, T_E) \cdot e^{-i\Phi(\mathbf{r}, T_E)} \quad (2.28)$$

where $M(\mathbf{r}, t)$ is the magnitude, and $\Phi(\mathbf{r}, T_E)$ is the phase of the signal at position \mathbf{r} and at the time T_E ³.

In an idealized imaging scenario, all protons would experience the same magnetic field \mathbf{B}_0 , and would therefore precess at the same rate. Each voxel's signal would thus have identical phase

$$\Phi(\mathbf{r}) = \omega \cdot T_E \quad (2.29)$$

with ω being the Larmor precession rate (2.2).

However, as discussed in the previous section, there are changes in the local magnetic field resulting from susceptibility differences between tissues. There are also other causes for field inhomogeneities, such as imperfect gradient performance, object motion or air-tissue interfaces. [18] For venous vessel imaging purposes, the susceptibility effects are of interest.

The local field changes affect the magnitude and the phase of the MR signal. If the field changes rapidly, i.e. the field varies significantly within a voxel, the spins within the voxel

³ T_E corresponds to the echo time in a gradient echo sequence, which represents the time interval between the RF pulse and readout of the maximum gradient echo signal.

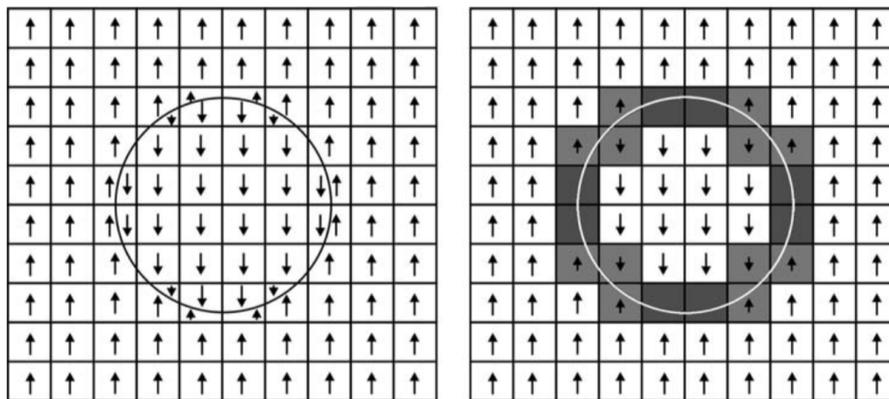


Figure 2.8: Illustration of the field inhomogeneity and spin dephasing effects to be expected near the interface of a vein and the surrounding tissue. The arrows represent the direction of the transverse magnetization vector. In this example the vein is 180° out of phase with the surrounding tissue. The voxels at the boundaries experience both field resulting in the dephasing of the spins within the voxel and thus an attenuation of the signals magnitude (represented by the voxel appearing darker). (Image reproduced from Ref. [18])

dephase and the magnitude of the signal gets gets attenuated. [18] For field variations that are of lower spatial frequency, the field within a voxel can still be regarded as being constant. The precession frequency of the entire spin ensemble within the voxel however will be altered. This becomes apparent in the signal's phase. The phase of two voxels experiencing different, but constant local magnetic fields differs by: [18]

$$\Delta\Phi = -\gamma \cdot \Delta B_0 \cdot T_E. \quad (2.30)$$

Figure 2.8 shows a schematic illustration of the effects of susceptibility-induced field changes on the phase and magnitude of a gradient echo signal.

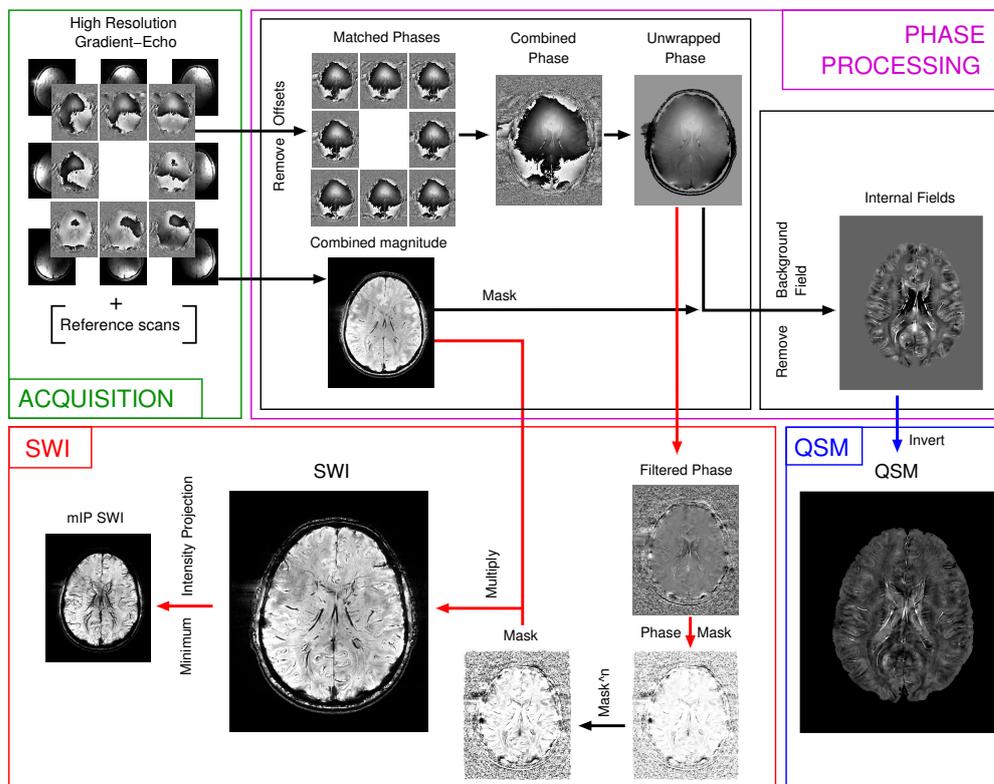


Figure 2.9: Processing steps to form a Susceptibility Weighted Image (SWI): The data from a high resolution Gradient-Echo sequence is combined to form magnitude and phase images. The phase image is unwrapped and high-pass filtered. Then a phase mask is generated which is multiplied into the magnitude image (a number of times) to enhance the magnitude's contrast between tissues with different susceptibility. The processing steps for Quantitative Susceptibility Mapping (QSM) are also shown. QSM will be discussed in Section 2.5. (Courtesy of Simon Robinson)

2.4 Susceptibility Weighted Imaging

As discussed in the previous sections, both the magnitude and the phase of gradient echo scans are sensitive to susceptibility. By using the phase information in combination with the magnitude image, one can generate a susceptibility weighted image (SWI) that shows enhanced contrast between tissues with different susceptibilities. [25, 18] Figure 2.9 gives an overview of the processing steps typically used to form a SWI. In the following, the processing steps for generating a susceptibility weighted image are described in more detail.

2.4.1 Phase Unwrapping

The phase signal, as it is measured ($\Phi_{measured}$), is limited to values ranging from 0 to 2π . So whenever the actual phase value (Φ_{act}) lies outside of that range, it still gets represented by a value within the interval. That process is called phase wrapping or aliasing and leads sharp borderlines appearing in the original phase image (Figure 2.10a). In order to obtain the real range of the phase variations, the phase needs to be unwrapped, which is equivalent to adding multiples of 2π at the appropriate positions:

$$\Phi_{act}(\mathbf{r}) = \Phi_{measured}(\mathbf{r}) + 2\pi \cdot n(\mathbf{r}) \quad (2.31)$$

where n is an integer (positive or negative).

The original SWI approaches [27, 25] did not explicitly unwrap the phase, but instead used a homodyne filter [28] to remove phase variations of low spatial frequencies. Homodyne

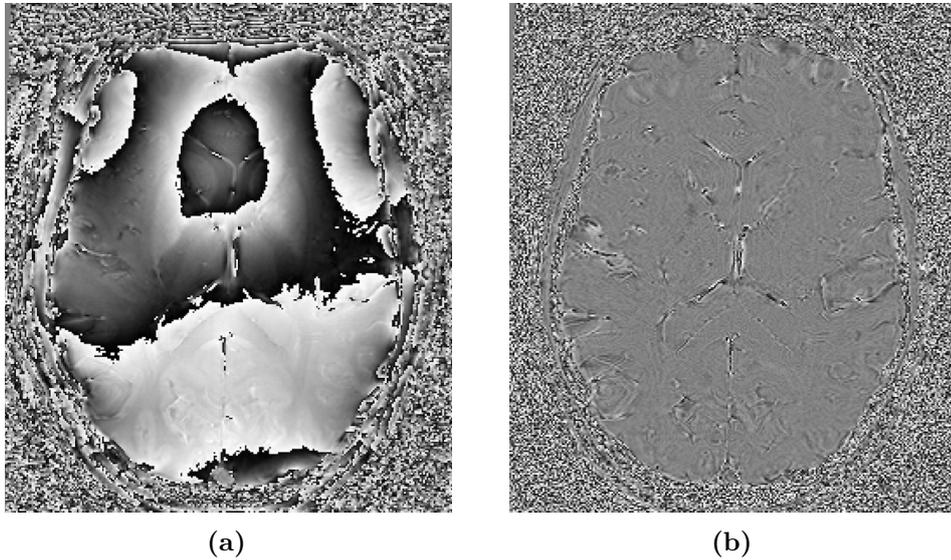


Figure 2.10: (a) Raw phase image showing phase wraps and (b) corresponding unwrapped phase image. Unwrapping was performed with a Laplacian-based algorithm that is implemented in TGV-QSM [26].

filtering is performed by smoothing (low-pass filtering) the complex image and dividing the smoothed version into the original complex signal to create a high-pass filter effect. The obtained phase image only contains high frequency spatial components, which are associated with susceptibility changes, while the low frequency components (background field inhomogeneities) are filtered out. While homodyne filtering alone reduces the amount of phase wraps, it does not eliminate them completely. [29]

An alternative to homodyne filtering is unwrapping the true phase by aiming to determine $n(\mathbf{r})$ in equation (2.31). Several unwrapping techniques have been successfully applied to MR Phase data, such as Φ UN [30], PRELUDE [31] and Laplacian-based algorithms [32]. Figure 2.10b shows the unwrapped version of Figure 2.10a. The unwrapped phase images contains all spatial frequencies. However, in SWI, only the local field changes that result from susceptibility changes are of interest. To remove contributions from background field inhomogeneities, the unwrapped phase image can be high-pass filtered.

2.4.2 Combination of Magnitude and Phase

After the phase image has been unwrapped and high-pass filtered, a phase mask is generated which is then combined with the corresponding magnitude image by voxel-wise multiplication. The phase mask can be designed to enhance positive phase changes (Fig. 2.11a), negative phase changes (Fig. 2.11b), or both (Fig. 2.11c). In order to enhance positive phase changes, for example, all voxels with $\Phi < 0$ are set to unity, while the voxels with positive phase values are scaled linearly from 1 to 0:

$$\Phi_{\text{masked}} = \begin{cases} 1 - \frac{\Phi}{\Phi_{\text{max}}}, & \text{if } \Phi > 0 \\ 1, & \text{otherwise.} \end{cases} \quad (2.32)$$

The phase mask can be multiplied with the magnitude a number of times, each time further reducing the magnitude signal in voxels with positive phase.

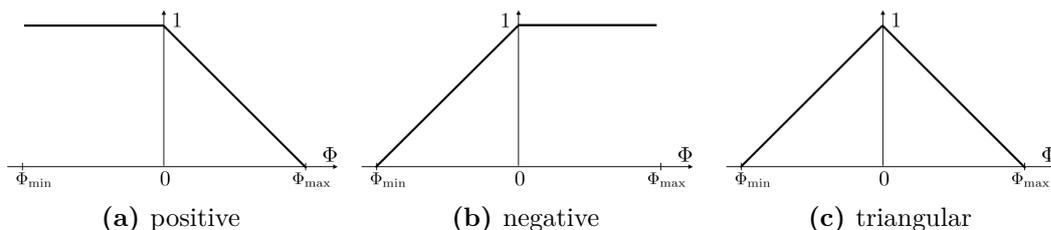


Figure 2.11: Functions for creating the phase mask. (a) enhances positive phase changes, (b) enhances negative phase changes, and (c) enhances both.

2.4.3 Minimum Intensity Projection

To visualize the venous vasculature in SWI, a minimum intensity projection (mIP) from the fully processed SWI is typically used. [18] A mIP combines multiple adjacent slices into one image by taking the minimum value at every position and displaying it. This

method for vein visualization relies on the assumption that the vein has smaller values than any other point along the projection. [18] While this assumption is often true for slices through the middle of the brain, mIP visualization of veins near the brain’s surface is impractical because of the low magnitude values of voxels outside of the brain.

2.5 (Quantitative) Susceptibility Mapping

As discussed in Section 2.3 the phase images depict the field changes resulting from the susceptibility distribution inside a sample. These field changes, however, can be spatially distant to the source of susceptibility changes and depend on the geometry and orientation to the main field [5]. Susceptibility mapping aims to map these field changes back to their source: the local susceptibility distribution. The basic concept behind (quantitative) susceptibility mapping is described in the following:

2.5.1 Basic Concept

Assuming a main magnetic field in z-direction, and susceptibility values $\chi \ll 1$, the magnetic field perturbation due the susceptibility distribution $\chi(\mathbf{r})$, can be written as

$$\Delta B_z(\mathbf{r}) = B_0 \cdot \chi(\mathbf{r}) * G(\mathbf{r}) \quad (2.33)$$

where $*$ denotes a convolution and $G(\mathbf{r})$ is the point dipole response function given by:

$$G(\mathbf{r}) = \frac{1}{4\pi} \frac{3 \cos^2 \theta - 1}{r^3} \quad (2.34)$$

with θ being the angle between \mathbf{r} and \mathbf{z} . [5] Using the convolution theorem in Fourier space (2.33) simplifies to

$$\mathcal{F}[B_z(\mathbf{r})] = B_0 \cdot \mathcal{F}[\chi(\mathbf{r})] \cdot \mathcal{F}[G(\mathbf{r})]. \quad (2.35)$$

The Fourier transform of $G(\mathbf{r})$ is given by [5]

$$\mathcal{F}[G(\mathbf{r})] = G(\mathbf{k}) = \begin{cases} \frac{1}{3} - \frac{k_z^2}{|\mathbf{k}|^2}, & \text{for } k \neq 0 \\ 0, & \text{for } k = 0 \end{cases} \quad (2.36)$$

where k_x , k_y and k_z are the coordinates in k-space and $|\mathbf{k}|^2 = k_x^2 + k_y^2 + k_z^2$, and the Fourier transform of the susceptibility distribution $\chi(\mathbf{r})$ is $\chi(\mathbf{k})$. The magnetic field perturbation ΔB_z is then found by computing the inverse Fourier transform of (2.35):

$$\Delta B_z(\mathbf{r}) = B_0 \cdot \mathcal{F}^{-1}[G(\mathbf{k}) \cdot \chi(\mathbf{k})] \quad (2.37)$$

To compute susceptibility maps from the phase images, (2.37) needs to be solved as an inverse problem. However, due to zero values of $G(\mathbf{k})$ along the magic angles, some spatial frequencies are under-sampled, which means the inverse problem is ill-posed [5]. Several algorithms have been proposed for solving the ill-posed inverse problem, a review can be found in [5].

2.5.2 Venous Vessel Appearance in QSM

Figure 2.12 shows the appearance of venous vessels in phase images and QSM images. In the phase images, the dipole effect is clearly visible with the phase appearing bright above and below the vein, and appearing dark to the left and right. Using the conventional SWI processing scheme, would incorrectly enhance the outside of veins perpendicular to the field [18]. Susceptibility mapping removes the dipolar appearance and shows the veins as continuous bright structures. There are however some streak artifacts in the QSM image which appear along the magic angles.

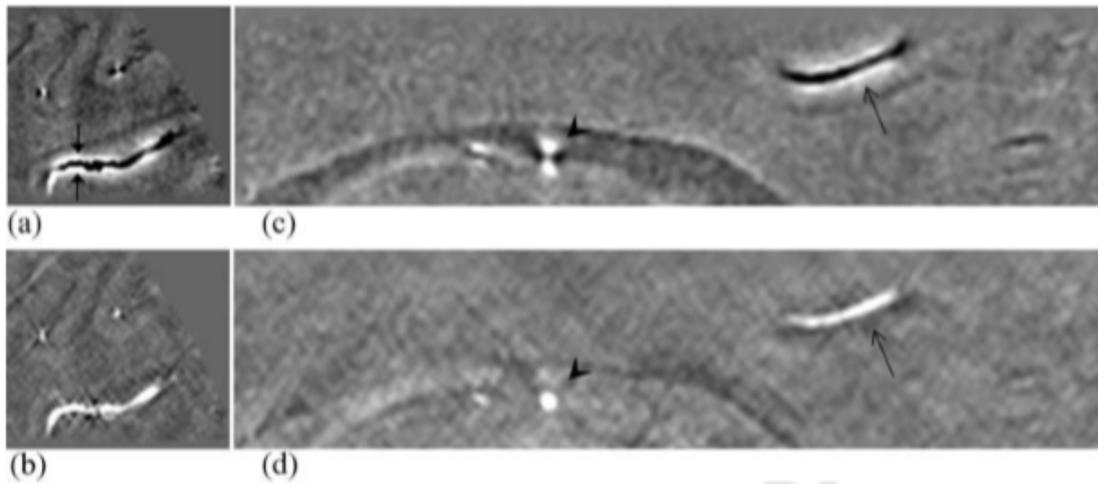


Figure 2.12: Comparison of vessel appearance in phase images (a) and (b) vs QSM images (c) and (d). (Image reproduced from [18, ch.25])

This page intentionally left blank.

Background: Vessel Extraction

In this chapter, the basics for the vessel extraction and segmentation are laid out. In Section 3.1, two popular vessel enhancing techniques, Frangi vesselness filtering and Vessel Enhancing Diffusion are described. Then Section 3.2 defines the term segmentation and discusses methods to perform segmentation. Finally, Section 3.3 deals with the Random Forest (RF) classifier and explains how it can be applied for segmentation and feature importance measurements.

3.1 Hessian-Based Vessel Enhancement

This section describes two methods which are often used for vessel detection and have also been used in susceptibility based images (See Table 1.1). The two methods are the multi-scale vessel enhancing filter developed by Frangi et al. [12], and the Vessel Enhancing Diffusion (VED) filter by Manniesing et al. [13]. Both are built upon the assumption that in 3D, vessels have a cylindrical structure, and use concepts of scale space theory to search for such structures within images [12, 13].

Frangi's multi-scale vessel enhancing filter examines the local structure of an image by analyzing the Hessian matrix of each voxel at multiple scales. A vesselness measure is obtained that represents the likelihood of each voxel to be part of a tubular (vessel-like) structure. Vessel Enhancing Diffusion [13] utilizes the vesselness measure to steer an anisotropic diffusion scheme, which aims at improving vein continuity and suppressing non-vessel structures.

In the following, the methods from Frangi et al. and Manniesing et al. are described in detail.

3.1.1 Vesselness Filter (Frangi)

Hessian matrix and Scale Space

The Hessian matrix is a matrix that contains the second order partial derivatives in x , y and z direction. Before we can go into detail about the analysis of the Hessian for vessel detection, we need to define differentiation in an imaging context.

In scale space theory the (normalized) differentiation of an image I at position \mathbf{r} is defined as a convolution with derivatives of Gaussians [12]:

$$\frac{\partial}{\partial x} I(\mathbf{r}, \sigma) = \sigma^\gamma I(\mathbf{r}) * \frac{\partial}{\partial x} G(\mathbf{r}, \sigma) \quad (3.1)$$

where the D -dimensional Gaussian is defined as [12]:

$$G(\mathbf{r}, \sigma) = \frac{1}{\sqrt{(2\pi\sigma^2)^D}} e^{-\frac{\|\mathbf{r}\|^2}{2\sigma^2}} \quad (3.2)$$

with standard deviation σ . The size of the standard deviation determines the scale of the features that are being looked for. σ^γ in (3.1) is a normalization factor that is necessary to allow a fair comparison of the features at multiple scales. When no scale is preferred γ is set to unity [12].

The Hessian therefore is defined as a convolution of the original image with the the partial second derivatives of a Gaussian. For a 3D image $I(\mathbf{r})$ the Hessian (of scale σ) reads as

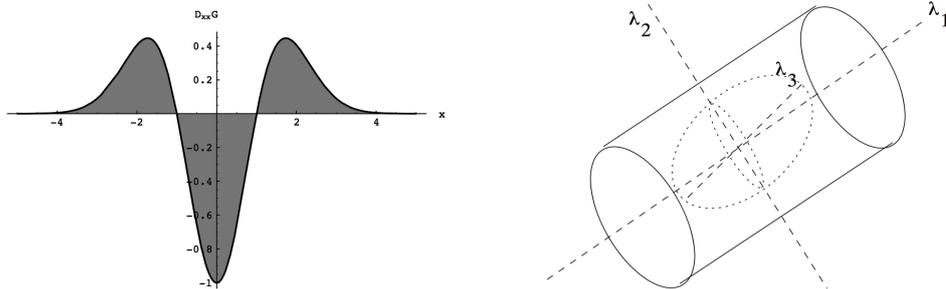
$$H(\mathbf{r}, \sigma) = \begin{pmatrix} I_{xx} & I_{xy} & I_{xz} \\ I_{yx} & I_{yy} & I_{yz} \\ I_{zx} & I_{zy} & I_{zz} \end{pmatrix} \quad (3.3)$$

with the partial second derivatives $I_{xy} = \frac{\partial^2}{\partial_x \partial_y} I(\mathbf{r}, \sigma)$ being calculated as defined in (3.1).

Hessian Eigenvalue Analysis

The second derivative of a Gaussian kernel at position \mathbf{r} at scale σ generates a kernel that measures the contrast around \mathbf{r} in the range $(-\sigma, \sigma)$, as illustrated in Figure 3.1a. A subsequent eigenvalue analysis of the Hessian gives the principal directions (eigenvectors $\hat{\mathbf{v}}_i$) in which the local second order structure can be decomposed. The three orthogonal eigenvectors, scaled by their eigenvalues (λ_j) span an ellipsoid that describes the 2nd order structure around a specific voxel (see Figure 3.1b). The shape of this ellipsoid, more specifically the relations between the eigenvalues (see Table 3.1), allows one to determine the geometrical structure that the voxel belongs to.

In 3D, an idealized vessel has a tubular structure. A voxel belonging to a vessel region will, therefore, be signaled by λ_1 being small and λ_2 and λ_3 of a large magnitude and equal sign (positive sign for dark vessels on bright background and negative sign for



(a) Second order derivative of a Gaussian kernel that measures the contrast inside/outside of the range $(-\sigma, \sigma)$ (here $\sigma = 1$).

(b) Second order ellipsoid that describes the local principal directions of the curvature embedded in a tubular (vessel-like) structure.

Figure 3.1: Illustrations relating to Vesselness Filtering (reproduced from [12]).

bright vessels):

$$|\lambda_1| \approx 0 \quad (3.4)$$

$$|\lambda_1| \ll |\lambda_2| \quad (3.5)$$

$$\lambda_2 \approx \lambda_3. \quad (3.6)$$

The eigenvector $\hat{\mathbf{v}}_1$ will point in the direction of vessel.

Dissimilarity measures

To calculate the likelihood for each voxel to be part of a tubular (vessel-like) structure Frangi et al. use three dissimilarity measures and combined them in a probability-like estimation of vesselness.

λ_1	λ_2	λ_3	Structure
L	L	L	noisy, no preferred structure
L	L	H-	plate like structure (bright)
L	L	H+	plate like structure (dark)
L	H-	H-	tubular structure (bright)
L	H+	H+	tubular structure (dark)
H-	H-	H-	blob-like structure (bright)
H+	H+	H+	blob-like structure (dark)

Table 3.1: Eigenvalues of the Hessian matrix and corresponding 3D image structure [12]. The eigenvalues are ordered $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$. (L=low absolute value, H=high absolute value, +/- indicate the sign of the eigenvalue.)

The first dissimilarity measure differentiates between plate and line like structures:

$$A = \frac{|\lambda_2|}{|\lambda_3|}. \quad (3.7)$$

A will be 1 for line-like structures and go to zero for a plate-like pattern. Only values between 0 and 1 are possible because λ_3 will always be greater than λ_2 or equal in magnitude (they are sorted that way).

The second dissimilarity measure is a measure of the deviation from blob-like structures:

$$B = \frac{|\lambda_1|}{\sqrt{|\lambda_2||\lambda_3|}} \quad (3.8)$$

This ratio attains its maximum when all three eigenvalues are about equal in magnitude (blob like structures) and vanishes when $\lambda_1 \approx 0$.

In an idealized (noiseless) scenario the features A and B would be enough to measure vessel-like structures, but in reality, medical images contain random noise fluctuations. To prevent these fluctuations from producing unpredictable filter responses, a third dissimilarity measure is needed. A distinguishing property of background pixels is that the magnitude of the derivatives (and therefore the eigenvalues) is small [12]. Therefore Frangi et al. proposed the use of the (Frobenius) norm of the Hessian as the third dissimilarity measure:

$$S = \|H\|_F = \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2} \quad (3.9)$$

In the background, where no structure is present, S will be low. For regions, which are part of a structure (see Table 3.1), at least one eigenvalue will be large. Therefore S will be higher in those regions than in background regions. While the first two measures, A and B , are grey-level invariant (i.e. they remain constant under intensity re-scaling), S is sensitive to the absolute intensity of the image. [12]

Frangi Vesselness Function

All three measures (Eqs. (3.7) to (3.9)) are combined to define a single vesselness function. In case of dark vessels on bright background the (Frangi-) vesselness is defined as:

$$\mathcal{V}_O(\sigma) = \begin{cases} 0, & \text{for } \lambda_1 < 0 \text{ or } \lambda_2 < 0 \\ \left(1 - e^{-\frac{A^2}{2\alpha^2}}\right) e^{-\frac{B^2}{2\beta^2}} \left(1 - e^{-\frac{S^2}{2\gamma^2}}\right), & \text{otherwise} \end{cases} \quad (3.10)$$

where σ represents the current scale, and α , β and γ are parameters to adjust the influence of each individual measure A , B and S (Eqs. (3.7) to (3.9)).

To detect all vessels, small and big, the vesselness response is calculated at multiple scales σ by computing the Hessian with Gaussian derivatives at multiple scales. At every

voxel location, the vesselness output with the highest response is selected to obtain a final estimate of vesselness:

$$\mathcal{V}_F = \max_{\sigma_{\min} \leq \sigma \leq \sigma_{\max}} \mathcal{V}_O(\sigma) \quad (3.11)$$

where σ_{\min} and σ_{\max} are the maximum and minimum scales at which relevant structures are expected to be found. [12]

3.1.2 Vessel Enhancing Diffusion

Vessel Enhancing Diffusion (VED), described by Manniesing et al. in [13], is an iterative process that combines a smoothed vesselness filter with a nonlinear anisotropic diffusion scheme. During each iteration the vesselness measure is calculated. Then a diffusion scheme of varying (an)isotropy is applied. Based on the response of the vesselness filter, the diffusion tensor is constructed to achieve:

- *isotropic diffusion* for voxels with low vessel likelihood to reduce background noise and
- strong *anisotropic diffusion* for vessel structures *in the direction of the vessel* to improve vein continuity.

Smoothed Vesselness Filter (Manniesing Vesselness)

VED uses a vesselness filter that is based on Frangi's vesselness filter (see Section 3.1.1). However, the vesselness function defined in (3.10) is not smooth at the origin and can thus not be used directly to construct a vesselness diffusion equation. [13] To cope with that limitation, Manniesing et al. extended (3.10) with a smoothing term that resembles a Gaussian function with its argument inverted [13]:

$$e^{-\frac{2c^2}{|\lambda_2|\lambda_3^2}} \quad (3.12)$$

The parameter c should be chosen small to only have influence around the origin [13].

Multiplying $\mathcal{V}_O(\sigma)$ with the smoothness term results in a smoothed vesselness function

$$\mathcal{V}_S(\sigma) = \begin{cases} 0, & \text{for } \lambda_1 \leq 0 \text{ or } \lambda_2 \leq 0 \\ \left(1 - e^{-\frac{A^2}{2\alpha^2}}\right) e^{-\frac{B^2}{2\beta^2}} \left(1 - e^{-\frac{S^2}{2\gamma^2}}\right) e^{-\frac{2c^2}{|\lambda_2|\lambda_3^2}}, & \text{otherwise} \end{cases} \quad (3.13)$$

with A , B and S being the dissimilarity measures defined in Eqs. (3.7) to (3.9) and α , β and γ their respective weighting factors. Note that the domain definition has changed: $\mathcal{V}_M(\sigma)$ is set to zero for for $\lambda_{\{2,3\}}$ less or equal to zero oppose to $\mathcal{V}_O(\sigma) = 0$ for $\lambda_{\{2,3\}} < 0$.

Similar to (3.10) a multiscale approach is achieved by taking the maximum response of all scales for each voxel:

$$\mathcal{V}_M = \max_{\sigma_{\min} \leq \sigma \leq \sigma_{\max}} \mathcal{V}_S(\sigma) \quad (3.14)$$

Nonlinear, Anisotropic Vessel Enhancing Diffusion

In scale space theory, different scale versions of an image can be generated according to the diffusion equation

$$I_t = \nabla(D\nabla I) \quad (3.15)$$

where I is the original image and D is a diffusion tensor that controls the amount and direction of the blurring or smoothing.[13]

In order to guide the diffusion process to enhance vessel like structures, Manniesing et al. defined the diffusion tensor D as follows

$$D = Q\Lambda'Q^T \quad (3.16)$$

with Q being a matrix that represents the eigenvectors of the Hessian $Q = [\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \hat{\mathbf{v}}_3]$ and Λ' a diagonal matrix that contains the following elements:

$$\lambda'_1 = 1 + (\omega - 1) \cdot \mathcal{V}_M^{(\frac{1}{s})} \quad (3.17)$$

$$\lambda'_2 = \lambda'_3 = 1 + (\epsilon - 1) \cdot \mathcal{V}_M^{(\frac{1}{s})} \quad (3.18)$$

The parameters ω and ϵ control the anisotropic diffusion for vessel structures and the parameter s can be used to control the sensitivity to the vesselness response. ω should be a large value (may be larger than one) as it steers the diffusion in vessel direction. ϵ should be chosen very small, but $\epsilon > 0$ to ensure positive definiteness of the tensor. For non-vessel structures ($\mathcal{V}_M \rightarrow 0$), the diffusion is high and isotropic and background noise is reduced.

3.2 Segmentation

Image segmentation can be defined as the process of extracting one or more objects of interest from an image [33]. For the objective of this thesis, we are interested in distinguishing venous vessels (*Foreground*) from the rest of the image (*Background*). The goal is thus to assign a discrete label $l(\mathbf{x})$ to each voxel \mathbf{x} in the following way

$$l(\mathbf{x}) = \begin{cases} 0, & \text{if background} \\ 1, & \text{if venous vessel.} \end{cases} \quad (3.19)$$

3.2.1 Thresholding

The simplest way do such a binary segmentation is by setting a threshold T and assigning the label according to the voxel intensity $I(\mathbf{x})$ in relation to T :

$$l(\mathbf{x}) = \begin{cases} 1, & \text{if } I(\mathbf{x}) < \text{ or } > T \\ 0, & \text{otherwise.} \end{cases} \quad (3.20)$$

The threshold T is often determined manually by interactively varying T until the result is considered satisfactory by the human operator. A well-known method for automatically determining T is Otsu’s method [34], which chooses T such that the intra-class variance is minimal.

In SWI, veins appear hypo-intense compared to their surrounding. A quick way to segment veins could therefore be thresholding the magnitude image or the processed SWI. There are however other structures that appear hypo-intense, such as the iron-rich basal ganglia or the skull surrounding the brain. Grey-value based thresholding on the magnitude image or SWI would falsely classify those as veins.

A more accurate approach is to threshold vessel-enhanced versions of the images, which are produced by applying one of the methods described in Section 3.1. This approach has been used for venous vessel segmentation from fully processed SWI [35] and magnitude images [8, 6]. However, these Hessian-based filters are based on idealized geometrical assumption and tend to undersegment the vasculature when those geometrical assumptions are not met.

3.2.2 Classifier

Segmentation can also be treated as a voxel-wise classification problem. Each voxel is then represented as a point in an p -dimensional feature space with p being the number of features that are used to describe the voxel. Such features can for example be the voxel intensity or, as also used in this thesis, the vesselness measure that is computed for this voxel (see Section 3.1.1). The features space can also be derived from multiple images, like multi-echo MR-sequences, multiple MR-sequences or differently processed versions of an image. Figure 3.2a shows an example of a 2 dimensional feature space.

Classifiers partitions the feature space into sub-populations by using data with known labels. [36] Figure 3.2b illustrates this partitioning. New, i.e. unsegmented, voxels are then labeled according to their position in the partitioned feature space.

Classifier-based segmentation pipelines are, in machine-learning terminology, supervised methods, because they require training data that has been manually segmented and is then used as reference for automatically segmenting new data. [36]

A wide range of classification algorithms can be used for segmentation tasks [36], some popular examples are K-nearest neighbors (KNN) [37], Support Vector Machines (SVM) [38] and Artificial Neural Networks. In this thesis, a Random Forest (RF) [14] classifier is used, which constructs a multitude of decision trees and segments the image based on the cumulative votes of the individual trees. A more detailed description of Random Forests can be found in Section 3.3.

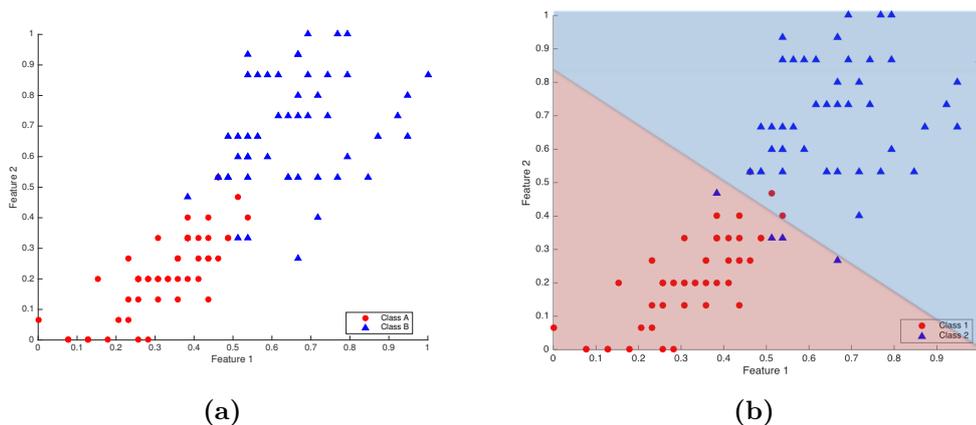


Figure 3.2: Illustrative example of the classification procedure. (a) Training dataset in a 2-dimensional feature space. For a segmentation task each point in this space would represent a voxel, and the color would represent the class label (e.g. red circle = vein, blue triangle = background). (b) shows the partitioning of the feature space after the classifier has been trained. New samples (unsegmented voxels) are classified according to this partitioning.

3.3 Random Forest

Random Forest (RF) is an ensemble classifier that was originally described by Leo Breiman and Adele Cutler in [14]. It is based on constructing multiple decision trees (a ‘forest’) from randomly selected subsamples of the initial training data. Classification is then performed by aggregating the individual tree’s vote and deciding for the statistical mode. The averaging over multiple decision trees and the introduction of randomness in the tree building process (see Section 3.3.1) improves the predictive accuracy and controls over-fitting to the training data, i.e. it does not excessively adapt to the training and fail to classify new data [14, 39].

Additionally, the RF algorithm provides a built-in measurement of the ‘importance’ of each feature for the classification [14]. Details on the variable importance measurement are discussed in Section 3.3.2.

3.3.1 The Random Forest Algorithm

In the following the key concepts the RF algorithm are provided with a focus on the application for segmentation tasks. A more general descriptions of the RF algorithm can be found in [14, 39].

Random forests is a supervised learning algorithm. This means that for training the algorithm, a labeled training set needs to be available. The training set consists of a feature matrix $\mathbf{X}_{\text{Train}} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\} \in \mathbb{R}^{n \times p}$ where each of the n training voxels is represented in the p -dimensional feature space, and the corresponding output set

$\mathbf{y}_{\text{Train}} = \{y_1, y_2, \dots, y_n\}$. In the segmentation case $\mathbf{y}_{\text{Train}}$ contains the manually annotated class labels of the training set.

From the provided training data, the RF prediction model is generated by the following algorithm:

For each tree T_b , $1 \leq b \leq ntree$:

1. Draw m cases at random with replacement to create a subset of the training dataset. This procedure is often referred to as bootstrapping.
2. Recursively grow a decision tree T_i to the drawn sub-set by applying following steps on each node:
 - a) Randomly sample $mTry$ feature variables from the p features available
 - b) Choose the best feature variable / split-point from those $mTry$ features. The best split is defined as the split that results in the biggest reduction of impurity.
 - c) Split the node at into two sub-nodes
 - d) Repeat the recursive steps for each sub-node until the node size falls below a given minimum node size.

Breiman suggests the following parameters for classification tasks: $mTry = \sqrt{p}$, a bootstrap size m of about two thirds of the training data and a minimum node size of one. [14]

Node Splitting and Impurity Measure

Ideally a node is split in such a way that each of the sub-nodes only consists of samples of one class, i.e. each sub-node is pure. Hence, to determine the quality of a potential split point, a measure of purity or impurity is required. In practice, the *Entropy* or the *Gini index* are commonly used as impurity measures [39].

With n_t denoting the total number of samples in a node t , n_k the number of those samples belonging to class k and

$$p(k|t) = \frac{n_k}{n_t} \quad (3.21)$$

representing the proportion of class k in the node t , the Gini impurity function $i_G(t)$ and Entropy $i_E(t)$ are defined as:

$$i_G(t) = \sum_{l \neq k} p(k|t)p(l|t) = \sum_{k=1}^K p(k|t)(1 - p(k|t)) = 1 - \sum_{k=1}^K p^2(k|t) \quad (3.22)$$

$$i_E(t) = - \sum_{k=1}^K p(k|t) \log_2 p(k|t) \quad (3.23)$$

With those impurity measures, one can now determine the reduction of impurity for a binary split of node t into two sub-nodes t_L and t_R by:

$$\Delta i(t) = i(t) - P_l \cdot i(t_L) - (1 - P_l) \cdot i(t_R) \quad (3.24)$$

where P_l denotes the proportion of samples that are assigned to the left sub-node when splitting the node at point s of the feature variable j :

$$P_l = P_l(j, s) = P(x_j \leq s) = \frac{n_{t_L}}{n_t} \quad (3.25)$$

The best split then determined by finding the split that results in the biggest reduction of impurity.

Segmentation / Prediction

Segmenting new data with a RF-model is done by predicting the class label \hat{y} for each of the voxels \mathbf{x}_{new} . Note that \mathbf{x}_{new} is the feature space representation of a voxel that is to be segmented. The features need to be extracted in the exact same way as the features for the trainings set $\mathbf{X}_{\text{train}}$ have been extracted.

\mathbf{x}_{new} is run down all of the trees of the RF-model and each tree ‘votes’ for one of the class labels. In a segmentation, i.e. classification task, the predicted class label is then the one that receives the majority of the votes.

3.3.2 Feature Importance

A useful property of the RF algorithm is that it also allows an estimate of the importance of each feature for the classification task. The author of RF proposes two measures for feature importance [14]. A detailed characterization of those of those can be found in [40].

The first method referred to as *Gini importance* and is based on the Gini impurity measurement (3.22). As discussed in the previous subsection, the Gini impurity function i_G can be used in as criteria to determine the split at each node. During the creation of the RF model, one can now simply keep track of all the times a feature variable is chosen as splitting criteria and average over the impurity reduction Δi_G (computed via (3.24)) for each of those splits. This mean decrease of Gini impurity then serves as a measurement of feature importance.

The second method for measuring feature importance in RF is referred to as *permutation importance*. It measures the prediction strength of each variable by using the Out-Of-Bag (OOB) samples. The OOB samples of each tree are those samples that are not selected by the bootstrapping procedure, i.e are not used for the training of the tree. The permutation importance is computed by the following procedure [40]:

For each tree T_b , $1 \leq b \leq ntree$:

1. Identify the OOB samples. (the samples that were not used for growing the tree T_b)
2. Predict the class label for the OOB samples using T_b and count the correct predictions.
3. for each feature $j = 1, \dots, p$:
 - a) Randomly permute the OOB samples' values of feature j
 - b) Predict the class labels for the modified OOB samples
 - c) Calculate the difference in the accuracy between the prediction of the unmodified OOB samples and the prediction of the permuted OOB samples

The decrease in accuracy resulting from the permutation of a feature j is averaged over all trees and yields the features' permutation importance.

According to [40] and [39] the Gini importance and the permutation importance correlate well as long as the features are of the same type. Although the absolute values of the two feature importance measurements differ, the most predictive features rank amongst the top features in both.

Strobl et al. [41] showed that, when used for a mixture of continuous and categorical types of features, the Gini importance measurements are biased towards continuous features. They propose to use a modified version of permutation accuracy in those cases [41].

This page intentionally left blank.

Materials and Methods

4.1 Acquired Data

The datasets that are used in this thesis were acquired at a 7 Tesla MR whole body scanner (Magnetom, Siemens Healthcare, Erlangen, Germany) with a 32-channel head coil (Nova Medical, Wilmington, USA) at the High Field MR Centre at the Medical University of Vienna.

A total number of nine datasets from five different healthy volunteers are used. The subjects are referred to as s1, s3, s4, s5 and s6¹. For each subject, a high resolution multi-echo 3D gradient echo scan was performed. Resolutions and echo times are listed in Table 4.1. Subjects s5 and s6 were scanned with identical multi-echo sequences using an isotropic resolution of 0.7 mm and echo times of $TE = 7, 14, 20$ ms. The resolutions and echo times of subjects s1, s3 and s4 are slightly different.

The data from the 32 individual coils was combined using either the Virtual Reference Coil approach [42] (for s1, s3, and s4), or the COMPOSER approach [43] (s5 and s6). This results in one combined magnitude, and one combined (but not yet unwrapped) phase image for each of the three echoes of each dataset. These magnitude and phase images posed the starting point for the analysis done in this thesis.

For subject s6, two additional single-echo datasets were acquire in the same scan session. The echo time for those single-echo sequences was set to $TE = 7$ ms to match the first echo time of the multi-echo dataset. The resolution was set to isotropic 0.7 mm and isotropic 0.5 mm respectively. These single-echo datasets were used to test the performance of the proposed Random Forest segmentation method on short single echo sequences, which are faster in acquisition.

¹The reason that there is no s2 is because there were problems in the reconstruction of the s2 dataset and the same person was scanned again later (subject s5) with different acquisition parameters). Thus the dataset s2 is not used.

Table 4.1: Overview of the datasets that are used in this thesis, listing voxel dimensions, echo times, Repetition Time (TR), Flip Angle (FA), Acquisition Time (TA) and reconstruction method. All datasets are acquired at a field strength of 7T. The datasets for subjects s1, s3, and s4 are combined using the Virtual Reference Coil (VRC) [42] method and s5 and s6 are combined using the COMPOSER [43] approach.

Subj.	voxdim [mm·10 ⁻¹]	Echotimes [ms]			TR [ms]	FA [°]	TA	Reconstr. Method
		TE1	TE2	TE3				
s1	4x4x7	8	14	21	26	8	10 min 17 sec	VRC
s3	5x5x5	7.5	13	19.5	? ^a	? ^a	? ^a	VRC
s4	5x5x12	8	12.3	17.5	27	15	5 min 37 sec	VRC
s5	7x7x7	7	14	20	32	10	11 min 25 sec	COMPOSER
s6	7x7x7	7	14	20	32	10	11 min 25 sec	COMPOSER
	7x7x7	7	-	-	17	8	06 min 04 sec	COMPOSER
	5x5x5	7	-	-	17	10	09 min 51 sec	COMPOSER

^aThe original data and scanning parameters were lost due to a harddisk failure.

4.2 Performance Metrics

In SWI vessel segmentation literature, evaluation of segmentation performance is often performed visually [8, 6, 7] by overlaying the segmentation result on to minimum intensity projections of the SWI data or visualizing 3D renderings of the segmentation. These visualizations are then judged in terms of vessel-like appearance and connectivity of the vessels. However, in order to assess the quality of the obtained segmentations quantitatively and to enable a comparison of different segmentation approaches, quantitative performance metrics are needed. Table 4.2 provides a list of such performance metrics, which enable a quantitative evaluation of automatic segmentation results compared with the “ground truth”. Since an actual “ground truth” is often not available, “ground truth” refers to an expert’s interpretation of the image in form of a manually drawn label map.

All the performance metrics in Table 4.2 work on a voxel-by-voxel basis, i.e. they compare the voxel labeling obtained by the automatic segmentation with the manually assigned label. Although the metric “Accuracy” (the ratio of correctly classified voxels) is very intuitive, it is not particularly informative for unbalanced test samples, as in such cases it is biased towards the bigger class. Dice Similarity Coefficient and Matthews Correlation Coefficient (MMC) can both handle unbalanced test data. In image segmentation literature the Dice Similarity Coefficient is commonly used as performance score. Therefore it was also used as main metric for evaluation in this thesis to simplify comparison with other approaches. However, scores regarding the other performance metrics are provided as well in order to enable detailed analysis of the segmentation performance.

Table 4.2: Performance Metrics for quantitative comparison of the automatic segmentation with the “ground truth”, which is represented by a manually drawn label map.

Notation	Name	Description
TP	True Positives	Number of voxels that are correctly classified as veins.
FP	False Positives	Number of voxels that are falsely classified as veins.
TN	True Negatives	Number of voxels that are correctly classified as background.
FN	False Negatives	Number of voxels that are incorrectly classified as background.
N		Total number of (annotated) voxels.

Metric	Equation	Description
Precision	$\frac{TP}{TP+FP}$	Ratio that measures how many of the cases that are classified as positives are actually true positives.
Recall	$\frac{TP}{TP+FN}$	Ratio that measures how many of the positives are found by the classifier.
Accuracy	$\frac{TP+TN}{N}$	Ratio that measures how many of the tested cases (voxels) are classified correctly.
Dice	$\frac{2*Precision*Recall}{Precision+Recall}$	Dice Similarity Coefficient: Harmonic Mean between Precision and Recall. Also often referred to as F1-Score.
AUC-ROC	$\int_0^1 ROC$	Area under the Receiver Operating Characteristics (ROC) Curve. ROC is a plot of Sensitivity vs (1 - Specificity)
MMC	$\frac{TP*TN-FP*FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$	Matthews Correlation Coefficient [44]

This page intentionally left blank.

Analysis

In this chapter, the image processing and analysis that was performed within this thesis is presented in detail. Starting with the raw multi-echo magnitude and phase images that are acquired using a high-resolution, multi-echo Gradient Echo sequence (see Section 4.1), the first steps are to unwrap the phase, generate QSM images, and remove intensity

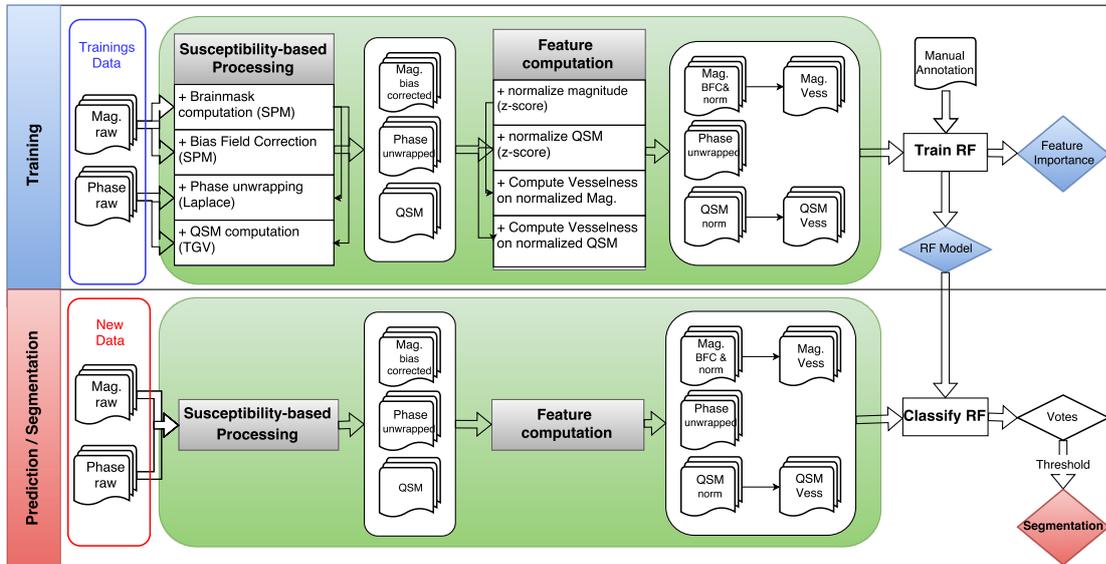


Figure 5.1: Flowchart of the Random Forest (RF) vein segmentation pipeline showing the necessary steps for training an RF model on manually annotated training data and using this RF model to segment new data. The susceptibility-based processing and feature extraction steps that are highlighted in green are identical for training and segmentation.

inhomogeneities. Section 5.1 describes these processing steps. Section 5.2 deals with the analysis of vessel extraction from either a single magnitude or QSM image, or from a combination of multiple features using the Random Forest (RF) algorithm. Figure 5.1 illustrates the RF vein segmentation pipeline that was developed and implemented in this thesis.

5.1 Susceptibility Based Image Processing

5.1.1 Brainmask generation

In order to limit computation and analysis to the brain region, a binary mask was computed, which indicates voxels within the brain as 1, and voxels outside the brain as 0. Two methods of brainmask generation were tested in this thesis. The first was the *Brain Extraction Tool* (BET) [45] from the *FMRIB Software Library*¹ (FSL) and the second a custom made MATLAB routine using the *Statistical Parametric Mapping*² suite (SPM).

The procedure for generating the brainmask with SPM was automated in MATLAB using SPMs matlabbatch system and comprises of the following steps:

1. Load the magnitude images into SPM (all three echoes of the multi-echo scans are used)
2. Use the SPM Segmentation module to:
 - a) perform bias field correction (see Section 5.1.2)
 - b) generate 5 tissue probability maps, each representing voxel-wise probability of belonging to one of the following tissues respectively: (p_1) grey matter, (p_2) white matter, (p_3) CSF, (p_4) bone, and (p_5) soft tissue. A detailed description of SPMs segmentation procedure can be found in the user manual³ to SPM12.
3. Generate an initial brainmask \mathcal{BM}_{init} from the probability maps by thresholding the sum of the the probabilities for grey matter, white matter and CSF :

$$\mathcal{BM}_{init}(\mathbf{x}) = \begin{cases} 1, & \text{if } p_1(\mathbf{x}) + p_2(\mathbf{x}) + p_3(\mathbf{x}) > 0.95 \\ 0, & \text{otherwise.} \end{cases} \quad (5.1)$$

A threshold of 0.95 was used for all the datasets.

4. Close the gaps inside \mathcal{BM}_{init} with MATLAB's `imclose` function.

¹FSL can be downloaded from: <http://fsl.fmrib.ox.ac.uk/>

²SPM12 can be downloaded from <http://www.fil.ion.ucl.ac.uk/spm/software/>

³The SPM12 user manual can be found at: <http://www.fil.ion.ucl.ac.uk/spm/doc/manual.pdf>

5.1.2 Biasfield Correction

MR images often show a smoothly varying intensity inhomogeneity across the dataset which has no anatomical relevance but is instead caused by factors such as magnetic field inhomogeneities, non-uniform reception coil sensitivity, eddy currents, or electrodynamic interactions with the measured object [46]. This intensity non-uniformity is referred to as bias field and, although not always visible for a human observer, can obstruct automated processing of the images [47]. Thus the bias field needs to be removed to perform robust automatic segmentation.

In this work, bias field correction of the magnitude images was performed with SPM's Segmentation module (which was already used in the process of creating brainmasks as described in Section 5.1.1).

5.1.3 Phase Processing and QSM computation

All processing steps for generating a Quantitative Susceptibility Map (QSM) from the combined phase image (see Sections 2.4 and 2.5) were performed using a very recent algorithm using Total Generalized Variation (TGV) developed by Langkammer et al.[26].

TGV-QSM⁴ takes magnitude image, phase image and a brainmask as input and performs phase unwrapping[32], background field removal, and dipole inversion to generate a QSM image [26]. In the original version, the brainmask is eroded to ensure the exclusion of the noisy regions outside of the brain. This erosion, however removes significant parts of the surface vasculature. Since, in this thesis, we are interested in the surface veins, the PYTHON script was modified not to perform this erosion. Due to this omission of the erosion step, the brainmask is required to be more accurate than in the original version. The default parameters of 1000 iterations and a TGV-regularization-parameter-ratio of $\frac{\alpha_0}{\alpha_1} = 1$ are used for all QSM computations in this thesis.

5.2 Vessel Extraction

This section describes how different methods for segmenting veins from susceptibility based images were performed and how they were analyzed. Section 5.2.1 describes the manual annotation of a part of the datasets, which then served as “ground truth” for training and evaluation of the automatic segmentation methods. Section 5.2.3 describes the parameter optimization and evaluation of automatic vein segmentation via vesselness filtering. Sections 5.2.4 and 5.2.5 explain the way the Random Forest algorithm was used for vein segmentation and analysis of the feature importances.

⁴The program can be obtained from <http://www.neuroimaging.at/pages/research/quantitative-susceptibility-mapping.php>.

5.2.1 Manual Annotation

Two datasets were manually annotated to serve as a reference for training the Machine Learning algorithm and as gold standard for quantitative evaluation of the segmentation performance. Annotation is the manual assignment of a label $L(\mathbf{x}) \in \{0, 1, 2, 3\}$ to individual voxels. Because of the huge number of voxels within each dataset, annotation of the whole 3D dataset would take a disproportionate amount of time. Thus only a portion of the voxels within two datasets (s1 and s5) was annotated with either a vein, or a background label. ITK-SNAP⁵ was used as annotation tool and both, a magnitude and a QSM image were displayed simultaneously to help the identification of veins. Figure 5.2 shows a screenshot of the view during the annotation process.

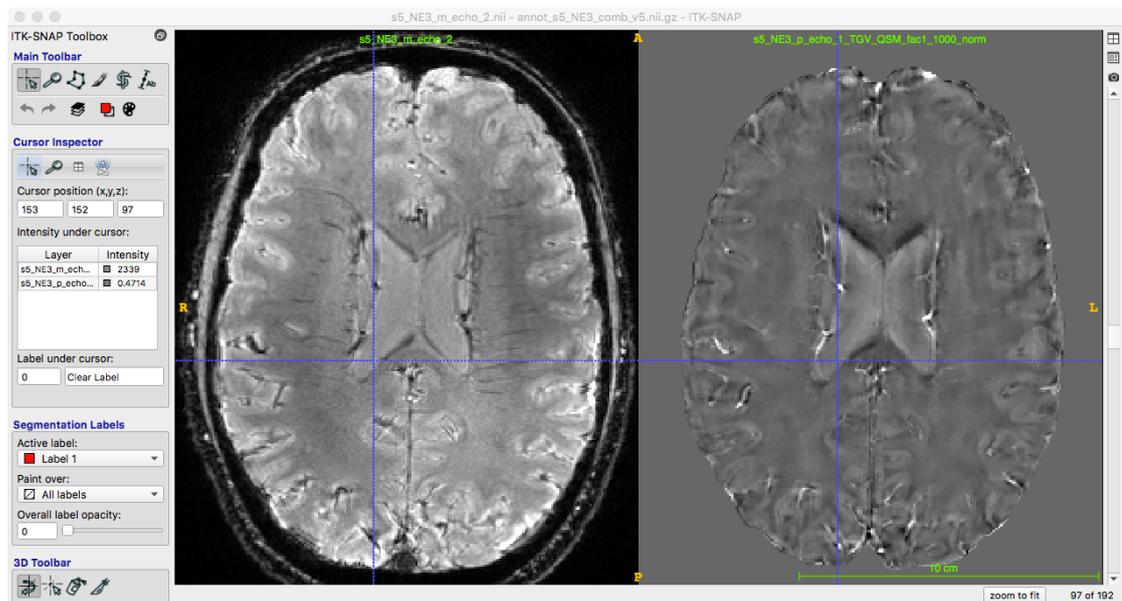


Figure 5.2: Screenshot of the ITK-Snap window demonstrating the view that was used in the annotation process. The magnitude image of echo 2 and the QSM image of echo1 are displayed next to each other to help identify venous vessels.

When only looking at individual slices it is difficult to distinguish between veins and background. Therefore the annotation strategy was to identify veins by scrolling through slices and following the veins, drawing the annotation voxel by voxel. Because not all veins were annotated, it is unreasonable to treat all voxels that have not been annotated as background. Thus a number of regions within the brain that do not contain venous vessels were marked with a background label, using the paintbrush tool. Only voxels with either a vein- or a background label were considered in training and for quantitative evaluation.

⁵<http://www.itksnap.org>

Table 5.1: Number of manually annotated Voxels within dataset s1 and s5. The label called “Sinuses” represents the superior sagittal sinus and the straight sinus which are not considered for training and validation. The color of the label corresponds with the colors in the visualization of the annotations (Fig. 5.3).

	Voxelcount		% of Brainvolume	
	s1	s5	s1	s5
Vein	76751	23227	0.64%	0.53%
Background	173224	48786	1.45%	1.12%
Sinuses	65779	-	0.55%	-
# Voxels within Brainmask	11973907	4361245	-	-

In order to capture the variety of different tissues in the background region, background labels were drawn in multiple brain regions, focusing on regions that are challenging for automatic vein segmentation such as the cerebral surface and basal ganglia. Figure 5.3 shows 3D renderings of the manual annotations of s1 and s5. In dataset s1 the superior sagittal sinus and the straight sinus were initially marked as veins as well. However these two sinuses are much bigger than the typical cerebral vein and thus would get a disproportionately high importance in quantitative evaluation. The labeling of those two sinuses were thus changed to a separate “Sinuses” label and was neither used for training nor for quantitative evaluation. In dataset s5 those two sinuses were not annotated. Table 5.1 shows the number of labeled voxels for each of the two annotated datasets.

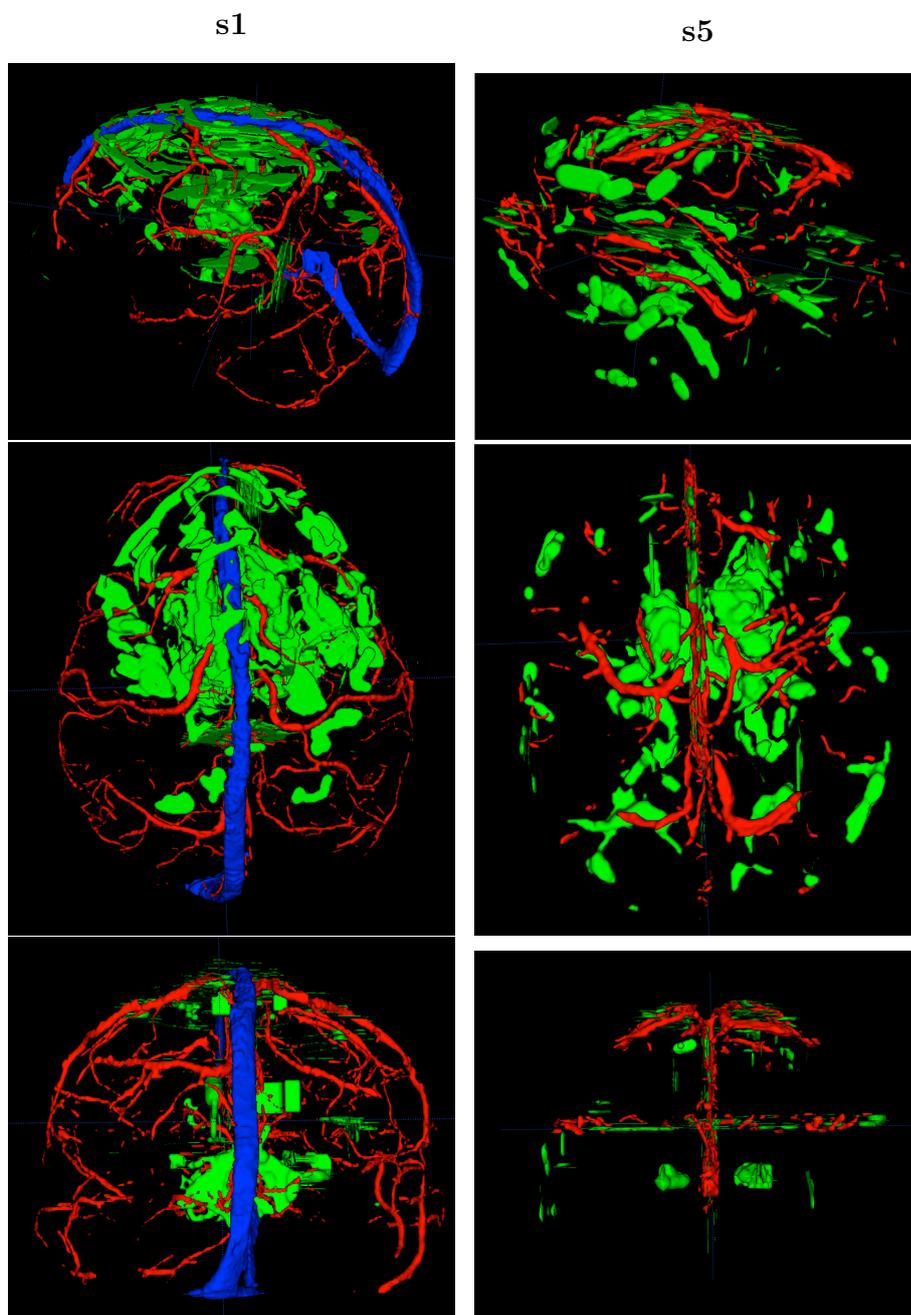


Figure 5.3: 3D views of the manual annotations for datasets s1 (*left*) and s5 (*right*). Venous vessels annotations are displayed in red and background labels are green. In s1, the superior sinus sagittalis and the straight sinus are labeled in blue.

5.2.2 Data normalization

In MR images, the scale of the signal intensity varies between subjects. To get a uniform intensity distribution across different datasets the intensity values were normalized using the z-score:

$$I'(\mathbf{x}) = \frac{I(\mathbf{x}) - \hat{\mu}}{\hat{\sigma}} \quad (5.2)$$

where $\hat{\mu}$ and $\hat{\sigma}$ are the mean and standard deviation of the intensity values. For each image, $\hat{\mu}$ and $\hat{\sigma}$ were computed from all the voxels within the brainmask. The magnitude and QSM images of each dataset were normalized in this way.

5.2.3 Frangi Vesselness

Frangi vesselness filtering was performed on the normalized magnitude and normalized QSM images using MATLAB⁶. The vesselness computation can be tuned with 7 parameters: FrangiAlpha, FrangiBeta and FrangiC, which represent the α , β and γ parameters from Eq. (3.10); the scale parameters ScaleRange (minimum and maximum scale), NumScales (number of scales) and ScaleType (linear or exponential distribution of the scales); and the BlackWhite option, which was set to True to detect vessels that appear dark with respect to their surrounding and False to detect bright vessels.

An exhaustive grid search was performed on the magnitude and QSM images on the two annotated multi-echo datasets s1 and s5 to optimize the parameters for the cerebral vein segmentation task. This parameter iteration was computed for the images of each echo individually to investigate if and how the echo time influences the performance

⁶An implementation of a Frangi vesselness filter can be found at: <https://www.mathworks.com/matlabcentral/fileexchange/24409-hessian-based-frangi-vesselness-filter>

Table 5.2: Parameters for Frangi Vesselness Parameter Iteration on the normalized magnitude and QSM images.

Parameter	Value	Reason
FrangiAlpha	0.5	Default parameters as described in Frangi’s original paper [12].
FrangiBeta	0.5	
FrangiC	varied	Varied from 0.1 to 2. Step size 0.1
ScaleRange	varied	Tried 3 Ranges: [0.1 1]mm as in Koopmans paper [6], [0.5 2.5]mm like Beriault [9], and [0.1 3]mm.
NumScales	7	Empirically determined. More scales did not improve accuracy.
ScaleType	exp	
BlackWhite	True (Mag)	Veins appear dark in magnitude images.
	False (QSM)	Veins appear bright in QSM images.

of the vesselness filter. The parameter settings that were used in the grid search are summarized in Table 5.2.

In order to evaluate the vesselness filtering output, the output was turned into a binary segmentation by thresholding. For each image, the threshold was optimized with respect to the maximum Dice overlap with the manually annotated vein and background labels. In addition to the Dice overlap score, multiple other performance metrics (see Table 4.2) were computed.

5.2.4 Random Forest Segmentation

In order to segment venous vessels on the basis of information from the multi-echo magnitude, phase, and QSM images, a Random Forest algorithm⁷(RF) was used to classify each voxel within the brain as either vein or background. This section provides a detailed explanation of how the RF algorithm was used, starting with the features that are used. Then the processes of building an RF model and using this model for vein segmentation are described. Finally the validation methods that were used are elaborated.

Features

The features that are used for the RF segmentation are computed from the magnitude, phase and QSM images of a multi-echo scan. Since all the images of each subject were acquired in a single multi-echo scan, these images share a common space and do not need to be coregistered. The following features were used in this thesis:

Mag	Voxel intensity of the biasfield-corrected and normalized magnitude image
MagVess	Frangi vesselness computed on the biasfield-corrected and normalized magnitude image. The following parameters are used: FrangiAlpha=FrangiBeta=0.5; FrangiC=1; 7 exponentially spaced scales in the range of [0.1 1]mm; BlackWhite=True
Ph	Unwrapped phase value (computed with the laplace unwrapping method implemented in TGV-QSM)
QSM	Voxel intensity of the normalized QSM image
QSMVess	Frangi vesselness computed on the normalized QSM image. Parameters used: FrangiAlpha=FrangiBeta=0.5; FrangiC=1; 7 exponentially spaced scales in the range of [0.1 1]mm; BlackWhite=False

The vesselness parameters for the features MagVess and QSMVess were selected after the vesselness parameter iteration that is described in Section 5.2.3. As can be seen in

⁷The RF implementation used in this thesis can be downloaded from <http://code.google.com/p/randomforest-matlab/>

Section 6.2 and Appendix A, the selected parameter sets were amongst those that yielded the highest performance scores on all three echoes of both annotated datasets.

For each voxel within the brainmask \mathcal{BM} , a feature vector $\mathbf{f}(\mathbf{x})$ was generated containing either a subset or all of the features described above. The full feature vector using all features from a 3 echo scan thus contained 15 features:

$$\begin{aligned} \mathbf{f}_{\text{full, ME}}(\mathbf{x}) = \{ & \text{Mag}(\text{TE1}, \mathbf{x}), \text{Mag}(\text{TE2}, \mathbf{x}), \text{Mag}(\text{TE3}, \mathbf{x}), \\ & \text{MagVess}(\text{TE1}, \mathbf{x}), \text{MagVess}(\text{TE2}, \mathbf{x}), \text{MagVess}(\text{TE3}, \mathbf{x}), \\ & \text{Ph}(\text{TE1}, \mathbf{x}), \text{Ph}(\text{TE2}, \mathbf{x}), \text{Ph}(\text{TE3}, \mathbf{x}), \\ & \text{QSM}(\text{TE1}, \mathbf{x}), \text{QSM}(\text{TE2}, \mathbf{x}), \text{QSM}(\text{TE3}, \mathbf{x}), \\ & \text{QSMVess}(\text{TE1}, \mathbf{x}), \text{QSMVess}(\text{TE2}, \mathbf{x}), \text{QSMVess}(\text{TE3}, \mathbf{x}) \} \end{aligned} \quad (5.3)$$

The feature vector was used for training a RF model and segmentation of new cases.

Training a RF Model

In order to obtain a segmentation model for a given feature subset, the RF needs to be trained on this feature set. For this purpose, a representative number of the manually annotated voxels (\mathcal{A}) are used to form a training set \mathcal{T} :

$$\mathcal{T} = \{(\mathbf{f}(\mathbf{x}), y(\mathbf{x})) \mid \mathbf{x} \in (\text{ROI}^{\text{train}} \cap \mathcal{A})\} \quad (5.4)$$

that contains the feature vector $\mathbf{f}(\mathbf{x})$ and the corresponding vein, or background label $y(\mathbf{x})$ of all voxels within the training region $\text{ROI}^{\text{train}}$ that have been annotated as either vein or background. Details on how the data was divided up into test- and training set are given in the Validation subsection below.

The influence of different values for the RF parameters $n\text{trees}$ and $m\text{Try}$ was investigated in an initial experiment using all annotated voxels of s5 and the full featureset $\mathbf{f}_{\text{full, ME}}$. In accordance with that experiment (see Section 6.3.1 for details) and the suggestions from Breiman et. al. [14] number of trees was set to $n\text{trees} = 200$ and the number of features tried at each split was set to the square root of the number of features $m\text{Try} = \sqrt{n\text{Feat}}$ in all experiments.

Segmentation with RF model

To segment veins in new data, each voxel is classified as either vein or background by the trained RF model using the same feature set as in the training step (e.g. $\mathbf{f}_{\text{full, ME}}$). The classification is based on the RF votes, where each tree in the RF model votes for the voxel to either belong to the background class ($\hat{y}_i = 0$) or the vein class ($\hat{y}_i = 1$). The ratio of trees that vote for the vein class $\frac{1}{n\text{trees}} \sum_{i=1}^{n\text{trees}} \hat{y}_i$ can be seen as the RF's confidence that the voxel belongs to a vein. The ratio for each voxel is stored in a votes map which is then thresholded to obtain the segmentation. The default threshold in the RF algorithm is 0.5, which is equivalent to a majority vote.

Validation

To evaluate the segmentation performance of the proposed multi-echo RF method the manually annotated data was split into training and test set. The annotated voxels within the training set were used for training the Random Forest, and the annotated voxels within the test set were used for quantitative evaluation using performance metrics described in Section 4.2.

Two different ways of splitting the manually annotated voxels of datasets s1 and s5 into test and train set were performed:

- To test the performance of the RF when training and test set originate from the same dataset a 4-fold cross-validation technique was used for each of the two annotated datasets. Each dataset was split evenly into 4 non-overlapping regions of interest (ROIs). Three of the four regions (i.e. all the annotated voxels within these three regions) were used for training an RF model. This RF model was then used to segment the fourth region. That whole procedure was repeated 4 times, each time leaving a different region as test region and using the other 3 regions for training. With this procedure the whole brain gets segmented without ever using the same data for training and prediction. The segmentation generated in this way is referred to as **RF-intra**.
- The second way was to use all of the annotated voxels of one annotated dataset (either s1 or s5) to train an RF model and then segment the other dataset with that RF model. The segmentations obtained in that way are referred to as **RF-inter**.

The performances of both, the RF-intra and RF-inter segmentations, were compared to the segmentations obtained via thresholding the magnitude, phase, or QSM images directly or via using a the Frangi vesselness filter on either the magnitude, or QSM images. For the segmentation performance of the vesselness on the individual magnitude and QSM images, the best scores obtained in the vesselness parameter iteration (Section 5.2.3) were taken. The thresholding segmentation performances were also maximized w.r.t. the highest Dice score. In other words, we compared the RF segmentation with the best obtainable segmentation when using vesselness or thresholding on individual images.

Visual assessment of the segmentation performance in different brain regions was performed by overlaying the segmentation onto a magnitude or QSM image. Also 3D renderings of the segmented surface veins are shown.

Other datasets with varying acquisition parameters (see Table 4.1) were segmented as well using RF-models trained on either s1 or s5. Because no annotations was available for those datasets, those segmentations are evaluated only visually.

5.2.5 Feature Importance

The Gini importance values, which are obtained when training a RF (see Section 6.4) were used to assess the importance of the multi-echo magnitude, phase and QSM features for the venous vessel segmentation task.

Additionally, in order to evaluate the quality of the Feature importance ranking, RF training and prediction was performed using only subsets of features. Table 5.3 shows the feature subsets that were used. Segmentation performance measurements were computed in the same way the full featureset was evaluated.

Table 5.3: Name and feature-list of the feature subsets that are used for evaluating the feature importance.

Subset	List of features
$\mathbf{f}_{\text{full,TE1}}$	{Mag(TE1), MagVess(TE1), Ph(TE1), QSM(TE1), QSMVess(TE1)}
$\mathbf{f}_{\text{full,TE2}}$	{Mag(TE2), MagVess(TE2), Ph(TE2), QSM(TE2), QSMVess(TE2)}
$\mathbf{f}_{\text{full,TE3}}$	{Mag(TE3), MagVess(TE3), Ph(TE3), QSM(TE3), QSMVess(TE3)}
$\mathbf{f}_{\text{MagOnly}}$	{Mag(TE1), Mag(TE2), Mag(TE3), MagVess(TE1), MagVess(TE2), MagVess(TE3)}
$\mathbf{f}_{\text{QSMOnly}}$	{QSM(TE1), QSM(TE2), QSM(TE3), QSMVess(TE1), QSMVess(TE2), QSMVess(TE3)}
$\mathbf{f}_{\text{NoPhase}}$	{Mag(TE1), Mag(TE2), Mag(TE3), MagVess(TE1), MagVess(TE2), MagVess(TE3), QSM(TE1), QSM(TE2), QSM(TE3), QSMVess(TE1), QSMVess(TE2), QSMVess(TE3)}
$\mathbf{f}_{\text{Top3-Gini}}$	{MagVess(TE2), QSM(TE1), QSMVess(TE1)}

This page intentionally left blank.

Results

This chapter presents the obtained results in the following way: In Section 6.1 the visual appearance of venous vessels in the magnitude, phase and QSM images of different echo times is compared and the influence of the brainmask on the surface vein appearance in QSMs is shown. Section 6.2 provides a quantitative summary of the vesselness parameter iteration on magnitude and QSM images of different echo times. In Section 6.3 the performance of the proposed Random Forest (RF) segmentation method is evaluated and compared with the segmentations obtained with vesselness filtering and thresholding. These comparisons are performed quantitatively using the manually annotated datasets and visually. Finally the feature importance rankings are presented in Section 6.4.2.

6.1 Susceptibility Based Image Processing

For the computation of Quantitative Susceptibility Maps (QSM), a brainmask is needed to avoid regions outside of the brain where the phase values are noisy and unreliable. Brainmasks that are computed with FSL-BET, which are typically used in QSM computations[5], turned out to be too inaccurate when one is interested in depicting the surface veins in QSM. Brainmasks that are computed by the method described in Section 5.1.1 using the Statistical Parameter Mapping Tool (SPM) were more accurate on our datasets and led to less artifacts in the surface regions of the QSMs. This is demonstrated in Fig. 6.1, which shows one QSM image computed using a BET-mask and one QSM image using the SPM-mask next to a magnitude image of the same slice. One drawback of the SPM-mask is that some parts of the superior sagittal sinus, and potentially some smaller parts of other surface veins, are cut away by the tight SPM-mask. However as we will see in the 3D renderings of the segmented surface veins in Section 6.3, there is enough of the veins left to allow good visualizations. For all the analysis going forward, only voxels within the SPM mask are considered.

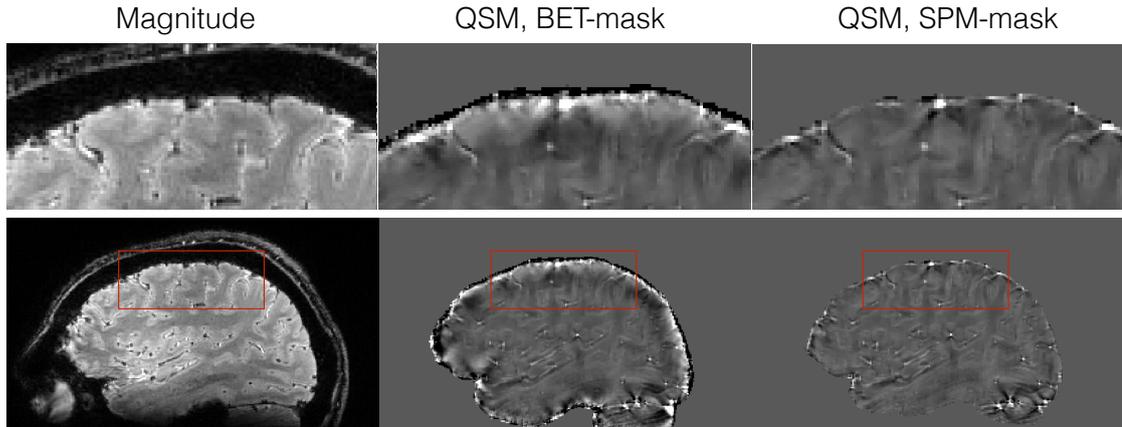


Figure 6.1: Surface vein appearance in a magnitude image (*left*) and two QSM images. Both QSM images are computed using TGV-QSM with the same parameters, once with a brainmask computed using FSL-BET (*middle*) and once with a brainmask computed using SPM (*right*). QSM computation with the larger BET-mask results in bright artifacts in regions close to the brain surface, which hampers the identification of surface veins.

The use of multi-echo datasets enables us to investigate the influence of different echo times (TE) on the depiction of veins on magnitude, phase and QSM images. Figure 6.2 shows images of the same slice in the mid-sagittal brain region for TEs of 7, 14 and 20 ms. In the magnitude images, the contrast for veins increases for higher echo times. Also the veins appear bigger in higher TE magnitude images than in the lower ones due to increased signal loss around the veins by the susceptibility-induced fields. The phase images, which have been unwrapped with the Laplace unwrapping method that is implemented in the TGV-QSM program, show unwrapping errors in the 14 ms version and even more at TE = 20 ms. The unwrapping errors are mostly located in and around larger venous vessels. Therefore the QSM images generated from these later phase images do not show the veins as clearly as the QSM image created from the 7 ms phase image.

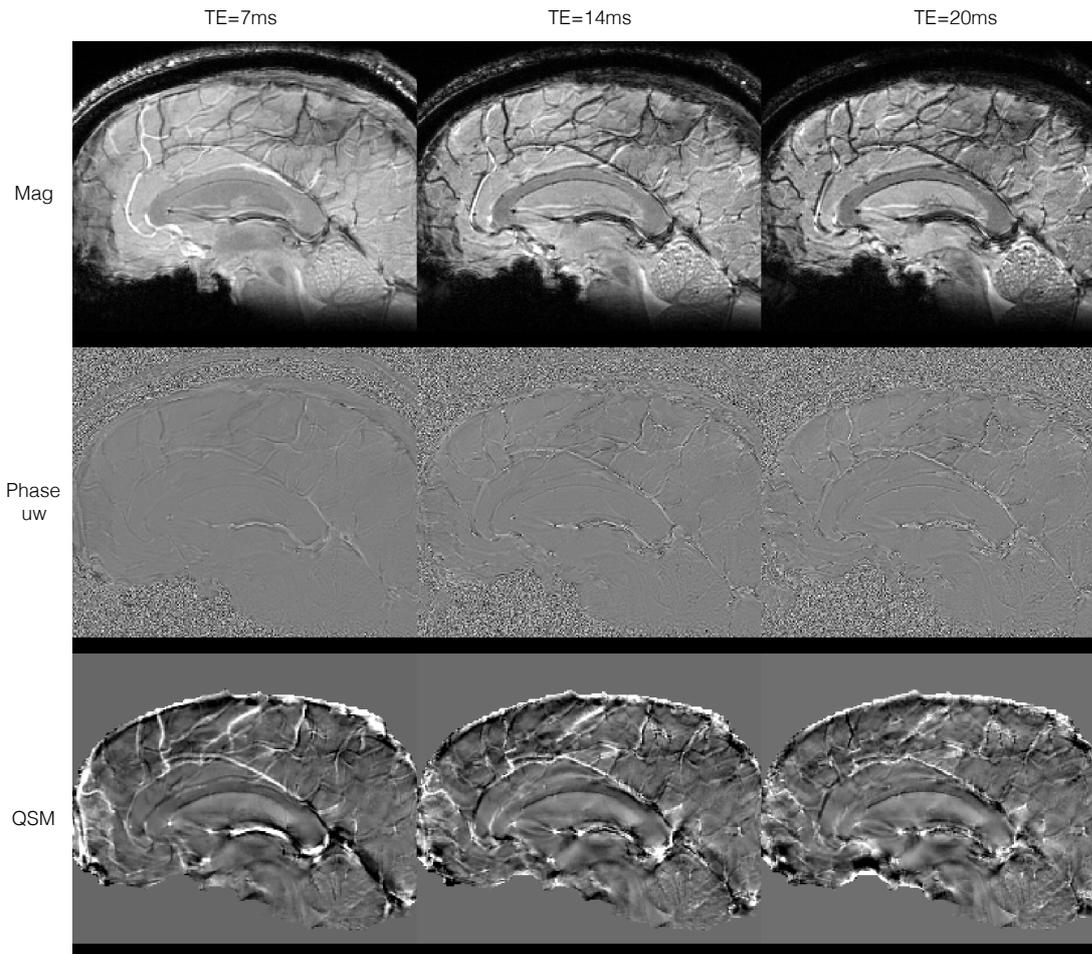


Figure 6.2: Magnitude, phase and QSM images at three different echo times acquired in a single scan, showing the change in venous vessel appearance for progressing Echo Times (TE). The vein contrast in the magnitude image increases with higher TE, while for QSM, the contrast is best for the shorter TE.

6.2 Frangi Vesselness Parameter Iteration

As described in Section 5.2.3 we performed a grid search to optimize the Frangi Vesselness parameters with respect to the maximum Dice overlap score with the manual annotations. The parameter iteration was performed on the normalized magnitude and normalized QSM images of each echo individually. Table 6.1 shows those parameter sets that performed best. In all but one case a ScaleRange of 0.1 to 1 mm yielded the highest score. In Appendix A, figures visualizing the performance of all parameter-sets on each individual image are provided.

Table 6.1: Frangi vesselness parameters for the normalized Magnitude and normalized QSM images of s1 and s5 that yielded the highest Dice score after parameter iteration. The scores regarding various other performance metrics can be found in Tables 6.2 and 6.3.

Subj.	Method	TE [ms]	Dice [%]	FrangiC	ScaleRange	
					min [mm]	max [mm]
s5	Mag. Frangi	7	65.70	0.8	0.1	1
		14	70.67	0.9	0.1	1
		20	69.39	1	0.1	1
	QSM Frangi	7	71.96	1.1	0.1	1
		14	65.04	1.1	0.1	1
		20	61.92	1	0.5	2.5
s1	Mag. Frangi	8	73.69	0.9	0.1	1
		14	80.94	1	0.1	1
		21	81.07	1.1	0.1	1
	QSM Frangi	8	86.24	1.6	0.1	1
		14	80.98	1.5	0.1	1
		21	75.80	1.3	0.1	1

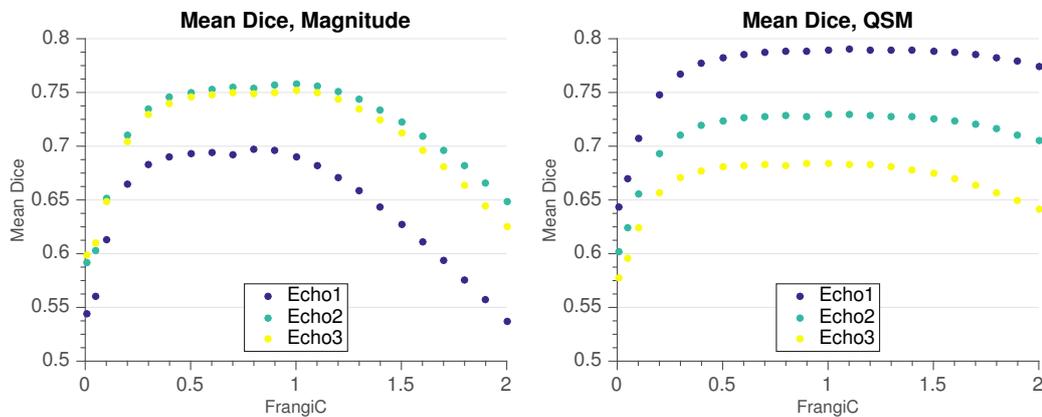


Figure 6.3: Mean Dice score for Frangi Vesselness segmentation with a ScaleRange of 0.1 to 1 computed on magnitude images (*left*), and QSM images (*right*) for each echo. The individual plots for each image and all scale ranges can be found in Appendix A.

Figure 6.3 shows the mean Dice score of the individual echoes for a scale range of 0.1 to 1 over the different FrangiC values. For the magnitude images, echo2 (14 ms) and echo3 (~20ms) performed almost equal, both better than echo1 (~7ms). However for vesselness filtering on the QSM images, echo1 yielded best results. This indicates that a multi-echo acquisition could be beneficial for vein segmentation.

In Section 6.3 the best individual Frangi vesselness segmentations are compared quantitatively and visually with segmentations obtained by the proposed multi-echo Random Forest (RF) approach. For the Frangi vesselness features in the RF a scale range of 0.1 to 1 mm and a FrangiC value of 1 is chosen for all the magnitude and the QSM images of each echo.

6.3 Segmentation Performance of RF using all Features (multi-echo)

6.3.1 Random Forest Parameter Optimization

To test the influence of the RF parameters $mTry$ and $nTrees$, multiple RFs were trained on all annotated voxels of s5 using all 15 features. Figure 6.4 shows the error rate of the Out Of Bag samples for a growing number of trees using $mTry$ values ranging from 2 to 6. The error rate for all the curves drops exponentially until it stabilizes at approximately 150 to 200 trees. The parameter $mTry$ does not seem to have a big influence on the error rate for this classification task. For all the following RF segmentations in this thesis, the RF models are trained with $nTrees = 200$ and the default value of $mTry = \sqrt{\text{number of features}}$ (3 for the full featureset of 15).

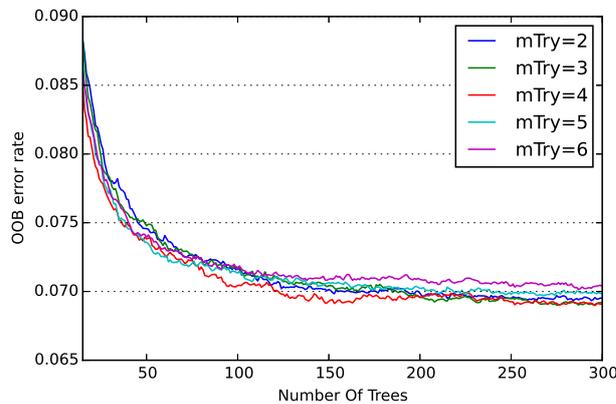


Figure 6.4: Out Of Bag (OOB) error rate while training a Random Forest classifier on s5 using all features (5 features for each of the 3 echoes = 15) with increasing number of trees for different $mTry$ values ranging from 2 to 6.

6.3.2 Quantitative Segmentation performance compared to Manual Annotation

The bar plots in Fig. 6.5 shows the Dice Similarity Coefficient between different automatic segmentations and the manually annotated voxels of datasets s1 and s5. A more detailed evaluation using multiple performance metrics is provided in Tables 6.2 and 6.3. For each of the magnitude, QSM and phase thresholding segmentations, the threshold which yielded the highest Dice score was chosen. In the “PhaseThreshPos” experiments, all voxels that have a bigger phase value in the unwrapped images than the threshold are classified as vein, similar to a positive phase mask in SWI (see Fig. 2.11). In the “PhaseThreshNeg” voxels with a smaller value are classified as veins. For the vesselness segmentations, the best values from the vesselness parameter iteration (Section 6.2) are chosen. The way the annotated voxels are split into training and test dataset for intra- and inter-dataset validation of the Random Forest is described in Section 5.2.4.

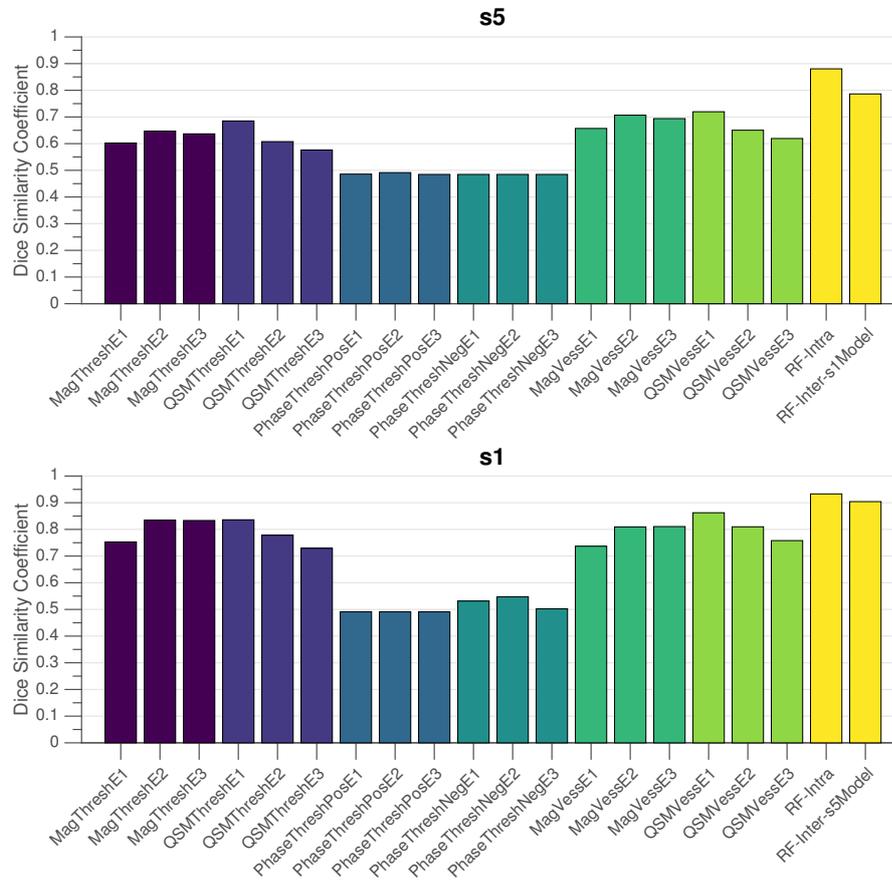


Figure 6.5: Dice similarity coefficient of different segmentation methods compared to manual annotation on s1 and s5.

6.3. Segmentation Performance of RF using all Features (multi-echo)

Table 6.2: Detailed evaluation of segmentation performance on **s5** compared to the manual annotation. The operating point was set to the maximum Dice similarity coefficient.

Method	TE [ms]	Dice [%]	Accuracy [%]	Precision [%]	Recall [%]	AUC-ROC [%]	MCC [%]
MagThresh	7	60.24	75.76	63.37	57.41	74.92	56.65
	14	64.71	74.33	57.75	73.58	80.46	58.51
	20	63.67	73.09	56.03	73.73	77.89	57.45
QSMThresh	7	68.52	80.66	71.47	65.80	83.30	63.34
	14	60.76	74.75	60.40	61.14	74.28	56.76
	20	57.63	70.02	52.58	63.75	70.99	53.91
PhaseThreshPos	7	48.59	71.05	56.23	42.78	51.77	48.96
	14	49.14	73.54	63.78	39.97	50.58	49.06
	20	48.47	31.99	31.99	100.00	50.39	-
PhaseThreshNeg	7	48.47	31.99	31.99	100.00	48.23	-
	14	48.47	31.99	31.99	100.00	49.42	-
	20	48.47	31.99	31.99	100.00	49.61	-
MagVess	7	65.70	81.43	80.30	55.60	75.89	61.64
	14	70.67	83.09	79.39	63.67	79.70	65.73
	20	69.39	82.91	81.25	60.55	78.22	64.77
QSMVess	7	71.96	82.81	75.26	68.94	81.61	66.51
	14	65.04	79.01	69.61	61.03	76.42	60.50
	20	61.92	72.90	56.23	68.90	73.51	56.72
RF-Intra	ME	88.04	92.45	89.26	86.86	97.18	83.85
RF-Inter-s1Model	ME	78.64	86.26	78.20	79.08	92.76	72.84

Table 6.3: Detailed evaluation of segmentation performance on **s1** compared to the manual annotation. The operating point was set to the maximum Dice similarity coefficient.

Method	TE [ms]	Dice [%]	Accuracy [%]	Precision [%]	Recall [%]	AUC-ROC [%]	MCC [%]
MagThresh	8	75.29	84.91	80.59	70.64	88.14	69.79
	14	83.43	89.52	85.97	81.04	92.83	78.31
	21	83.34	89.26	84.14	82.56	93.04	78.05
QSMThresh	8	83.61	89.83	87.97	79.66	90.36	78.66
	14	77.82	86.99	87.41	70.13	86.23	72.72
	21	72.92	83.41	77.79	68.61	82.65	67.42
PhaseThreshPos	8	49.11	32.55	32.55	100.00	42.75	-
	14	49.11	32.55	32.55	100.00	43.98	-
	21	49.11	32.55	32.55	100.00	47.22	-
PhaseThreshNeg	8	53.14	73.15	61.51	46.78	57.25	51.78
	14	54.75	76.64	74.08	43.43	56.02	52.93
	21	50.23	72.39	60.77	42.80	52.78	49.85
MagVess	8	73.69	85.15	87.01	63.91	80.71	68.80
	14	80.94	88.38	86.85	75.79	86.66	75.78
	21	81.07	88.30	85.66	76.94	86.96	75.81
QSMVess	8	86.24	91.48	90.91	82.04	90.26	81.78
	14	80.98	88.41	86.91	75.81	86.62	75.82
	21	75.80	85.23	81.18	71.10	83.43	70.31
RF-Intra	ME	93.31	95.65	93.56	93.06	98.74	90.52
RF-Inter-s5Model	ME	90.41	93.85	91.74	89.12	97.83	86.75

In all of the performance metrics, the segmentation obtained by combining the features using Random Forests yielded higher scores than the segmentations obtained by using the features individually. Segmenting veins by only thresholding the unwrapped phase values seems to be problematic. In some cases the highest Dice overlap score is yielded by a threshold so low that all voxels are classified as veins. This is probably caused by the nonlocal character of the phase.

It has to be emphasized that because quantitative evaluation requires a “ground truth” to test against, these scores are only calculated from voxels that have manual annotations. In s5 the annotated voxels are primarily focused on regions that are challenging for automatic segmentation (see Section 5.2.1). Therefore the performance scores on this dataset are generally lower than in s1.

6.3.3 Qualitative Segmentation Performance

Figures 6.6 to 6.8 illustrate the Random Forest segmentation of datasets s1 and s5 in comparison with the those Frangi Vesselness segmentations that yielded the highest Dice scores in Section 6.2. Fig. 6.6 shows minimum/maximum intensity projections of a Magnitude image and the corresponding QSM image overlaid with the segmentations to demonstrate the performance inside the brain, while Fig. 6.8 shows the 3D renderings of the surface vein segmentations of the same subject. In Fig. 6.7, the 3D rendered surface vein segmentations for s1 are shown.

With the **magnitude vesselness** segmentation approach, most of the veins inside the brain are recognized. In the surface region however, the vesselness filter struggles to detect veins because there, due to the hypointense appearance of both the veins and the skull, the idealized vessel shape is not given. Because a lot of the manually annotated voxels are on the surface (this is the region we are most interested in) and the segmentation parameters are optimized for those annotations, the threshold for the magnitude vesselness to classify a voxel as vein is forced to be rather low. This leads to a slight tendency for oversegmentation inside the brain while still not detecting the surface veins sufficiently.

The **QSM vesselness** performs well on the surface veins. When compared with the 3D rendering of the manually annotated surface veins of s1 (Fig. 6.7), one can see that all the surface veins that are annotated are picked up by the QSM vesselness as well, and even some that were not annotated. Inside the brain, the QSM vesselness segments most of the veins as well, however it does miss some of the veins in the mid-sagittal region while still having some obvious misclassifications.

Because the **Random Forest** model has information from the magnitude vesselness filter as well as the QSM vesselness filter, alongside the intensity values, the RF segmentation performs well both inside the brain and for surface veins.

In order to test generality, other datasets with slightly varying acquisition parameters (see Table 4.1) are segmented with the RF models trained on either s5 or s1. The 3D surface vein renderings of those segmentations are shown in Fig. 6.9 and Fig. 6.10 shows the segmentations inside the brain.

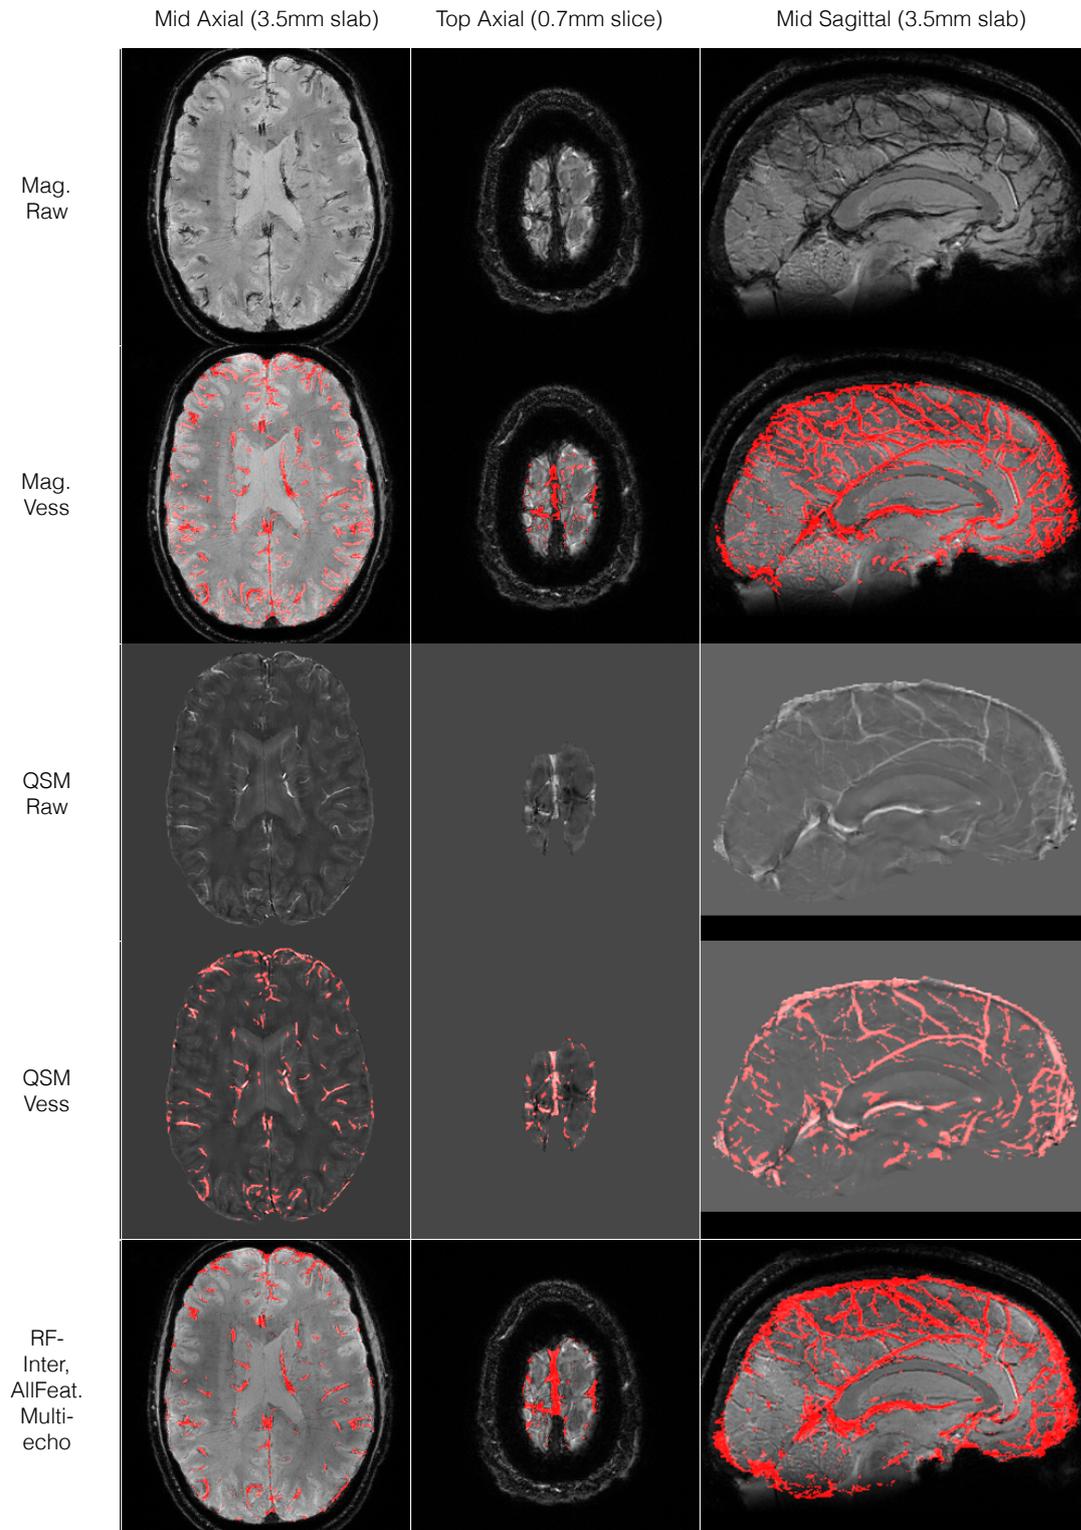


Figure 6.6: Comparison of binary vein segmentation of s5 using Frangi vesselness filtering on the magnitude image (Echo2), Frangi vesselness on the QSM image (Echo1) and the proposed RF. The RF was trained on a different dataset (s1), which has different acquisition parameters (acquisition parameters are listed in Table 4.1).

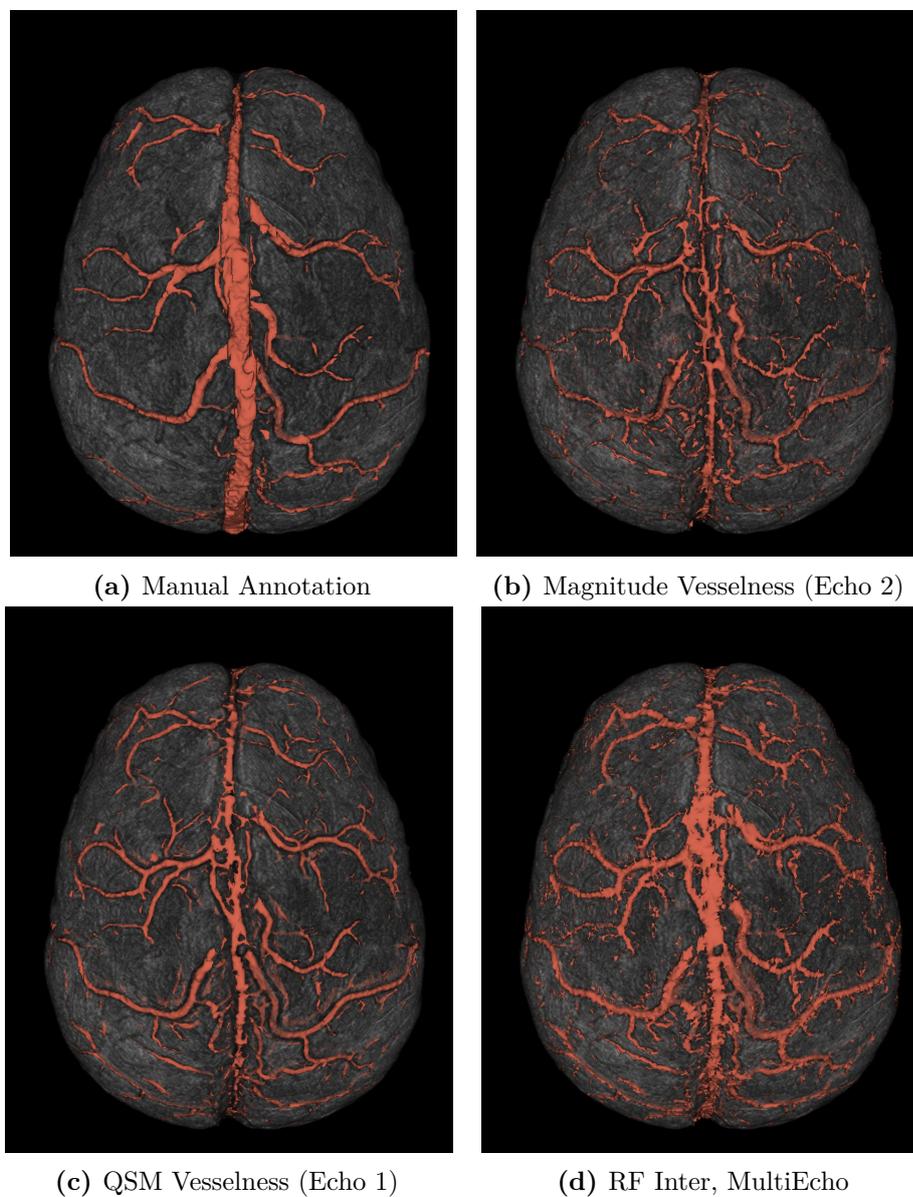


Figure 6.7: Comparison of three automatic segmentation methods with the manual annotation on the same subject (s1) in 3D. (b) and (c) show the best segmentations obtained by thresholding the output of the Frangi vesselness filter on the magnitude and QSM. (d) shows the segmentation obtained by using the proposed Random Forest approach. (The RF model was trained on s5.)

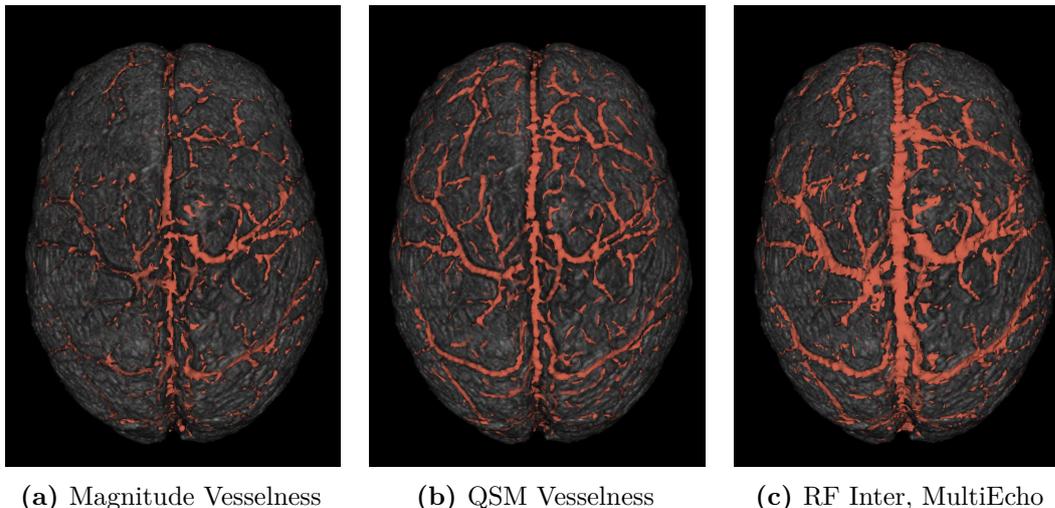


Figure 6.8: Visualization of the surface vasculature of s5 after automatic segmentation. (a) and (b) show the best segmentations obtained by thresholding the output of the Frangi vesselness filter on the magnitude (Echo2) and QSM (Echo1). (c) shows the segmentation obtained by using the proposed Random Forest approach. (The RF model was trained on s1.)

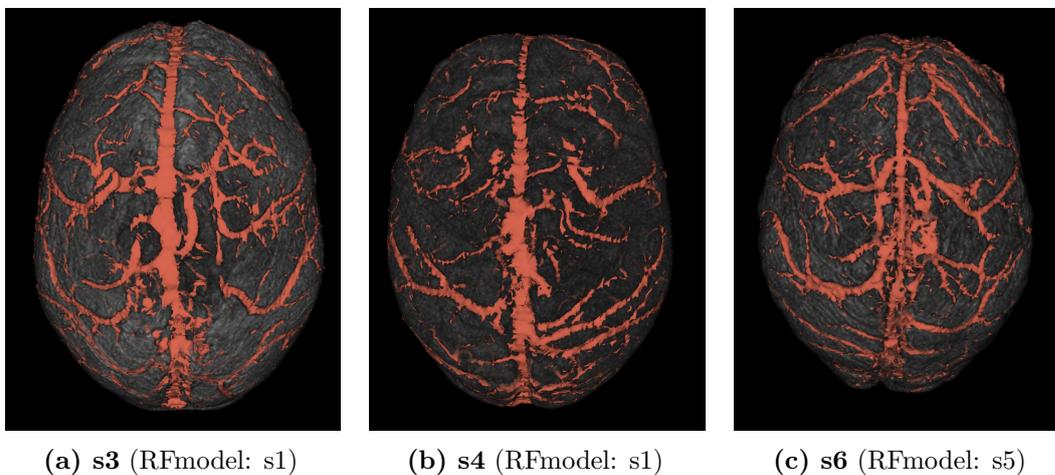
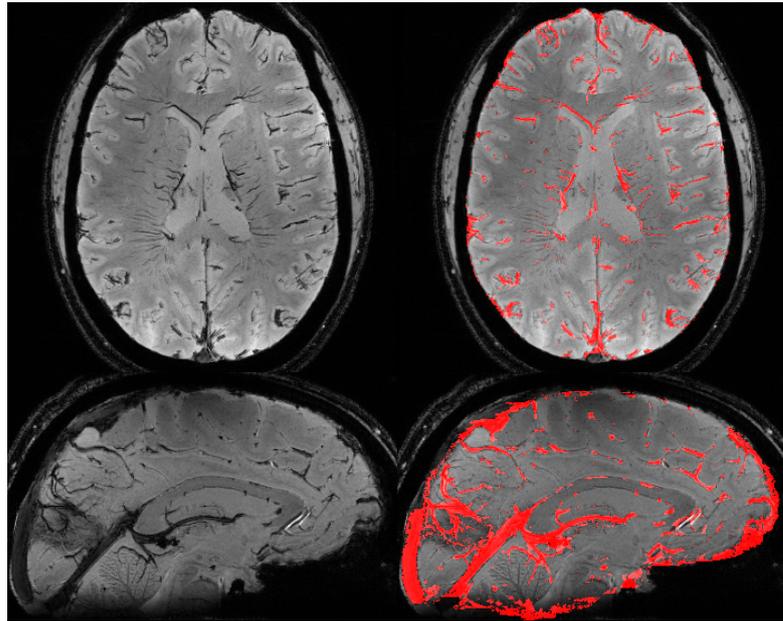
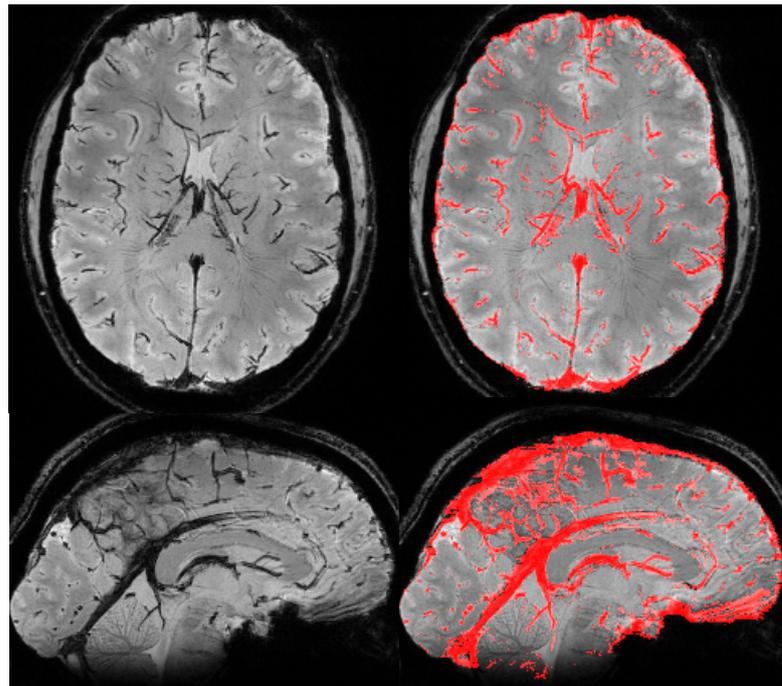


Figure 6.9: 3D surface vein rendering of the automatic segmentation of other datasets using the proposed RF approach. s3 and s4 have slightly different acquisition parameters than the datasets on which the RF models have been trained (see Table 4.1).



(a) s3 (RF-model: s1)



(b) s6 (RF-model: s5)

Figure 6.10: Automatic vein segmentation of other subjects using the RF approach overlaid on minimum intensity projections of the Magnitude image.

6.4 Feature Ranking

6.4.1 Gini Importance

As described in Section 3.3 the mean decrease of Gini impurity for a feature in the Random Forest training step can be used as a measure for the feature’s importance. The bar plots in Fig. 6.11 show these Gini importance values obtained by training a RF with the full multi-echo feature set on all annotated data of s5 or s1 respectively. In both cases the vesselness value computed on the (normalized) magnitude of the second echo and the vesselness and intensity value of the QSM image of the first echo show the highest importance values. Furthermore it is apparent from these importance values that the first echo contains most of the information for the vein segmentation task.

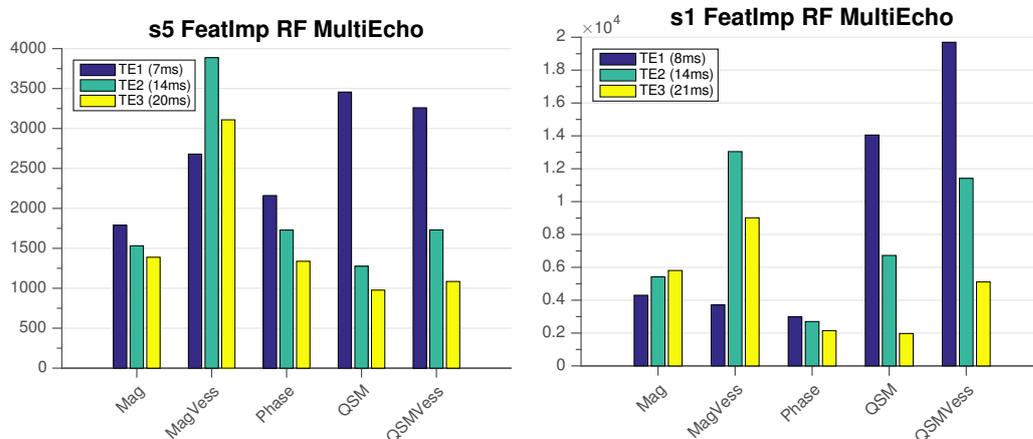


Figure 6.11: Feature importances for a Random Forest trained on all features of the multi-echo datasets s5 (*left*), and one trained on s1 (*right*).

6.4.2 Segmentation Performance of Feature Subsets

In order to further investigate the importance of different features for venous vessel segmentation, the RF training and segmentation was performed using only subsets of features for training and segmentation. Tables 6.4 and 6.5 show the performance scores of those feature-subsets for both annotated datasets, obtained by intra dataset cross-validation (see Section 5.2.4). When comparing these performance scores, one can see that including the phase feature additionally to the Mag. and QSM features does not add a lot of value, with Dice scores of 93.31% with the full feature-set versus 92.59% without the phase for s1, and 88.04% versus 86.44% for s5. The performance scores of using only the magnitude features (intensity and vesselness of each echo) are very similar to using only the QSM intensity and vesselness features.

When only using the features from one particular echo, the first echo yields the highest performance scores, closely followed by the second echo. This result further supports the

findings from the Gini importance measurement, that most of the information necessary for vein segmentation is contained in the shorter echoes.

In Figs. 6.12 and 6.13 the performance of single-echo Random Forest segmentations of the surface veins is demonstrated. Figure 6.12 shows 3D renderings of three segmentations of the same (multi-echo) dataset: a RF segmentation using features from all three echoes, a RF segmentation using only the features of the first echo, and manually annotated veins. In Fig. 6.12 the multi-echo RF segmentation of subject s6 is compared with single-echo RF segmentations performed on different datasets of the same subject that are acquired with single-echo sequences.

In both figures, the surface vein segmentations of the single-echo RF looks similar to the multi-echo RF segmentations. This indicates that, for surface vein segmentation purposes, it is sufficient to only use a single, short echo time.

Table 6.4: Segmentation performance of feature-subsets for RF-intra classification on dataset **s5**. The features that are used for each subset are listed in Table 5.3

Featureset	Dice [%]	Accuracy [%]	Precision [%]	Recall [%]	AUC-ROC [%]	MMC [%]
$\mathbf{f}_{\text{full,TE1}}$	84.46	90.21	85.82	83.14	94.78	79.54
$\mathbf{f}_{\text{full,TE2}}$	83.34	89.43	83.99	82.71	94.47	78.18
$\mathbf{f}_{\text{full,TE3}}$	79.91	87.22	80.34	79.49	92.56	74.31
$\mathbf{f}_{\text{MagOnly}}$	80.01	87.36	80.94	79.09	91.46	74.46
$\mathbf{f}_{\text{QSMOnly}}$	79.17	86.90	80.56	77.83	91.89	73.60
$\mathbf{f}_{\text{NoPhase}}$	86.44	91.45	87.69	85.23	96.21	81.89
$\mathbf{f}_{\text{Top3-Gini}}$	79.23	87.48	84.39	74.67	89.48	74.00
$\mathbf{f}_{\text{full,MultiEcho}}$	88.04	92.45	89.26	86.86	97.18	83.85

Table 6.5: Segmentation performance of feature-subsets for RF-intra classification on dataset **s1**. The features that are used for each subset are listed in Table 5.3.

Featureset	Dice [%]	Accuracy [%]	Precision [%]	Recall [%]	AUC-ROC [%]	MMC [%]
$\mathbf{f}_{\text{full,TE1}}$	91.89	94.77	92.75	91.06	97.94	88.66
$\mathbf{f}_{\text{full,TE2}}$	91.16	94.27	91.54	90.79	97.82	87.67
$\mathbf{f}_{\text{full,TE3}}$	89.33	93.06	89.36	89.31	97.10	85.29
$\mathbf{f}_{\text{MagOnly}}$	89.25	93.08	90.28	88.25	96.27	85.25
$\mathbf{f}_{\text{QSMOnly}}$	89.59	93.37	91.59	87.68	96.82	85.74
$\mathbf{f}_{\text{NoPhase}}$	92.59	95.23	93.69	91.52	98.40	89.60
$\mathbf{f}_{\text{Top3-Gini}}$	89.59	93.37	91.71	87.56	95.83	85.74
$\mathbf{f}_{\text{full,MultiEcho}}$	93.31	95.65	93.56	93.06	98.74	90.52

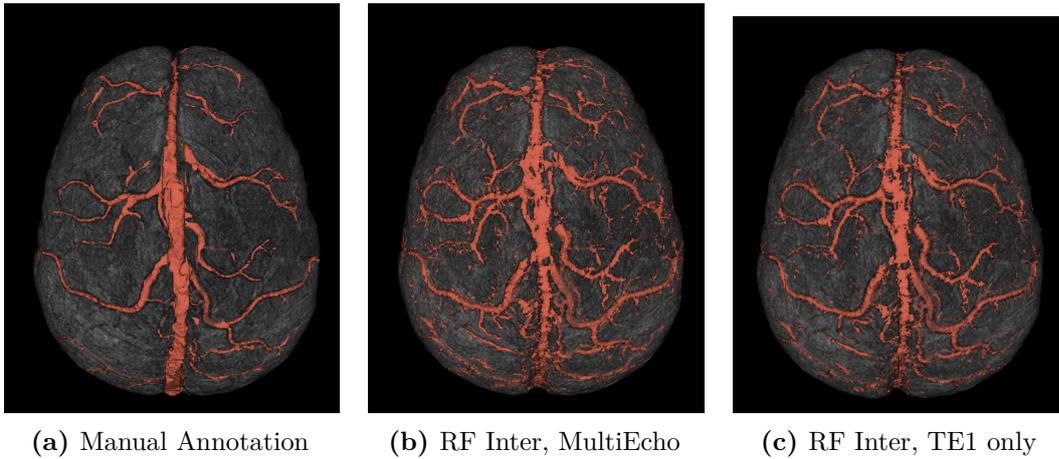


Figure 6.12: Comparison of the surface vein segmentation of s_1 obtained by using only the features of the first echo (c) with the segmentation obtained by using the full multi-echo feature set (b) and the manual annotation of the same subject (a). Both, the multi-echo and the single-echo RF models are trained on the annotated data of s_5 .

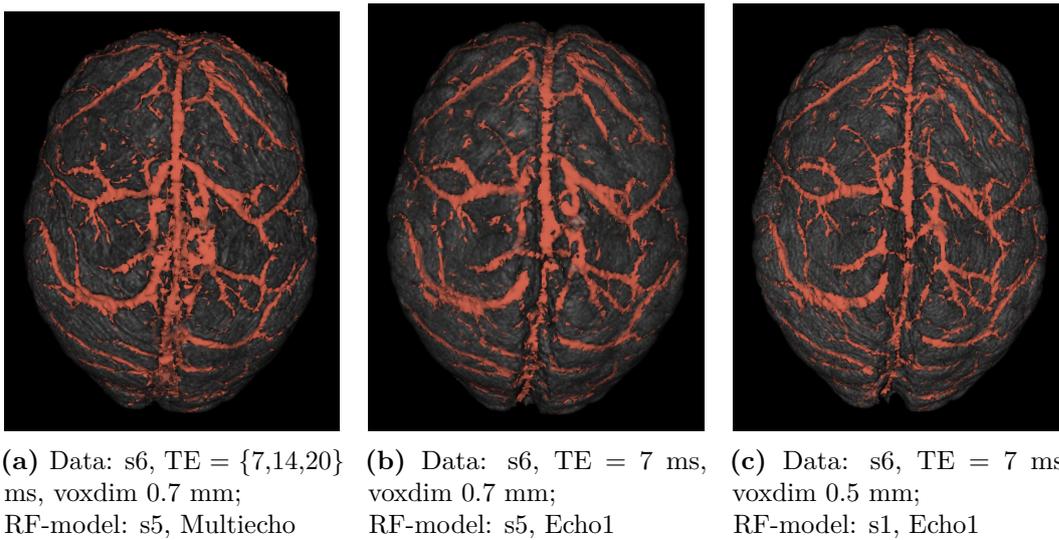


Figure 6.13: Comparison of RF surface vein segmentations of subject s_6 using data acquired in 3 different scans (see Table 4.1 for acquisition details): (a) is the image from Fig. 6.9c using multi-echo acquisition and a RF-model trained on all features of s_5 . (b) and (c) show the surface veins of the same subject segmented from datasets acquired with single-echo sequences using a short echo time of 7 ms and a voxel size of 0.7 mm and 0.5 mm respectively. The RF-models that are used for segmenting (b) and (c) are trained using only features of the first echo of s_5 (b) or s_1 (c).

This page intentionally left blank.

Discussion and Conclusion

Patient-specific visualization and modeling of the cerebral vasculature can provide valuable information for image-guided neurosurgical procedures and can thereby help to minimize the risk of post-operative neurological deficits. In particular, 3D models of the surface vasculature have the potential to be used as landmarks during surgery. In this project, a method for segmentation of cerebral veins from ultra-high field susceptibility-based MR imaging was developed. The presented method uses a Random Forest (RF) classifier to segment veins from a combination of appearance and shape features that are computed separately from magnitude images, phase images, and QSM images of a multi-echo T_2^* -weighted Gradient Echo scan. The RF approach was demonstrated to enable automatic segmentation of veins inside the parenchyma and, more importantly, on the brain's surface, where standard methods such as vesselness filtering [8, 6] tend to undersegment the veins. The results further indicate that most of the information that is needed for (surface) vein segmentation is already contained in the first echo, which potentially enables quicker data acquisition.

In the following sections the results that were reported in the previous chapter are discussed and compared with the literature. First the susceptibility-based image processing and the vesselness parameter iteration are discussed in Sections 7.1 and 7.2 respectively. Sections 7.3 and 7.4 then deal with the main contributions of the thesis: the segmentation performance of the proposed Random Forest approach and the importance of the individual features for venous vessel segmentation.

7.1 Susceptibility-based image processing

QSM computation In the original TGV-QSM publication [26], BET was used to create an initial brainmask, which is then eroded to ensure the exclusion of noisy phase regions outside the brain. This approach, however, cuts away the surface vasculature.

The SPM-based brainmask generation method was more accurate than BET. QSMs generated with SPM-based brainmask showed surface veins more clearly.

A limitation of the SPM-based brainmask computation is, however, that the tight brainmask cuts off significant parts of the Superior Sagittal Sinus (SSS) and can in some cases cut off parts of other surface veins. This can be mitigated by adjusting the disk-size parameter of MATLABs *imclose* function which is used in the creation of the SPM mask. Using other methods of background field removal in QSM such as [48] might potentially enable the visualization of the full venous vasculature of the cerebral surface including the SSS.

Influence of echo time on venous contrast in magnitude and QSM images

A visual comparison of the venous contrast in magnitude and QSM images from the different echo times was performed. In the magnitude images, the contrast for veins was best developed in the second and third echoes (14 ms and 20 ms). This is in good agreement with the findings from Koopmans et. al. in [6], who reported optimal venous contrast in magnitude images at 15 ms for 7T acquisitions. For the calculated QSM images, the veins were most clearly visible at the shortest echo time (7 ms). This is consistent with [5], where, for imaging veins at 7 T, Haacke et. al. suggest an echo time of “6 ms instead of the usual 10 ms used for SWI to reduce phase aliasing effects, in order to reduce streaking and improve the accuracy of susceptibility mapping”. This difference in echo time for optimal vein contrast in magnitude vs. QSM images indicates that a multi-echo acquisition strategy might be beneficial for imaging veins.

7.2 Vesselness Parameter Iteration

The performance of venous vessel segmentations that are obtained by vesselness filtering either magnitude images or QSM images was evaluated. All results presented for that segmentation approach have to be considered as upper bound since the vesselness parameters and the segmentation threshold were optimized for each image individually using the manual annotations.

Vesselness filtering on the magnitude images Even given those optimized parameters, the vesselness filtering approach tends to undersegment surface veins when used on magnitude images. Issues of the vesselness filter failing to accurately detect surface veins from magnitude and SWI images have also been reported by Koopmans et. al. in [6] and by Bériault et. al. in [9, 11].

Vesselness filtering on the QSM images In QSM images that are computed with the tight SPM brainmask, the vesselness filtering approach did segment the surface veins quite well, as the comparison with the manual annotation in Fig. 6.7 showed. Inside the brain, however, the contrast for veins in QSM was not as good as in magnitude images. Therefore, vesselness filtering the QSM did miss some of the veins inside the parenchyma,

while also generating some obvious false positives. To the author’s knowledge, the only publication that also reported vesselness filtering on QSM (at 3 T) is the ISMRM abstract by Ward et. al. [10]. They reported low performance scores for vesselness filtering (DICE=0.27 for QSM and DICE = 0.46 for SWI). Supposedly, one reason for the low scores is that they used Otsu thresholding [34] to transform the vesselness output to a binary segmentation. They did not report the vesselness parameters that were used.

Influence of the echo time The vesselness filtering based segmentation approach was performed on each of the three echoes ($\{7, 14, 20\}$ ms for s5 and $\{8, 14, 21\}$ ms for s1) individually. For vesselness filtering on the magnitude images, the two longer echo times (14 and 20/21 ms) yielded the higher performance scores, while for vesselness filtering QSM images, the shorter echo time (7/8 ms) performed best.

7.3 Segmentation Performance of the Multi-Echo RF-Approach

The segmentation approach that was developed within this thesis uses a Random Forest algorithm to combine appearance and shape features from magnitude images, phase images, and QSM images of a multi-echo T_2^* weighted gradient echo scan. In all of the performance metrics that were used within our experiments, the RF approach yielded higher scores than either of those features used individually. Specifically, the RF approach outperformed the common vesselness filtering approach in all similarity measures that were computed against the manual annotations. Visual assessment of 3D renderings of the surface veins confirmed that the segmentations obtained by the RF approach did look very similar to renderings of the manual annotations. Generality of the approach was shown by segmenting other datasets, some of which were acquired with slightly differing acquisition parameters.

It has to be noted that quantitative evaluation that was performed within this thesis is limited to regions that are manually annotated. This means that misclassifications in regions that do not have manual annotations are not reflected in these measurements. However, the used manual annotations are focused on regions that are of particular interest for presurgical planning and regions that are reportedly difficult for automatic segmentation such as the brain surface, the mid sagittal region, and the basal ganglia.

Furthermore the the manual annotations were performed by the author of this diploma. To validate the quality of these manual annotation a second annotator (ideally a radiologist) could annotate the same parts of the dataset.

One limitation of the proposed RF segmentation approach is that segmentation is performed on a voxel-by-voxel basis. While some amount of information about each voxel’s neighborhood is included in the vesselness features, including further neighborhood information could help to smooth the segmentations and get rid of isolated false positives. This could for example be done with Markov Random Fields or Conditional Random

Fields, as done in [9, 11, 10]. Another possibility would be to further enforce tubularity by performing Vessel Enhancing Diffusion [13] on the probabilistic output of the Random Forest.

Another obvious limitation is that the Random Forest classification is a supervised machine-learning approach, i.e. that the RF segmentation model is learned from manually labeled training samples. Although we successfully segmented datasets, that were acquired with slightly different scan parameters than the dataset on which the RF model was trained on, the trained RF models are expected to fail at other field strengths or when the acquisition parameters differ significantly. Further experiments would need to be performed to determine the degree of accuracy loss that is introduced when varying the data acquisition. Ideally the training set is acquired with the exact same sequence as the datasets that are to be segmented.

7.4 Feature Importance

The Gini importance measurement, which is integrated in the RF algorithm, was used to investigate the importance of each feature from the multi-echo dataset. Since the Gini importance values are a byproduct of the training step, these measurements could only be performed on the two annotated datasets. In both datasets the vesselness of the second echo magnitude image and the vesselness and intensity value of the first echo QSM showed the highest importance values. The results further indicated that most information necessary for venous vessel segmentation is already contained in the first echo. It's important to keep in mind though, that these results represent those regions of the brain that are used for training the RF, i.e. those voxels that were manually annotated¹. Since special emphasis was put on the segmentation of the surface veins, surface voxels are represented to a disproportionately high degree in the manual annotations. The measured feature importance thus mainly describes the importance for surface vein segmentation.

Finally, we investigated the segmentation performance of RF models that only use subsets of the features. The results indicate that the first echo does in fact contain enough information to accurately segment the surface vasculature. This could potentially lead to shorter measurement times as using a short echo time allows the repetition time to be reduced.

7.5 Conclusion

In this thesis, a method for automatic segmentation of cerebral veins from susceptibility-based MR imaging was developed. The combination of features from magnitude, phase, and QSM images has been shown to be more accurate than the standard vesselness filtering approach, especially for the challenging task of surface vein segmentation. The

¹For the feature importance measurements, RF models were trained on all annotated voxels.

proposed segmentation framework was implemented in MATLAB, which enables fully automatic whole-brain venous vessel segmentation.

This page intentionally left blank.

Vesselness Parameter Iteration

In the following pages plots of the quantitative evaluation scores for vesselness filtering with different parameters on each of the echoes of the two annotated datasets s1 and s5 are provided. A.1 shows plots for the Dice similarity coefficient; A.2 shows the plots for accuracy; A.3 the Matthews Correlation Coefficient (MMC); and A.4 the Area Under the Receiver Operating Characteristics Curve (AUC_ROC).

In order to compare the the vesselness filtering output with the manually annotated labels, the filter output is turned into a binary segmentation by thresholding. For each trial the threshold was optimized w.r.t. the respective evaluation score (Dice, Accuracy, or MMC).

A.1 Dice

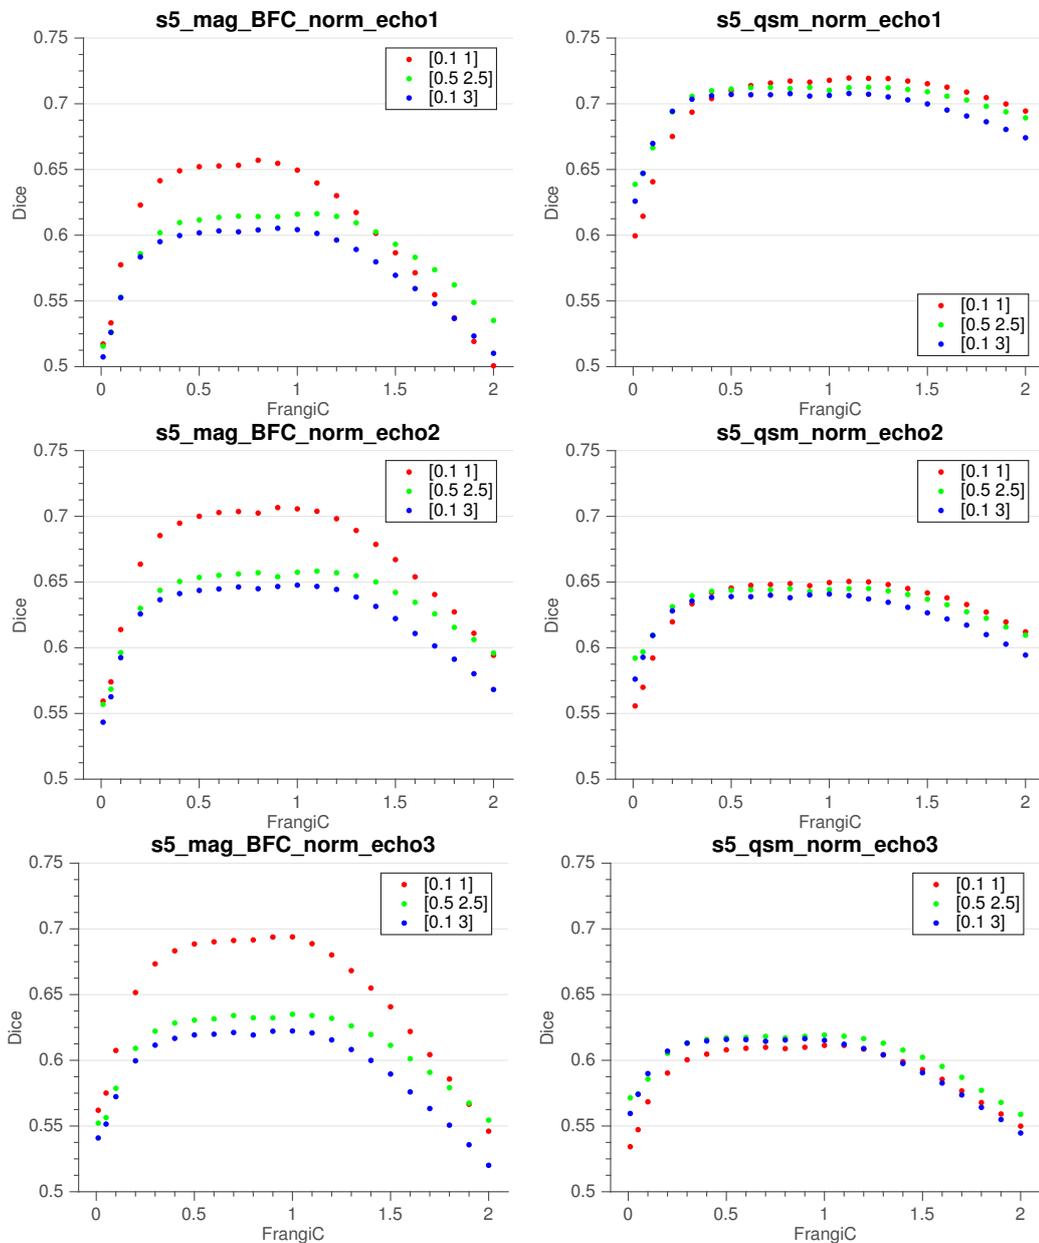


Figure A.1: Dice Scores for Vesselness Parameter Iterations on s5

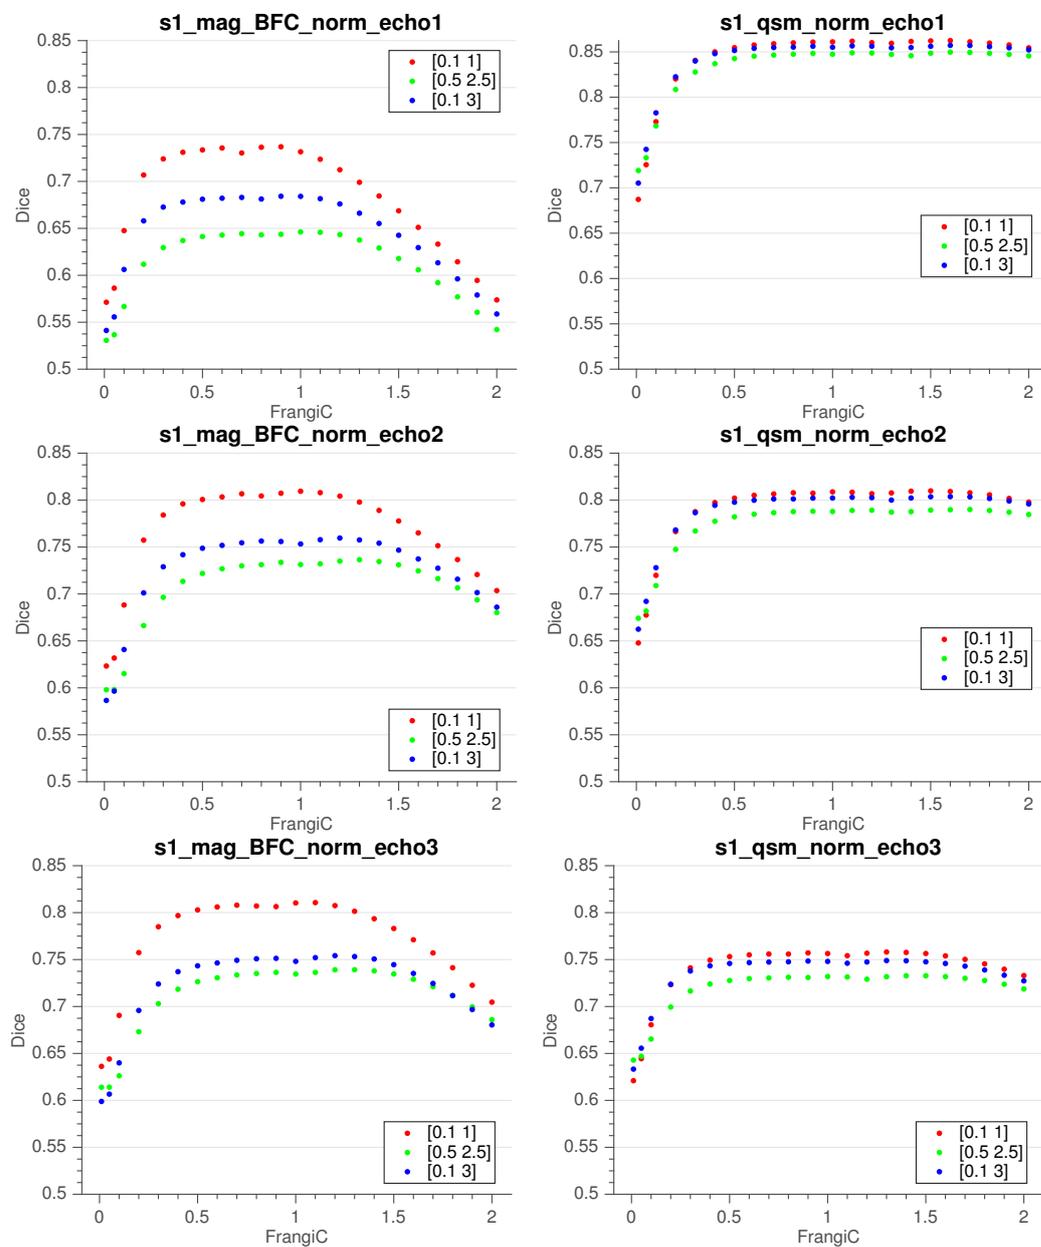


Figure A.2: Dice Scores for Vesselness Parameter Iterations on s1

A.2 Accuracy

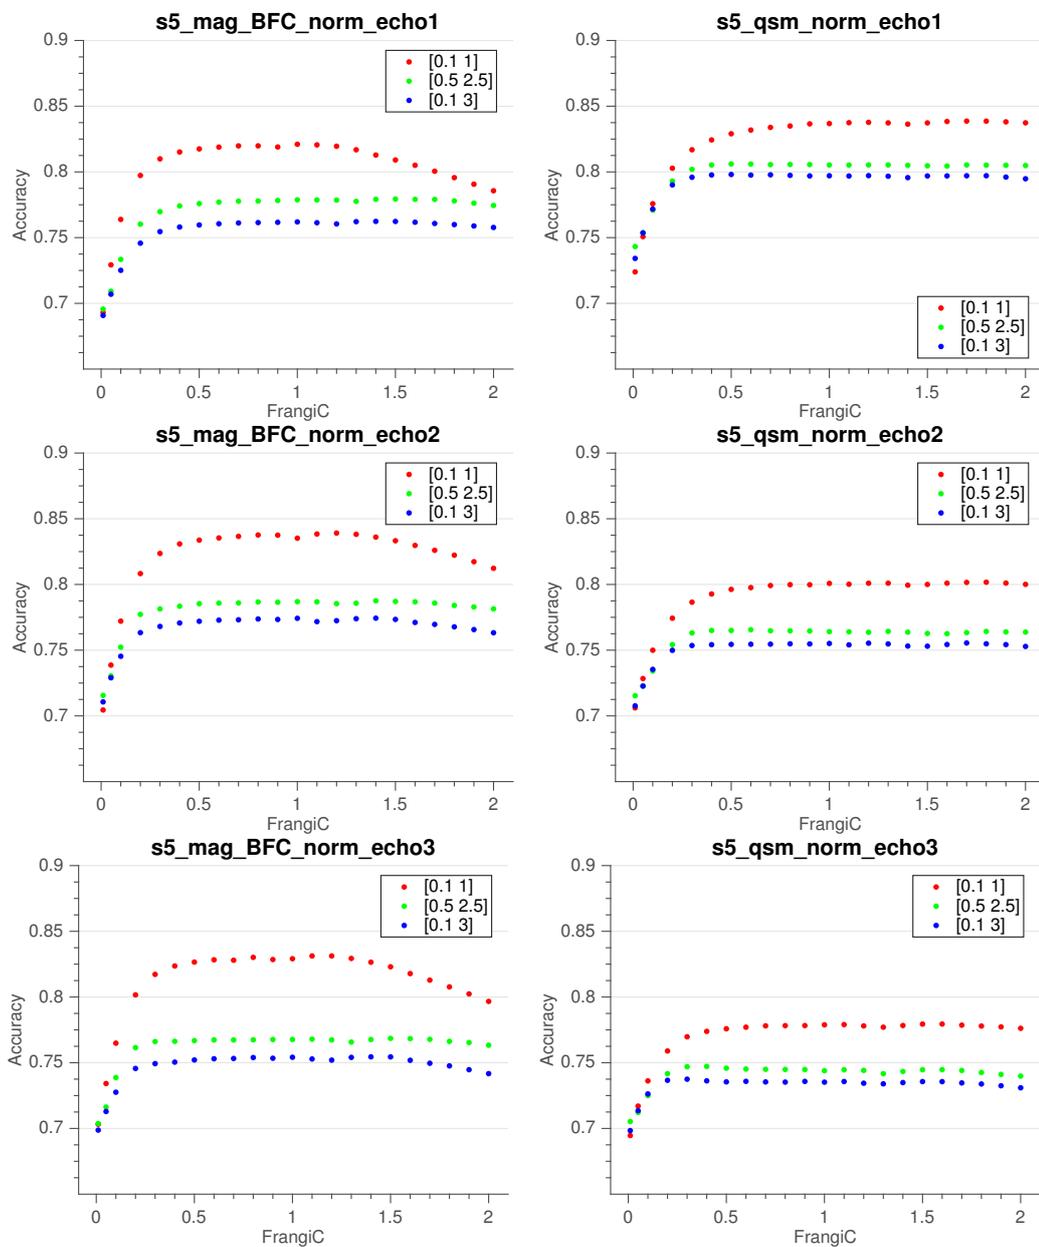


Figure A.3: Accuracy for Vesselness Parameter Iterations on s5

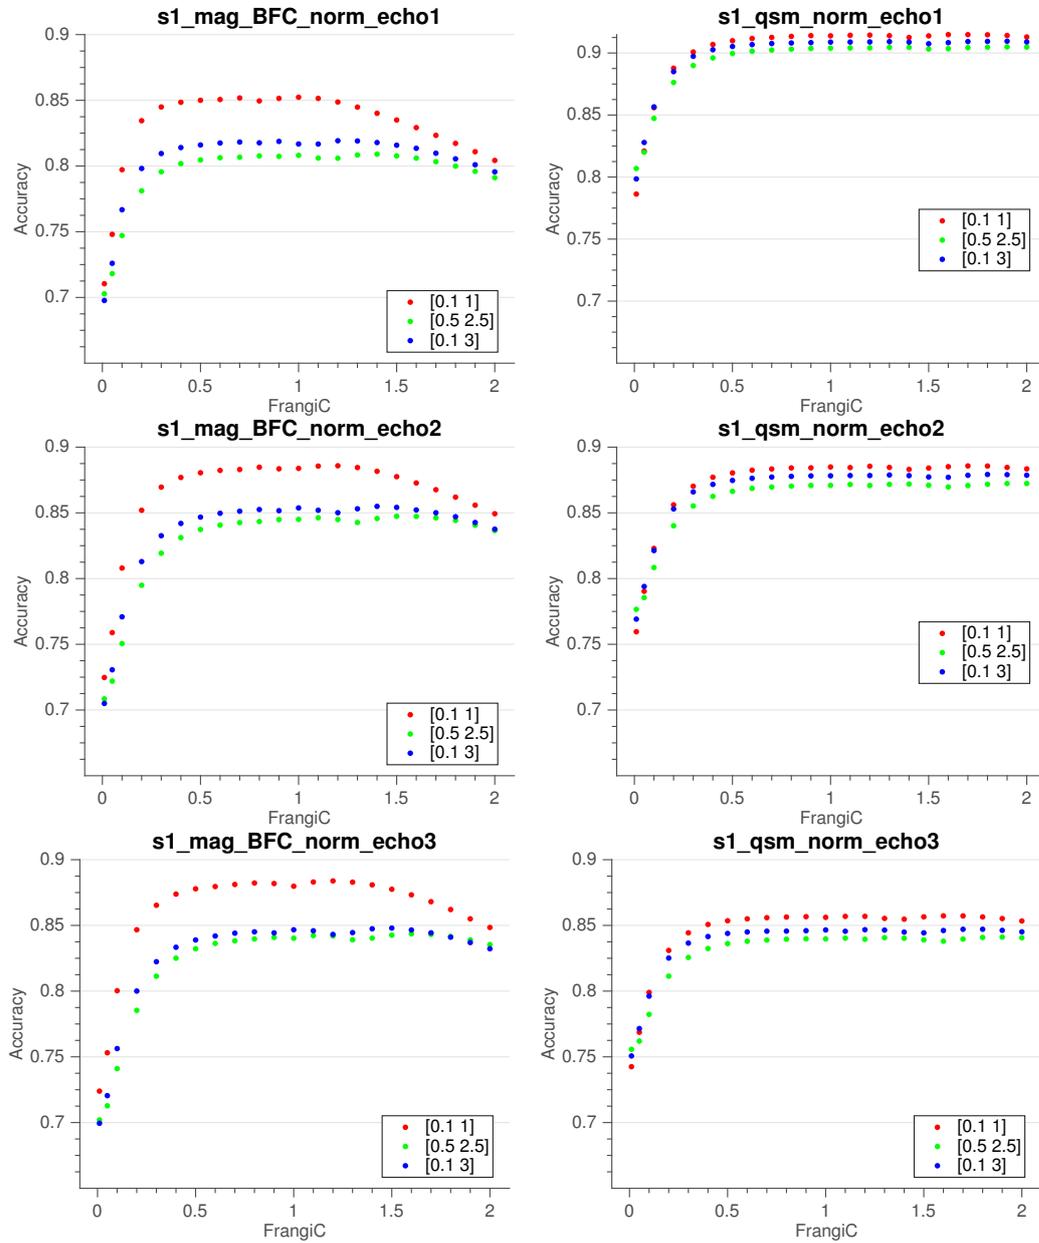


Figure A.4: Accuracy for Vesselness Parameter Iterations on s1

A.3 Matthews Correlation Coefficient (MMC)

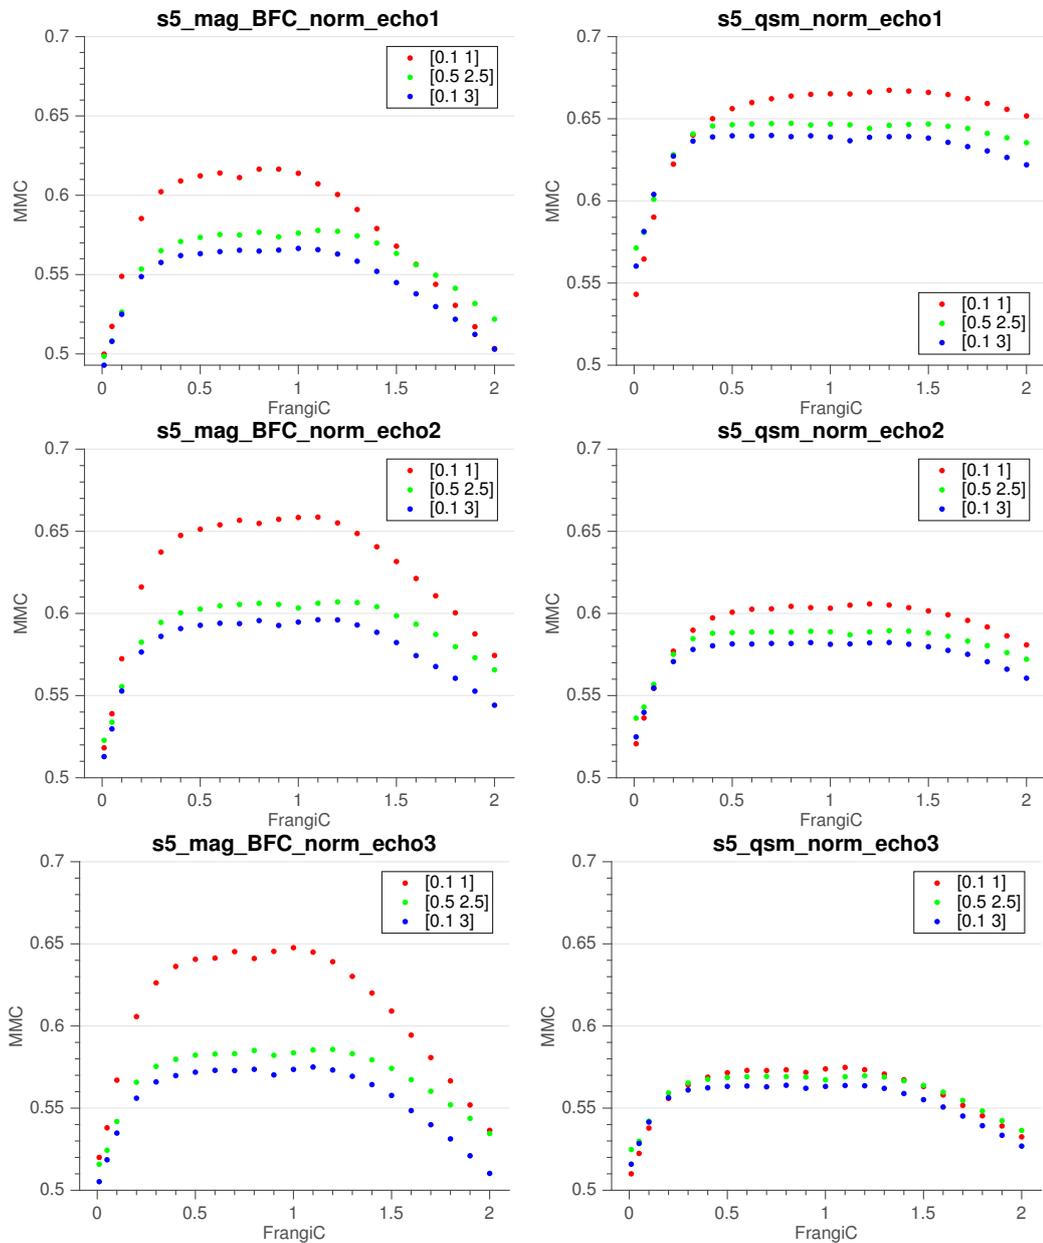


Figure A.5: Matthews Correlation Coefficient (MCC) for Vesselness Parameter Iterations on s5

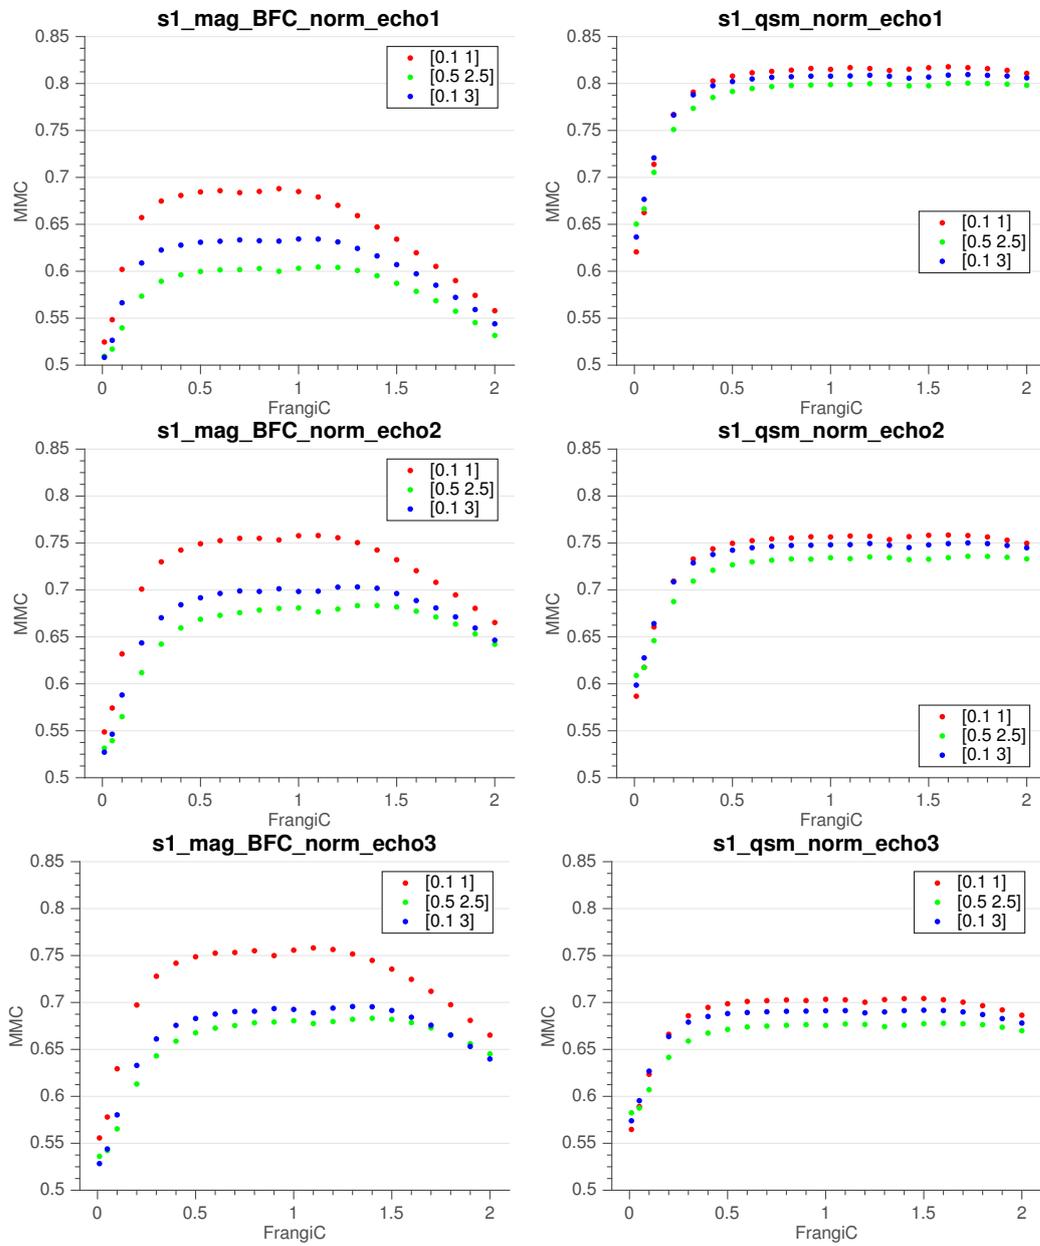


Figure A.6: Matthews Correlation Coefficient (MCC) for Vesselness Parameter Iterations on s1

A.4 Area under the Curve: Receiver Operating Characteristics

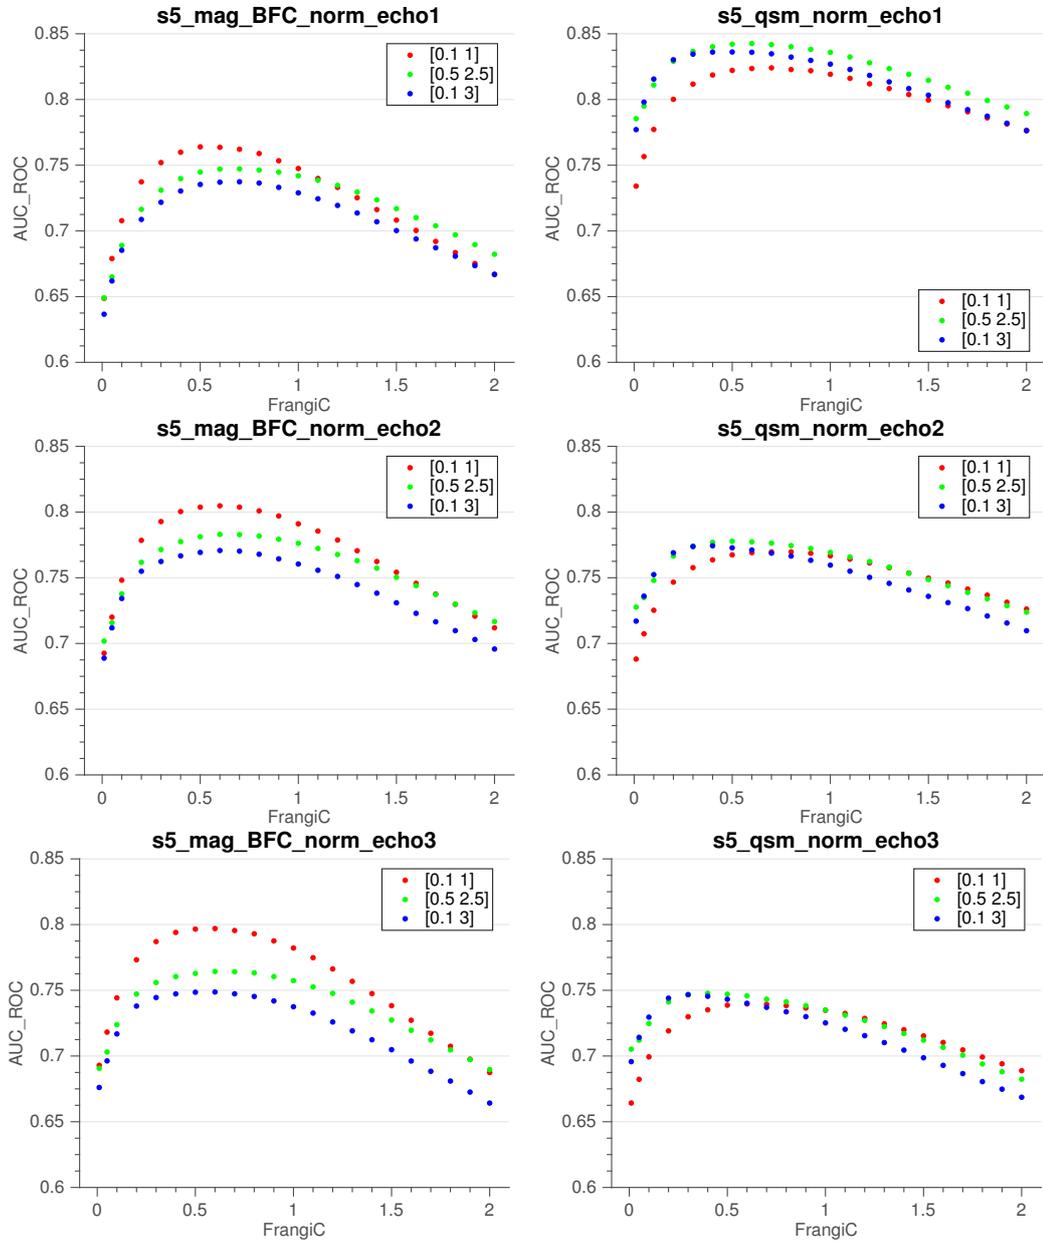


Figure A.7: Area under the Receiver Operating Characteristics Curve (AUC-ROC) for Vesselness Parameter Iterations on s5

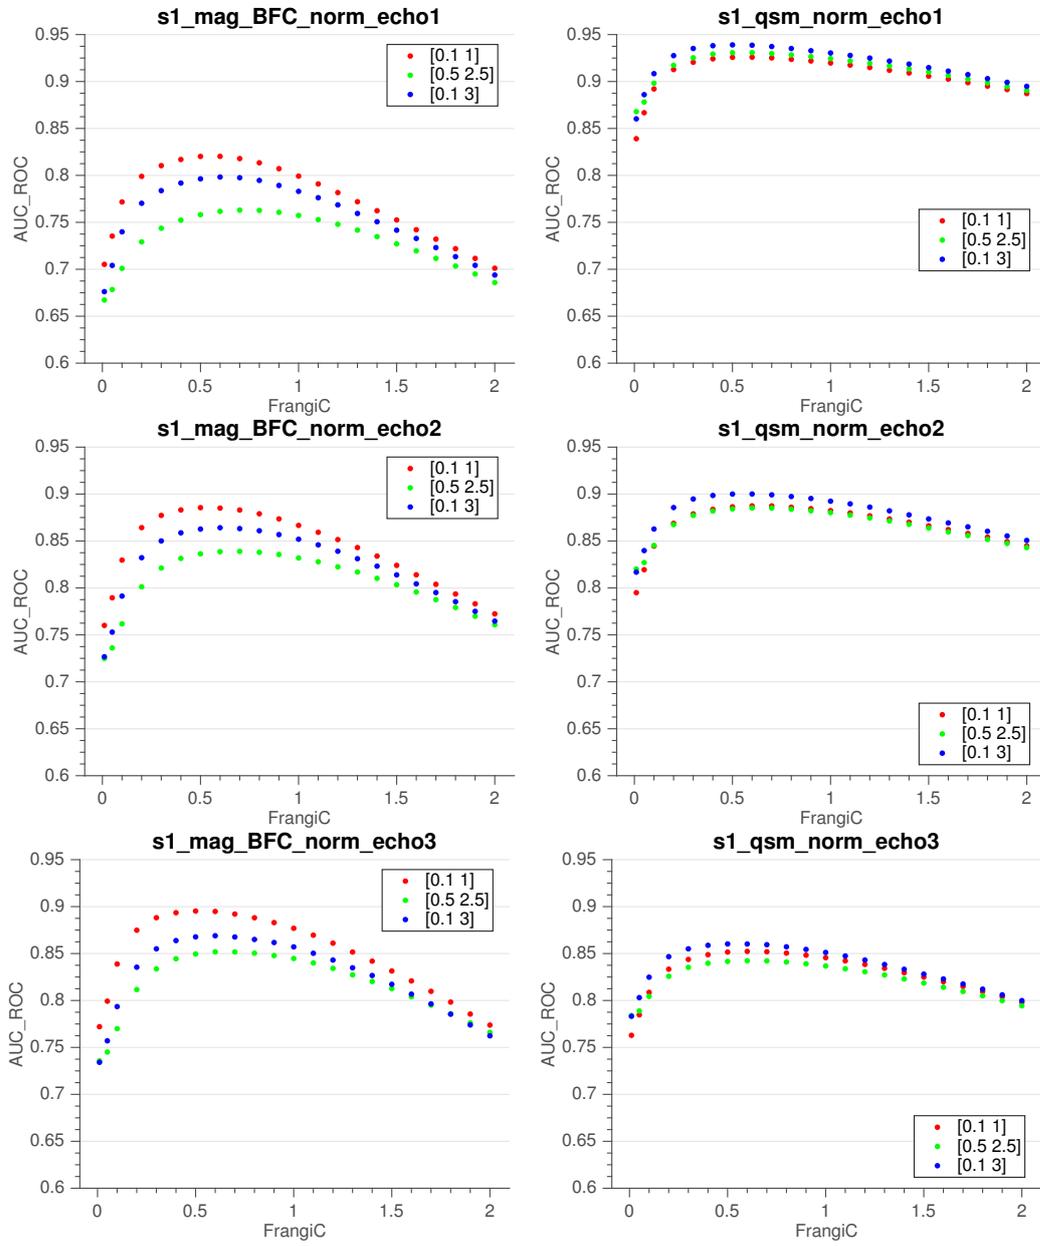


Figure A.8: Area under the Receiver Operating Characteristics Curve (AUC-ROC) for Vesselness Parameter Iterations on s1

List of Figures

2.1	Magnitude and Phase of the Magnetization	8
2.2	Spins within B_0 . Microscopic and Macroscopic Picture	9
2.3	Effect of a 90° pulse on a spin ensemble	10
2.4	Image Acquisition process. Fourier Transform	13
2.5	2D Gradient Echo Sequence	14
2.6	Magnetic susceptibility values of different materials and tissues	16
2.7	Magnetic Field Perturbation of Cylindrical Structure	17
2.8	Effects of field inhomogeneities on the spins at Vein-Tissue interface	18
2.9	SWI processing steps: overview	19
2.10	Original (wrapped) and unwrapped phase	20
2.11	Functions for creating a Phase Mask	21
2.12	Vein appearance in phase images vs QSM images	23
3.1	Illustrations relating to Vesselness Filtering	27
3.2	Example of a 2D Feature Space Classification	32
5.1	Flowchart: Vein Segmentation Pipeline Overview	41
5.2	Screenshot of view during the manual annotation process.	44
5.3	Manual Annotations of datasets s1 and s5	46
6.1	Surface vein appearance in a magnitude image and two QSM images computed with two different brainmasks.	54
6.2	Comparison of Magnitude, Phase and QSM images at different Echo Times	55
6.3	Mean Dice score for Frangi vesselness Parameter Iteration.	56
6.4	RF Parameter Optimization	57
6.5	Quantitative Comparison of Segmentation Methods	58
6.6	Visualization of vein segmentation of s5.	61
6.7	Surface Vein Segmentation of s1	62
6.8	Surface Vein Segmentation of s5	63
6.9	Surface Vein Segmentation of other subjects.	63
6.10	Automatic vein segmentation of other subjects using the RF approach overlaid on minimum intensity projections of the Magnitude image.	64
6.11	Feature Importance	65
6.12	Surface Vein Segmentation of s1 using only the first echo	67

6.13	Comparison of single echo vs multi echo segmentation of s6	67
A.1	Dice Scores for Vesselness Parameter Iterations on s5	76
A.2	Dice Scores for Vesselness Parameter Iterations on s1	77
A.3	Accuracy for Vesselness Parameter Iterations on s5	78
A.4	Accuracy for Vesselness Parameter Iterations on s1	79
A.5	Matthews Correlation Coefficient (MCC) for Vesselness Parameter Iterations on s5	80
A.6	Mathews Correlation Coefficient (MCC) for Vesselness Parameter Iterations on s1	81
A.7	Area under the Receiver Operating Characteristics Curve (AUC-ROC) for Vesselness Parameter Iterations on s5	82
A.8	Area under the Receiver Operating Characteristics Curve (AUC-ROC) for Vesselness Parameter Iterations on s1	83

List of Tables

1.1	Previous Work	3
2.1	Magnetic Field inside and outside Sphere and Cylinder	17
3.1	Eigenvalues of the Hessian matrix and corresponding 3D image structure	27
4.1	Available Data	38
4.2	Performance Metrics	39
5.1	Number of manually annotated Voxels within dataset s1 and s5.	45
5.2	Parameters for Frangi Vesselness Parameter Iteration on the normalized magnitude and QSM images.	47
5.3	List of Feature Subsets used for evaluation.	51
6.1	Best Frangi Parameters	56
6.2	Quantitative Performance Metrics for s5	59
6.3	Quantitative Performance Metrics for s1	59
6.4	Performance of RF with feature-subsets on s5.	66
6.5	Performance of RF with feature-subsets on s1.	66

Bibliography

- [1] I.-F. Talos, L. O'Donnell, C.-F. Westin, S. K. Warfield, W. Wells III, S.-S. Yoo, L. P. Panych, A. Golby, H. Mamata, S. S. Maier, *et al.*, "Diffusion tensor and functional MRI fusion with anatomical MRI for image-guided neurosurgery," *Medical Image Computing and Computer-Assisted Intervention*, vol. 2878, pp. 407–415, 2003.
- [2] A. Nabavi, P. McL. Black, D. T. Gering, C. F. Westin, V. Mehta, R. S. Pergolizzi, M. Ferrant, S. K. Warfield, N. Hata, R. B. Schwartz, W. M. Wells, R. Kikinis, and F. A. Jolesz, "Serial intraoperative magnetic resonance imaging of brain shift," *Neurosurgery*, vol. 48, no. 4, pp. 787–798, 2001.
- [3] J. R. Reichenbach, R. Venkatesan, D. J. Schillinger, D. K. Kido, and E. M. Haacke, "Small vessels in the human brain: MR venography with deoxyhemoglobin as an intrinsic contrast agent.," *Radiology*, vol. 204, no. 1, pp. 272–277, 1997.
- [4] K. Shmueli, J. A. de Zwart, P. van Gelderen, T.-Q. Li, S. J. Dodd, and J. H. Duyn, "Magnetic susceptibility mapping of brain tissue in vivo using MRI phase data.," *Magnetic Resonance in Medicine*, vol. 62, no. 6, pp. 1510–1522, 2009.
- [5] E. M. Haacke, S. Liu, S. Buch, W. Zheng, D. Wu, and Y. Ye, "Quantitative susceptibility mapping: current status and future directions," *Magnetic Resonance Imaging*, vol. 33, no. 1, pp. 1–25, 2015.
- [6] P. J. Koopmans, R. Manniesing, W. J. Niessen, M. A. Viergever, and M. Barth, "MR venography of the human brain using susceptibility weighted imaging at very high field strength," *Magnetic Resonance Materials in Physics, Biology and Medicine*, vol. 21, pp. 149–158, 2008.
- [7] Z. Jin, L. Xia, and Y. P. Du, "Enhancement of Venous Vasculature in Susceptibility Weighted Images of the Brain Using Multi-Scale Vessel Enhancement Filtering," in *International Conference on Biomedical Engineering and Informatics*, vol. 3, pp. 226–230, IEEE, 2010.
- [8] A. Deistung, M. Kocinski, P. Szczypinski, A. Materka, and J. R. Reichenbach, "Segmentation of Venous Vessels using Multi-Scale Vessel Enhancement Filtering in Susceptibility Weighted Imaging," *International Society for Magnetic Resonance in Medicine*, vol. 14, p. 1948, 2006.

- [9] S. Bériault, M. Archambault-Wallenburg, A. Sadikot, D. Louis Collins, and G. Bruce Pike, “Automatic Markov Random Field Segmentation of Susceptibility-Weighted MR Venography,” in *Clinical Image-Based Procedures*, vol. 8361 of *Lecture Notes in Computer Science*, pp. 39–47, Springer International Publishing, 2014.
- [10] P. G. D. Ward, N. J. Ferris, A. C. L. Ng, D. G. Barnes, D. L. Dowe, G. F. Egan, and P. Raniga, “Venous segmentation using Gaussian mixture models and Markov random fields,” in *Proceedings of the ISMRM*, vol. 23, p. 3759, 2015.
- [11] S. Bériault, Y. Xiao, D. L. Collins, and G. B. Pike, “Automatic SWI Venography Segmentation Using Conditional Random Fields,” *IEEE Transactions on Medical Imaging*, vol. 34, no. 12, pp. 2478–2491, 2015.
- [12] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, “Multiscale vessel enhancement filtering,” *Medical Image Computing and Computer-Assisted Intervention - MICCAI’98. Lecture Notes in Computer Science, vol 1496*, vol. 1496, pp. 130–137, 1998.
- [13] R. Manniesing, M. A. Viergever, and W. J. Niessen, “Vessel Enhancing Diffusion: A Scale Space Representation of Vessel Structures.,” *Medical Image Analysis*, vol. 10, no. 6, pp. 815–825, 2006.
- [14] L. Breiman, “Random Forests,” *Machine Learning*, vol. 45, pp. 5–32, 2001.
- [15] M. A. Brown and R. C. Semelka, *MRI: Basic Principles and Applications*. John Wiley & Sons, 4 ed., 2010.
- [16] R. A. De Graaf, *In Vivo NMR Spectroscopy: Principles and Techniques*. John Wiley & Sons, 2007.
- [17] I. P. Arlart, G. M. Bongartz, and G. Marchal, *Magnetic Resonance Angiography*. Springer-Verlag Berlin Heidelberg, 2002.
- [18] E. M. Haacke and J. R. Reichenbach, *Susceptibility weighted imaging in MRI: basic concepts and clinical applications*. John Wiley & Sons, 2014.
- [19] E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic Resonance Imaging: Physical Principles and Sequence Design*, vol. 1. Wiley, 1999.
- [20] H. K. Song, “Physics of Magnetic Resonance,” in *Diagnostic Radiology Physics: A Handbook for Teachers and Students*, ch. 14, International Atomic Energy Agency, 2014.
- [21] M. A. Bernstein, K. F. King, and X. J. Zhou, *Handbook of MRI Pulse Sequences*. Elsevier, 2004.
- [22] E. M. Haacke, S. Mittal, Z. Wu, J. Neelavalli, and Y.-C. N. Cheng, “Susceptibility-weighted imaging: technical aspects and clinical applications, part 1.,” *American Journal of Neuroradiology*, vol. 30, no. 1, pp. 19–30, 2009.

- [23] J. F. Schenck, “Safety of strong, static magnetic fields.,” *Journal of Magnetic Resonance Imaging*, vol. 12, no. 1, pp. 2–19, 2000.
- [24] R. Bowtell, “Magnetic susceptibility of biological tissues.” Presentation during *ESMRMB Lecture on MR: Susceptibility weighted imaging and quantitative mapping*, Graz, 2014.
- [25] E. M. Haacke, Y. Xu, Y.-C. N. Cheng, and J. R. Reichenbach, “Susceptibility weighted imaging (SWI).,” *Magnetic Resonance in Medicine*, vol. 52, no. 3, pp. 612–618, 2004.
- [26] C. Langkammer, K. Bredies, B. A. Poser, M. Barth, G. Reishofer, A. P. Fan, B. Bilgic, F. Fazekas, C. Mainero, and S. Ropele, “Fast quantitative susceptibility mapping using 3D EPI and total generalized variation,” *NeuroImage*, vol. 111, pp. 622–630, 2015.
- [27] J. R. Reichenbach, M. Barth, E. M. Haacke, M. Klarhöfer, W. A. Kaiser, and E. Moser, “High-resolution MR venography at 3.0 Tesla.,” *Journal of Computer Assisted Tomography*, vol. 24, no. 6, pp. 949–957, 2000.
- [28] D. C. Noll, D. G. Nishimura, and A. Macovski, “Homodyne Detection in Magnetic Resonance Imaging,” *IEEE Transactions on Medical Imaging*, vol. 10, no. 2, pp. 154–163, 1991.
- [29] A. Rauscher, M. Barth, J. R. Reichenbach, R. Stollberger, and E. Moser, “Automated unwrapping of MR phase images applied to BOLD MR-venography at 3 tesla,” in *Journal of Magnetic Resonance Imaging*, vol. 18, pp. 175–180, 2003.
- [30] S. Witoszynskyj, A. Rauscher, J. R. Reichenbach, and M. Barth, “Phase unwrapping of MR images using Φ UN—a fast and robust region growing algorithm,” *Medical Image Analysis*, vol. 13, no. 2, pp. 257–268, 2009.
- [31] M. Jenkinson, “Fast, automated, N-dimensional phase-unwrapping algorithm,” *Magnetic Resonance in Medicine*, vol. 49, no. 1, pp. 193–197, 2003.
- [32] M. A. Schofield and Y. Zhu, “Fast phase unwrapping algorithm for interferometric applications.,” *Optics letters*, vol. 28, no. 14, pp. 1194–1196, 2003.
- [33] M. Rajchl, *Rapid Segmentation Techniques for Cardiac and Neuroimage Analysis*. PhD thesis, The University of Western Ontario, 2014.
- [34] N. Otsu, “A threshold selection method from gray-level histograms,” *IEEE Transactions on Systems, Man and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.
- [35] Z. Al-Rekabi, K. McPhee, K. Fitzpatrick, B. Bjornson, and A. Rauscher, “Presurgical Visualization of Cerebral Surface Veins with Susceptibility Weighted Imaging,” in *Proceedings of the ISMRM*, vol. 17, p. 44, 2009.

- [36] D. L. Pham, C. Y. Xu, and J. L. Prince, “Current Methods in Medical Image Segmentation,” *Annual Review of Biomedical Engineering*, vol. 2, pp. 315–337, 2000.
- [37] N. S. Altman, “An introduction to kernel and nearest-neighbor nonparametric regression,” *The American Statistician*, vol. 46, no. 3, pp. 175–185, 1992.
- [38] C. Cortes and V. Vapnik, “Support-Vector Networks,” *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [39] T. Hastie, R. Tibshirani, and J. Friedman, “The Elements of Statistical Learning,” *Elements*, pp. 337–387, 2009.
- [40] K. J. Archer and R. V. Kimes, “Empirical characterization of random forest variable importance measures,” *Computational Statistics and Data Analysis*, vol. 52, no. 4, pp. 2249–2260, 2008.
- [41] C. Strobl, A.-L. Boulesteix, A. Zeileis, and T. Hothorn, “Bias in random forest variable importance measures: Illustrations, sources and a solution,” *BMC bioinformatics*, p. 25, 2007.
- [42] D. L. Parker, A. Payne, N. Todd, and J. R. Hadley, “Phase reconstruction from multiple coil data using a virtual reference coil,” *Magnetic Resonance in Medicine*, vol. 72, no. 2, pp. 563–569, 2014.
- [43] S. D. Robinson, B. Dymerska, W. Bogner, M. Barth, O. Zaric, S. Goluch, G. Grabner, X. Deligianni, O. Bieri, and S. Trattnig, “Combining phase images from array coils using a short echo time reference scan (COMPOSER),” *Magnetic Resonance in Medicine*, 2015.
- [44] B. W. Matthews, “Comparison of the predicted and observed secondary structure of T4 phage lysozyme,” *Biochimica et Biophysica Acta (BBA)-Protein Structure*, vol. 405, no. 2, pp. 442–451, 1975.
- [45] S. M. Smith, “Fast robust automated brain extraction,” *Human Brain Mapping*, vol. 17, no. 3, pp. 143–155, 2002.
- [46] J. G. Sled and G. B. Pike, “Understanding intensity non-uniformity in MRI,” in *MICCAI*, vol. 1496 of *Lecture Notes in Computer Science*, pp. 614–622, 1998.
- [47] K. Van Leemput, F. Maes, D. Vandermeulen, and P. Suetens, “Automated model-based bias field correction of MR images of the brain,” *IEEE transactions on medical imaging*, vol. 18, no. 10, pp. 885–896, 1999.
- [48] R. Topfer, F. Schweser, A. Deistung, J. R. Reichenbach, and A. H. Wilman, “SHARP edges: Recovering cortical phase contrast through harmonic extension,” *Magnetic Resonance in Medicine*, vol. 73, pp. 851–856, 2015.