

Assessment of Collateral Blood Flow in the Brain using Magnetic Resonance Imaging



Thomas W. Okell

Physics Group, FMRIB Centre,
Nuffield Department of Clinical Neurosciences
and Wolfson College

University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Supervisor: Professor Peter Jezzard

Hilary Term 2011

Assessment of Collateral Blood Flow in the Brain using Magnetic Resonance Imaging

Thomas W. Okell

Wolfson College, University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Hilary Term 2011

Abstract

Collateral blood flow is the compensatory flow of blood to the tissue through secondary channels when the primary channel is compromised. It is of vital importance in cerebrovascular disease where collateral flow can maintain large regions of brain tissue which would otherwise have suffered ischaemic damage. Traditional x-ray based techniques for visualising collateral flow are invasive and carry risks to the patient. In this thesis novel magnetic resonance imaging techniques for performing vessel-selective labelling of brain feeding arteries are explored and developed to reveal the source and extent of collateral flow in the brain non-invasively and without the use of contrast agents.

Vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL) allows the selective labelling of blood water in different combinations of brain feeding arteries that can be combined in post-processing to yield vascular territory maps. The mechanism of VEPCASL was elucidated and optimised through simulations of the Bloch equations and phantom experiments, including its sensitivity to sequence parameters, blood velocity and off-resonance effects.

An implementation of the VEPCASL pulse sequence using an echo-planar imaging (EPI) readout was applied in healthy volunteers to enable optimisation of the post-labelling delay and choice of labelling plane position. Improvements to the signal-to-noise ratio (SNR) and motion-sensitivity were made through the addition of background suppression pulses and a partial-Fourier scheme. Experiments using a three-dimensional gradient and spin echo (3D-GRASE) readout were somewhat compromised by significant blurring in the slice direction, but showed potential for future work with a high SNR and reduced dropout artefacts.

The VEPCASL preparation was also applied to a dynamic 2D angiographic readout, allowing direct visualisation of collateral blood flow in the brain as well as a morphological and functional assessment of the major cerebral arteries. The application of a balanced steady-state free precession (bSSFP) readout significantly increased the acquisition efficiency, allowing the generation of dynamic 3D vessel-selective angiograms. A theoretical model of the dynamic angiographic signal was also derived, allowing quantification of blood flow through specified vessels, providing a significant advantage over qualitative x-ray based methods.

Finally, these methods were applied to a number of patient groups, including those with vertebro-basilar disease, carotid stenosis and arteriovenous malformation. These preliminary studies demonstrate that useful clinical information regarding collateral blood flow can be obtained with these techniques.

Acknowledgements

There are a huge number of people who have provided support and insightful discussion to improve the work presented in this thesis. I would like to thank everyone involved for their help and support, and for making these past three years so enjoyable.

In particular, I would like to thank my supervisor, Peter Jezzard, for his guidance and support, for inspiring and exciting discussions on our latest work, for seemingly knowing a huge amount about almost every aspect of MR physics and beyond, for working late or at weekends to help me out with a deadline and for always finding time to answer my questions, even if he had several other urgent tasks. He has been a most excellent supervisor and mentor.

I also owe a lot to my friend and colleague, Rob Tijssen, for perhaps being one of the most patient and consistently good-natured people you are ever likely to meet. I have enjoyed greatly our discussions on MRI and he has many times saved me hours of work by pointing out an easier way to do something or providing me with a paper which answers the question I was posing. He also, without exception, tolerated my many interruptions to his work without complaint for “sanity checks”, to ask about some aspect of our work or moan about my latest problem. His good humour and company have made my D.Phil. experience greatly enjoyable. He must also go down in history for using the immortal line “let’s jump on the echo train to outer k -space”, as well as his first presentation to the physics group on BS fMRI.

One of my most fruitful collaborations has been with Michael Chappell, who also kindly tolerates my many questions and emails about his sophisticated analysis methods, as well as proof-reading parts of this thesis and other work. Without his help and support the vessel-encoded techniques described here would be at a significant disadvantage, and his advice on modelling and Bayesian methods was most useful in the development of the angiographic quantification method described herein.

I have also been greatly aided by a number of other collaborators, including Matthias Günther, David Feinberg, Xiaoming Bi and Peter Schmitt who have kindly provided excellent pulse sequences for us to modify. On the clinical side, Ursula Schulz has been critical in running the vertebro-basilar patient study, with help from Jingyi Xie, enlightening me on topics related to cerebrovascular disease and for proof reading the patient related chapter of this thesis. Natalie Voets has also been most helpful in enabling the application of the vessel-encoded tech-

niques to epilepsy patients, including those with arteriovenous malformation, and answering my naive questions about neuroscience.

FMRIB is full of people who seem to know a lot about nearly everything, and who have saved me huge amounts of time by pointing out a relevant Matlab function or describing clearly some of the more complex concepts of MR physics or image analysis. In particular, I should thank Adrian Groves for his advice on shell-scripting, Matlab and any other programming language you would care to use (and who had the misfortune of sitting near my desk and being plagued with questions), Karla Miller, for her ability to describe the most complex concepts in an easily understandable way and providing valuable guidance on SSFP and other topics, Dan Gallichan, for advice on anything related to image reconstruction or the intricacies of the MR systems we are using, Matt Robson, for his intricate knowledge of pulse sequence programming, James Meakin, for knowing far too much about MR and ASL for someone only a year into their D.Phil. and to all the other members of the physics group and FMRIB in general, who have contributed significantly to my understanding of MRI and functional imaging, as well as making FMRIB a great environment to work in.

Finally, I would like to thank my friends and family for their support during this time, but I owe the greatest debt to my wonderful wife, Helen, who has supported me through this D.Phil. (and everything leading up to it), put up with my late nights before deadlines and taken up the slack when I've had a heavy workload, for making me smile even when things weren't going well and for helping to keep everything in perspective. Your love, patience and support have kept me going, and made life outside work so enjoyable. I am certainly one of the luckiest men in the world to have you. And to my daughter Emily, who makes it impossible to keep a straight face and continuously amazes us with her latest tricks, unconventional use of English or pseudo-Greek and enthusiastic renditions of nursery rhymes on the journey to work, I thank you for making me laugh every day.

Contents

Preface	vii
Declaration of the Contributions of Others	vii
List of Publications Arising from this Thesis	vii
1 Introduction	1
1.1 Motivation	1
1.2 Thesis Outline	3
2 Background	6
2.1 Cerebral Vasculature	6
2.2 Cerebrovascular Disease	9
2.3 Principles of Magnetic Resonance Imaging	13
2.4 Arterial Spin Labelling	25
2.5 Conclusions	34
3 Current Methods for Assessment of Collateral Flow	35
3.1 Conventional Angiographic Methods	35
3.2 Indirect Tissue Perfusion Methods	44
3.3 Vessel-Specific Arterial Spin Labelling	45
3.4 Vessel-Encoded Arterial Spin Labelling	49
3.5 Conclusions	54
4 Assessment and Optimisation of Vessel-Encoded ASL	55
4.1 Introduction	55
4.2 Simulation Methods	56
4.3 Sequence and Phantom Design	61
4.4 Inversion Mechanism	63
4.5 Velocity Dependence	68
4.6 Parameter Choice for Maximal ASL Contrast	71
4.7 Spatial Modulation of Inversion Efficiency	73
4.8 Off-Resonance Effects	75
4.9 Vessel Angulation Through the Labelling Plane	77
4.10 Effect on Static Magnetisation	78
4.11 Development of Vessel-Encoded ASL Analysis Methods	81
4.12 Conclusions	98

5	Vessel-Encoded Tissue Perfusion Imaging	101
5.1	Introduction	101
5.2	Implementation with a 2D Multi-Slice EPI Readout	102
5.3	Optimisation of 2D Multi-Slice EPI	116
5.4	Moving to a 3D-GRASE Readout	126
5.5	Normal Variants of the Cerebral Vasculature	133
5.6	Conclusions	135
6	Vessel-Encoded Dynamic Angiography	136
6.1	Introduction	136
6.2	Spoiled Gradient Echo Readout	137
6.3	Readout Optimisation Using Balanced Steady-State Free Precession	163
6.4	Conclusions	181
7	Quantification of Dynamic Angiography	183
7.1	Introduction	183
7.2	Theory	185
7.3	Methods	198
7.4	Preliminary Results	200
7.5	Discussion	207
7.6	Conclusions	212
8	Patient Studies	213
8.1	General Methods	214
8.2	Vertebro-basilar Disease	215
8.3	Carotid Stenosis	221
8.4	Arteriovenous Malformation	221
8.5	Conclusions	225
9	Summary and Future Work	226
9.1	Thesis Summary	226
9.2	VEPCASL Tissue Perfusion or Angiography?	232
9.3	Future Research	233
	References	240

Preface

Declaration of the Contributions of Others

I would like to acknowledge the contribution of Dr. Michael Chappell, who developed the Bayesian analysis method for separating vascular components in the vessel-encoded data (see Section 4.11.4), in collaboration with myself. In particular, Dr. Chappell wrote the appendix on this analysis method in the referenced paper (Okell *et al.*, *Magn Reson Med*, 64 (2010), 698-706), which is paraphrased and acknowledged in this thesis.

List of Publications Arising from this Thesis

Journal Articles

1. **T. W. Okell**, M. A. Chappell, M. W. Woolrich, M. Günther, D. A. Feinberg, and P. Jezzard, "Vessel-encoded dynamic magnetic resonance angiography using arterial spin labeling," *Magn Reson Med*, vol. 64, pp. 698–706, Sep 2010.
2. M. A. Chappell, **T. W. Okell**, P. Jezzard, and M. W. Woolrich, "A general framework for the analysis of vessel encoded arterial spin labeling for vascular territory mapping," *Magn Reson Med*, vol. 64, pp. 1529–39, Nov 2010.
3. K. L. Miller, R. H. Tijssen, N. Stikov, and **T. W. Okell**, "Steady-state MRI: methods for neuroimaging," *Imaging in Medicine*, vol. 3, no. 1, pp. 93–105, 2011.

Conference Proceedings

1. M. A. Chappell, **T. W. Okell**, P. Jezzard, and M. W. Woolrich, "Vascular territory image analysis using vessel encoded arterial spin labeling," *Med Image Comput Comput Assist Interv*, vol. 12, no. Pt 2, pp. 514–21, 2009.
2. M. A. Chappell, **T. W. Okell**, P. Jezzard, and M. W. Woolrich, "A probabilistic model-based approach for the analysis of vessel encoded arterial spin labelling," in *Proceedings 15th Annual Meeting of the Organization for Human Brain Mapping*, San Francisco, p. 353, 2009.

3. **T. W. Okell**, M. Günther, D. Feinberg, and P. Jezzard, "A new method for selective dynamic MRI angiography using arterial spin labeling," in *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu*, p. 1330, 2009.
4. M. A. Chappell, **T. W. Okell**, P. Jezzard, and M. W. Woolrich, "Resolving arterial contributions in vessel encoding dynamic angiography," in *Proceedings 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Stockholm*, p. 4422, 2010.
5. M. A. Chappell, **T. W. Okell**, P. Jezzard, and M. W. Woolrich, "A general framework for the analysis of vessel encoded arterial spin labelling," in *Proceedings 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Stockholm*, p. 675, 2010.
6. **T. W. Okell**, M. A. Chappell, M. Günther, and P. Jezzard, "A comparison of 3D-GRASE and EPI for vessel-encoded arterial spin labeling," in *Proceedings 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Stockholm*, p. 4071, 2010.
7. **T. W. Okell**, M. A. Chappell, U. Schulz, and P. Jezzard, "Vessel-encoded dynamic angiography: Initial experience in healthy volunteers and patients," in *Proceedings 16th Annual Meeting, British Chapter of the International Society for Magnetic Resonance in Medicine, Nottingham*, p. 11, 2010.
8. **T. W. Okell**, P. Schmitt, X. Bi, M. A. Chappell, R. H. Tijssen, K. L. Miller, and P. Jezzard, "4D vessel-encoded arterial spin labeling angiography," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, p. 4034, 2011.
9. M. A. Chappell, **T. W. Okell**, P. Jezzard, M. W. Woolrich, and S. J. Payne, "Fast analysis of vessel encoded ASL perfusion and angiographic images for clinical applications," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, p. 366, 2011.
10. **T. W. Okell**, M. A. Chappell, U. Schulz, and P. Jezzard, "Quantification of vessel-encoded arterial spin labeling dynamic angiography with auto-calibration," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, p. 3405, 2011.

Chapter 1

Introduction

Collateral blood flow in the brain is the subject of considerable clinical interest due to its importance in stroke, one of the major causes of mortality and morbidity in the developed world [1]. This D.Phil. thesis describes the development of magnetic resonance imaging (MRI) methods to study collateral blood flow in the brain, both in the blood vessels and the resulting perfusion of tissue. In this chapter the motivation for this work is described, and an outline of the remainder of the thesis provided.

1.1 Motivation

A significantly large reduction in the blood supply to a brain region (ischaemia) can cause a neurological deficit (stroke). If blood supply is not restored, this may lead to death of brain tissue (infarction) and permanent neurological damage. Such a reduction in blood flow can be caused by a number of conditions, including thrombosis, embolism, stenosis, occlusion and haemorrhage [2].

Liebeskind [3] defines the cerebral collateral circulation as “the subsidiary network of vascular channels that stabilise cerebral blood flow when principal conduits fail”. This network is of vital importance if one or more principal arteries are

compromised and is often observed to be recruited in ischaemic patients [3]. The level of collateral flow may be important in determining the rate of ischaemic injury [4] and thus tissue fate [5], irrespective of subsequent recanalisation [6], with a number of studies showing that increased collateral flow is well correlated with better clinical outcome following a stroke [7, 8]. The presence or absence of collateral flow is also of use in therapeutic planning [4, 9]. For example, in patients with bilateral carotid stenoses, revascularization may be a better option for patients with poor collateral circulation but medical, rather than surgical, treatment more suitable for those with good collateral flow [10]. Assessment of collateral flow is therefore important for the diagnosis, prognosis and aiding of therapeutic decisions for stroke patients and those with cerebrovascular disease.

Current clinical assessment of collateral flow within the cerebral vessels is generally limited to digital subtraction angiography (DSA), which makes use of x-ray contrast agents injected directly into the artery of interest using a catheter [11]. This technique provides excellent temporal and spatial resolution, but is limited by its invasive nature, the requirement for exposure to ionising radiation, multiple acquisitions required to assess multiple arteries, and the lack of information about the perfusion that results from the collateral flow. In addition, these procedures carry a risk of contrast agent reaction, silent ischaemia [12], or even stroke [13]. The injection pressure may modify the normal flow patterns or distort the appearance of distal vessels [3], leading to misrepresentation of the patient's condition. Other angiographic methods, such as conventional magnetic resonance angiography (MRA), are not able to easily distinguish the source of collateral flow [3].

Collateral flow can also be inferred from measurements of perfusion using a variety of modalities such as positron emission tomography (PET), computed tomography (CT) with contrast agent, and non-selective arterial spin labelling

(ASL). These methods allow assessment of regional perfusion, but do not show the source of any collateral flow present [3] or its contribution to perfusion in a particular area. A noninvasive alternative, such as one making use of MRI, is therefore desirable.

A number of other clinical areas may benefit from vessel specific information. For example, arteriovenous malformation (AVM), in which new blood vessels form, creating connections between arteries and veins that bypass the capillary bed. Therapy of this disease is aided by knowledge of which vessels are supplying the AVM, which is often not clear on a standard angiographic image. Such techniques may also be of use in determining the main feeding arteries to a tumour.

1.2 Thesis Outline

This thesis describes the development of non-invasive MRI methods to study collateral blood flow in the brain. Chapter 2 gives an overview of the background relevant to the work of this thesis, including cerebrovascular anatomy, collateral pathways, the physical principles underlying MRI and explanations of non-invasive perfusion measurements using arterial spin labelling techniques.

In Chapter 3 current methods of assessing collateral flow are discussed along with their relative merits and limitations. These include conventional methods within a number of imaging modalities (such as x-ray DSA, computed tomography, transcranial Doppler ultrasonography and conventional magnetic resonance angiography), indirect methods such as regional perfusion imaging from which collateral flow can be inferred in some cases, novel techniques for selective labelling of individual vessels using MRI and finally vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL), which is the method used exten-

sively in this thesis.

The assessment and optimisation of VEPCASL through simulations of the Bloch equations and experiments in a flow phantom are described in Chapter 4. In particular, the mechanism by which magnetisation can be inverted and the spatial modulation of inversion efficiency using this technique are explored, along with its sensitivity to a variety of factors such as blood velocity and off-resonance frequency. A number of methods for the subsequent separation of vascular components from vessel-encoded data are then tested and compared.

Experiments in healthy volunteers begin in Chapter 5, where the VEPCASL preparation is attached to an echo planar imaging (EPI) readout method to generate vascular territory maps weighted by regional tissue perfusion. Optimisation of the labelling plane location, post labelling delay and EPI readout technique are discussed, followed by a comparison with a three-dimensional gradient and spin echo (3D-GRASE) readout method. Variants of the cerebrovasculature identified in some healthy volunteers using these methods are shown, demonstrating the ability to identify abnormal flow patterns.

In order to visualise collateral flow directly in the blood vessels, the VEPCASL preparation was attached to a 2D dynamic (“*cine*”) angiographic readout in Chapter 6. Methods for the acquisition and processing of such data are described, along with example results in healthy volunteers. Preliminary work using a much more efficient balanced steady-state free precession readout is also described.

In Chapter 7 a novel modelling and analysis technique for the quantification of vessel-encoded dynamic angiographic data is derived and tested, yielding blood volume flow rates in the major vessels around the circle of Willis which are consistent with the literature. This is a significant advantage over many other techniques for the assessment of collateral flow, which provide qualitative information only.

The application of the VEPCASL tissue perfusion and angiographic methods

in a number of patient groups is described in Chapter 8. Preliminary results are shown for patients with vertebro-basilar disease, carotid stenosis and arteriovenous malformation, demonstrating the ability of these methods to provide clinically useful information on collateral flow in patients.

Finally, a summary of this work is given in Chapter 9, including a discussion of the relative merits of the VEPCASL tissue perfusion and angiographic methods previously mentioned. Prospects for future research in this area are also discussed, including vessel-specific cerebral blood flow quantification for the tissue perfusion measurements and acceleration techniques which would be beneficial for angiographic acquisitions.

It should be noted that references for all chapters can be found at the end of this thesis.

Chapter 2

Background

In this chapter the underlying principles behind the ideas discussed in the remainder of the thesis are summarised. These include cerebrovascular anatomy and collateral pathways in the brain (Section 2.1), cerebrovascular disease (Section 2.2), the theory behind magnetic resonance imaging (MRI) (Section 2.3) and arterial spin labelling (ASL) (Section 2.4). Current methods for the assessment of collateral blood flow in the brain are discussed in the next chapter.

2.1 Cerebral Vasculature

In this section the main cerebral arteries are described along with potential routes for collateral blood flow which are commonly recruited when one or more of the main arteries are compromised. More detailed information can be found in the literature [14, 3, 10].

The main arteries that supply the brain originate from the aorta, just superior to the heart. The right and left common carotid arteries (CCAs) split (bifurcate) into the internal and external carotid arteries (ICAs and ECAs). Under normal circumstances the ICAs supply the brain while the ECAs supply the face and scalp. The right and left vertebral arteries (VAs) originate at the subclavian ar-

teries which themselves feed off the aorta. The VAs supply the brain whilst the subclavian arteries feed the arms and upper thorax.

Fig. 2.1 shows the cerebral vasculature from just above the carotid bifurcation. Just below the brain the VAs fuse to form the basilar artery (BA). The ICAs and BA come together in a circular structure at the centre of the brain known as the “circle of Willis” before branching off into smaller arteries which go on to feed the majority of the brain: the right and left anterior cerebral arteries (ACAs), middle cerebral arteries (MCAs) and posterior cerebral arteries (PCAs).

The cerebellum, a structure at the base of the brain towards the posterior, is fed by smaller arteries (not visualised in Fig. 2.1) which branch off the posterior circulation, each with a right and left counterpart: the posterior inferior cerebellar arteries (PICAs) branch off the ipsilateral VAs; the anterior inferior cerebellar arteries (AICAs) branch off the BA just above the point where the two VAs fuse; the right and left superior cerebellar arteries (SCAs) originate at the BA just before it bifurcates into the PCAs. Similarly, the brainstem, which is the most inferior brain structure linking the brain to the spinal cord, is fed by a number of small arteries which branch off the VAs and BA.

At the circle of Willis there are small “communicating” arteries which allow collateral flow between the BA and two ICAs. The anterior communicating artery (AComA) allows flow between the two ICAs and the right and left posterior communicating arteries (PComAs) allow flow between the BA and the relevant ICA. This redundancy provides a useful backup system, reducing the chance that any brain area will be left without blood supply following failure of the usual feeding artery. For example, if the RICA were to become occluded then, in the absence of collateral flow, the brain tissues supplied by the RMCA and RACA would become infarcted. However, collateral flow from the LICA across the AComA or from the BA across the RComA could continue to maintain these tissues and

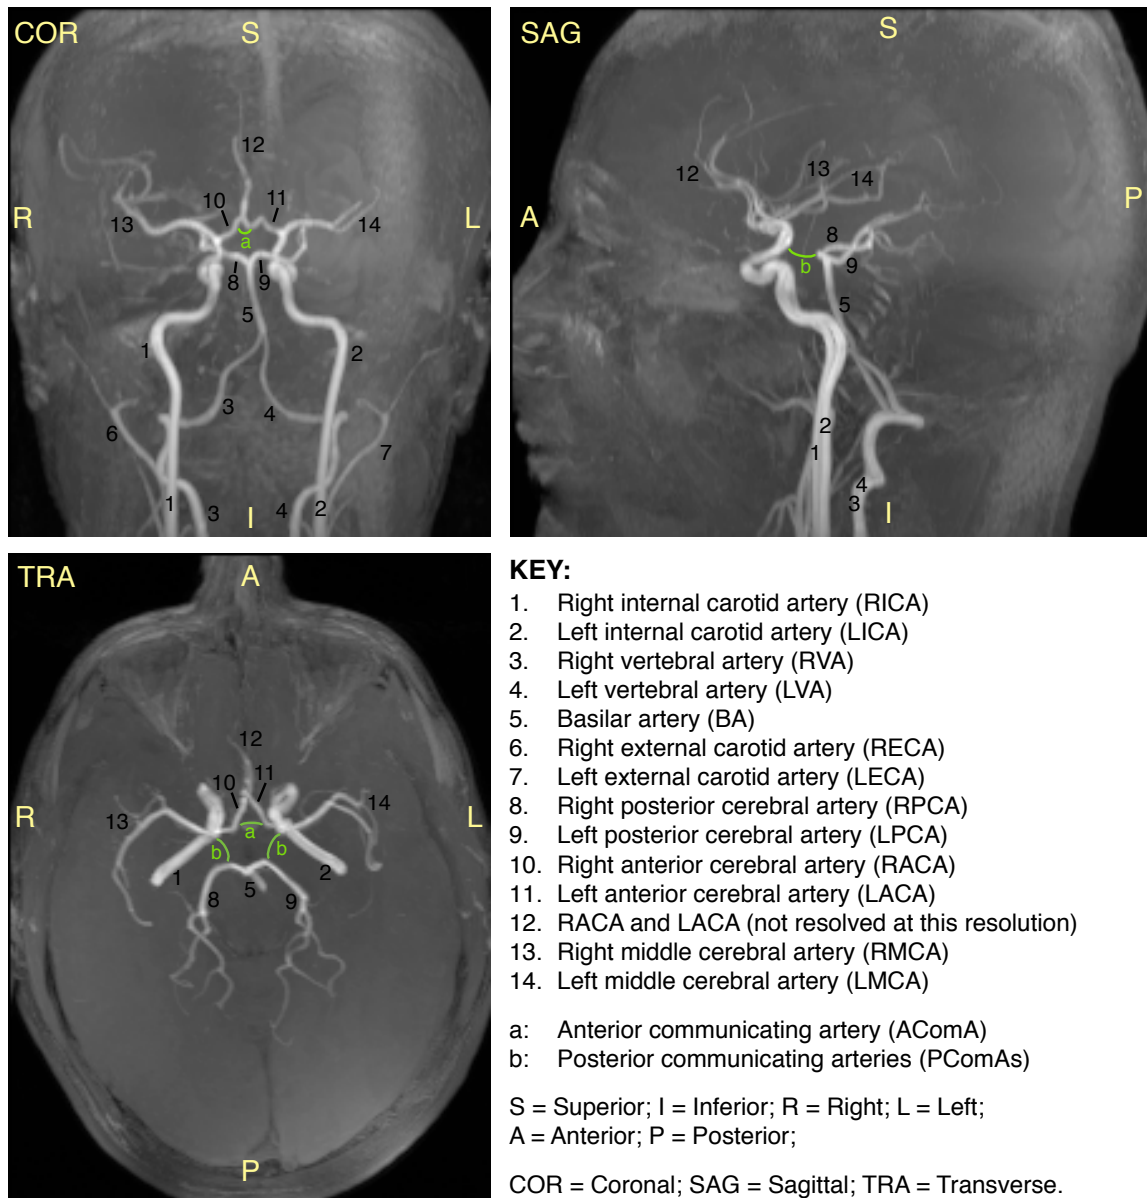


Figure 2.1: Cerebral vasculature illustrated using maximum intensity projections from a time-of-flight magnetic resonance angiography data set in a healthy volunteer. The main vessels of interest below and above the circle of Willis are marked. In this healthy volunteer the main collateral routes around the circle of Willis (the AComA and PComAs) are not visible since there is no flow across them, so their approximate positions are marked in green. Note that to display the circle of Willis more clearly the transverse maximum intensity projection excludes the inferior and superior vessels. Also, there is significant overlap of the vessels on the left and right in the sagittal view. Images are displayed according to radiological convention (right of image is left of subject).

prevent a catastrophic stroke (defined in the next section). Note that not all of the communicating arteries are present in all individuals due to natural variation. In fact, only about 25% of patients have a complete circle of Willis [14].

In addition to the circle of Willis, there are a number of other potential collateral pathways. Extracranial sources include branches of the ECA via the ophthalmic artery (which feeds the eye and connects to the ICA). Intracranial sources include a number of small pathways in the membrane covering the brain which are known as “pial” or “leptomeningeal” vessels. These routes are described in more detail by Liebeskind [3].

2.2 Cerebrovascular Disease

There is a large number of diseases which can affect the cerebral vasculature. A brief overview of the main diseases relevant to this thesis is provided here, but the reader is referred to the literature for a more detailed discussion [15, 16].

A “stroke” or “brain attack” occurs when the flow of blood into the brain is compromised, leading to ischaemia (lack of blood) to regions of brain tissue. This reduced blood supply prevents oxygen and vital nutrients reaching the cells and prevents them from functioning normally. In addition, waste products normally taken away by blood flow may build up and cause cellular damage. In the complete absence of blood flow, these cells cannot survive for more than a few minutes [1]. Below a critical threshold (about 20 ml blood per 100 g of brain tissue per minute [17]), neuronal cells cease to be active but if the blood supply is sufficient they may remain viable, such that if normal blood flow can be restored quickly enough then these cells can regain their normal function.

The hypoperfused tissues can be categorised into three main types [17]:

- Core: tissues with severely reduced blood supply that will inevitably die;

- Oligaemia: tissues with only minor reduction in blood supply that are likely to survive;
- Penumbra: tissues that may either die or survive and are therefore potentially salvageable.

If blood supply is not restored the penumbral tissues are gradually recruited into the core. Collateral flow is likely to be important in maintaining the penumbra [3], allowing time for reperfusion therapy to be administered. The reduction in blood supply to the brain can be caused by a number of factors which are split into three broad groups below.

It should also be noted that some patients experience minor, reversible stroke symptoms known as “transient ischaemic attacks” (TIAs), which last less than 24 hours and result in no radiographic evidence of stroke [16]. These are clinically important as patients who experience a TIA are at higher risk of having a full stroke [16].

2.2.1 Steno-Occlusive Disease

Stroke is often caused by the complete blockage (occlusion) or narrowing (stenosis) of brain-feeding arteries. This is commonly due to a blood clot which forms there (thrombosis), or which is produced elsewhere and traverses the vascular system before becoming lodged in a smaller downstream vessel (embolism). A common treatment for such cases is the administration of a thrombolytic agent, such as tissue plasminogen activator (tPA), to dissolve the clot and restore normal blood flow. However, this treatment must be given within about three hours of stroke onset [4], otherwise there is a danger that the vessels distal to the clot, having been deprived of blood supply for a significant period, will have begun to necrose. When the clot is dissolved, there is significant risk that the inflowing

blood will rupture these weakened vessels, causing haemorrhage (see below) and further damage to the brain [4]. In addition, prior to administration of such an agent, it must be ascertained that the cause of stroke was not a haemorrhage, since the administration of a thrombolytic agent in such cases is detrimental to the patient.

Another common cause of steno-occlusive disease is atherosclerosis, in which fatty deposits (plaques) develop on the wall of an artery. The addition of further fatty deposits and the resulting cellular proliferation around them can lead to significant reduction in blood flow through the artery [2]. In addition, the plaque can catalyse the formation of blood clots, leading to thrombosis or embolism. After some time the plaques can cause the walls of the artery to become more rigid, increasing the chances of haemorrhage (defined below).

There are a number of treatments for atherosclerosis, including angioplasty, in which an empty balloon-like device is inserted adjacent to the plaque using a catheter under x-ray guidance. Once in position, the balloon is inflated, widening the vessel to allow more blood to pass. After this widening, a stent may also be placed. This is a small tube-like metallic device which expands to fill the artery, helping to maintain blood flow through the vessel. Alternatively, an attempt may be made to remove the plaque surgically in a procedure known as endarterectomy. The diseased vessel is clamped to prevent blood loss during the procedure, the vessel wall is dissected, the plaque removed and the vessel wall repaired. During this procedure the patient is therefore reliant on collateral flow to maintain perfusion in the brain. These surgical procedures can be highly beneficial to the patient but also carry risks, including ischaemia whilst the vessels are clamped in endarterectomy, rupture of the hardened vessels or creation of an embolus from the plaque material which travels higher into the vascular tree and occludes a smaller vessel there, causing a stroke. Medical treatments, such as the

use of statins, are also available.

2.2.2 Haemorrhage

Haemorrhage is the rupture of a blood vessel which can damage the brain by increasing the intracranial pressure, damaging delicate brain tissue, and by reducing the blood flow to distal tissues. In addition, bleeding into the brain tissues may lead to delayed neurological damage [18]. Approximately 12% of strokes are due to haemorrhage, with the remaining 88% due to ischaemia [19].

Haemorrhage can follow after a weakening of the arterial wall, causing it to distend (aneurysm) and eventually rupture, or after traumatic injury. Risk of haemorrhage is also increased in atherosclerosis, as mentioned above.

2.2.3 Other

There is a number of other diseases which can lead to reduced blood supply to the tissues. One of particular interest in this thesis is arteriovenous malformation (AVM), in which a series of abnormal blood vessels form links between arteries and veins directly (anastomoses), bypassing the tissue which can then become ischaemic. This lesion can cause severe headache, stroke, haemorrhage and epilepsy [20]. Therapy options include radiotherapy, surgical removal and embolization, in which the main vessels feeding the AVM are blocked with a glue-like substances injected via a catheter [21]. Methods developed for studying collateral flow may also be of use in assessing the main feeding arteries to the AVM for planning of embolization procedures.

Another disease of interest in the study of collateral flow is “moyamoya” disease, in which there is progressive stenosis of some of the cerebral arteries. To compensate for this, a series of small vessels are formed to provide collateral flow

to the affected regions [22] as well as via more common collateral routes (see Section 2.1). Collateral flow is therefore very significant in these patients and methods for improving visualisation of this phenomenon are likely to be useful in the study of this disease.

2.3 Principles of Magnetic Resonance Imaging

In this section the main principles underlying magnetic resonance imaging (MRI) are described, including the generation of net magnetisation using a large external magnetic field, the Bloch equation governing the motion of the magnetisation, methods for spatial localisation and k -space formalism used to describe image formation. For a more comprehensive discussion of this wide topic, the reader is referred to the literature [23, 24, 25, 26, 27, 28, 29, 30], from which the information in this section was derived.

2.3.1 Magnetisation

When the nucleus of a hydrogen atom is placed within an external magnetic field, \mathbf{B} , the quantum mechanical Zeeman effect dictates that the component of the spin parallel to this field can only take two possible values: one aligned with the field (the lower energy state) and the other against it (the higher energy state). The difference in energy between these states, ΔE , is given by:

$$\Delta E = \hbar\gamma B \tag{2.1}$$

where B is the magnitude of the vector \mathbf{B} , \hbar is Planck's constant divided by 2π and γ is the gyromagnetic ratio of hydrogen. Boltzmann statistics can be used to determine the relative populations of these two energy states:

$$\frac{n_{\uparrow}}{n_{\downarrow}} = \exp\left(\frac{\Delta E}{k_B T}\right) = \exp\left(\frac{\hbar\gamma B}{k_B T}\right) \quad (2.2)$$

where n_{\uparrow} and n_{\downarrow} are the numbers of nuclei with spin components parallel and anti-parallel to \mathbf{B} respectively, k_B is Boltzmann's constant and T is the temperature in degrees Kelvin. For typical magnetic field strengths encountered in MRI and at room temperature $k_B T \gg \hbar\gamma B$ and thus:

$$\frac{n_{\uparrow}}{n_{\downarrow}} \approx 1 + \frac{\hbar\gamma B}{k_B T} \quad (2.3)$$

Thus there is a slight excess of nuclei with spin components parallel to the field over those anti-parallel to the field (10 parts per million at 3T and room temperature). This is a minute number, but given the enormous quantity of hydrogen nuclei present in the water of our bodies, this is sufficient to generate an observable signal.

Parallel spins have a magnetic moment $\mu_{\uparrow} = \gamma\hbar/2$ in the direction of the external field and anti-parallel spins have $\mu_{\downarrow} = -\gamma\hbar/2$. In a large ensemble of nuclei the components of μ perpendicular to \mathbf{B} cancel out, giving a net magnetisation vector, \mathbf{M} , parallel to \mathbf{B} , with magnitude, M , given by:

$$M = n_{\uparrow}\mu_{\uparrow} + n_{\downarrow}\mu_{\downarrow} = (n_{\uparrow} - n_{\downarrow})\frac{\gamma\hbar}{2} \approx n_{\downarrow}\frac{\hbar\gamma B}{k_B T}\frac{\hbar\gamma}{2} \approx N\left(\frac{\hbar\gamma}{2}\right)^2\frac{B}{k_B T} \quad (2.4)$$

where $N = n_{\uparrow} + n_{\downarrow} \approx 2n_{\downarrow}$ is the total number of nuclei in the ensemble. It is the motion of this magnetisation which is observable in MRI experiments.

Transitions between the two energy levels can occur with the absorption or emission of a photon with energy E_{ph} , equal to the difference in energy between the two states:

$$E_{\text{ph}} = \Delta E = \hbar\gamma B \quad (2.5)$$

The energy of a photon is related to its angular frequency, ω , by $E_{\text{ph}} = \hbar\omega$ and thus:

$$\omega = \gamma B \quad (2.6)$$

At common MRI field strengths, this corresponds to electromagnetic radiation in the radio-frequency band. This equation is central to the theory behind MRI since it shows that the frequency of radiation absorbed and emitted by the nuclei, known as the Larmor frequency, is directly proportional to the magnetic field they experience. This allows spatial localisation of the signal, as discussed below.

2.3.2 The Bloch Equation

Bloch *et al.* [23, 24] showed that the motion of the magnetisation vector, \mathbf{M} , resulting from a large number of nuclei, can be accurately described using a classical model:

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B} - \frac{1}{T_2} \begin{pmatrix} M_x \\ M_y \\ 0 \end{pmatrix} + \frac{1}{T_1} \begin{pmatrix} 0 \\ 0 \\ M_0 - M_z \end{pmatrix} \quad (2.7)$$

where M_x , M_y and M_z are the components of \mathbf{M} in each direction, it is assumed that B is approximately parallel to the z direction, M_0 is the equilibrium magnetisation, T_2 is the spin-spin relaxation time and T_1 is the spin-lattice relaxation time.

Ignoring the relaxation terms, the Bloch equation (Eq. 2.7) describes the precession of \mathbf{M} about \mathbf{B} with angular frequency equal to the Larmor frequency $\omega = \gamma B$, in a left-handed sense, as shown in Fig. 2.2.

The motion of the magnetisation is somewhat clearer when viewed in a frame




Figure 2.2: Precession of the magnetisation, \mathbf{M} , about the magnetic field, \mathbf{B} , in a left-handed sense. The phase of the magnetisation as measured in the standard laboratory frame, ϕ_{lab} , is defined as the angle made between the x axis and the transverse magnetisation (i.e. the component in the x - y plane), also in a left-handed sense. Relaxation effects are neglected here.

of reference rotating about the z axis at the Larmor frequency, $\omega_0 = \gamma B_0$, which corresponds to the main magnetic field of the scanner, \mathbf{B}_0 . The Bloch equation thus becomes:

$$\frac{d\mathbf{M}'}{dt} = \gamma \mathbf{M}' \times (\mathbf{B} - \mathbf{B}_0) - \frac{1}{T_2} \begin{pmatrix} M'_x \\ M'_y \\ 0 \end{pmatrix} + \frac{1}{T_1} \begin{pmatrix} 0 \\ 0 \\ M_0 - M_z \end{pmatrix} \quad (2.8)$$

where the prime marks represent quantities observed in the rotating frame. Note that quantities pointing purely along the z axis are unaffected by observation in the rotating frame. The rapid precession caused by \mathbf{B}_0 is effectively removed, and thus only motion due to other influences is apparent. For example, consider a radiofrequency (RF) pulse applied at the Larmor frequency which leads to a magnetic field \mathbf{B}_1 , applied in the transverse plane (i.e. x - y plane) such that:

$$\mathbf{B}_1 = B_1 \begin{pmatrix} \cos(\omega_0 t + \phi_1) \\ \sin(\omega_0 t + \phi_1) \\ 0 \end{pmatrix} \quad (2.9)$$

where t represents time and ϕ_1 is the phase of the RF pulse. In the standard laboratory frame of reference, it is difficult to see the effect this will have on the magnetisation. However, in the rotating frame, this transforms to:

$$\mathbf{B}'_1 = B_1 \begin{pmatrix} \cos \phi_1 \\ \sin \phi_1 \\ 0 \end{pmatrix} \quad (2.10)$$

Ignoring the relaxation terms, the Bloch equation becomes:

$$\frac{d\mathbf{M}'}{dt} = \gamma \mathbf{M}' \times \mathbf{B}'_1 \quad (2.11)$$

In this frame the magnetisation simply precesses about the axis defined by ϕ_1 , which tips it into the transverse plane, as shown in Fig. 2.3. The total angle, θ , through which the magnetisation is rotated is given simply by:

$$\theta = \int \omega dt = \int \gamma B_1 dt \quad (2.12)$$

Thus, by appropriate control of the duration and amplitude of the RF pulse, any arbitrary flip angle can be achieved. It is also important to note that an RF pulse played out away from the resonant frequency of the magnetisation will not be static in the rotating frame of reference, but will rotate at angular velocity given by the difference in frequency between the RF pulse and the Larmor frequency. This will cause the direction of motion of the magnetisation due to the RF pulse to




Figure 2.3: The tipping of magnetisation into the transverse plane by an RF pulse, illustrated in the rotating frame of reference. The initial magnetisation, M_i , is aligned along the z axis. The RF pulse creates a magnetic field, B_1 , with phase ϕ_1 , which rotates the magnetisation by an angle θ , into its final position, M_f , in the transverse plane. Relaxation effects are neglected here.

change constantly, leading to very little overall effect. Thus, only RF pulses which are at the resonant frequency of the magnetisation will have any significant effect on it, which is consistent with the quantum mechanical view.

Once the magnetisation has a component in the transverse plane, in the lab frame it will precess at a frequency $\omega = \gamma B$. This continuous change in the magnetisation can be detected using a radiofrequency coil, yielding a voltage which is proportional to the net magnetisation within the sample. It is this signal which is recorded and forms the basis of MRI images.

Thus far we have neglected the relaxation terms involving T_1 and T_2 . The equilibrium position of the magnetisation is aligned with the z axis with magnitude M_0 , since at this point $dM/dt = 0$. In the rotating frame of reference and in the absence of RF pulses or other magnetic fields, these relaxation terms cause magnetisation which is not at its equilibrium position to decay exponen-

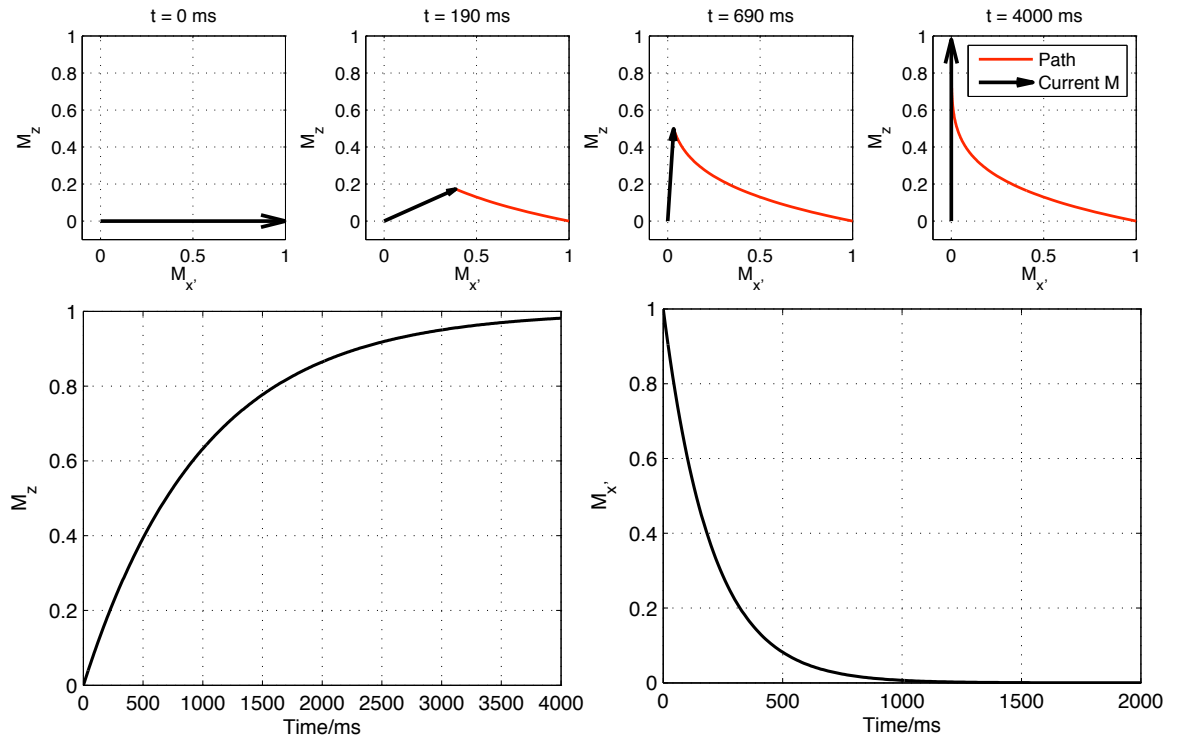


Figure 2.4: Relaxation processes in the rotating frame in the absence of RF pulses and other magnetic fields. Plots are shown of a magnetisation vector which begins aligned with the x' axis (e.g. just after a 90° RF pulse). The top row shows the path of the magnetisation vector in the x' - z plane at a number of time points, showing regrowth of the longitudinal magnetisation (M_z) and the simultaneous decay of the transverse magnetisation ($M_{x'}$). The bottom row shows these processes plotted against time for longitudinal magnetisation (left) which recovers with time constant T_1 (here set to 1000 ms) and transverse magnetisation (right) which decays with time constant T_2 (here set to 200 ms). In these plots the magnetisation is normalised against M_0 .

tially back towards it. The transverse magnetisation, and thus the detected MRI signal, decays towards zero with time constant T_2 and the longitudinal magnetisation (aligned with z) recovers back towards M_0 with time constant T_1 . This process is illustrated in Fig. 2.4.

2.3.3 Spatial Localisation

In the previous section, the behaviour of the magnetisation under the influence of magnetic fields was discussed, along with the use of RF pulses to flip the magnetisation into the transverse plane, allowing a signal to be detected. However,

this signal will be representative of the total magnetisation in the sample or subject. In this section methods for spatial localisation of this signal are outlined. These make use of (almost exclusively) three gradient coils within the MRI scanner that create linear variation in the z component of \mathbf{B} across the three different directions. The gradient strength is conveniently represented as a vector quantity, \mathbf{G} :

$$\mathbf{G} = \begin{pmatrix} G_x \\ G_y \\ G_z \end{pmatrix} = \begin{pmatrix} dB_z/dx \\ dB_z/dy \\ dB_z/dz \end{pmatrix} \quad (2.13)$$

Note that the x and y gradients *do not* generate magnetic fields along x or y , they merely modify the component of the magnetic field in the z direction, B_z , along x and y . With modern hardware these gradient fields can be switched on and off very rapidly, allowing fast manipulation of the magnetisation, as we shall see later.

The local magnetic field experienced by the magnetisation at a position within the sample, \mathbf{x} , in the presence of gradient fields is:

$$\mathbf{B} = \mathbf{B}_0 + \mathbf{G} \cdot \mathbf{x} \quad (2.14)$$

In the rotating frame of reference the angular precession frequency of magnetisation thus varies at different points in space, given by:

$$\omega' = \gamma \mathbf{G} \cdot \mathbf{x} \quad (2.15)$$

It is this variation in precession frequency with position that allows the spatial localisation of the signal.

Perhaps the simplest method of spatial localisation is slice selection, whereby a




Figure 2.5: Slice selective excitation of magnetisation: applying a RF pulse with centre frequency ω'_c and bandwidth $\Delta\omega$ in the presence of a gradient (here along z) leads to the selective excitation of magnetisation with precession frequency within this range, which corresponds to a slice centred at z_c with thickness Δz as shown on the left. An ideal RF pulse for this application has sharp edges and uniform power within the defined bandwidth (top right). The time-domain representation of this RF pulse can be approximately calculated by Fourier transform (FT) for small flip angle pulses. In this case the pulse is a sinc function (bottom right).

single slice of the sample, with finite thickness, is tipped into the transverse plane without disturbing the magnetisation elsewhere in the sample. This is achieved by playing out an RF pulse with a finite bandwidth in the presence of a gradient field (e.g. along z), as illustrated in Fig. 2.5. The gradient results in magnetisation at different positions precessing at different frequencies. Magnetisation outside the range of frequencies present in the RF pulse is relatively unaffected, whilst that inside is excited. The slice position can be moved by modifying the central frequency of the RF pulse, and the slice width can be changed by varying the RF bandwidth or gradient strength.

Once a slice of magnetisation has been excited, the distribution of magnetisation within this sample needs to be determined. In one dimension (1D) it can be

straightforwardly assessed by measurement of the signal during the application of a gradient pulse in the plane of the excited slice (e.g. along x). This causes magnetisation and different locations along the x axis to precess at different frequencies. The measured signal is the sum of all of these frequencies, weighted by the amount of magnetisation at each point. A simple Fourier transform (FT) separates out these frequency components, resulting in a 1D image of the distribution of magnetisation along x . This process is commonly referred to as “frequency encoding”, with x defined as the “readout” direction.

This process can be generalised to multiple dimensions by a process described as “phase encoding”, where a small gradient is applied just after the excitation pulse along y to create a variation in the phase of the magnetisation across y , before applying the x gradient to perform frequency encoding, as before. The mechanism by which this allows 2D image formation is clearer when considering this process in Fourier space, also known as k -space, discussed below.

2.3.4 k -space Formalism

In MRI the signal measured is based entirely upon the transverse magnetisation (i.e. in the x - y plane). A convenient mathematical construct to represent such 2D vectors is complex numbers. The transverse magnetisation, M_+ , can be written:

$$M_+ = M_x + iM_y \quad (2.16)$$

where $i = \sqrt{-1}$. The total signal, S , measured by a homogeneous receive coil is proportional to the sum of all the magnetisation vectors in the sample:

$$S \propto \iiint dx dy dz M_+(x, y, z) \quad (2.17)$$

The magnitude of the transverse magnetisation is proportional to the density

of hydrogen nuclei at this point, ρ , and its phase, ϕ , in the rotating frame of reference:

$$M_+(x, y, z) \propto \rho(x, y, z) \exp(-i\phi(x, y, z, t)) \quad (2.18)$$

Note that a number of effects which modulate the magnitude and phase of M_+ are neglected here. ϕ can be calculated by the time integral of the angular frequency of the magnetisation at this point, which we know is related to the local magnetic field it experiences:

$$\phi(x, y, z, t) = \int dt \omega'(x, y, z, t) \quad (2.19)$$

$$= \int dt \gamma \mathbf{G} \cdot \mathbf{x} \quad (2.20)$$

For simplicity we define:

$$\mathbf{k}(t) = \int dt \gamma \mathbf{G} \quad (2.21)$$

Thus $\phi(x, y, z, t) = \mathbf{k}(t) \cdot \mathbf{x}$, which simplifies the signal equation to:

$$S(\mathbf{k}) \propto \iiint dx dy dz \rho(x, y, z) \exp(-i\phi(x, y, z, t)) \quad (2.22)$$

$$\propto \iiint dx dy dz \rho(x, y, z) \exp(-i\mathbf{k}(t) \cdot \mathbf{x}) \quad (2.23)$$

We identify this with the Fourier transform (FT) equation, relating the measured signal to the spin density:

$$S(\mathbf{k}) \propto \text{FT}(\rho(\mathbf{x})) \quad (2.24)$$

Therefore, \mathbf{k} represents spatial frequencies of the spin density and depends only upon the time integral of the gradients which can be simply controlled. Thus, the gradients can be used to move around in this Fourier, or k -space, and collect samples along the way. Once sufficient data have been acquired, a multi-dimensional Fourier transform can be performed to create an image proportional to the spin density.

The frequency encoding mentioned previously corresponds to sampling in the presence of a constant gradient which moves us through k -space at constant speed, effectively filling in a single line of data. Phase encoding corresponds to jumping to a different position in k -space before acquiring another line. Thus, a combination of frequency and phase encoding allow the coverage of 2D k -space.

This is the basis of imaging using MRI and is equally valid for generating 2D images following slice-selective excitation or for 3D acquisitions where phase encoding occurs along two directions. Methods for covering k -space in an efficient manner, generating contrast based on T_1 , T_2 or other parameters, and requirements for sampling density and extent in k -space are discussed in greater detail in the literature (e.g. [29]) or where relevant later in this thesis.

2.3.5 Effective Magnetic Field

For some applications, it is convenient to define a total effective magnetic field, \mathbf{B}_{eff} , in the rotating frame, which incorporates the effect of the RF field, \mathbf{B}_1 , the gradient fields, \mathbf{B}_G , and any other sources of magnetic field along z which lead to off-resonance precession with angular frequency ω_{OR} , \mathbf{B}_{OR} :

$$\mathbf{B}_{\text{eff}} = \mathbf{B}_1 + \mathbf{B}_G + \mathbf{B}_{\text{OR}} \quad (2.25)$$

$$\text{where } \mathbf{B}_1 = B_1(\cos \phi_1, -\sin \phi_1, 0), \quad (2.26)$$

$$\mathbf{B}_G = (0, 0, \mathbf{G} \cdot \mathbf{x}) \quad (2.27)$$

$$\text{and } \mathbf{B}_{\text{OR}} = (0, 0, \omega_{\text{OR}}/\gamma) \quad (2.28)$$

This simplifies the Bloch equation in the rotating frame to:

$$\frac{d\mathbf{M}'}{dt} = \gamma\mathbf{M}' \times \mathbf{B}_{\text{eff}} - \frac{1}{T_2} \begin{pmatrix} M'_x \\ M'_y \\ 0 \end{pmatrix} + \frac{1}{T_1} \begin{pmatrix} 0 \\ 0 \\ M_0 - M_z \end{pmatrix} \quad (2.29)$$

Thus, once the effective magnetic field is calculated, it is straightforward to visualise the magnetisation as rotation about \mathbf{B}_{eff} along with relaxation processes.

2.4 Arterial Spin Labelling

The methods developed for this thesis are based upon a technique for the non-invasive measurement of cerebral tissue perfusion known as arterial spin labelling. A brief discussion of the general principles and various implementations of non-vessel-selective ASL are given, but more detailed discussion can be found in the literature, on which this section is based [31, 32, 33, 34, 35, 36, 37, 38, 39, 40]. Discussion of vessel-selective modifications of ASL which can be used to more directly measure collateral flow are discussed in the next chapter (Sections 3.3 and 3.4).

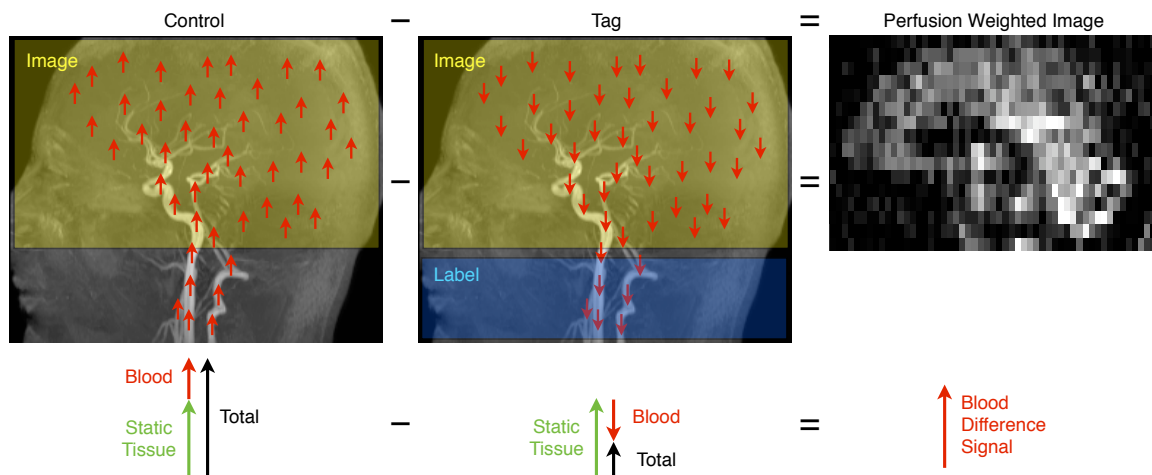


Figure 2.6: Illustration of the idea behind arterial spin labelling (ASL): a “control” image is acquired with no modification of the blood magnetisation (left). Thus the blood magnetisation adds to the static tissue signal to give the total measured signal (bottom row). A “tag” image is also acquired (middle) in which the blood water magnetisation in the neck is inverted, followed by a delay to allow it to flow into the brain and then image acquisition. In this case, the inverted blood water magnetisation subtracts away from the static tissue signal, yielding a reduced total signal. If the tag image is subtracted from the control image (right), all the static tissue signal is removed, leaving an image weighted by the blood signal which has accumulated in the brain and therefore tissue perfusion. Note that here T_1 decay of the inverted blood water is neglected for illustrative purposes.

2.4.1 General Principles

Arterial Spin Labelling (ASL) techniques use the water in arterial blood as an endogenous tracer for perfusion imaging. The protons in arterial water are inverted in the neck and after a certain delay, referred to as the inversion time (TI) or post labelling delay (PLD), the brain is imaged. The experiment is repeated without the labelling process, and the subtraction of the two images removes any static tissue components to create a perfusion weighted map (see Fig. 2.6). There are a number of different implementations of ASL which are discussed briefly in the following sections.

2.4.2 Pulsed Arterial Spin Labelling

In pulsed ASL (PASL), the inversion of blood water in the neck is achieved by applying a (typically) single 180° inversion pulse in a 10-15 cm slab below the brain. Care must be taken to ensure that any effects of the labelling pulse on the static tissue in the imaging region are reproduced in the control image, otherwise erroneous signal will appear in the subtracted images unrelated to perfusion. One particular effect is magnetisation transfer (MT). Labelling RF pulses applied in the neck appear off-resonance to magnetisation in the imaging region, and so are not directly excited. However, hydrogen nuclei bound to macromolecules have a broad range of resonance frequencies and thus some of these nuclei within the imaging region are affected and become saturated (i.e. the magnetisation magnitude is reduced significantly). Some magnetisation from free water is transferred to this pool of nuclei [29], reducing the free water signal and creating a difference between tag and control images not due to perfusion [35].

In the echo-planar MR imaging and signal targeting with alternating radio frequency (EPISTAR) technique [41], labelling is performed using a single 180° adiabatic inversion pulse. MT effects are better matched if the control condition consists of two similar 180° RF pulses, each with half the RF power of the labelling pulse. This leads to 360° rotation of the blood magnetisation, leaving it relatively unaffected, but MT effects are matched in the tag and control images and therefore subtract away without artefact [42].

In flow alternating inversion recovery (FAIR) [43, 44], the control condition is prepared using a slice-selective inversion pulse applied to the imaging region. In the tag condition, a non-selective inversion pulse is used. In both cases the static tissue is inverted, but for the tag case the inflowing blood is also inverted and for the control case it is not, giving ASL contrast when the images are subtracted. MT effects are minimised by the use of the same RF pulse for both tag and control

images.

2.4.3 Continuous Arterial Spin Labelling

Rather than invert all the magnetisation proximal to the imaging region using a single, relatively short RF pulse as in PASL, in continuous ASL (CASL) a long constant RF waveform is played out in the presence of a constant (usually negative) gradient in the z direction. This defines a labelling plane at the point where the RF pulse is on-resonance within the gradient field. Blood magnetisation passing through this labelling plane on the way to the brain experiences an effective magnetic field which changes slowly over time. When the blood is some distance proximal to the labelling plane, the gradient component dominates \mathbf{B}_{eff} , causing it to point approximately along the z axis. The magnetisation is thus relatively unaffected at this point. As the blood passes through the labelling plane and beyond, the gradient component to \mathbf{B}_{eff} reduces to zero then becomes negative, causing \mathbf{B}_{eff} to sweep from approximately parallel to $+z$ to approximately parallel to $-z$. The constant RF pulse ensures \mathbf{B}_{eff} always has a fixed transverse component. Since \mathbf{M} precesses about \mathbf{B}_{eff} , a slow change in \mathbf{B}_{eff} causes \mathbf{M} to follow its direction, resulting in its inversion. This effect is shown diagrammatically in Fig. 2.7.

This flow driven adiabatic inversion causes all blood flowing through the labelling plane to be inverted for as long as the RF pulse and gradient are switched on, allowing the creation of boluses of labelled blood water with very long temporal duration, giving a higher signal change in the brain and thus better SNR than PASL techniques. In fact, the original implementation of ASL used this methodology [32, 31].

MT effects are even more significant for CASL than PASL due to the long continuous RF pulse and need to be carefully controlled for. In the original im-

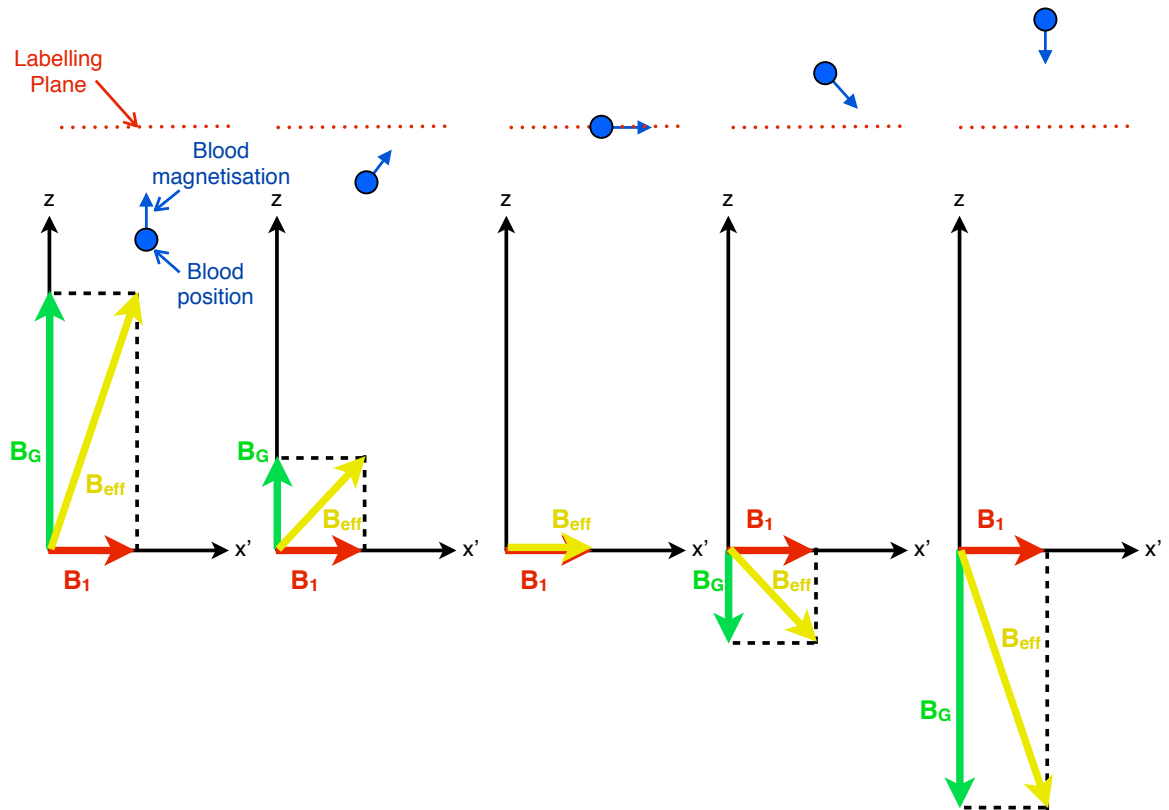


Figure 2.7: Schematic diagrams showing how the effective magnetic field, B_{eff} , changes as the blood passes through the labelling plane in CASL, and the effect this has on the blood's magnetisation. The diagrams on the top row show the position of the blood relative to the labelling plane and the direction of its magnetisation vector. The diagrams on the bottom row show the contributions to B_{eff} at each point. The contribution from the RF pulse, B_1 , is constant throughout. However, the contribution from the negative z gradient, B_G , is initially large and positive, but as the blood moves towards the labelling plane it decreases in magnitude, passes through zero at the labelling plane and then regrows in the negative direction. Thus B_{eff} sweeps from pointing approximately towards $+z$ to approximately towards $-z$ and since this movement is slow, the magnetisation tends to follow it, leading to inversion of the blood water.

plementation [32, 31], the control image was acquired after an equivalent labelling preparation performed distal to the imaging region. Thus no blood is labelled, but the MT effects in the tissue are similar. However, this is only valid for single slice acquisitions. Multi-slice imaging is made possible by modulating the RF waveform with a sinusoidal function in the control condition, effectively creating two labelling planes which first invert, then un-invert the blood [45]. This gives similar MT effects in the brain in tag and control cases. However, the main disadvantage of CASL is that the generation of very long duration RF pulses is often not possible using clinical scanners without additional hardware.

2.4.4 Pseudo-Continuous Arterial Spin Labelling

Pseudo-continuous ASL (PCASL) [46, 40] combines the SNR advantage of CASL (even exceeding it [47]) with the ability to run the sequence using standard hardware, as in PASL. The PCASL module consists of a train of RF pulses and associated z gradient pulses in rapid succession (approximately every 1 ms), as shown in Fig. 2.8. The time averaged B_1 field and gradient fields are comparable to CASL and in the same way magnetisation is inverted in a flow-driven adiabatic fashion as it passes through a defined labelling plane. However, understanding the mechanism of this inversion requires a little more thought.

The repeated application of RF pulses and gradients leads to steady-state behaviour, such that a spin in the steady state orientation mid-way between two consecutive RF pulses returns to the same position after the next pulse. This is identical to the behaviour of the magnetisation in steady-state free precession sequences [49]. However, here it is assumed that the magnetisation experiences the RF pulses for only a short time while it is near the labelling plane, during which relaxation effects can be ignored. The steady state orientation depends on the phase shift experienced between RF pulses and can therefore be influenced by




Figure 2.8: Pseudo-continuous ASL pulse sequence (left). The control condition is realised by alternating the phase of every other RF pulse between 0 and π (dashed lines). The net gradient need not be zero in order to achieve an efficient control condition [48]. The z component of the steady-state magnetisation position is shown as a function of phase accrual between RF pulses, ϕ , in the absence of T_1 and T_2 decay (right). Here the RF flip angle was assumed to be 20° . The red and black lines both represent valid steady-state positions.

gradients. The z component of the steady-state magnetisation as a function of phase shift between RF pulses, ϕ is given by [46]:

$$M_{z,ss} = \frac{\pm M_0 \sin \alpha \sin \phi/2}{\sqrt{(1 - \cos \alpha)^2 + \sin^2 \alpha \sin^2 \phi/2}} \quad (2.30)$$

This function is plotted in Fig. 2.8. In a similar manner to CASL, if the phase accrual between RF pulses is slowly varied by the spin moving through the z gradient field, then the steady-state position of the magnetisation slowly inverts. The blood magnetisation will tend to follow this steady-state position and thus is also inverted. The control condition is realised by using RF pulses with a phase that alternates between 0 and π , flipping the magnetisation back and forth through a small angle, but not perturbing it significantly. Dai *et al.* also modify the gradient pulses to produce zero net gradient in the control condition [40]. However, Wong showed that this is not necessary to produce an efficient control [48], and keeping the gradient pulses identical between tag and control reduces the possibility of differences in eddy currents produced by the sequences which could result in unwanted differences between the tag and control images.

It should be noted that the gaps between RF pulses allow rephasing gradients along z . This means that the mean gradient can be kept to a low value comparable to those used in CASL experiments, leading to an efficient adiabatic inversion, but at the same time the gradient during RF pulses can be relatively large. This greatly reduces the potential MT effects in the imaging region which is much further from resonance in PCASL than CASL. In addition, PCASL also has the advantage of having gaps between RF pulses which can be exploited to selectively label individual arteries, as discussed in the next chapter.

The inversion mechanism and sensitivity of PCASL to various factors such as blood velocity and off-resonance frequency are explored through simulations of the Bloch equations and phantom experiments in Chapter 4.

2.4.5 Blood Flow Quantification

A lot of work has been done to allow the quantification of cerebral blood flow (CBF) in units of ml of blood delivered per 100 g of brain tissue per minute through ASL measurements. These range from the initial studies [32, 31] through to a model free approach [50]. Such quantification approaches have the potential to make ASL a more useful clinical tool and better suited for multi-centre trials.

Perhaps the most commonly used approach is the general kinetic model of Buxton *et al.* [33]. This uses a single-compartment kinetics assumption that water is very rapidly exchanged between vascular and tissue compartments at the capillary bed to predict the ASL signal, defined here as ΔM , as a function of measurement time (see Fig. 2.9). It should be noted that beyond the peak in this curve (i.e. at times greater than the arrival time of blood in the voxel plus the labelling duration), ΔM is approximately proportional to CBF but before this time, the measurement is confounded by blood arrival time. If only a single time point is to be sampled it is therefore critical to choose a post labelling delay (PLD) long




Figure 2.9: Variation in the ASL difference signal, ΔM , as a function of measurement time using the general kinetic model of Buxton *et al.* [33] for a CASL/PCASL type sequence. Initially the difference signal is zero before the front of the labelled bolus has arrived in the voxel of interest. As the bolus washes in, the exchange of labelled blood water into tissue causes the label to accumulate, thereby increasing ΔM . After the bolus has passed through the voxel, the difference signal decays with the T_1 of tissue. Here it was assumed that the arrival time of blood in the voxel, $\Delta t = 500$ ms, the labelling duration, $\tau = 1400$ ms and other values (e.g. the T_1 of blood and gray matter) were taken from the literature. The time at which the image is acquired, t_{acq} , for CASL or PCASL is equal to the labelling duration, τ , plus the post labelling delay (PLD, set to 1000 ms here).

enough to reach this regime. However, even if this is not the case, a reasonable ΔM is likely to be achieved, giving some indication that perfusion is occurring, except for extremely delayed transit where the arrival time is greater than the labelling duration plus the PLD.

In order for absolute quantification to be possible, a number of calibration and correction steps (e.g. for coil sensitivity profiles) must be undertaken, but these are beyond the scope of this thesis. The reader is referred to the literature for a discussion of these topics [33].

2.5 Conclusions

In this chapter cerebrovascular anatomy was discussed along with how collateral pathways can maintain brain tissue despite relatively severe cerebrovascular disease. The concepts underlying magnetic resonance imaging and arterial spin labelling, which are used heavily throughout this thesis, were also outlined. In the following chapter current methods for the assessment of collateral blood flow in the brain are discussed. In particular, how arterial spin labelling can be modified to make it sensitive to the blood flow from individual arteries and thus reveal the presence of collateral flow in a non-invasive manner.

Chapter 3

Current Methods for Assessment of Collateral Flow

A number of methods is available for the assessment of collateral blood flow in the brain across a range of imaging modalities. In this chapter a brief summary of the available methods is presented, along with discussion of their relative benefits and pitfalls. These include conventional angiographic methods (Section 3.1), techniques to assess regional tissue perfusion which allow indirect assessment of collateral flow in some cases (Section 3.2), the adaptation of arterial spin labelling (ASL) to target specific vessels (Section 3.3) and finally the method used extensively in this thesis, vessel-encoded ASL (Section 3.4).

3.1 Conventional Angiographic Methods

Collateral blood flow can be assessed through the visualisation of blood vessels (angiography). Ideally, the images would reflect the blood flow due to each brain-feeding artery to allow clear visualisation of collateral flow. Dynamic images of blood flow, even if not vessel-selective, may provide information on flow through

common collateral vessels. For example, early arriving blood appearing in the right internal carotid artery (RICA) followed by filling of the right posterior communicating artery (RComA) and then the right posterior cerebral artery (RPCA) shows that collateral flow is occurring from the RICA to the RPCA rather than the other way around. Methods which assess flow direction or velocity provide similar information.

Even looking at the vessel morphology itself may provide clues as to the level of collateral flow. In the example above, a dilated RComA is suggestive of collateral flow, although whether this is from the RPCA to the RICA or vice versa is not known. However, it can perhaps be inferred from other information. For example, if it is known that there is a severe stenosis in the basilar artery (BA), then it is highly likely that the collateral flow is from RICA to RPCA to compensate for the reduced blood flow from the BA.

The methods described below provide one or more of these aspects (vessel-specificity, dynamic information and/or vessel morphology) and each has distinct advantages and disadvantages.

3.1.1 Conventional X-ray Angiography

One of the most common methods for the assessment of collateral flow in clinical practice is x-ray digital subtraction angiography (DSA) [11], and this is often considered to be the gold standard against which other modalities are compared. A catheter is inserted into the arterial system, typically at the femoral artery [13], and guided into the artery of interest. Once there, a contrast agent containing radio-opaque dye (such as iodine) is injected and images acquired either fluoroscopically (giving dynamic information) or as higher quality individual radiographs. A “mask” image is acquired at some point during the procedure which is subtracted from the images with contrast to remove overlying anatomy and show

only the passage of the contrast agent through the arterial system. The x-ray system is often mounted on a gantry, allowing observation of the patient from a range of angles. Images are commonly viewed in 2D, although 3D reconstructions can be performed [51].

This procedure results in vessel-specific dynamic data sets with high spatial and temporal resolution. An example image is shown in Fig. 3.1 in a patient with collateral flow from the RICA to the RPCA via the RPCoA. The collateral flow is clearly shown, as are small distal vessels, which allows this technique to visualise collateral flow from small leptomeningial vessels [7]. An additional advantage of this technique is that it can be easily used alongside an interventional procedure such as angioplasty.

However, x-ray DSA has a number of disadvantages: it is an invasive procedure, causing discomfort to the patient and requiring a number of experienced staff members to be present, resulting in a relatively high economic cost. There are also significant risks to the patient from the procedure, including exposure to ionising radiation, allergic reaction to the contrast agent, haematoma (bleeding into the tissues, commonly around the site of catheter insertion), silent (asymptomatic) ischaemia [12], or even permanent neurological damage (stroke) [13]. The pressure required for the injection of the contrast agent may modify the normal flow patterns or distort the appearance of distal vessels [3], leading to misrepresentation of the patient's condition. Multiple contrast injections and catheter repositioning are required to get a complete picture of collateral flow from all feeding arteries [3].

In addition, the technique does not give quantitative information on blood flow rates through the arteries, making the assessment of vessel function subjective, and only the status of arteries is assessed, not the resulting perfusion of tissue [10]. The occlusion of a small distal vessel may not be conspicuous in such




Figure 3.1: X-ray digital subtraction angiography (DSA) of the right common carotid artery showing collateral flow into the right posterior cerebral artery (arrow) via the right posterior communicating artery. The subject's orientation is marked on the image: anterior (A), posterior (P), superior (S) and inferior (I). Image courtesy of Dr. Ursula Schulz.

an angiographic technique, and ultimately it is the tissue perfusion which is critical to the patient's condition. Due to the physical mechanism of image formation in x-ray DSA it is generally not possible to obtain a transverse view in 2D imaging, which depicts collateral flow around the circle of Willis most clearly, because this view is obstructed by the patient's body. Finally, as with most of the techniques described in this chapter, this method is sensitive to motion between the acquisition of mask and contrast data sets which can lead to poor subtraction of static tissue, resulting in image artefacts.

3.1.2 Computed Tomography Angiography

Computed tomography (CT) allows the generation of 3D maps of x-ray attenuation within the patient by acquiring images at a large number of angles around the patient. Vessels are not conspicuous in standard CT images, so computed tomography angiography (CTA) is performed by injection of a contrast agent, usually intravenously. Following a slow continuous injection a single 3D image is normally acquired, with the vessels appearing bright due to the high attenuation caused by the contrast agent. Dynamic scanning after fast bolus injection is also possible, giving extra dynamic information [52].

CTA has been shown to be able to identify the presence of collateral vessels [53] and is considerably less invasive than x-ray DSA. However, it does not give vessel-specific information, results in ionising radiation exposure to the patient, carries a risk of contrast agent reaction, has limited spatial resolution (leading to poor visualisation of distal occlusions [54]) and often no temporal information is available.

3.1.3 Transcranial Doppler Ultrasonography

Transcranial Doppler (TCD) ultrasonography makes use of the reflection of ultrasound waves from within the body to generate an image. When these waves are reflected from flowing blood their frequency changes in a manner dependent on the velocity of the blood. Collateral flow can be inferred by the measurement of flow direction in vessel segments [55, 56]. The more sophisticated transcranial color-coded duplex sonography (TCCD) allows the generation of velocity maps rather than measuring velocity at a single location. These techniques are non-invasive, require no ionising radiation, no contrast agent injection and give useful information about flow direction and speed.

However, they also suffer from a number of drawbacks. The bone of the skull attenuates the ultrasound waves, so reasonable signal can only be obtained at certain cranial positions, limiting the coverage of the vascular tree. This becomes more difficult in older patients, meaning that an examination may not be able to take place in up to 30% of cases [56]. Performance and interpretation of the images and/or measured velocities are also subject to considerable variability [3] and the spatial resolution is limited.

3.1.4 Contrast-Enhanced Magnetic Resonance Angiography

A variety of MRI techniques have been developed to generate angiographic contrast. Some make use of gadolinium-based contrast agents which dramatically shorten the T_1 of blood, giving signal enhancement on T_1 weighted images [57]. These are commonly performed by using a short interval between one RF excitation pulse and the next, referred to as the repetition time (TR). Spoiling is also commonly used to dephase any residual transverse magnetisation prior to the next RF pulse, typically by using a large gradient pulse in combination with

variation of the phase of the RF pulses, meaning that the signal measured is proportional only to the longitudinal magnetisation prior to the RF pulse. In such short TR sequences, the longitudinal magnetisation does not have time to fully recover before the next RF pulse, giving higher signal intensities for tissues with a short T_1 , thereby highlighting areas where the contrast agent has reached.

Following the intravenous contrast injection images can be acquired dynamically to give timing information, or as a single image with the vasculature filled with contrast. Both of these acquisitions are typically very fast (10-30 s [57]). The high SNR of these images allows the use of acceleration techniques such as parallel imaging [58, 59] and temporal undersampling [60] to be used with high acceleration factors whilst still maintaining diagnostic image quality. This results in high spatial and temporal resolution images of the cerebral vasculature.

However, along with the discomfort to the patient of having an intravenous injection and the lack of vessel-specific angiographic information there is increasing concern about the link between gadolinium based contrast agents and nephrogenic systemic fibrosis (NSF) in patients with poor kidney function [61]. Non-contrast MRA techniques are therefore preferable where possible.

3.1.5 Non-Contrast Magnetic Resonance Angiography

One of the most commonly used non-contrast techniques is time-of-flight (TOF) magnetic resonance angiography (MRA) [62]. The TOF acquisition is also a short TR spoiled gradient echo sequence which yields high signal intensity for inflowing blood because it has not experienced many previous RF pulses and therefore still has a large longitudinal magnetisation. Static tissue sees all of the pulses and becomes saturated, giving good contrast between flowing blood and tissue. Flow-compensated gradients (see Bernstein *et al.* [29] for a detailed explanation) prevent signal loss due to dephasing of flowing blood during gradient pulses.

Superior saturation slabs destroy the venous signal so that only arteries appear bright in the resulting images. These 3D data sets are often displayed as maximum intensity projections (MIPs). For example, to produce a coronal projection, the maximum voxel intensity along lines parallel to the anterior-posterior direction are calculated and used to form a 2D image. Since the vessels are much brighter than static tissue they dominate the MIP, allowing visualisation of the entire vessel in one image. An example of such a data set is given in Fig. 3.2.

TOF data sets are relatively fast to acquire, are completely non-invasive, show relatively high resolution images of the vasculature (although not comparable to x-ray DSA) and are less sensitive to motion than techniques that rely on image subtraction. However, static tissue is not perfectly suppressed so smaller vessels can often be obscured, particularly in MIPs. No dynamic or vessel-selective information is provided, so the inference of collateral flow can only be made by observation of signal enhancement and/or increased diameter in collateral vessels, suggestive of increased flow. Even if such vessels can be identified, the flow direction is ambiguous. In addition, since the technique relies on inflowing blood, it is insensitive to collateral routes where slow flow occurs [56].

Another popular non-contrast technique is phase-contrast (PC) MRA [63]. In this technique additional gradient pulses are added just after the excitation pulse which cause a phase shift in moving magnetisation in proportion to its velocity. Quantitative 3D velocity maps can be generated in this manner [64], allowing the estimation of volume flow rates in various vessel segments as well as flow direction. However, as for TOF MRA, this technique is insensitive to slow flowing blood [56] and has been shown to underestimate collaterals in some cases [10]. Although flow direction information is obtained, the relative contributions to downstream vessels from each of the feeding arteries is not directly visualised. For example, if there are complex flow patterns around the circle of Willis with




Figure 3.2: Time-of-flight (TOF) magnetic resonance angiography (MRA) in a healthy volunteer. Maximum intensity projections (MIPs) are shown in transverse (TRA), coronal (COR) and sagittal (SAG) views, along with an example single transverse slice from the 3D data set, taken at the position shown by the dashed red line in the COR and SAG MIPs. Subject orientation is marked on the images: R = right, L = left, A = anterior, P = posterior, S = superior and I = inferior.

multiple feeding arteries supplying a distal vessel it would not be possible to determine exactly the contribution of each as it is with vessel-selective techniques. It is also somewhat hampered by long acquisition times [65].

It has been shown that the arterial spin labelling (ASL) technique (see Section 2.4) is capable of generating angiographic contrast if images are acquired rapidly following the ASL preparation pulses rather than waiting for the labelled blood to exchange into tissue [66, 67, 68, 69, 70, 71]. This allows the generation of dynamic angiograms with good suppression of static tissue (see example in Fig. 3.3). Even full 4D (i.e. dynamic 3D) acquisitions are possible [72, 73], allowing the generation of MIP angiograms in any desired orientation. These techniques do not rely on fast flow to generate contrast so are likely to be more able to detect lower levels of collateral flow.

However, the SNR is not as high as some other MRA techniques, and motion between the tag and control images causes unwanted background signal which can obscure smaller vessels. The spatial and temporal resolution are lower than x-ray DSA and some other contrast-enhanced MRA techniques. In addition, visualisation of very distal vessels is compromised by the T_1 decay of the ASL label, particularly if there is delayed transit, for example due to extensive collateral flow. Despite these disadvantages, there are a number of methods for generating vessel-selective ASL contrast, giving it a distinct advantage over other methods (see Sections 3.3 and 3.4).

3.2 Indirect Tissue Perfusion Methods

Angiographic methods allow the assessment of the brain feeding arteries and collateral pathways directly, but ultimately it is the perfusion of the tissue which is vital for the patient. Therefore, methods that can assess tissue perfusion directly




Figure 3.3: Non-selective ASL dynamic angiography using a FAIR labelling scheme and balanced steady-state free precession readout [72, 73]. Selected time frames are shown of the transverse MIP after subtraction of tag and control data sets. The time between the labelling pulse and the acquisition of each frame is displayed under each image. Clear inflow of arterial blood is shown and distal vessel visibility is good. Only the transverse MIP is shown here but this 4D acquisition can also be viewed in other orientations. Note that this FAIR labelling scheme also gives contrast to venous inflow (red arrow) and there is some residual artefact due to eye motion (green arrow).

have some advantage. There are a large number of methods for the measurement of tissue perfusion across a wide range of modalities [74, 75]. These include computed tomography perfusion (CTP) [76], single photon emission computed tomography (SPECT) [77], positron emission tomography (PET) [78], contrast-enhanced MR perfusion [79] and arterial spin labelling (see Section 2.4). Collateral flow may be inferred from the late arrival of blood to an area or vascular enhancement in some imaging modalities [80, 81]. However, the information obtained from such studies is limited as the source of the collateral flow remains unknown.

3.3 Vessel-Specific Arterial Spin Labelling

One of the big advantages of ASL over other potential imaging methods for collateral flow is its ability to be adapted to label specific arteries of interest. In a similar manner to x-ray DSA, the ultimate destination of the blood from the

labelled artery can be visualised, providing unambiguous evidence of collateral flow. In this section methods for labelling a single artery at a time are discussed for each of the main ASL categories: continuous ASL (CASL), pulsed ASL (PASL) and pseudo-continuous ASL (PCASL). These are in contrast to vessel-encoded methods, where multiple arteries are labelled simultaneously and separated using image analysis techniques, as described in the next section. Many of these methods are discussed in more detail in the review by Paiva *et al.* [82].

3.3.1 Continuous ASL

One method for vessel-selectivity with continuous ASL (CASL) involves using a separate labelling coil placed close to the artery to be labelled. The limited spatial coverage of this coil combined with an oblique CASL labelling plane leads to the inversion of blood water only in arteries close to the coil [83]. This approach requires extra hardware, is not able to accurately target specific arteries and is limited to vessels close to the skin. A similar approach can be taken using a standard head transmit/receive coil and exploiting the drop in sensitivity at the edges of the coil [84]. However, this approach is limited to separating the left and right circulation and prohibits the use of a separate receive head coil, as is often used in clinical practice.

Another method which can be implemented with standard hardware is continuous artery-selective spin labelling (CASSL) [85]. The CASL labelling plane is oriented obliquely and rotated with constant angular speed such that there is only a single point which remains constantly within this rotating labelling plane. At this point the inversion efficiency is reasonably high (about 80% relative to standard CASL), and elsewhere the blood is generally saturated. The control condition is achieved by sinusoidal amplitude modulation of the RF pulse, as described in Section 2.4.3, also leading to saturation of spins away from the targeted

vessel. This technique has been shown to be able to label small intra-cranial vessels above the circle of Willis [86]. However, the rotation of the gradient leads to artefacts in superior slices and therefore imaging cannot occur close to the site of labelling. Given the specificity of the labelling point it is also highly sensitive to subject motion and, like standard CASL, many standard clinical scanners are unable to generate sufficiently long RF pulses without additional hardware.

3.3.2 Pulsed ASL

There are two main approaches which make use of individual inversion pulses to generate ASL contrast in specific vessels. The first makes use of multi-dimensional RF pulses to target specific arteries [87]. Selective perfusion images are obtained, but this technique is limited by the presence of side lobes in the RF inversion profile which must be placed outside the head to ensure no unwanted effects occur in the imaging region. The resulting full width at half maximum (FWHM) of the inversion profile is limited to about 10 mm, so closely spaced vessels cannot be easily separated. In addition, the inversion efficiency averaged across the labelled bolus is relatively low.

The second approach makes use of selective inversion slabs to target specific arteries of interest. Although not a conventional ASL technique since the blood water is saturated rather than inverted, one of the earliest examples of such a method was described by Furst *et al.* [88]. Oblique saturation slabs were used to null the signal coming from all the vessels except those of interest. This technique was combined with a TOF MRA acquisition to generate selective angiograms of the carotid arteries and vertebral arteries to study collateral flow.

This idea has been extended with the use of inversion rather than saturation slabs oriented with the aid of MRA images (such as TOF) to invert blood water in a specific artery [89, 90]. Such an approach has been shown to provide

similar information to x-ray DSA regarding collateral flow when combined with TOF MRA in patients with steno-occlusive disease [91]. Typically this preparation is followed by a delay and then the acquisition of images which have tissue perfusion weighting for the labelled artery.

Angiographic contrast can also be obtained with this method. Van Osch *et al.* [71] used a Look-Locker readout strategy [92], which provides multiple images at different inversion times after a single labelling block, to generate vessel-selective dynamic angiograms of the circle of Willis. A single transverse imaging slab 25 mm thick was used to visualise the whole circle of Willis.

However, all of these slab-selective methods are limited to arteries that can be separated using such slabs (e.g. the two ICAs and combined VAs and BA) and some contamination is difficult to avoid given their tortuous nature. It would therefore be difficult to separate the relative contributions from the two VAs for example, or label above the circle of Willis. They are also very sensitive to subject motion between the MRA and ASL acquisitions since for optimal labelling the inversion slabs have to be placed close to other vessels. The planning of such acquisitions to minimise contamination between arteries is also relatively time-consuming.

3.3.3 Pseudo-continuous ASL

Recently, two groups have identified methods similar to the CASSL technique described above, but applied to pseudo-continuous ASL (PCASL). Dai *et al.* [93] added transverse gradients between the PCASL RF pulses, with a direction that rotates over time. The RF phase is also modulated to maintain coherence with the targeted vessel but not at other locations. They found that amplitude modulation of the transverse gradients reduced the labelling efficiency away from the targeted vessel. In contrast to CASSL, the transverse gradients are applied only when RF

pulses are not being played out, so the labelling plane remains constant, rather than rotating with time. This eliminates the problem with artefacts in superior slices found with CASSL. This method has recently been used to generate vessel-selective angiograms capable of showing collateral flow around the circle of Willis and beyond [94].

Helle *et al.* [95] devised a similar method with a range of transverse gradient patterns alongside the rotating gradient method used by Dai *et al.* [93], but without the amplitude modulation. Similar results were obtained, although there is considerably higher labelling efficiency some distance from the targeted vessel, leading to unwanted partial labelling of adjacent vessels, presumably due to the lack of amplitude modulation. Despite this, relatively selective images of small vessels distal to the circle of Willis were demonstrated.

As for CASSL, both of these methods are fairly sensitive to motion between the vessel-localisation scan and the selective ASL acquisition. In addition, the labelling efficiency is reduced relative to non-selective PCASL and there is still considerable signal contamination from arteries up to about 3 cm away from the targeted point, making it hard to disambiguate low levels of collateral flow from partial labelling of adjacent vessels. For example, in the study of Robson *et al.* [94] this resulted in contaminated signals from other vessels which averaged $12 \pm 7\%$, even when labelling was performed in the large vessels below the circle of Willis which are relatively well separated.

3.4 Vessel-Encoded Arterial Spin Labelling

The vessel-specific methods mentioned in the previous section have the advantage that they are generally relatively straightforward to set up since only a single vessel is targeted per acquisition. However, for a complete assessment of collateral

flow it is important to have knowledge of the blood flow from all of the feeding arteries. Labelling a single artery at a time is quite inefficient, since the total amount of labelled blood is considerably lower in each acquisition than if multiple vessels are labelled simultaneously.

To overcome this problem, a number of methods have been devised in which multiple arteries are labelled in different combinations (“encoded”) over a number of acquisitions (“encoding cycles”). The vascular components can then be separated out in post-processing (“decoded”). If there are an equal number of vessels labelled and controlled on each encoding cycle, then the SNR will be comparable to a non-selective acquisition, but with the addition of vessel-selective information.

The first use of this idea was described by Zimine *et al.* [96], where labelling was performed using the angled labelling slab PASL approach. Three encoding cycles were used:

- non-selective control;
- labelling of the RICA and posterior circulation (VAs and BA)
- labelling of the LICA and posterior circulation (VAs and BA)

Vascular territories of the two ICAs and the posterior circulation were calculated by first subtracting each of the label images from the control image and then performing various additions and subtractions of the resulting images. However, this method implicitly assumes that there is no overlap of the territories (i.e. areas supplied by more than one artery), which is not ideal for the study of collateral flow where such patterns are expected.

Günther *et al.* [97] proposed a similar approach to generate vascular territory maps of the two ICAs and BA by using four encoding cycles. One was a non-selective control and the other three cycles involved the labelling of two

feeding arteries in different combinations. This allowed the calculation of the contributions of all three vessels to the signal in each voxel, allowing for overlap of vascular territories. However, it was found that there was still some significant contamination of signals between arteries, perhaps due to slight differences in labelling efficiency between the different encoding cycles. Independent component analysis resulted in better separation of the territories, although it is not clear how such analysis would perform in cases where there is significant overlap of the vascular territories [97].

Both of these approaches are more SNR efficient than the single artery labelling strategies, and somewhat simplify the prescription of labelling slabs to cover the arteries of interest. However, they still suffer from contamination of components related to imperfect positioning of the labelling slab or variation in inversion efficiency with encoding cycle, and thus cannot well represent areas supplied by multiple arteries. They also have lower SNR than a method based on CASL or PCASL due to the limited temporal duration of the labelled bolus.

Wong [48] introduced several significant advances to this field to overcome some of the problems mentioned previously. His method, known as vessel-encoded pseudo-continuous ASL (VEPCASL) makes use of a modified PCASL pulse train, which has the advantage of high SNR, low MT effects and lower demands on hardware than other ASL methods. The addition of transverse gradients between RF pulses and cycling of the RF phase allows the manipulation of the inversion efficiency across the labelling plane in any given direction (see Fig. 3.4). The tag and control locations and the length scale of this profile can be tuned by varying the transverse gradient strength and RF phase. This allows a number of vessel-encoded cycles to be performed with encoding in different directions, labelling different combinations of vessels each time.

As for the vessel-specific PCASL based methods mentioned previously, la-

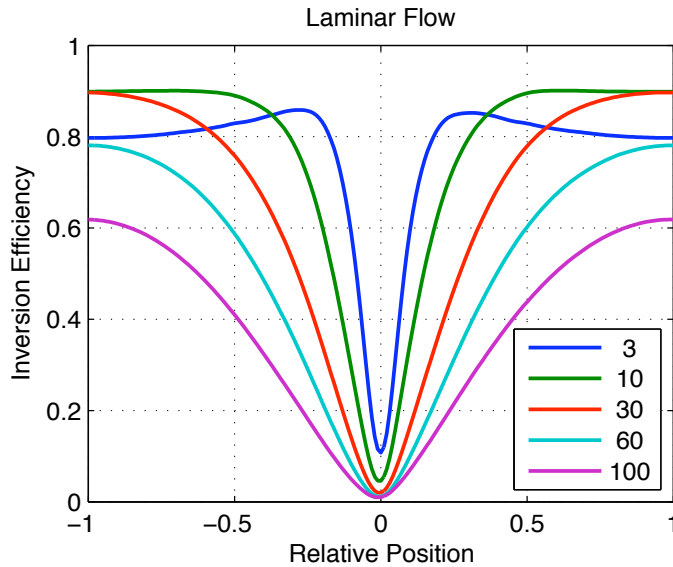


Figure 3.4: The spatial modulation of inversion efficiency across the direction defined by the transverse gradient pulses in vessel-encoded pseudo-continuous ASL [48]. An inversion efficiency of 100% corresponds to perfect inversion (i.e. $M_z = -M_0$ just beyond the labelling plane). A variety of flow speeds are shown, with the assumption of laminar flow and average velocity given in the legend in cm/s. Note that this profile is periodic across space. This graph was derived from simulations of the Bloch equation described in Chapter 4.

bellings occurs only within the defined labelling plane. Thus, prescription of the encoding cycles to effectively label different combinations of vessels is relatively straightforward compared to PASL techniques where an oblique labelling slab must cover the entire vessel whilst excluding other vessels.

The inversion efficiency in each vessel can be measured directly from the data [48] which aids the analysis used to separate the measured signals into vascular components (see Section 4.11). This separation, if done correctly, effectively removes signal contamination that plagues some of the other methods described above. In addition, the analysis can be used to estimate patient motion between vessel-localisation scans and the VEPCASL acquisition, making it less sensitive to subject motion than vessel-specific techniques [98].

VEPCASL results in well defined vascular territory maps below the circle of Willis and has been shown to correlate well with x-ray DSA in patients with ICA

or MCA stenosis [99]. It can also separate out the territories of the two vertebral arteries [100], which has not been shown with other techniques. Application above the circle of Willis is also possible [101], although in this case the analysis used made the assumption that each voxel is fed only by a single vessel, limiting its utility for observation of collateral flow. However, the use of a greater number of encoding cycles could remove this restriction.

In summary, VEPCASL has many advantages over other methods for the assessment of collateral flow, providing highly vessel-selective, potentially dynamic tissue perfusion information (if multiple post-labelling delays are used). The SNR efficiency is high relative to other vessel-selective ASL techniques, it is completely non-invasive, is not limited by hardware constraints and does not result in artefacts in the imaging region. As a result, this is the method implemented to study collateral flow for the remainder of this thesis.

However, it does have a number of disadvantages. As with all ASL techniques, the SNR is limited and in cases of severely delayed blood arrival T_1 decay of the labelled blood will reduce the SNR further. The analysis techniques required are a little more complex than some other methods which label a single artery (see Section 4.11), so it may not be suitable for clinical use until the separation of vascular components can be performed robustly on the scanner console. This complexity increases considerably when labelling large numbers of feeding arteries. In addition, when collecting data weighted by tissue perfusion, no morphological information is obtained on the feeding vessels themselves and collateral flow across the circle of Willis or other pathways is not visualised directly. However, it is possible to combine the VEPCASL preparation with an angiographic readout to obtain such information (see Chapter 6).

3.5 Conclusions

In this chapter current methods for the assessment of collateral blood flow in the brain were described. The current clinical gold standard, x-ray digital subtraction angiography, carries significant risks to the patient and is expensive. A non-invasive method which provides similar information is therefore desirable. Many of the alternative methods are hampered by the need for contrast agent injections, or the lack of vessel-specific or dynamic information about blood flow. Arterial spin labelling techniques avoid many of these problems, with vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL) being the most SNR efficient of these methods. For this reason VEPCASL was selected as the method to study collateral flow for the remainder of this thesis. In the following chapter simulations and phantom experiments are conducted to assess and optimise the VEPCASL pulse sequence, along with discussion of methods for the post-processing of vessel-encoded data, allowing the robust separation of vascular components.

Chapter 4

Assessment and Optimisation of Vessel-Encoded ASL

4.1 Introduction

The mechanism by which the pseudo-continuous arterial spin labelling (PCASL) [40] pulse train inverts magnetisation when in “tag” mode and leaves the magnetisation relatively unperturbed when in “control” mode is not a simple one. The addition of transverse gradients to perform vessel-encoding [48] adds extra complexity. Before implementing such a sequence and testing it in healthy volunteers, it is important to understand how the magnetisation responds to this sequence of RF and gradient pulses and how this behaviour is modified by various sequence parameters or external influences.

In order to do this, simulations of the Bloch equations [23], which govern the motion of the magnetisation (see Section 2.3.2), were performed to see how the magnetisation responds to the vessel-encoded PCASL (VEPCASL) pulse train under a range of conditions. Results of the simulations were validated using a flowing water phantom and an implementation of the VEPCASL sequence with a

simple gradient echo readout.

The simulations and experiments were designed to help understand:

- the mechanism by which the magnetisation vectors are inverted or controlled;
- the sequence parameters that optimise the ASL signal;
- the effect of blood velocity and off-resonance frequency on labelling efficiency;
- the spatial modulation of the inversion efficiency for vessel-encoded acquisitions;
- the sensitivity of the inversion efficiency to vessel angulation through the labelling plane; and
- the effect of the sequence on static magnetisation near the labelling plane.

By understanding these effects an optimal protocol for VEPCASL can be established.

Following these experiments, methods for separating out vascular components from the vessel-encoded data are described and tested (Section 4.11), ranging from simple but naive approaches to a full Bayesian framework developed by Dr. Michael Chappell.

4.2 Simulation Methods

Simulations of the Bloch equations (defined in Section 2.3.2) were performed for a single magnetisation vector, M , (also referred to as a “spin” for brevity) at a time. The gradients, G , radio-frequency (RF) pulses, B_1 , (with amplitude, B_1 ,

and phase, ϕ_1), spin position, \mathbf{x} , and off-resonance frequency, ω_{OR} , are all defined at discrete time points, t , separated by an interval, Δt .

At each time point, the effective local magnetic field, \mathbf{B}_{eff} (see Section 2.3.5), experienced by the spin was determined in a frame of reference rotating at the Larmor frequency, $\omega_0 = \gamma B_0$, thus enabling calculation of the rotation matrix used to act upon this spin.

To recap, \mathbf{B}_{eff} has contributions from:

- the transmitted RF pulse, \mathbf{B}_1 ;
- precession about the z axis due to the spin's position in the gradient field, \mathbf{B}_G ; and
- precession about the z axis due to the spin's off-resonance frequency, \mathbf{B}_{OR} , caused by local field inhomogeneity.

Thus \mathbf{B}_{eff} can be written in the rotating frame as:

$$\mathbf{B}_{\text{eff}} = \mathbf{B}_1 + \mathbf{B}_G + \mathbf{B}_{OR} \quad (4.1)$$

where

$$\mathbf{B}_1 = B_1 \begin{pmatrix} \cos \phi_1 \\ -\sin \phi_1 \\ 0 \end{pmatrix}, \quad \mathbf{B}_G = \begin{pmatrix} 0 \\ 0 \\ \mathbf{x} \cdot \mathbf{G} \end{pmatrix} \quad \text{and} \quad \mathbf{B}_{OR} = \begin{pmatrix} 0 \\ 0 \\ \omega_{OR}/\gamma \end{pmatrix} \quad (4.2)$$

These vector quantities are shown graphically in Fig. 4.1 (left). Ignoring relaxation processes temporarily, the Bloch equations (see Section 2.3.2) in the rotating frame reduce to:

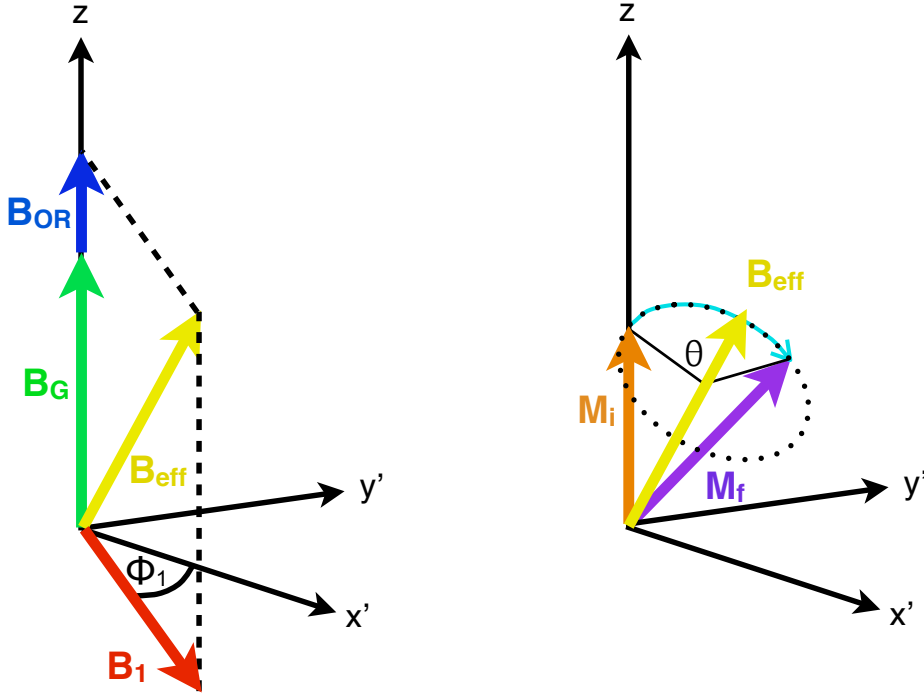


Figure 4.1: Left: contributions to the effective magnetic field, \mathbf{B}_{eff} , (yellow) from the RF pulse, \mathbf{B}_1 , (red, with phase ϕ_1), the spin's position in the gradient field, \mathbf{B}_G , (green), and the local off-resonance frequency due to field inhomogeneities, \mathbf{B}_{OR} , (blue). Note that the RF phase is defined such that $\phi_1 = 0$ corresponds to \mathbf{B}_1 aligned with the x' axis and an increase in ϕ_1 rotates \mathbf{B}_1 about the z axis according to the left hand rule. Right: the effect of \mathbf{B}_{eff} on an initial magnetisation vector, \mathbf{M}_i , (orange): rotation about the axis of \mathbf{B}_{eff} through an angle θ according to the left hand rule, following the cyan path to the final magnetisation vector, \mathbf{M}_f (purple).

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B}_{\text{eff}} \quad (4.3)$$

Thus in time interval Δt , \mathbf{M} rotates about \mathbf{B}_{eff} (according to the left hand rule) by an angle θ given by:

$$\theta = \gamma |\mathbf{B}_{\text{eff}}| \Delta t \quad (4.4)$$

The advantage of this approach is that it is accurate when θ is large (e.g. if the spin is far from isocentre during a large gradient), since the rotations due to all the influences are considered together. If the rotation due to each component

of \mathbf{B}_{eff} were performed separately, this would lead to significant errors when θ is large since such rotation matrices do not commute. For example, a spin far from isocentre in a large gradient which is also subjected to a \mathbf{B}_1 field would see an effective magnetic field which is approximately parallel to the z axis. Thus, if it were initially aligned along z , then the magnetisation vector would not change significantly during a small simulated time interval. However, if the \mathbf{B}_1 field were allowed to act upon the magnetisation, tipping it away from the z axis before the precession due to the large gradient were considered, the final magnetisation would not be aligned with z as it should be.

Using the effective magnetic field approach thus allows the use of longer simulated time intervals, Δt , provided that there are only small changes to the RF, gradient or off-resonance influences during this time.

The simulation also models both T_2 and T_1 decay. Following the application of the rotation matrix the transverse components are attenuated by a factor $\exp(-\Delta t/T_2)$ and the difference between M_z and M_0 shrinks by a factor $\exp(-\Delta t/T_1)$.

In these simulations it was assumed that $T_1 = 1932$ ms and $T_2 = 275$ ms: values typical for blood at 3T [102]. These are somewhat higher than other literature values, although the effect on the results of simulations is minimal for small changes in these values.

All simulations were performed using Matlab[®] (Mathworks[™], Natick, Massachusetts, USA). For these simulations it was important to have a realistic estimate of blood flow velocities in the arteries supplying the brain. Values quoted in the literature vary considerably. Wong [48] assumes values between 5 and 40 cm/s in simulations. Most other studies report values in the range 15-70 cm/s [103, 104, 105, 106]. Internal carotid artery (ICA) systolic velocities up to 250 cm/s in those with significant stenosis (at the level of the stenosis) have been reported

[107]. These simulations will therefore consider velocities between 5 and 80 cm/s as a representative sample unless otherwise specified.

Sequence parameters (such as the gradient strength, RF pulse duration etc.) are given in Table 4.1. These are identical to those used by Wong [48] unless otherwise specified, with the exception of the RF pulse shape: Gaussian rather than Hanning pulses were used to match the sequence implemented on Siemens Healthcare (Erlangen, Germany) MRI systems (see Section 4.3). These simulations used the numerical output from the Siemens sequence simulator to ensure an accurate comparison with phantom data.

To be able to compare performance of the sequence under different conditions, the inversion efficiency, E_I , is calculated. This accounts for T_1 decay following the inversion of the spin by extrapolating the z component of the magnetisation at the end of the simulation, $M_{z,\text{final}}$, back to the time at which the spin crossed the labelling plane and normalising to the case of an optimal inversion:

$$E_I = \frac{M_0 - M_{z,\text{final}}}{2M_0} \exp\left(\frac{t_{\text{final}} - t_{\text{label}}}{T_1}\right) \quad (4.5)$$

where t_{final} is the time at the end of the simulation and t_{label} is the time at which the spin crosses the labelling plane. For a perfect inversion, $E_I = 1$, and for a perfect control, $E_I = 0$. It is also useful to define the contrast, C , between tag and control conditions:

$$C = E_{I,\text{tag}} - E_{I,\text{control}} \quad (4.6)$$

This accounts for any perturbation in the magnetisation caused by the control sequence and is proportional to the ASL difference signal which would be measured experimentally.




Figure 4.2: One subunit of the sequence implemented for phantom experiments, showing the VEPCASL pulse train, a spoiler to remove any residual transverse magnetisation (on G_z) and a simple one line readout followed by another spoiler. The VEPCASL labelling pulse train consists of a series of radio-frequency (RF) pulses and associated z gradient (G_z) pulses to invert spins passing through an axial labelling plane. In the control condition, consecutive RF pulses have a phase difference of π relative to each other. Transverse gradients (shown here on G_x) are applied along the desired vessel encoding direction only when selective labelling is required. In this case the RF phase is also modulated to maintain phase coherence with magnetisation at the tag location. Three blue dots represent continuation of the same repeating pattern, omitted for clarity. Note that the z gradient active during the labelling RF pulses is negative, due to an automatic rotation matrix used by the scanner, but this has no bearing on its ability to label the blood. (ADC = analogue to digital converter; RF = radiofrequency pulses; G_x = vessel encoding/phase encoding gradient; G_y = slice selection gradient; G_z = labelling/readout gradient).

4.3 Sequence and Phantom Design

The VEPCASL preparation was implemented following the design set out by Wong [48] as a self-contained module, allowing it to be easily included in a number of different pulse sequences. For these initial phantom experiments a simple single-line gradient echo readout was used (see Fig. 4.2) to avoid potential artefacts which can arise from more complex readouts. Sequence parameters are given in Table 4.1, which were used by default unless otherwise specified.

In order to test the results of the simulations and ensure the implemented

<i>Group</i>	<i>Parameter</i>	<i>Value</i>
VEPCASL Pulse Train	VEPCASL pulse train ("tag") duration	500 ms
	RF Pulse Type	Gaussian
	RF Flip Angle	20°
	RF Duration	600 μ s
	RF Separation	960 μ s
	Mean Tagging Gradient	0.8 mT/m
	Tagging Gradient Amplitude*	6 mT/m
	Post label delay	10 ms
Readout	Imaging Plane	Coronal
	Phase Encode Direction	Perpendicular to the tubes
	Field of View	Approx 250 x 62.5 mm
	Matrix size	Approx 256 x 64
	Slice Thickness	10 mm
	TR	3000 ms
	TE	3.6 ms
Other	Excitation Flip Angle	90°
	Average water velocity	10 cm/s

Table 4.1: Default sequence parameters used for phantom experiments, which closely match those of Wong [48]. *Refers to gradient amplitude during RF pulses in the VEPCASL pulse train.

VEPCASL sequence was performing as expected, a simple flowing water phantom was constructed. It consisted of a long plastic tube with an internal diameter of approximately 5 mm attached to a pump which had been calibrated to allow accurate setting of the average velocity. The tube was passed through the bore of the scanner, a loop of several metres left outside the scanner before being passed back through the bore in the opposite direction. This would allow visualisation of the labelling of spins flowing in both directions. The two sections of the tube were attached to a wooden board to prevent motion during the scan. A standard bottle phantom containing NiCl_2 solution was placed just above the tubes to allow assessment of the effect of the sequence on static magnetisation, provide sufficient loading for the coil and allow accurate scanner frequency adjustment prior to the scan.




Figure 4.3: Time-course of a magnetisation vector flowing through the labelling plane at 10 cm/s for the CASL (left) and PCASL (right) sequences. The curves are very similar, although the refocussing gradients on G_z in the PCASL sequence cause the magnetisation to oscillate quickly about the z axis. Note that the inversion occurs in a small region close to the labelling plane (i.e. within 100 ms = 1 cm).

4.4 Inversion Mechanism

The PCASL sequence can be thought of as a discrete approximation to the continuous ASL (CASL) regime, as discussed in Section 2.4.4. In order to understand how these sequences cause inversion of magnetisation, a single spin flowing through the labelling plane was simulated for both CASL and PCASL sequences using identical mean B_1 and mean G_z values to allow comparison. The PCASL sequence used here was simply the VEPCASL sequence described by Wong [48] in the non-selective “tag all” mode. Fig. 4.3 shows the time course of the components of the magnetisation and Fig. 4.4 shows a 3D view of the same magnetisation vector just before, at, and after the labelling plane.

As discussed in Section 2.4.3, as the spin flows through the labelling plane in the CASL sequence, the effective magnetic field, \mathbf{B}_{eff} , rotates from pointing towards $+z$ to pointing towards $-z$ due to the varying contribution from the applied gradient and constant contribution from the RF pulse. Since this rotation of \mathbf{B}_{eff} is slow, this causes the magnetisation vector, which is precessing about the

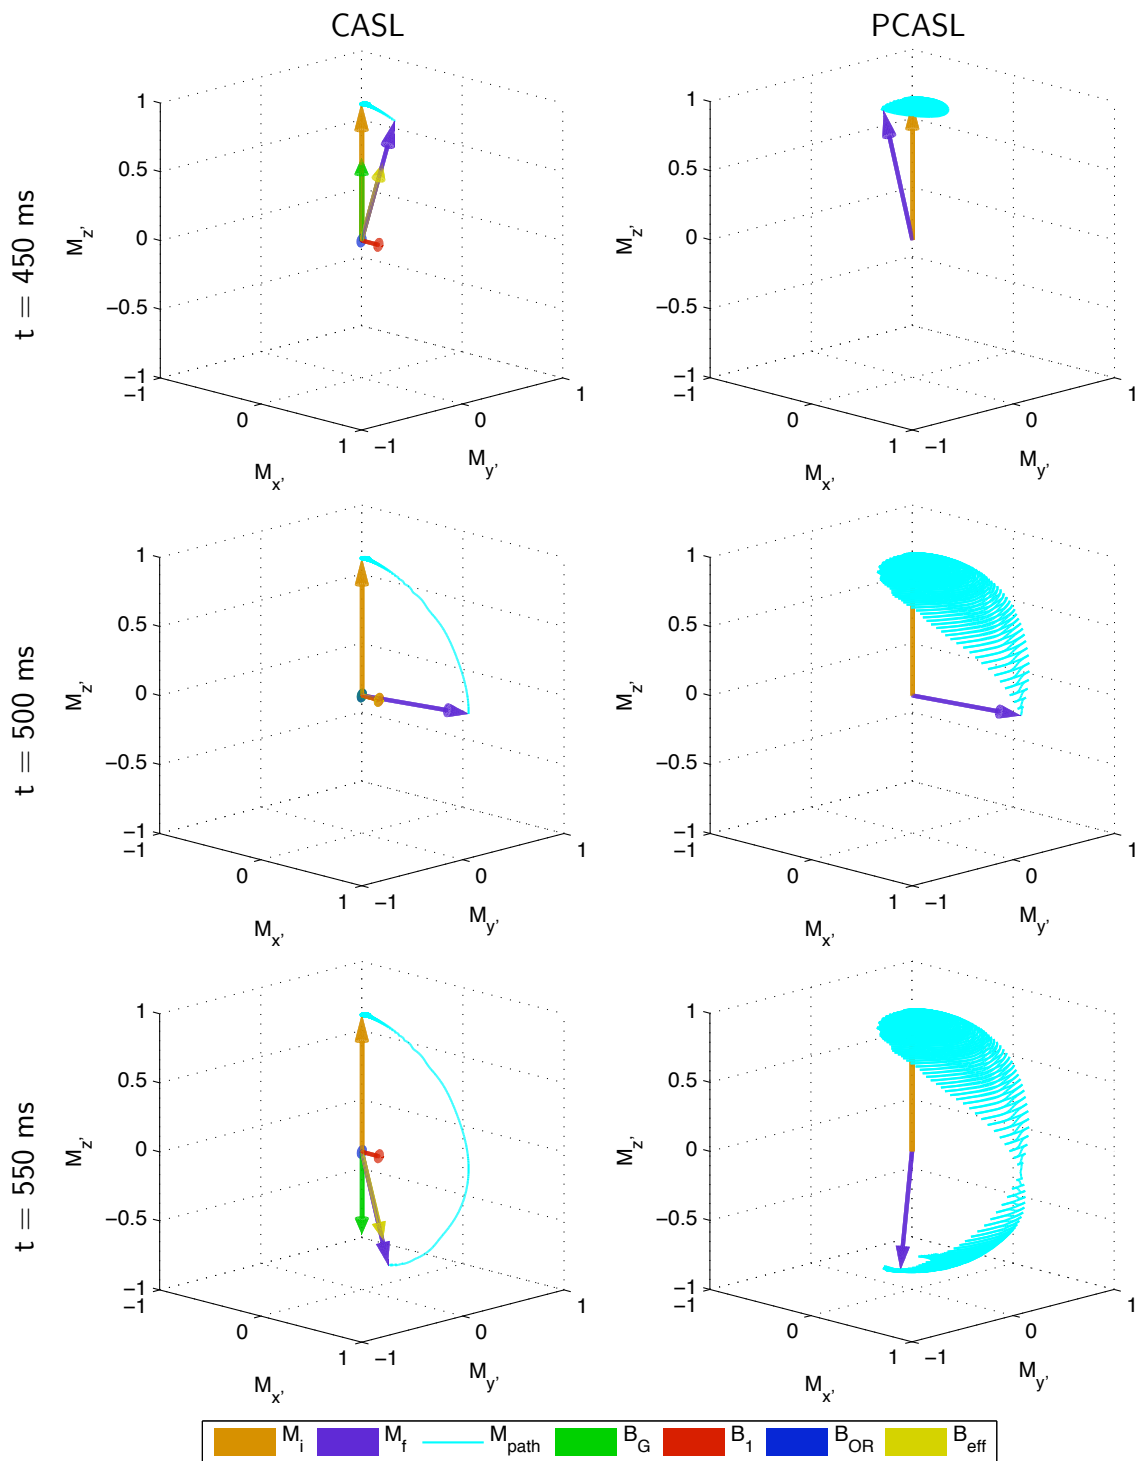


Figure 4.4: 3D view of the magnetisation vector flowing through the labelling plane at 10 cm/s for the CASL (left) and PCASL (right) sequences when the spin is before (top), at (middle), and past (bottom) the labelling plane. The times shown correspond to those in Fig. 4.3. The colours are the same as those in Fig. 4.1, representing the initial magnetisation, M_i , the final magnetisation, M_f , and the path of the magnetisation, M_{path} . The components of the effective field, B_{eff} , due to gradients (B_G), RF pulses (B_1) and off-resonance (B_{OR} , zero here) are shown for the CASL case where the magnetisation tracks B_{eff} .




Figure 4.5: Final M_z (left) and the calculated inversion efficiency (right) of a set of spins with different starting locations along the z axis, moving at 10 cm/s and subjected to a PCASL pulse train of duration 500 ms. This leads to a well defined bolus with high inversion efficiency.

effective magnetic field, to also invert. In the PCASL sequence, the spin follows not the effective field itself, but the steady state position of a static spin in that position, as discussed in Section 2.4.4. In a manner analogous to CASL, this steady state vector inverts as the spin passes through the labelling plane, leading to adiabatic inversion of the spin. Both sequences demonstrate efficient inversion of the blood magnetisation.

If the PCASL pulse train is played out for a significant time interval it leads to a well defined bolus of labelled blood, as shown in Fig. 4.5, with a consistent inversion efficiency of approximately 90% within the bolus.

It is also important to establish that the PCASL control condition does not significantly perturb the magnetisation so that the contrast between the tag and control signals is maximised. In “control all” mode the RF phase is alternated between zero and π such that the magnetisation near the labelling plane is flipped back and forth, but not inverted. The time course of blood magnetisation flowing through the labelling plane is shown in Fig. 4.6, showing that the control condition produces minimal perturbation of the z component of the magnetisation.

To confirm these results in the flow phantom, “tag all” and “control all” cycles

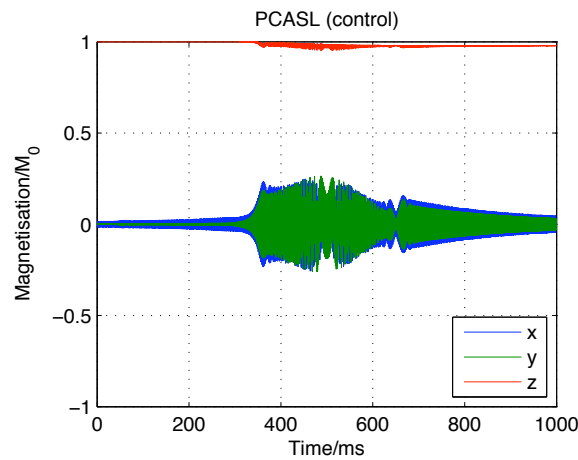


Figure 4.6: Efficiency of the PCASL control condition. The time-course of the magnetisation is shown for the same situation described in Fig. 4.3 but with the sequence in “control all” mode, such that the RF phase is alternated between zero and π . This leads to the magnetisation being flipped back and forth, resulting in very little perturbation of the z component of the magnetisation and thus an effective control condition.

were played out using a long tag duration (2000 ms) and average flow speed of 10 cm/s. Complex subtraction was performed and the resulting images normalised to twice the intensity of the control image. Once corrected for T_1 decay this gives an approximate measure of contrast, as defined in Eq. 4.6. The calculation was performed only within a mask created from the control image to exclude low signal regions. Example images and results are shown in Fig. 4.7.

If the absolute value of the complex subtraction were to be taken then the resulting values would be somewhat biased by noise in areas where the contrast is low since the result is always positive. In order to avoid this problem, a phase correction step was undertaken before the complex subtraction. This involved subtracting the phase of the control image from both the tag and control images on a voxelwise basis before performing the complex subtraction and taking the real part of the result. This is equivalent to projecting the complex difference signal along the direction of the signal in the control image, thereby excluding noise which is orthogonal to this direction and allowing both positive and negative results to be obtained.

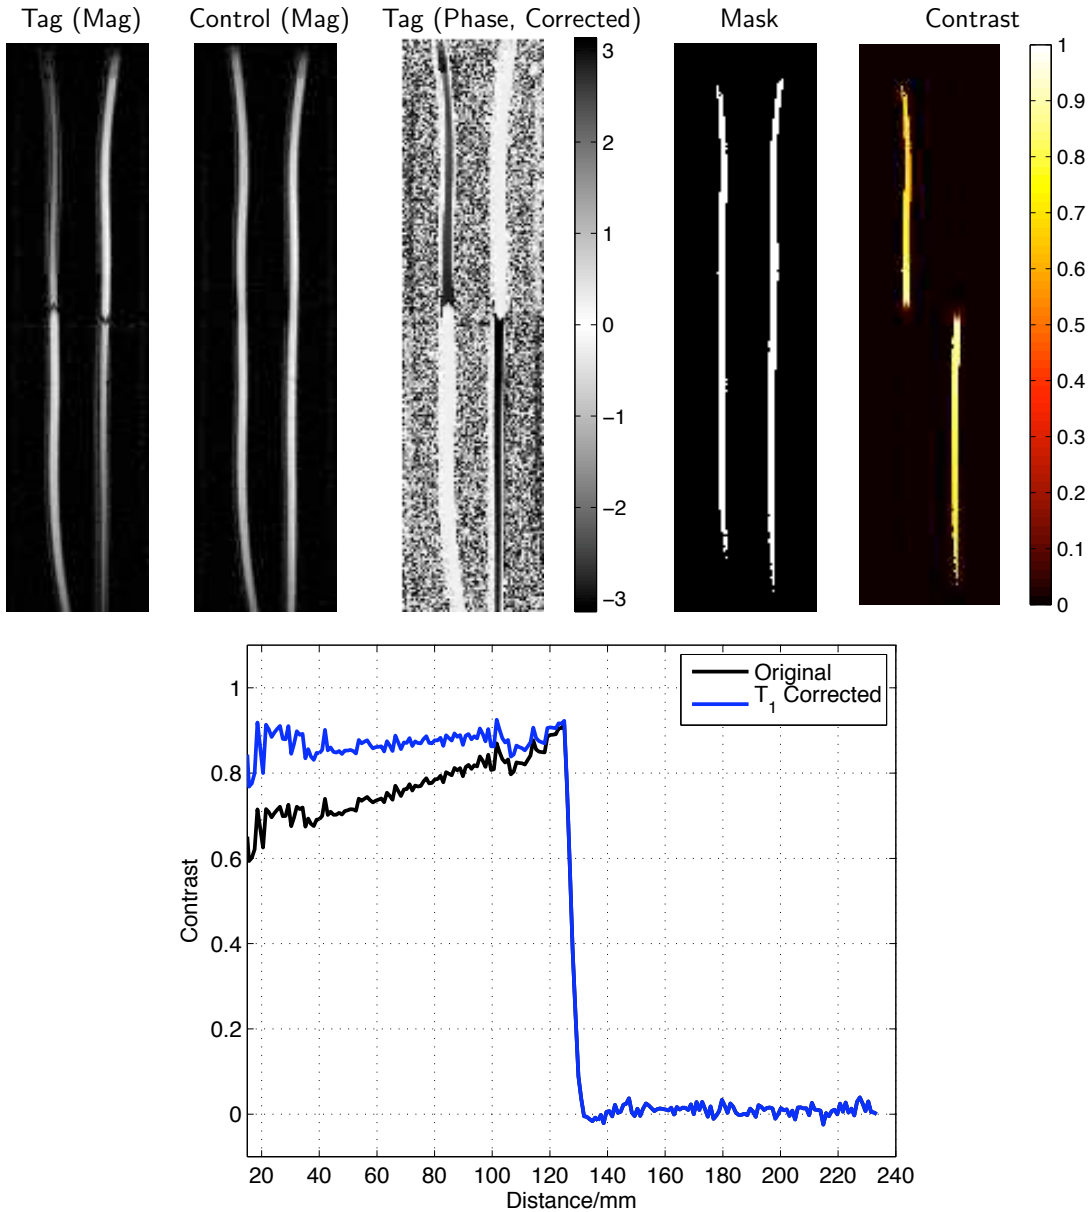


Figure 4.7: Assessment of ASL contrast in the flow phantom for non-selective PCASL. The images shown were collected using a 2000 ms tag duration and average flow speed of 10 cm/s. Water in the tube on the left of the image is flowing towards the top of the image. The opposite is true for the right hand tube. Magnitude (“Mag”) images are shown for both the non-selective tag and control cycles. The phase of the tag image after phase correction (see text) is also shown: the labelled bolus has a phase of π relative to the control condition, demonstrating clearly that inversion of the magnetisation has occurred. The mask created from the control image and used to exclude low signal regions is also shown along with the calculated ASL contrast within this mask (before correction for T_1 decay). The plot shows the mean contrast within the mask for the right hand tube, with and without correction for T_1 decay. Since the mask limits the calculations to a region near the centre of the tube where faster flowing water is present, a velocity of 20 cm/s was assumed for this correction which gives the expected flat contrast profile within the bolus.

The mean contrast within the mask at each point along the right hand tube in Fig. 4.7 was calculated and corrected for T_1 decay, assuming that the T_1 for water is 2200 ms at 3T [108]. The mean T_1 -corrected contrast within the mask after the labelling plane is $87.1 \pm 2.7\%$. The equivalent value calculated from simulations, assuming a laminar flow profile (see below) with an average velocity of 10 cm/s is 85.7%, giving a good agreement between theory and experiment.

4.5 Velocity Dependence

Given the potential variation in blood velocity between different arteries and subjects, it is important to determine how the PCASL sequence performs over a range of velocities. Simulations were performed assuming both uniform flow within the artery (i.e. constant velocity over the cross-section of the vessel) and laminar flow, where the velocity, v , varies quadratically with distance from the centre of the vessel, r :

$$v(r) = \begin{cases} 2v_{av} \left(1 - \left(\frac{r}{r_0}\right)^2\right) & r \leq r_0 \\ 0 & r > r_0 \end{cases} \quad (4.7)$$

where v_{av} is the average velocity and r_0 is the radius of the vessel. Laminar flow has been shown to be the more realistic assumption for blood flow in carotid arteries [109].

Fig. 4.8 shows that the inversion efficiency has a significant velocity dependence. At low velocities, the magnetisation spends longer in the transverse plane as it is inverted, leading to greater decay of the magnetisation via T_2 decay and thus reduced inversion efficiency. At high velocities the magnetisation passes through the labelling plane too quickly, preventing it from effectively following the steady-state position of the magnetisation and thus the efficiency of the inver-

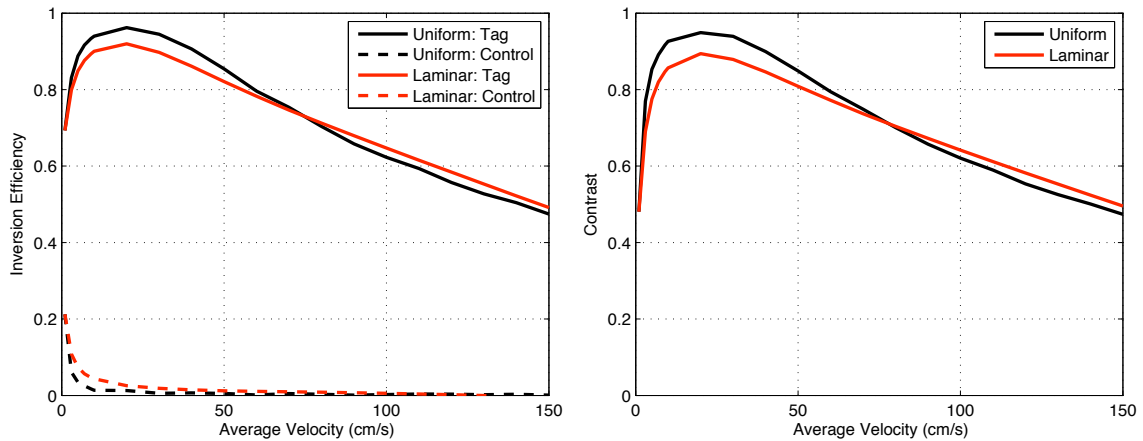


Figure 4.8: Inversion efficiency (left) for non-selective tag and control PCASL cycles and contrast (right) as a function of average blood velocity through the labelling plane. Results are shown assuming both uniform flow and laminar flow inside the vessel.

sion is also reduced. This dependence may result in an underestimation of blood flow if the velocity is unusually high or low unless this effect is accounted for. However, the inversion efficiency remains reasonably high over the intermediate range of velocities typically found in the cerebral arteries. The magnetisation is generally unaffected by the control sequence except at very low velocities when the blood spends a much greater time near the labelling plane, allowing T_2 decay to become significant. As might be expected, the more realistic assumption of laminar flow is a somewhat smoothed out version of the curve assuming uniform flow since laminar flow incorporates a range of velocities.

This velocity dependence was tested experimentally by varying the average flow velocity in the phantom using a tag duration of 500 ms. Since lower velocities were considered here, a significant proportion of the flowing water was sometimes unlabelled and could therefore bias calculations of the ASL contrast if the entire phantom were included in the calculations. To avoid this problem, masks were created for the labelled bolus by identifying regions which have non-negligible magnitude signal in the control case (to exclude air) and a phase difference close to π in the phase corrected tag image (to isolate inverted magneti-

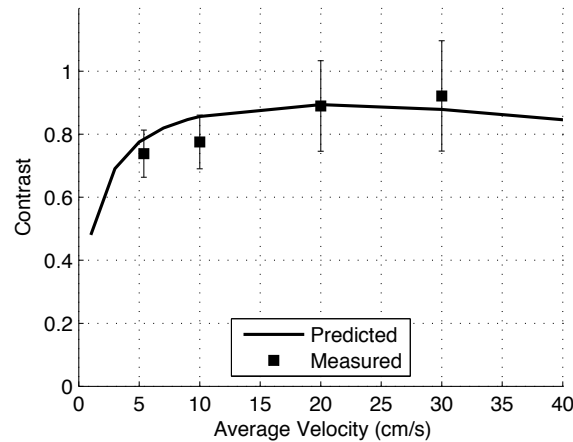


Figure 4.9: Variation in contrast with average water velocity. The mean and standard deviation were calculated within the labelled bolus after correction for T_1 decay and averaged between the two tubes of the phantom. Predicted values from simulations of the Bloch equations assuming a laminar flow profile fit the data well.

sation). The mean and standard deviation of the contrast were calculated within this mask and corrected for T_1 decay.

Results are shown in Fig. 4.9, and qualitatively fit well with the results from simulations assuming laminar flow. This can be tested more rigorously by calculating the statistic $\chi^2 = 1.1$. The null hypothesis is defined to be that the data were drawn from normal distributions centred about the predicted curve. The probability, p , that this value of χ^2 or greater were drawn from this distribution can then be calculated as $p = 0.89$. Thus, there is no evidence to reject the null hypothesis that there is any significant difference between the predicted and measured ASL contrast.

Note that due to the lack of flow compensation in the readout, there was some signal loss for the higher velocity experiments. This affects both tag and control images equally, so does not affect the calculated contrast, but it does lead to a decrease in SNR and therefore an increase in the error estimate of the contrast.




Figure 4.10: Contour plots showing how the contrast varies with \bar{G}_z and RF flip angle, α . The mean (left) and standard deviation (right) of the contrast are calculated over the velocity range 5 to 80 cm/s

4.6 Parameter Choice for Maximal ASL Contrast

There are a number of sequence parameters which can be varied in order to maximise the ASL contrast. The optimal parameter choice is influenced by the spin velocity, so here a range of velocities from 5 to 80 cm/s is simulated and the average contrast taken across this range. Figure 4.10 shows how the contrast varies with two of the main sequence parameters: the mean z gradient amplitude, \bar{G}_z , and the flip angle of each RF pulse, α .

Note the drop off in the contrast for \bar{G}_z above about 1 mT/m. This is due to aliased labelling planes, which occur at locations where the phase accrual between two RF pulses due to the mean gradient is equal to an integer multiple of 2π . As \bar{G}_z is increased, these aliased labelling planes move closer to the central labelling plane. Above some critical value, these aliased planes reach a point where the RF pulses have a significant effective flip angle, causing disruption to the magnetisation and inefficient labelling.

Regions of parameter space where the mean contrast is high are desirable to ensure maximal contrast over a range of flow velocities. In addition, regions

where the standard deviation of the contrast across the velocity range is low are less susceptible to errors in the case of unusually high or low blood velocity. It should be noted that the use of the standard deviation here does not imply that the distribution is Gaussian, but is merely a simple measure of variation across the velocity range of interest.

Finally, it is also important to use as small a flip angle as possible to minimise RF power deposition to the subject. These simulations therefore suggest that \bar{G}_z and α should be set to approximately 0.6 mT/m and 20° respectively. However, there is not a significant drop in contrast when increasing \bar{G}_z to 0.8 mT/m, so this was done in most of the experiments described in following chapters to match the parameters used by Wong [48] for comparative purposes.

These results were verified using the flow phantom by varying both α and \bar{G}_z whilst keeping the other fixed. Results are shown in Fig. 4.11, giving reasonable qualitative agreement with the simulations. This agreement was tested statistically as described in Section 4.5. For the variation with α this yielded $\chi^2 = 14.5, p = 0.006$, implying a significant difference. However, this calculation was dominated by the data point at $\alpha = 5^\circ$ where the simulated curve varies rapidly and is therefore very sensitive to small variation in the value of α at which the data point was measured. If it is assumed that α can vary by approximately 10% across the field of view [90], then repeating this calculation with the values of α reduced by 10% for the measured data points yields $\chi^2 = 7.5, p = 0.11$, implying there is no significant difference between the simulated and measured values. Although not statistically rigorous, this calculation demonstrates that the variation in α over the field of view is a reasonable explanation for the observed discrepancy. The equivalent statistical calculation for variation of the ASL contrast with \bar{G}_z gives $\chi^2 = 2.75, p = 0.60$, implying a good agreement between simulations and measurements in this case.




Figure 4.11: Variation in contrast with tagging RF flip angle, α , (left) and the mean z gradient, \bar{G}_z , (right) in the flow phantom. When varying α , \bar{G}_z was fixed at 0.8 mT/m. When varying \bar{G}_z , α was fixed at 20° . Predicted values are from simulations of the Bloch equations which agree well with the experimental data.

Simulations also show that there is very little variation of inversion efficiency with RF pulse duration, but the contrast is compromised for RF separations over $2000 \mu\text{s}$ and tagging gradient amplitudes of less than 5 mT/m. In keeping with Wong [48], the RF duration, RF separation and tagging gradient amplitude were set to $600 \mu\text{s}$, $960 \mu\text{s}$ and 6 mT/m respectively for future experiments.

4.7 Spatial Modulation of Inversion Efficiency

The transverse gradients of the VEPCASL sequence cause spins in vessels sited at the “tag” location to be inverted, whilst those at the “control” location are relatively unaffected. However, for a real vessel geometry it is seldom possible to arrange all vessels to be exactly at a tag or control location. In order to decode the contributions to the ASL signal from each of these vessels using the matrix inversion or Bayesian methods described in Section 4.11, it is essential to know how the inversion efficiency is modulated across space.

Results of simulations (Fig. 4.12) show that this inversion profile is periodic in nature. Magnetisation sited at locations which experience $2n\pi$ phase accrual




Figure 4.12: Spatial modulation of the inversion efficiency across the transverse gradient direction assuming both uniform flow (left) and laminar flow (right). The control location has a relative position of zero and the tag location is at 1 or -1, due to the periodic nature of this profile. Each line represents a different average blood velocity, given in cm/s in the legend.

(where n is an integer) due to the transverse gradient pulse effectively “see” RF pulses which all act in the same direction, equivalent to the PCASL tag condition. Magnetisation at locations which experience phase accrual of $(2n + 1)\pi$ due to the transverse gradient “see” RF pulses of alternating phase, equivalent to the PCASL control condition. Magnetisation at intermediate locations experiences partial inversion. The profiles show a strong velocity dependence, with a sharp trough at the control locations at low velocities and a more sinusoidal shape at higher velocities. Making the more realistic assumption of laminar flow modifies the profiles somewhat, particularly by decreasing the inversion efficiency at the tag locations and making the control trough narrower.

To confirm the simulated spatial modulation function, another phantom experiment was conducted to assess the labelling efficiency as a function of spatial position. To achieve this, the tag location was set to a number of locations at or between the two tubes of the phantom. The control position was altered in tandem to maintain a constant length scale of the inversion efficiency profile. The measured contrast relative to a non-selective inversion experiment is plotted in




Figure 4.13: Variation in contrast (relative to non-selective PCASL) with position of the phantom tube within the vessel encoding. Predicted values taken from simulations are shown for full laminar flow and for the maximum velocity only. A closer match to the maximum velocity trend is achieved, perhaps because the mask used covers the central portion of the tube where the flow is faster.

Fig. 4.13 along with predictions made by simulations. A much better fit to the measured data is achieved by assuming uniform flow with a velocity equal to the maximum expected within the laminar flow pattern ($\chi^2 = 16.2, p = 0.18$), rather than the full laminar profile ($\chi^2 = 252, p = 0$). This is likely to be due to the mask used which is based partly on magnitude data and thus the measured ASL contrast is dominated by the centre of the tubes where the flow is fastest.

4.8 Off-Resonance Effects

Local susceptibility effects and main magnetic field (B_0) inhomogeneities can lead to spins in different locations precessing with slightly different frequencies. The resulting phase accrual between PCASL RF pulses could disturb the inversion mechanism. This is particularly important because shimming is usually optimised

for the imaging region, leaving the neck potentially poorly shimmed.

Simulations showed that off-resonance frequencies above about 100 Hz have a significant impact on inversion efficiency (Fig. 4.14). Off-resonance has two effects on the PCASL labelling scheme. Firstly, there is a small shift to the position in z at which B_{eff} is entirely within the transverse plane during the RF pulses. Due to the high gradient strength during RF pulses, this shift is quite small (e.g. 0.8 mm for 200 Hz off-resonance and tagging gradient amplitude 6 mT/m).

The dominant effect is the shifting of the effective labelling plane location, also along z , which is the point at which there is zero phase accrual between RF pulses and thus the steady state position of the magnetisation is in the transverse plane. For example, for $\vec{G}_z = 0.8$ mT/m and 200 Hz off-resonance, the point at which there is zero phase accrual between RF pulses is 6 mm away from the nominal labelling plane location. At this point, the effective flip angle due to the RF pulse is considerably lower and may not be sufficient to cause the adiabatic inversion. This effect can also cause the control condition to become less effective. Another way to think of this effect is that in the extreme case where the off-resonance is high enough to cause a phase accrual of π between RF pulses then the control condition effectively becomes the tag condition and vice versa (for an RF separation of 960 μs this occurs at an off-resonance frequency of 520 Hz). This dependence of the inversion efficiency on phase accrual between RF pulses also underlies the ability to spatially modulate the inversion efficiency with transverse gradient pulses, as shown in Section 4.7.

It is noted that the inversion efficiency is not symmetric about the off-resonance frequency axis. This is because the shift in the effective labelling plane location due to off-resonance leads to a different degree of T_1 decay by the time the magnetisation reaches the tissue (see Fig. 4.15). This is particularly important for low velocities, leading to a greater degree of asymmetry. A subtle asymmetry is also




Figure 4.14: The effect of off-resonance frequency on PCASL labelling for a variety of spin velocities (as given in the legend, in cm/s). Left: Inversion efficiency of tag (solid lines) and control (dashed lines) conditions. Right: Contrast.

apparent in the spatial modulation function for lower velocities (see Fig. 4.12) for the same reason, although this time it is the transverse gradient which is the source of the phase accrual rather than off-resonance frequency.

4.9 Vessel Angulation Through the Labelling Plane

The carotid and vertebral arteries run approximately in the inferior-superior direction for most of their length. However, in subjects with atypical vasculature or for labelling in locations other than the neck it may not be possible to ensure the vessels run perpendicularly to the labelling plane. It is therefore important to establish how efficiently the VEPCASL scheme works for vessels running at an angle to the labelling plane.

Simulations were performed using non-selective and vessel-selective cycles since these are likely to have different sensitivities to vessel angulation if the spin has a velocity component in the direction of the transverse gradient. Fig. 4.16 shows that the sequence is very robust to vessel angulation, with significant decreases in contrast only evident at vessel angulations above about 60° . At high




Figure 4.15: Off-resonance asymmetry in inversion efficiency is due to the difference in T_1 decay resulting from the change in the effective labelling plane location. Shown here are two cases with off-resonance frequencies +100 Hz and -100 Hz for blood travelling at 3 cm/s. There is considerable difference in M_z between the two cases beyond the labelling plane.

velocities, moderate vessel angulation actually improves contrast because the velocity component perpendicular to the labelling plane is reduced.

The flow phantom was used to verify that there was no significant change in ASL contrast when the labelling plane was angled at 30° ($C = 0.79 \pm 0.06$) relative to the standard axial labelling plane ($C = 0.82 \pm 0.05$) using a t-test ($p = 0.38$).

4.10 Effect on Static Magnetisation

Any effects on the static magnetisation in the imaging region which vary between cycles of the VEPCASL sequence may result in artefacts in the calculated vascular territory images. The RF pulses must be of short duration so they can be closely spaced, which means they act over a significant region of space. In the case of the RF pulse derived from the scanner, the effective flip angle is non-negligible over a range of about 2 cm from the labelling plane when the RF pulse duration is set to $600 \mu\text{s}$ (Fig. 4.17). This was assessed by simulating a single RF pulse and recording the angle of the static magnetisation away from the z axis at the end of the simulation for a range of z positions.




Figure 4.16: Simulated variation of inversion efficiency and contrast due to vessel angulation through the labelling plane in the same direction as the transverse gradient. Left: Inversion efficiency for a spin moving at 20 cm/s for non-selective and vessel-selective cycles. Right: Vessel-selective contrast for a range of spin velocities (see legend, in cm/s). These results show that the VEPCASL sequence is robust to vessel angulations up to about 60° .




Figure 4.17: Simulated variation in the effective flip angle with z for the RF pulse shape as played out on the scanner, for RF duration of $600 \mu\text{s}$ and tagging gradient amplitude 6 mT/m.

As a result of the range of the RF pulse, there may be some perturbation of the static magnetisation away from the labelling plane. Static magnetisation was simulated during the playing out of the PCASL sequence in tag and control conditions (Fig. 4.18). Significant perturbations occur in the tag condition in positions where the phase accrual between RF pulses, ϕ , due to \bar{G}_z equals $2n\pi$, where n is an integer, since the spin experiences RF pulses in the same direction. We can calculate this position, z_p , by considering the rate of phase accrual due to the z gradient, $\Delta\omega$, acting over the time between RF pulses, T_{sep} :

$$\phi = \int \Delta\omega dt = \gamma z \int G_z dt = \gamma z \bar{G}_z T_{sep} \quad (4.8)$$

$$\text{so when } \phi = 2n\pi \quad (4.9)$$

$$\text{this yields } z_p = \frac{2n\pi}{\gamma \bar{G}_z T_{sep}} \quad (4.10)$$

For standard parameters this gives $z_p = 30.6$ mm which matches well with the simulation results. For the control condition, maximal perturbation occurs when the spin accrues $(2n + 1)\pi$ phase between RF pulses, since the RF phase is alternating between 0 and π , i.e. when:

$$z_p = \frac{(2n + 1)\pi}{\gamma \bar{G}_z T_{sep}} \quad (4.11)$$

When $n = 0$ this gives $z_p = 15.3$ mm for standard parameters, again matching well with the simulation results.

The magnitude of the perturbation drops off with distance from the labelling plane, as we would expect given the drop in the effective flip angle away from the labelling plane. These perturbations can be seen in static magnetisation in the phantom experiments and follow closely the predicted trends (see Fig. 4.19). Note that only the first two perturbations away from the labelling plane have




Figure 4.18: The effect of the PCASL sequence on static z magnetisation under tag all (left) and control all (right) conditions for $\bar{G}_z = 0.8$ mT/m.

been shown here. The next largest perturbation at $z = \pm 3\pi/\gamma\bar{G}_z T_{sep}$ is negligible since the difference between the static tissue magnetisation under tag and control conditions is less than 0.1% of the static tissue signal for the parameters used here. Once partial volume effects have been accounted for this effect will be even smaller, particularly if the static tissue signal is suppressed by background suppression techniques (see Section 5.3.1).

These results show that it is important to avoid imaging in the regions where these perturbations occur since they differ in the tag and control conditions which would lead to artefacts upon subtraction. Ideally a gap of more than 3 cm should be left between the labelling plane and imaging region for the default PCASL parameters.

4.11 Development of Vessel-Encoded ASL Analysis Methods

In the previous sections methods for optimisation of VEPCASL acquisitions have been discussed in order to ensure that the maximum ASL contrast is obtained.




Figure 4.19: Effect on the static magnetisation in a phantom due to the tag all (left) and control all (right) PCASL cycles along with the expected variation from simulations, normalised to the same intensity and averaged over the same size voxels as the data to allow for partial volume effects.

However, once vessel-encoded data have been acquired, a robust method for separating out each vascular component is essential. In this section the theory and methods for this separation are described, from a naive analysis assuming perfect inversions, methods for estimating the encoding matrix from the data and finally a Bayesian model developed in collaboration with Dr. Michael Chappell.

4.11.1 Theory

Using the mathematical formulation of Wong [48], the separation of vascular territories can be thought of as a matrix inversion problem. The signal measured in a particular voxel due to different cycles of the VEPCASL sequence can be described as:

$$\mathbf{y} = \mathbf{A}\mathbf{x} \quad (4.12)$$

where \mathbf{y} is a vector of signal intensities measured for each VEPCASL cycle, \mathbf{A} is the encoding matrix, and \mathbf{x} is a vector of contributions to the signal from

vascular and static tissue components. Note that the notation \mathbf{x} was previously used to represent spatial position, but it is repeated here to maintain consistency with Wong [48].

As an example, if R, L, B and S represent the signals in a particular voxel due to the RICA, LICA, BA and static tissue respectively, the encoding scheme of Günther [97] can be written as:

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} = \begin{pmatrix} -1 & 1 & -1 & 1 \\ 1 & -1 & -1 & 1 \\ -1 & -1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} R \\ L \\ B \\ S \end{pmatrix} \quad (4.13)$$

Note that the right hand column of the encoding matrix always has an entry of one, since static tissue should not be affected by the ASL preparation. Also note that the entries in the encoding matrix represent the modulation of the measured signal by the vessel-encoding preparation so are related to M_z (without T_1 decay) rather than the inversion efficiency directly.

The signal due to each artery and the static tissue can be recovered by inverting the matrix \mathbf{A} :

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{y} \quad (4.14)$$

or for non-square encoding matrices:

$$\mathbf{x} = \mathbf{A}^+\mathbf{y} \quad (4.15)$$

where \mathbf{A}^+ is the pseudo-inverse of \mathbf{A} . For the example given above, this results in the following solution:

$$\begin{pmatrix} R \\ L \\ B \\ S \end{pmatrix} = \frac{1}{4} \begin{pmatrix} -1 & 1 & -1 & 1 \\ 1 & -1 & -1 & 1 \\ -1 & -1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} \quad (4.16)$$

This type of encoding (Eq. 4.13) is known as Hadamard encoding and is SNR optimal because each vessel is labelled and controlled an equal number of times. This can be quantified by calculating the ratio of SNR of each vessel for the encoding scheme to the SNR of a simple non-selective experiment with the same number of acquired images, which yields [48]:

$$E_i = \frac{1}{\sqrt{N \sum_j A_{i,j}^+{}^2}} \quad (4.17)$$

where E_i is an index of efficiency for vessel i and N is the number of encoding cycles. For the encoding scheme of Günther [97], E_i is equal to one for all vessels. This can be contrasted with simple encoding of single vessels, as used by Hendrikse *et al.* [89], which gives significantly reduced efficiency of 0.7 for each vascular component.

If the encoding matrix can be estimated accurately, then it is possible to separate out the vascular components in post-processing. Various methods for this process are discussed below.

4.11.2 Simple Analysis

In simple situations where all the vessels line up exactly with the tag or control location, and we make the assumption that perfect inversion or control is achieved on each cycle then the encoding matrix can be populated in a simple fashion. To illustrate this process, data acquired in the flow phantom is used in which four

VEPCASL cycles were performed:

1. tag all;
2. control all;
3. tag right tube whilst controlling the left;
4. tag left tube whilst controlling the right.

In this case there are three sources of signal which constitute the vector \mathbf{x} : water which has been labelled in the right tube, R , water which has been labelled in the left tube, L , and water which is unlabelled or static magnetisation, S . Assuming perfect inversion we can populate the encoding matrix in the following manner:

$$\begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} = \begin{pmatrix} -1 & -1 & 1 \\ 1 & 1 & 1 \\ -1 & 1 & 1 \\ 1 & -1 & 1 \end{pmatrix} \begin{pmatrix} R \\ L \\ S \end{pmatrix} \quad (4.18)$$

As before, each row of the encoding matrix describes one of the four VEPCASL cycles (e.g. the top row is the tag all case: both the right and left tubes are inverted, giving an entry of -1, but static tissue is unaffected so has an entry of +1).

In order to determine the contributions of each vascular component (R , L and S in this case) to each voxel, Eq. 4.18 can be simply inverted to give:

$$\begin{pmatrix} R \\ L \\ S \end{pmatrix} = \begin{pmatrix} -1 & -1 & 1 \\ 1 & 1 & 1 \\ -1 & 1 & 1 \\ 1 & -1 & 1 \end{pmatrix}^{-1} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} = \frac{1}{4} \begin{pmatrix} -1 & 1 & -1 & 1 \\ -1 & 1 & 1 & -1 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{pmatrix} \quad (4.19)$$




Figure 4.20: Separation of vascular components using a simple matrix inversion method applied to flow phantom data to separate out (a) labelled water in the right tube, R , (b) labelled water in the left tube, L , and (c) static tissue and unlabelled water, S . These are shown as a combined colour map (d) where $R = \text{red}$, $L = \text{green}$ and $S = \text{blue}$. Note that flowing water which is unaffected by the labelling (e.g. upstream of the labelling plane) is calculated to be part of the S component since it does not change between cycles of the sequence.

where $^+$ represents the pseudo-inverse of a matrix. Note that this encoding matrix formalism does not attempt to model imperfect inversion or modification of the signal due to T_1 decay. Such factors can be accounted for at a later stage by simple multiplication if necessary. Applying this matrix to all voxels in the flow phantom image separately yields the separated components shown in Fig. 4.20.

The separation of signals in the two tubes is very clean since very targeted labelling could be performed in this simple case with only two vessels to label. In more complex situations such a simple approach will not be applicable. However, this approach can be generalised for vessels which do not sit exactly at tag or control locations. For a specified average velocity, the spatial modulation function described in Section 4.7 can be used to estimate all the entries in the encoding

matrix if the encoding prescription and vessel locations within the labelling plane are known. This approach can also be used to account for imperfect inversion efficiency in both selective and non-selective cycles.

For example, if there were three vessels, equally spaced at positions -1, 0 and 1 within the labelling plane, then assuming a velocity of 30 cm/s, the encoding matrix for the four cycle scheme described above can be calculated using the simulated inversion efficiency shown in Fig. 4.12:

$$\mathbf{A} = \begin{pmatrix} -0.81 & -0.81 & -0.81 & 1 \\ 0.96 & 0.96 & 0.96 & 1 \\ -0.81 & -0.53 & 0.96 & 1 \\ 0.96 & -0.58 & -0.81 & 1 \end{pmatrix} \quad (4.20)$$

Here the first and second rows of the encoding matrix describe non-selective tag and control conditions, but the imperfect inversion and control are accounted for. In the third cycle, the first vessel is tagged and the third vessel controlled, with the middle vessel partially inverted. In the fourth cycle the first vessel is controlled and the third vessel inverted, again with the middle vessel partially inverted.

This method is simple and very quick to compute. However, it is not capable of accounting for any effects which modify the entries of the encoding matrix. For example, if the subject moves between the vessel-localisation scan and the VEPCASL scan, then the entries in the encoding matrix will be incorrect. This can lead to poor separation of the vascular components and cross-contamination of their signals. Other effects, such as off-resonance at the labelling plane, are also not accounted for. Therefore, methods which can account for such changes are desirable.

4.11.3 Estimating the Encoding Matrix from the Data

Original Implementation

The method suggested in the original perfusion-based VEPCASL paper [48] for separating the vascular components was to estimate the entries in the encoding matrix using the data. By making the assumption that the majority of voxels are fed by a single vessel histograms of the signal in vessel-selective cycles relative to non-selective cycles reveal peaks which can be used to estimate the entries of the encoding matrix.

This method is illustrated here using the phantom data shown previously. The data are processed using the following pipeline:

1. The data are pre-processed by phase correcting each image to the control image (as described previously).
2. Pairs of encoding cycles are subtracted from each other (i.e. Cycle 2 - Cycle 1 to give a non-selective “perfusion” image and Cycle 4 - Cycle 3 to give a vessel-selective “perfusion” image: see Fig. 4.21 (a) and (b)). This removes static tissue to allow accurate calculation of the relative tagging efficiency.
3. A mask is created from the non-selective subtraction image to select the regions where labelled blood/water is present (Fig. 4.21 (c)). In this case the labelling plane is present in the image which perturbs the magnetisation of the bottle phantom near the labelling plane (as described in Section 4.10) so an additional constraint that the phase difference between tag and control images had to be close to π was also used to create the mask.
4. Within the mask, the quantity β is calculated by dividing the selective subtraction image by the non-selective subtraction image within the mask (Fig. 4.21 (d)).

5. A histogram of β within the mask is created and a Gaussian Mixture Model (GMM) implemented in Matlab[®] (Mathworks[™], Natick, Massachusetts, USA) is used to fit the two peaks (Fig. 4.21 (e)).
6. The centres of the fitted peaks to the histogram are used to estimate entries in the encoding matrix (see below).
7. Once the corrected encoding matrix has been constructed, it can be applied to the raw data to separate the vascular components.

This process is perhaps better understood by considering what β means mathematically. For simplicity here we assume that the non-selective cycles produce perfect tag and control conditions, but the vessel-selective cycles have unknown inversion efficiencies. The encoding matrix is therefore:

$$\mathbf{A} = \begin{pmatrix} -1 & -1 & 1 \\ 1 & 1 & 1 \\ a & b & 1 \\ c & d & 1 \end{pmatrix} \quad (4.21)$$

where a, b, c and d are unknown quantities, although from knowledge of the encoding set-up a and d are expected to be close to -1 and b and c to be close to +1. Generally the static component is of no interest so this problem can be simplified by subtracting pairs of cycles (in a manner analogous to the subtraction of images described above) to remove the static tissue component and leave:

$$\begin{pmatrix} y_{sub,ns} \\ y_{sub,s} \end{pmatrix} = \begin{pmatrix} 2 & 2 \\ c-a & d-b \end{pmatrix} \begin{pmatrix} R \\ L \end{pmatrix} = \begin{pmatrix} 2 & 2 \\ 2\beta_R & 2\beta_L \end{pmatrix} \begin{pmatrix} R \\ L \end{pmatrix} \quad (4.22)$$

where $y_{sub,ns}$ is the signal derived from subtracting the non-selective cycles




Figure 4.21: Example of estimating entries in the encoding matrix in the flow phantom, using the method described by Wong [48]. Pairs of encoding cycles are subtracted to give a non-selective (a) and selective (b) subtracted image. A mask (c) is calculated using the non-selective subtracted image to exclude regions which have no ASL contrast. Within this mask β is calculated (d) by dividing the selective by the non-selective subtracted image. A histogram of β (e) within the mask reveals two peaks corresponding to the right hand tube (R) and left hand tube (L) which are fitted using a Gaussian Mixture Model (GMM). The centres of these peaks are used to improve estimates of entries in the encoding matrix. Note that this histogram has been scaled to give the estimated probability density function (PDF) which has unit area.

and $y_{sub,s}$ is the signal derived from subtracting the selective cycles and $\beta_{R,L}$ are pre-emptively inserted here to reduce the number of variables.

Following the image processing pipeline, β is calculated by dividing the selective subtracted signal by the non-selective subtracted signal:

$$\beta = \frac{y_{sub,s}}{y_{sub,ns}} = \frac{2\beta_R R + 2\beta_L L}{2R + 2L} = \frac{\beta_R R + \beta_L L}{R + L} \quad (4.23)$$

For a voxel which is supplied entirely by the right tube of the phantom (i.e. $L = 0$) this gives:

$$\beta = \frac{\beta_R R}{R} = \beta_R \quad (4.24)$$

Similarly, for a voxel which is supplied entirely by the left tube of the phantom (i.e. $R = 0$) this gives:

$$\beta = \frac{\beta_L L}{L} = \beta_L \quad (4.25)$$

Since all of the voxels in this simple phantom are supplied either by one tube or the other then a histogram of β should contain two peaks, one at β_R and one at β_L . Note that in voxels which contain no labelled water β will be dominated by noise so a mask must be used to exclude such voxels. The search for these peaks in the histogram of β is initialised assuming perfect inversion or control of the two tubes. In this example, the GMM fit to the histogram yields $\beta_R = 0.90$ and $\beta_L = -0.92$. Thus the corrected encoding matrix with paired subtraction is:

$$\mathbf{A}_{sub,corr} = \begin{pmatrix} 2 & 2 \\ 1.80 & -1.84 \end{pmatrix} \quad (4.26)$$

The inverse of this matrix is applied to the subtracted image data to yield the two tube components, as shown in Fig. 4.22. Note that the static tissue




Figure 4.22: Components of the flow phantom calculated using the estimation of entries in the encoding matrix shown in Fig. 4.21. Components from the right tube (R) and left tube (L) are displayed separately and in a combined colour map with $R = \text{red}$ and $L = \text{green}$. Note that due to the disruption of the static magnetisation at and near the labelling plane ASL signal appears in these locations as well as in the labelled water.

component is not estimated in this case due to the subtraction of pairs of images. Since the estimated inversion efficiencies are close to ideal, there is not much difference between the components shown here and those calculated using the simple analysis method described in the previous section. However, if there had been any motion or off-resonance effects which modified the labelling efficiency in the vessel-selective images then such a method would be very useful.

Note that in Wong's paper [48] after estimation of the β values for each vessel, the original encoding matrix is reconstructed by "unsubtracting" the corrected subtracted encoding matrix. Since only a single β value for a given vessel and pair of cycles is obtained, the original matrix is restored by assuming that the entries in the encoding matrix are equidistant from zero. For example, it was found that $\beta_R = 0.90$, but two separate elements of the encoding matrix, a and c

need to be inserted (see Eq. 4.21). Remembering that $2\beta_R = c - a$ it is assumed that a and c are equidistant from zero so $a = -0.90$ and $c = +0.90$. However, if only the vascular components and not static tissue need to be reconstructed, then this step is not necessary.

This method is certainly an improvement over the simple analysis discussed in the previous section. However, it has a number of drawbacks. Firstly, if a larger number of vessels are encoded, it becomes significantly more difficult to pick out individual peaks in the 1D histogram of β values. Secondly, no use is made of spatial information since it is expected that vascular territories will be relatively contiguous. Thirdly, the maximum number of vessels which can be separated is limited by the rank of the encoding matrix and thus the number of independent encoding cycles.

Improvements to the Method

Some improvements to Wong's original method have been proposed. One idea is to extend the concept of β into multiple dimensions, also referred to as "tagging efficiency space" [101, 110]. For example, if four vessel-selective cycles are performed, two with left-right (LR) encoding and two with anterior-posterior (AP) encoding, then a 2D histogram of β can be constructed with β_{LR} on one axis and β_{AP} on the other axis. In this space, individual vessels can be identified by clustering methods (similar to the GMM used above). The cluster centres can be used to correct the encoding matrix or each voxel can be simply assigned wholly to the cluster it is closest to. This latter technique relaxes the constraint that the maximum number of vessels which can be separated is limited by the number of independent encoding cycles. However, it does mean that regions which are supplied by multiple arteries will not be well modelled. An alternative method assumes sinusoidal variation in the spatial modulation function [101]. The appli-

cation of sine and cosine modulations allow calculation of a “phase angle” which can be directly mapped back to the location of the supplying artery within the labelling plane. However, this approach also assumes each voxel is fed only by a single artery.

The approach taken here is to use a GMM to describe the clusters formed in tagging efficiency space. The cluster centres are then used to update the encoding matrix which can be applied to the data as before to calculate the vascular territories. Additionally, spatial information is incorporated into the clustering process to help better separate vessels which have been encoded in a similar manner, but are separated spatially (this was first suggested by Wong *et al.* [101] but not implemented at the time of writing). For example, if four vessel-selective cycles have been performed then after paired subtraction there are two dimensions to the tagging efficiency space. Each voxel is labelled with its β value for both of these encodings along with its (x, y, z) coordinate. The clustering therefore occurs in a five dimensional space. Additional steps were taken to remove outliers by iteratively performing the GMM fit and excluding voxels which lie outside 1.5 standard deviations from any cluster centre. This helps exclude small regions of mixed supply as can sometimes occur (e.g. in the anterior cerebral arteries). However, if significant mixing of blood occurs then separating the territories using this technique is generally not feasible (e.g. significant mixing of blood from the two vertebral arteries when they fuse to form the basilar artery before reaching the posterior territories of the brain).

A brief example of this technique is given here, although it has been superseded by the considerably more robust and flexible Bayesian model developed in collaboration with Dr. Michael Chappell and others, described in the next section.

A single slice of vessel-encoded perfusion data (described in Chapter 5) is used

as an example of the method (Fig. 4.23). Six VEPCASL cycles were used (two non-selective, two contrast left-right (LR) circulation and two contrast anterior-posterior (AP) circulation) with the aim of separating out the vascular territories of the right and left internal carotid arteries (RICA and LICA) and the basilar artery (BA). The GMM is applied to cluster the data in five dimensional space $(\beta_{LR}, \beta_{AP}, x, y, z)$. In this case the GMM correctly identifies the cluster centres both in tagging efficiency space and in the spatial domain and use of the corrected encoding matrix significantly reduces contamination of the ICA components into the BA territory. Applying the GMM in multi-dimensional space is much more robust than the original method of fitting to a 1D histogram since there is much more information available to separate out the different clusters. The method used to remove outliers also produces more robust estimates of the cluster centres.

However, this method still has a number of drawbacks. In the case shown here, ideally the vascular territories from the two vertebral arteries which fuse to form the basilar artery would be produced separately so the relative importance of the two vessels can be assessed. This clustering method only works if the majority of the image contains voxels supplied by a single artery and would therefore fail in this case where there is significant overlap of vascular territories. In addition, the rank of the encoding matrix still dictates how many vessels can be separated, making this method less applicable for greater numbers of vessels. In the next section a more sophisticated analysis method is described which overcomes these limitations.

4.11.4 Bayesian Analysis

In order to overcome some of the difficulties associated with the analysis methods described above, a general framework for the analysis of vessel-encoded ASL data was developed by Dr. Michael Chappell in collaboration with the current author

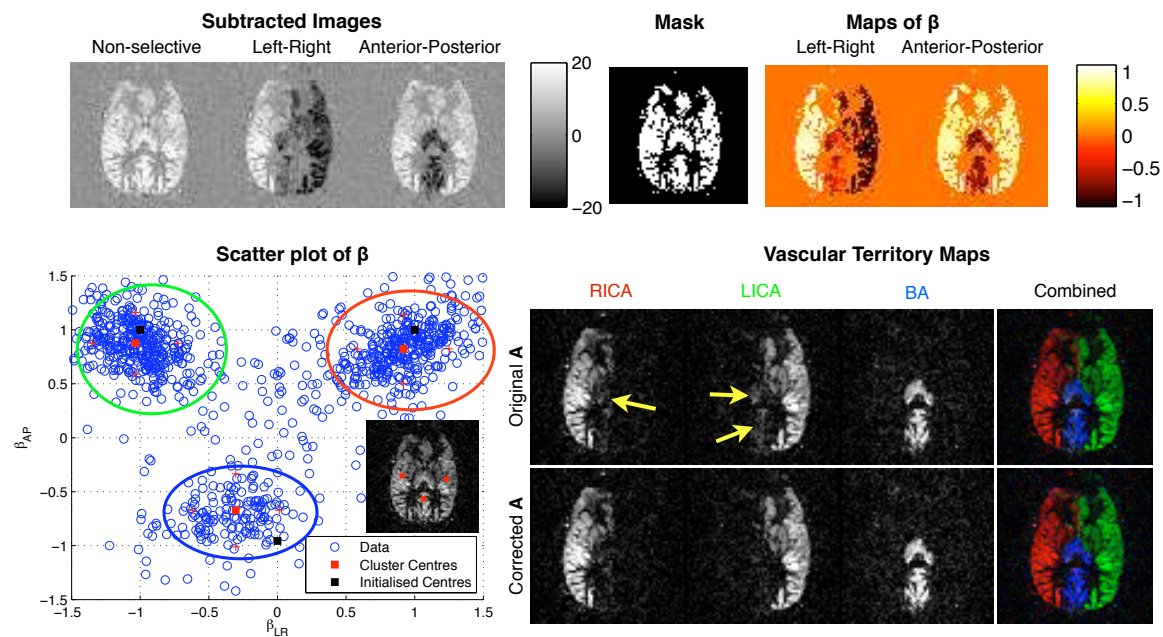


Figure 4.23: Improvements to the analysis method of Wong [48] to estimate encoding matrix entries from the data, applied to a single slice of vessel-encoded perfusion data with six VEPCASL cycles (two non-selective, two contrast left-right (LR) circulation and two contrast anterior-posterior (AP) circulation). Pairs of images are subtracted and a mask created from the non-selective subtracted image. Within the mask β is calculated for both LR and AP encodings. A 2D scatter plot of β_{AP} against β_{LR} is shown with clusters of voxels supplied by the same artery. The black squares represent the cluster centres which are estimated *a priori* using the spatial modulation function derived from Bloch simulations. These are used to initialise the GMM fit to the data, which yields cluster centres shown as red squares with standard deviations shown as red + symbols. Clustering is performed simultaneously in the spatial domain, yielding the cluster centres shown overlaid on the non-selective subtracted image in the insert. The clusters in the scatter plot represent the right internal carotid artery (red), left internal carotid artery (green) and basilar artery (blue) and are encircled in these colours for clarity. The vascular territory maps resulting from the original encoding matrix and the corrected encoding matrix are shown as separate components and a combined colour map. Note the contamination of RICA and LICA into the posterior territory (arrows) which is removed when the corrected encoding matrix is used. Images are displayed according to radiological convention (right of image is left of subject, top of image is anterior of subject).

and others. A brief overview of this method is given here along with a simple example of its use but the reader is referred to published work [111, 98] for a more detailed description.

This approach incorporates the spatial modulation function derived from the Bloch equation simulations described in Section 4.7, allowing the inversion efficiency for a given vessel location within the labelling plane to be estimated given the known encoding prescription in a similar manner to that described in Section 4.11.2. The advantages of both a clustering type approach, which can identify many more vascular territories than there are encoding cycles, and a matrix inversion method, which can properly model regions of mixed blood supply, are combined by introducing a classification scheme. Each class represents a subset of source vessels (typically two), since generally it is unlikely that more than two source vessels feed any given voxel. By considering these classes separately, the required rank of the encoding matrix is greatly reduced.

This model is encompassed within a Bayesian inference procedure, yielding probabilistic results which can account for multiple classes contributing to any given voxel and thus more feeding arteries than the set number of vessels per class. This inference procedure can also estimate the locations of the vessels within the labelling plane (in case there has been motion between the vessel-localisation scan and VEPCASL acquisition) and even the velocity of blood within the vessels. More recent developments allow restriction of vessel locations to rigid body transformations from their initial positions [112] (i.e. the relative positions of the vessels within the labelling plane are fixed, but global translations and rotations are allowed) which makes the inference more robust, particularly in cases where one or more of the arteries provides very little signal to the imaging region (e.g. due to severe stenosis or occlusion).

The inference can be performed either using Markov-Chain Monte Carlo meth-

ods (as in [98]) to attempt to sample the posterior probability distribution of the parameters, or the considerably faster maximum *a posteriori* (MAP) optimisation method (as in [112]) which aims to find the parameter set which maximises the posterior probability of the parameters given the data.

As an example of this method, the MAP implementation was applied to the single slice perfusion data set used above, this time attempting to identify the territories from both vertebral arteries (RVA and LVA) separately. Two vessels per class were used and an average flow velocity of 30 cm/s assumed. Fig. 4.24 shows that clean separation of the vascular components is achieved and the ability to represent the overlapping vertebral artery territories is a clear advantage of this method over clustering approaches and the method described above. From these images it is clear that both vertebral arteries contribute to the posterior circulation but the LVA is the greater contributor. Even the simple matrix inversion method without any assessment of encoding efficiency would fail here since the rank of the encoding matrix is four and there are four feeding arteries as well as static tissue, meaning that the pseudo-inverse of the encoding matrix is not well defined.

4.12 Conclusions

In this chapter, Bloch equation simulations and flow phantom experiments have been used to explore the inversion mechanism of VEPCASL, its sensitivity to various external influences and sequence parameters and its ability to encode vessels of interest. The results of simulations agree well with those in the original VEPCASL paper [48]. Small differences can be accounted for by the inclusion of T_1 decay in the simulations used here. The results from phantom experiments match well those of the simulations, validating the approach taken. This is of great use since simulations can be performed over a much greater range of parameters

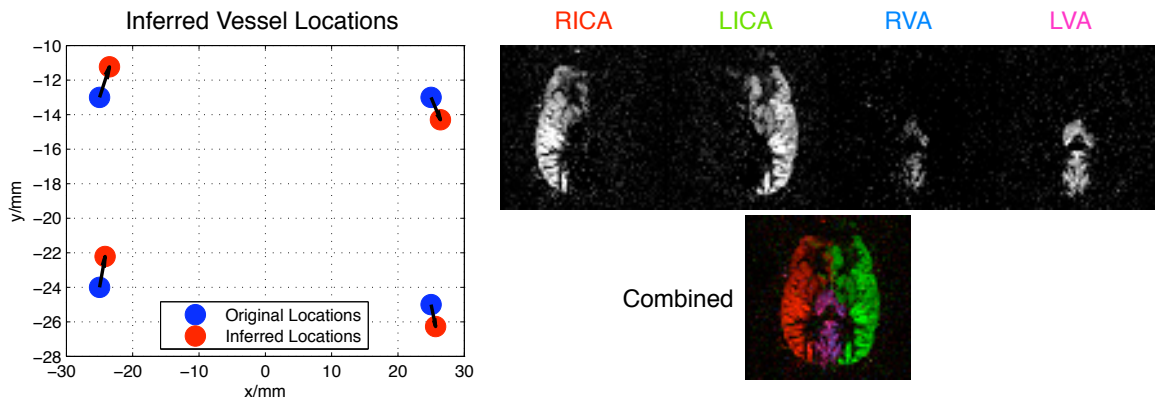


Figure 4.24: Example of the results of processing using the MAP implementation of the Bayesian inference model of Chappell *et al.* [98, 112]. In this case the vessel locations were kept fixed relative to each other but allowed to translate or rotate to fit the data. The original *a priori* locations and inferred locations are shown on the left. Each separate vascular component (both left and right internal carotids and vertebrals) are shown on the right as separate images and as a combined colour map (to represent four components here the colour scheme used is RICA = red, LICA = green, RVA = light blue, LVA = magenta). Note the excellent separation of territories with almost no contamination of components, and the ability to represent clearly overlapping territories in the posterior of the brain.

without using valuable scan time.

In particular, these simulations and experiments showed that PCASL can produce high levels of ASL contrast, at over 90% of the maximum achievable value. There is some sensitivity of the inversion efficiency to velocity, off-resonance and vessel angulation through the labelling plane, although these are relatively minor for the expected ranges to be found in most subjects. The sequence parameters used by Wong [48] are close to the optimum found in simulations so are used in future experiments to allow direct comparison. The simulated spatial modulation function produced in vessel-encoded acquisitions has a strong velocity dependence and this function can be used to predict entries in the encoding matrix to help calculate the vascular territories from vessel-encoded data. It was also shown that differences in the static magnetisation between tag and control conditions necessitate placing the labelling plane 3 cm or more from the imaging region (for the standard sequence parameters used here) to prevent erroneous ASL signal appearing the images.

Finally, a number of methods to separate out vascular components from vessel-encoded images were investigated, including a naive matrix inversion approach, methods for estimating entries in the encoding matrix from the data and finally a sophisticated Bayesian inference method developed in collaboration with Dr. Michael Chappell and others [98]. This latter approach was shown to be superior to the other methods since it can decode more vascular components than there are independent encoding cycles and can account for motion between vessel-localisation and VEPCASL acquisitions and variations in blood velocity. The resulting images are highly vessel specific and the maximum *a posteriori* (MAP) implementation of this method can be run in a matter of minutes, motivating its use in following chapters, where applications of the VEPCASL method to vessel-specific tissue perfusion imaging and dynamic angiography are explored.

Chapter 5

Vessel-Encoded Tissue Perfusion Imaging

5.1 Introduction

Vascular territory mapping is of great use in the study of collateral blood flow since it allows assessment of not only how much blood is reaching the tissue throughout the brain, but also the arterial source of this blood. Stenosis or occlusion of upstream vessels due to cerebrovascular disease will be apparent through reduced signal strength from the diseased artery and collateral flow can be inferred by the presence of signal from alternative sources. For example, collateral flow across the anterior communicating artery can be inferred via the presence of contralateral internal carotid artery signal in the anterior or middle cerebral artery territories usually fed by the ipsilateral internal carotid artery. The relative importance of each feeding artery can be assessed, which may be useful for diagnosis, prognosis and surgical planning.

In this chapter the design and optimisation of pulse sequences which create vascular territory images weighted by the rate of perfusion at the tissue level are

discussed. The vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL) module implemented for phantom experiments in the previous chapter is attached to a standard Echo Planar Imaging (EPI) [113] readout to allow the acquisition of whole brain vascular territory maps of the two internal carotid and two vertebral arteries. This sequence is applied in healthy volunteers to assess the optimum labelling plane location and post labelling delay.

Optimisation of this readout is then described, including the addition of background suppression pulses, partial Fourier and parallel imaging techniques. Preliminary data using a three dimensional gradient and spin echo (3D-GRASE) readout [114] are also presented and compared with the EPI method.

5.2 Implementation with a 2D Multi-Slice EPI Readout

5.2.1 Sequence Design

The VEPCASL preparation module described in the previous chapter was combined with a 2D multi-slice Echo Planar Imaging (EPI) readout [113]. This fast gradient echo sequence allows the acquisition of a single 2D slice in approximately 50 ms by rapidly rastering through k -space after a single 90° excitation pulse (see sequence diagram in Fig. 5.1). This high speed method allows a large number of slices to be acquired after each VEPCASL preparation, covering the entire brain with modest spatial resolution in approximately 1000 ms.

The VEPCASL pulse train is played out for a specified period (the “tag duration”), followed by a spoiler to dephase any residual transverse magnetisation and then a delay to allow the labelled blood water to traverse the cerebral vasculature and exchange into brain tissue. The time between the end of the VEPCASL

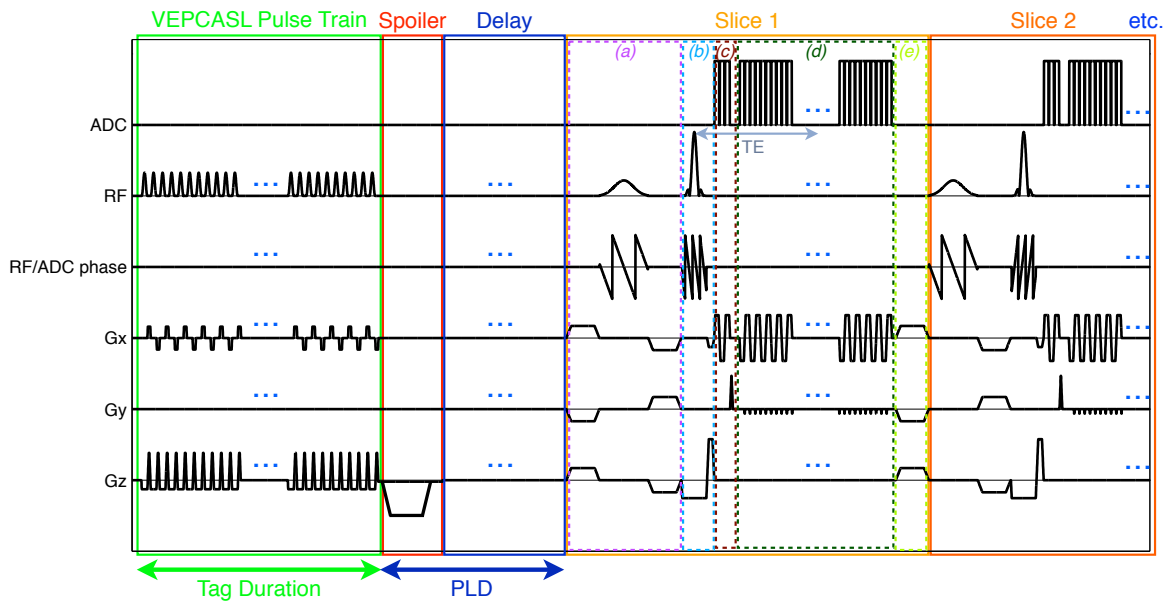


Figure 5.1: Sequence diagram for the VEPCASL module followed by an EPI readout. The VEPCASL pulse train is played out for a given time (the “tag duration”) with transverse gradients performing vessel encoding (shown along G_x here) only on the required encoding cycles. The RF phase is cycled to maintain coherence with the desired vessel (here at isocentre so no phase cycling is required) or alternated between 0 and π to achieve a control condition. Following this is a spoiler to dephase any transverse magnetisation and then a delay to allow the labelled blood to reach the brain tissue. The time from the end of the VEPCASL pulse train to the start of the readout is termed the “post label delay” (PLD). The 2D multi-slice EPI readout follows with slices being acquired in ascending order (inferior to superior, only two slices shown here). Within each slice, there are a number of components: (a) fat saturation; (b) excitation; (c) phase correction; (d) EPI readout; and (e) spoiler. The echo time (TE) is defined as the time between the centre of the RF excitation pulse and the time at which the centre of k -space is acquired. In order to display the sequence more clearly sections with repetitive patterns have been replaced by three blue dots. Note that the RF waveform represents only the magnitude of the RF pulses, with off-resonance frequency being represented as a linear ramp in the RF/ADC phase waveform.

pulse train and the acquisition of the first slice is referred to as the “post label delay” (PLD).

The acquisition of each slice involves a number of sequence components:

- Fat saturation: fat around the scalp can produce a significant signal in EPI images and due to the low effective bandwidth in the phase encode direction, this signal is subject to a strong chemical shift, and hence image displacement, causing some of it to overlap with brain tissue. Ideally this signal will be removed by the ASL subtraction process, but even a small amount of subject motion will cause this highly localised artefact to move and produce errors on subtraction. Thus, prior to each slice being acquired an off-resonance 90° pulse is applied at the fat resonant frequency before strong gradients to dephase the fat signal, preventing it from contributing to the final image.
- Slice selective excitation: a 90° RF pulse and associated slice selection gradient to excite the magnetisation and give the maximum possible signal strength.
- Phase correction lines: timing errors or eddy currents can lead to odd and even lines of k -space (which are acquired using gradients of opposite polarity, see below) being offset relative to each other, causing a ghost displaced from the main image by half the field of view in the phase encode direction. To correct for this, the central line of k -space is acquired three times (twice in the forward direction and once in the backward direction) just after the excitation pulse. This allows estimation of the zeroth and first order phase corrections which must be applied to alternate lines of k -space to remove this ghosting artefact [115, 29]. These corrections are performed automatically on the scanner prior to image reconstruction.

- EPI readout: initial gradients take the sampling point to the corner of k -space, after which each line is acquired consecutively by first acquiring a full line in the readout direction, followed by a small phase encode gradient (“blip”) to move up to the next line and then readout of the next line in the reverse direction. This pattern is continued until all the lines of k -space have been acquired.
- Spoiler: any residual transverse magnetisation is dephased using a strong gradient prior to the acquisition of the next slice.

Since blood tends to arrive earlier in inferior regions [114, 116], slices are acquired in ascending order such that the most superior slice is acquired last following each VEPCASL preparation. One volume of data is acquired for each cycle of the encoding scheme in turn before repeating the process to obtain further volumes for averaging, a requirement for any ASL technique with its inherently low SNR. Acquiring with this looping strategy ensures that any bias, such as scanner drift over the course of the acquisition, is shared between all the cycles and thus comparison of each of the averaged images is not greatly affected.

Typical parameter choices for these tissue perfusion experiments are given in Table 5.1. The main parameters of interest are the VEPCASL pulse train duration (“tag duration”) and post labelling delay (PLD). Clearly longer tag durations will lead to more labelled blood reaching the tissue and a greater perfusion signal. However, after each inverted spin exits the labelling plane, it undergoes T_1 decay, so the perfusion signal reaches an asymptote at large labelling durations [33]. Therefore, increasing the pulse train duration beyond about 1400 ms will not give a significant increase in perfusion signal, but will increase the required TR and thus total acquisition time. Many other studies have used a value close to 1400 ms and obtained high quality perfusion images [46, 48, 117], so this value is used in all experiments discussed in this chapter. The choice of PLD is fairly critical

<i>Group</i>	<i>Parameter</i>	<i>Value</i>
VEPCASL Pulse Train	VEPCASL pulse train duration ("tag duration")	1400 ms
	RF Pulse Type	Gaussian
	RF Flip Angle	20°
	RF Duration	600 μ s
	RF Separation	960 μ s
	Mean Tagging Gradient	0.8 mT/m
	Tagging Gradient Amplitude*	6 mT/m
	Post labelling delay (PLD)	1000 ms
Readout	Imaging Plane	Transverse/oblique
	Phase Encode Direction	Anterior-Posterior
	Field of View	220 x 220 mm
	Matrix size	64 x 64
	Number of Slices	Approx 18
	Slice Thickness	5 mm
	Slice Gap	2.5 mm
	Voxel size	Approx 3.4 x 3.4 x 5 mm
	Slice Ordering	Ascending (Inferior to Superior)
	TR	3500 ms
	TE	23 ms
	Bandwidth	2004 Hz per pixel
	Excitation Flip Angle	90°
Fat saturation	On	
Number of Averages	10 per VEPCASL cycle	
Other	Transmit Coil	Body
	Receive Coil	12 Channel Head
	Total Imaging Time	4 mins 40 s

Table 5.1: Default scanning parameters used for tissue perfusion experiments in healthy volunteers. *Refers to gradient amplitude during RF pulses.

in achieving accurate perfusion maps and is considered further in Section 5.2.3 below.

5.2.2 Methods

During sequence development and optimisation five healthy volunteers (four men and one women, mean age 28.0, age range 23-33) with no known neurological deficit were recruited and scanned under a technical development protocol agreed with local ethics and institutional committees. Three of these five volunteers were scanned twice on separate occasions. All data were acquired using a 3 Tesla TIM Trio Siemens MRI scanner (Siemens Healthcare, Erlangen, Germany). The acquisition schedule began with a localiser and multi-slab three-dimensional (3D) time of flight (TOF) angiography acquisition for labelling plane selection and vessel-localisation (see Section 3.1.5 for a description of the TOF technique). One or more VEPCASL perfusion acquisitions were then performed, as described in

the following sections.

If significant motion corruption was evident in the raw images then they were motion corrected using MCFLIRT [118] before being analysed using Matlab[®] (Mathworks[™], Natick, Massachusetts, USA). Images from each VEPCASL cycle were averaged before separation of vascular components using the maximum *a posteriori* (MAP) implementation of the Bayesian analysis method described in Section 4.11.4 [111, 98, 112].

Magnitude-only data were used for these analyses since, unlike the phantom experiments in the previous chapter, the labelled blood signal is always much smaller than the tissue signal by the time exchange has occurred. In addition, after the labelled blood water exchanges into tissue the phase of tissue and labelled water signals will be the same, so inverted magnetisation will simply subtract from the tissue signal, negating the need for complex signals to be considered. In addition, magnitude data is somewhat less susceptible to phenomena such as B_0 drift than full complex data, leading to reduced potential for artefacts in the ASL perfusion images.

5.2.3 Choice of Post Labelling Delay

In order for the ASL signal in each voxel to be approximately proportional to the regional cerebral blood flow (CBF), which describes the volume of blood delivered to the tissue per unit time per unit mass of tissue, the post labelling delay (PLD) must be greater than the blood transit time from the labelling plane to all voxels in the imaging region [119, 33]. However, using a longer PLD results in greater T_1 decay of the ASL contrast, reducing the SNR of the resulting perfusion images. Thus the optimum PLD is that which just exceeds the blood transit time to all voxels.

To determine the range of transit times present in healthy volunteers non-

selective PCASL perfusion images were acquired at various PLDs (see Fig. 5.2). These images have SNR and resolution comparable to those in the literature and are free from significant artefacts, showing that the PCASL labelling approach is working well *in vivo*. The small regions outside the brain with significant perfusion signal are likely to be branches of the external carotid artery which feed the face and scalp, or motion artefact around the scalp. These experiments suggest that using a PLD of 1000 ms is sufficient to capture most of the perfusion in normal volunteers since the majority of voxels have their peak signal within this time. This is in agreement with previous studies [120, 116].

However, it should be noted that for patient studies, the optimum PLD is likely to be higher due to increased transit times in diseased vessels or where the blood takes extended collateral routes to reach the tissue. In such cases the use of a longer PLD would increase the likelihood that late arriving blood has reached the tissue prior to imaging at the expense of SNR. Use of a range of PLDs would also help ensure that late arriving blood is not excluded whilst allowing some images to be obtained at shorter PLD values where the SNR is better. This would also facilitate the quantification of cerebral blood flow [33].

5.2.4 Labelling Plane Selection for Optimal Vessel-Encoding

A coronal TOF MIP in a typical healthy volunteer is shown in Fig. 5.3 with three potential locations for the labelling plane. These were selected by virtue of the main feeding arteries to the brain running approximately in the inferior-superior direction at these locations as well as being some distance from sharp turns in the vessels which could compromise the inversion efficiency of VEPCASL. The four main brain feeding arteries are the right internal carotid artery (RICA), left internal carotid artery (LICA), right vertebral artery (RVA) and left vertebral artery (LVA), which can be clearly seen in these images, allowing accurate prescription

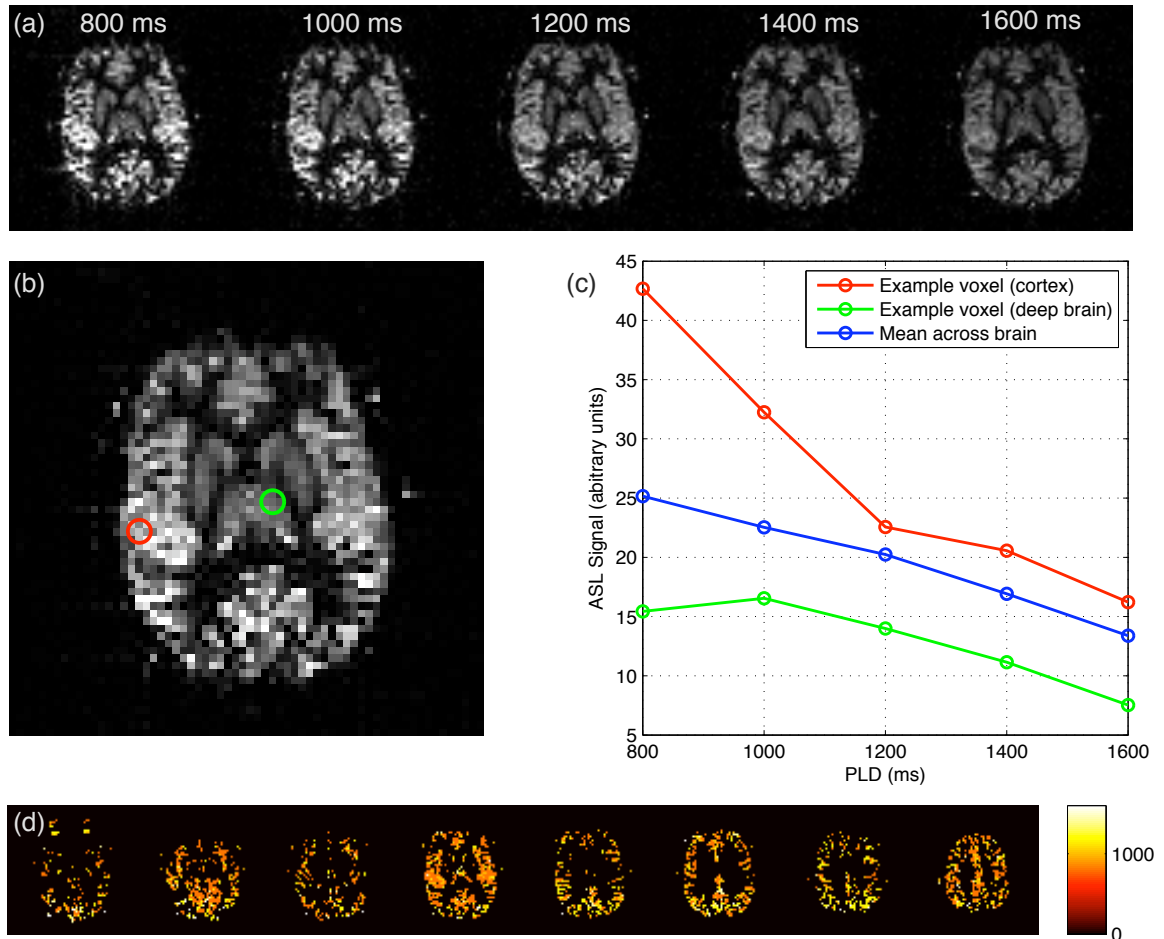


Figure 5.2: Variation in the ASL perfusion signal with post labelling delay: (a) non-selective ASL signal in one axial slice approximately in the centre of the brain for a variety of PLDs (as shown above each image); (b) mean perfusion signal in the same slice as (a) with two example voxels selected, one in cortex (red circle) and one in deep brain (green circle); (c) plots of the ASL signal against PLD for the two example voxels shown in (b) and the mean across the brain within a mask; (d) Maps of the PLD at which the perfusion signal peaks in each voxel for a number of slices (leftmost image is inferior, rightmost image is superior) within the mask. The mask used here is restricted to those voxels whose mean perfusion signal across PLDs is greater than half the 99th percentile perfusion signal across the brain. Images are displayed according to radiological convention (right of image is left of subject, top of image is anterior of subject).

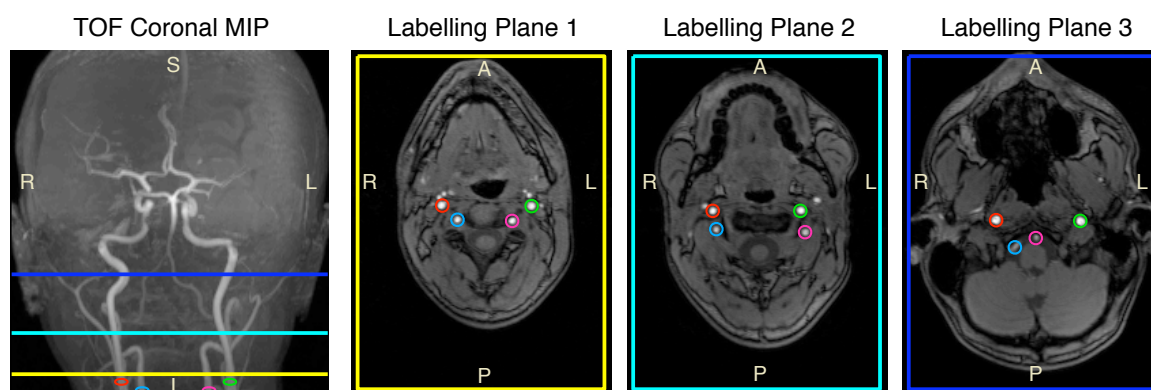


Figure 5.3: Selection of the labelling plane location. Three possible candidates where the main brain-feeding arteries run approximately in the inferior-superior direction are shown overlaid on a coronal maximum intensity projection (MIP) of a time-of-flight data set from a healthy volunteer. A transverse view of each plane is shown on the right in the corresponding colour. The four main arteries for which vascular territories are to be identified are highlighted (red = right internal carotid artery (RICA), green = left internal carotid artery (LICA), light blue = right vertebral artery (RVA) and magenta = left vertebral artery (LVA)). These colours are used within maps of vascular territories in the remainder of this thesis to allow multiple vascular components to be shown simultaneously. See text for a discussion of the relative merits of each candidate labelling plane. For clarity, anatomical labels have been added to these images (R = right, L = left, S = superior, I = inferior, A = anterior and P = posterior).

of the vessel-encoding cycles.

The first candidate labelling plane lies just above the bifurcation of the common carotid artery into the internal and external carotid arteries. The four vessels are well separated in both left-right (LR) and anterior-posterior (AP) directions. However, any simple encoding scheme will place some of the vessels in non-ideal locations where they will only be partially labelled or controlled, reducing the contrast which can be achieved. Better targeting of the individual vessels might be achieved by enforcing a VEPCASL spatial modulation function (see Section 4.7) with a short length scale. However, this is not desirable, firstly because such an acquisition would be much more sensitive to small amounts of subject motion and secondly because if the length scale of the modulation function approaches the diameter of the arteries, then only partial labelling will occur, again reducing the ASL contrast. In addition, this candidate labelling plane is the furthest from the brain, so longer PLDs would be required to allow the labelled bolus of blood

to reach the tissue, increasing the amount of T_1 decay and reducing the SNR.

In the second candidate labelling plane the arteries are arranged in a neat rectangular pattern. This allows simple encoding cycles to contrast the right and left circulation as well as anterior and posterior, with each vessel being in a tag or control condition on each cycle, resulting in high SNR. It is possible to locate such a plane in almost all subjects by locating a point midway between the two sharp twists in the vertebral arteries.

The third candidate plane is close to the point at which the two vertebral arteries fuse to form the basilar artery. Due to their close proximity, separating the contributions from the two would be difficult. In addition, this plane is quite close to the brain (part of the cerebellum can be seen in the TOF image here), so the disruption of the static magnetisation by the VEPCASL pulse train (see Section 4.10) will impinge considerably on the vascular territory maps, causing artefact in the lower brain slices.

Thus, a plane similar to the second plane shown in Fig. 5.3 is sought in all subjects. In such a plane it is possible to use an eight cycle encoding scheme which yields a full rank encoding matrix, better conditioning the separation of vascular components (see Section 4.11). The following cycles are performed, which are shown graphically in Fig. 5.4.

1. Tag all vessels;
2. Control all vessels;
3. Tag left circulation (LICA and LVA) whilst controlling right (RICA and RVA);
4. Tag right circulation whilst controlling left;
5. Tag anterior circulation (RICA and LICA) whilst controlling posterior (RVA and LVA);




Figure 5.4: Vessel-encoding in the optimal labelling plane location using eight cycles to attempt separation of vascular territories from the RICA (red), LICA (green), RVA (light blue) and LVA (magenta). The encodings are shown here overlaid on a TOF axial slice from a healthy volunteer. Cycles 1 and 2 are non-selective. Cycle 3: label at line A control at line B. Cycle 4: label B and control A. Cycle 5: label C and control D. Cycle 6: label D and control C. Cycle 7: label E and control F. Cycle 8: label F and control E. Note that there are two lines labelled F since the spatial modulation of inversion efficiency is periodic.

6. Tag posterior circulation whilst controlling anterior;
7. Tag diagonally, labelling the RICA and LVA whilst controlling the LICA and RVA;
8. Tag diagonally, labelling the LICA and RVA whilst controlling the RICA and LVA.

For the subject shown in Fig. 5.3 and assuming laminar flow with average speed 30 cm/s this yields the following encoding matrix:

$$\mathbf{y} = \mathbf{Ax} = \begin{pmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \end{pmatrix} = \begin{pmatrix} -0.81 & -0.81 & -0.81 & -0.81 & 1.00 \\ 0.96 & 0.96 & 0.96 & 0.96 & 1.00 \\ -0.81 & 0.96 & -0.80 & 0.75 & 1.00 \\ 0.96 & -0.81 & 0.83 & -0.80 & 1.00 \\ -0.80 & -0.81 & 0.91 & 0.88 & 1.00 \\ 0.88 & 0.91 & -0.81 & -0.80 & 1.00 \\ -0.81 & 0.73 & 0.73 & -0.81 & 1.00 \\ 0.96 & -0.80 & -0.80 & 0.96 & 1.00 \end{pmatrix} \begin{pmatrix} R_{\text{ICA}} \\ L_{\text{ICA}} \\ R_{\text{VA}} \\ L_{\text{VA}} \\ S \end{pmatrix} \quad (5.1)$$

where y_i is the signal measured in a given voxel in cycle i , $R(L)_{\text{ICA(VA)}}$ represents the contribution to the ASL signal in a given voxel from the right (left) ICA (VA) and S represents the signal of static tissue. The rank of this encoding matrix is five, allowing separation of all the vascular components and static tissue even with a simple matrix inversion approach. Its condition number, which represents its sensitivity to errors in the data, is 1.3, close to the ideal value of 1. The SNR efficiency, E_i , of the i^{th} vessel in this scheme can be assessed using the formulation of Wong [48], in which the SNR of the vessel-encoded scheme is calculated for each vessel relative to the SNR which would be achieved through simple averaging of the same number of non-selective ASL images:

$$E_i = \frac{\text{SNR}_{\text{encoded}}}{\text{SNR}_{\text{averaging}}} = \frac{1}{\sqrt{N \sum_j A_{i,j}^{+2}}} \quad (5.2)$$

where N is the number of encoding cycles and $A_{i,j}^{+2}$ is the component of the pseudo-inverse of \mathbf{A} in the i^{th} row and j^{th} column. Note that this formulation assumes perfect non-selective inversion and control. Applying this to the encoding matrix above (after scaling each entry to assume perfect inversion and control of

non-selective cycles) yields:

$$\mathbf{E} = \begin{pmatrix} E_{RICA} \\ E_{LICA} \\ E_{RVA} \\ E_{LVA} \\ E_S \end{pmatrix} = \begin{pmatrix} 0.99 \\ 0.96 \\ 0.94 \\ 0.96 \\ 1.00 \end{pmatrix} \quad (5.3)$$

This encoding scheme is therefore near-optimal, generating vascular territory maps with a SNR efficiency of approximately 95% compared to non-selective PCASL for the same number of acquired images. This might be expected since each vessel is close to the optimum tag or control location on all cycles, and is labelled or controlled an equal number of times. Another way of viewing this is that the encoding matrix consists of columns which are approximately equal to columns from a Hadamard matrix, as per the encoding strategies of Wong [48] and Günther [97].

It is worth noting that the Bayesian analysis approach used in this chapter to separate vascular components (see Section 4.11.4 and [111, 98]) may boost the SNR further by considering only sub-sections of the encoding matrix.

5.2.5 Application in Healthy Volunteers

Example vascular territory maps generated using this technique with optimal labelling plane location and PLD with eight encoding cycles are shown in Fig. 5.5. There is good separation of the vascular territories into the expected regions: the ICAs supply the anterior cerebral artery (ACA) and middle cerebral artery (MCA) territories and the VAs supply the cerebellum and posterior cerebral artery (PCA) territories. In addition, the SNR of this technique is high enough to draw clear conclusions from the data. Even in this healthy volunteer, non-standard flow




Figure 5.5: Vascular territory maps of a healthy volunteer calculated from vessel-encoded perfusion data using the Bayesian MAP method (see Section 4.11.4), which results in clean separation of the components for all four main brain feeding arteries. Each component is shown separately in selected transverse slices, and as a combined colour map (left). A 3D view of the same data set is also shown (right), with coronal (COR), sagittal (SAG) and transverse (TRA) views. The position of each slice is shown as a white line overlaid on the other two views. Features of interest include the ability to represent territories with mixed supply: in this case the right anterior cerebral artery territory is fed by both ICAs (yellow arrows), yielding a yellow/orange hue in the combined colour maps. The posterior cerebral artery territories are fed predominantly by the LVA (white arrows). However, the cerebellum is supplied mostly by the ipsilateral VA (orange arrows). Disadvantages of the EPI readout scheme include distortion and signal dropout near air/tissue interfaces (pink arrows).

patterns can be clearly seen: the RACA has contributions from both the RICA and LICA, from which it can be inferred that there is collateral flow from the LICA across the anterior communicating artery. In addition, the LVA appears to provide a much greater amount of blood to the brain than the RVA, particularly in the PCA territories. In the cerebellum the RVA makes more of a contribution on the ipsilateral side, as might be expected since the cerebellum is supplied by arteries which branch off the VAs before or just after they fuse to form the basilar artery (BA).

There is some residual artefact from eye motion present in these images, but

it does not appear to affect the signal present in the brain. The small regions of signal present just outside the brain are likely to be branches of the external carotid artery (ECA) or motion artefact arising near the sharp change in signal at the scalp. Since the ECAs have not been explicitly modelled in the analysis, their signal is assigned to any vessels that are encoded in a similar manner. Thus, interpretation of such images requires a little care. If significant collateral flow was present from the ECAs into the brain, then clearly these vessels must be modelled in the analysis and a scheme which efficiently encodes all the feeding arteries must be designed.

Once early technical difficulties were overcome and the optimum PLD and labelling plane location were established, reasonable vascular territory maps were generated using this technique. However, this method could be improved in a number of ways. Firstly, in patient populations where subject motion is generally more severe, methods for reducing sensitivity to motion are desirable. Secondly, increasing the SNR to improve image quality or reduce the required number of averages would be advantageous. Finally, reduction in distortion and dropout artefacts present in these EPI images would be beneficial. Optimisation of the 2D multi-slice EPI readout to help address these issues is discussed in the next section.

5.3 Optimisation of 2D Multi-Slice EPI

In this section a number of methods with the potential to improve the EPI readout mentioned above are discussed in turn. Example data using these methods are given after the discussions. The results shown here were generated as part of a study to compare EPI with a 3D Gradient and Spin Echo (3D-GRASE) technique (see Section 5.4), which formed the basis of a conference abstract [121].

5.3.1 Background Suppression

Background suppression pulses have been used by a number of groups to attenuate the signal in static tissue in ASL acquisitions [122, 114, 40]. This reduces the sensitivity to motion or other instabilities (such as physiological fluctuations which scale with static tissue signal [123]) or to reduce the total signal strength in 3D imaging to prevent digitisation errors becoming significant in the ASL maps. Such schemes generally consist of pre-saturation of the imaging region followed by a number of non-selective adiabatic inversion pulses [124] played out at specific times prior to image acquisition. This suppresses static tissue signal in the imaging region but leaves the inflowing blood signal relatively unaffected. Using a higher number of inversion pulses allows better suppression of a wider range of T_1 species. However, if the inversion efficiency of these pulses is not 100% then larger numbers of pulses can significantly degrade the ASL contrast [125].

For these tissue experiments a background suppression scheme similar to that used by Günther *et al.* [114] was implemented, consisting of a water suppression enhanced through T_1 effects (WET) pre-saturation scheme, similar to that of Golay *et al.* [90], for efficient pre-saturation of the imaging region, followed by the VEP-CASL module and then two global adiabatic hyperbolic secant inversion pulses [126] played out during the post labelling delay. A schematic of this sequence structure and the effect this has on static tissue and blood is shown in Fig. 5.6. As per Günther *et al.* [114] the inversion pulses are played out at times which allow perfect nulling of two species with T_1 equal to $T_{1,\text{opt}}$ and $2T_{1,\text{opt}}$, where $T_{1,\text{opt}}$ is a parameter which can be varied.

Due to the considerable duration of the VEP-CASL pulse train the placement of these inversion pulses is somewhat restricted, meaning that for a tag duration of 1400 ms and a post labelling delay of 1000 ms the maximum $T_{1,\text{opt}}$ achievable is 400 ms. However, species with longer T_1 are still subject to reasonable suppression

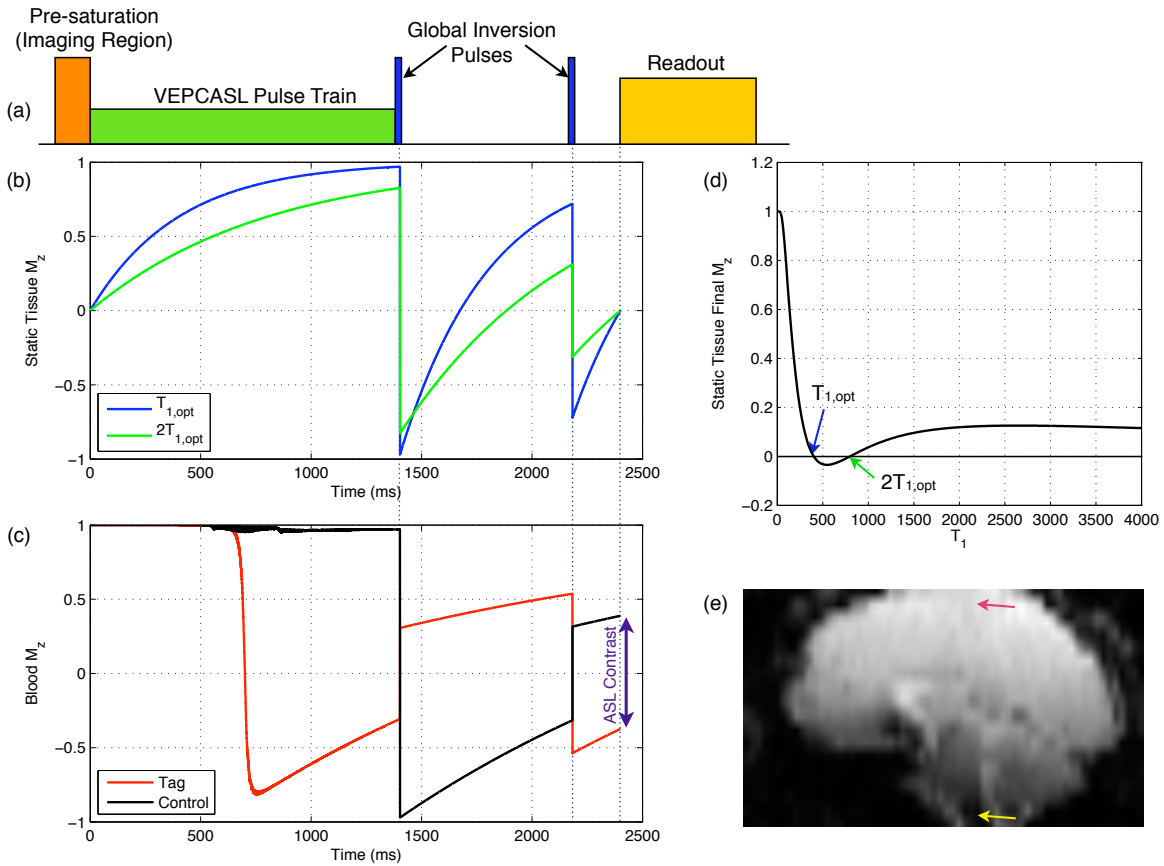


Figure 5.6: The addition of background suppression to the VEPCASL tissue perfusion sequence: (a) Schematic of the pulse sequence diagram showing the addition of a pre-saturation module (applied to the imaging region only) and two global adiabatic inversion pulses played out during the post labelling delay; (b) The effect of the pre-saturation and two global inversion pulses on static tissue in the imaging region (two species are shown with T_1 equal to $T_{1,opt}$ and $2T_{1,opt}$), calculated using the Bloch equation simulations described in Section 4.2; (c) The effect of the VEPCASL pulse train and global inversion pulses on blood travelling through the labelling plane in both “tag all” and “control all” conditions, showing the ASL contrast achieved which is unaffected by this scheme for perfect inversion pulses; (d) Graph showing the M_z of static tissue at the end of the preparation period (just prior to readout) as a function of T_1 (note the perfect suppression at $T_{1,opt}$ and $2T_{1,opt}$ as expected); (e) Raw EPI data in a sagittal view showing that good static tissue suppression is achieved in the first acquired, most inferior slice (yellow arrow) but significantly poorer in superior slices which are acquired later (red arrow). The magnetisation magnitudes shown here are scaled relative to M_0 .

(approximately 15% of M_0) as shown in Fig. 5.6 (d).

It should also be noted that this scheme is only aimed at suppression of static tissue at a single time point, which in this case is set to the start of the readout period, coinciding with the excitation pulse of the first acquired slice. Superior slices are acquired at later time points, giving the static tissue magnetisation time to recover via T_1 processes, leading to poorer suppression, as shown in Fig. 5.6(e). However, this level of background suppression is still an improvement on the previous implementation of the sequence.

5.3.2 Partial Fourier

As with all gradient echo sequences the signal strength dies away as the echo time (TE) is increased with time constant T_2^* , which is of the order of tens of milliseconds in brain tissue. Therefore any method which can reduce TE will significantly boost the SNR. One such method which is compatible with EPI readouts is partial Fourier, in which k -space is sampled asymmetrically, typically in the phase encode direction, to reduce the TE and total time for the readout. The reduced readout time gives an additional SNR boost for later slices by reducing the post labelling delay at which they are acquired. However, it should be noted that the reduction in the amount of k -space data acquired theoretically reduces the SNR somewhat, but this is likely to be a minor effect at the edges of k -space relative to the reduction in TE.

Fig. 5.7 compares a fully sampled trajectory with a 6/8 partial Fourier acquisition where only this fraction of k -space is acquired. If the missing lines are simply replaced with zeros (“zero-padding”) before the standard Fourier transform reconstruction is performed to form the images, this can be described as multiplying fully sampled k -space by a step-like modulation function, which is in turn equivalent to convolving the fully-sampled image by a kernel with finite

width and numerous side-lobes. This can lead to some blurring and ringing in the reconstructed images. To mediate the ringing artefacts the default reconstruction used on the scanner multiplies the data near the start of the trajectory by a smooth function.

More complex methods for estimating the unacquired data are available (e.g. see [127, 29]), reducing the ringing and blurring artefacts at some cost to SNR. However, for the purposes of such ASL acquisitions where the SNR is already relatively low and resolution is not at a premium, such sophisticated methods are not generally necessary.

It is worth noting that using partial Fourier techniques to reduce TE is also likely to reduce dropout artefacts (for reasonably well shimmed scanners) which are caused by dephasing of spins within a voxel during TE. However, since the effective bandwidth in the phase encode direction is the same, distortion artefacts will not be improved.

5.3.3 Parallel Imaging

Another method to reduce TE and the time to acquire each slice to boost the SNR is parallel imaging. In these techniques lines of k -space are skipped at regular intervals. This has the effect of creating an aliased image in the same way as reducing the field of view to a size smaller than the object being imaged. However, when a multi-channel coil is used to receive the signal and each channel has a different spatial sensitivity profile, this extra information can be used to “undo” the aliasing and recover an un-aliased image. The parallel imaging method used for the data in this thesis, performed automatically on the scanner, is generalized autocalibrating partially parallel acquisitions (GRAPPA) [59]. A full description of parallel imaging is beyond the scope of this thesis, but more information is available in the literature [58, 29, 128].

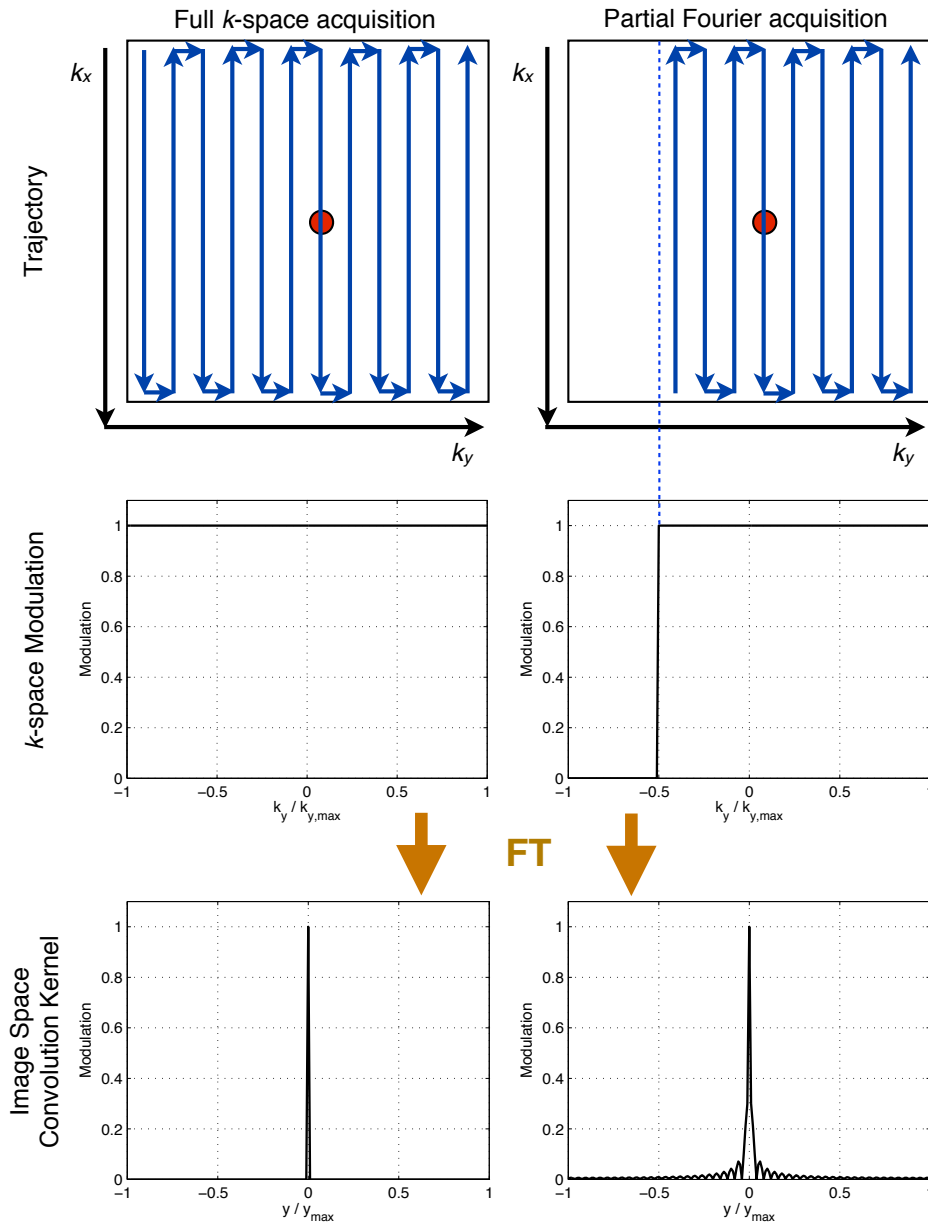


Figure 5.7: Partial Fourier with an EPI readout. The top row shows two trajectories: a fully sampled k -space EPI trajectory (left) and a partial Fourier acquisition (right) in which the first quarter of k -space is not acquired, dramatically reducing the TE (time to reach the centre of k -space, shown with a red circle). Only a small number of lines are shown here for clarity. In the middle row is the modulation function which can be used to describe the acquisitions: in the fully sampled trajectory the modulation function is one for all k_y points; the partial k -space trajectory is equivalent to multiplying the lines not acquired by zero, giving a step-like modulation function. The bottom row shows the equivalent process in image space: convolution with a kernel which is a delta function for the fully sampled case but which has finite width and lobes which can lead to blurring and ringing artefacts in the final image if a zero-padded reconstruction is used. The convolution kernels are calculated by Fourier Transform (FT) of the k -space modulation function. Note that only the magnitude of the convolution kernels are shown here.

The increase in speed of the acquisition by using parallel imaging is described by an acceleration factor, R . For example, using $R = 2$ means skipping every other line of k -space, increasing the acquisition speed by a factor of two. For EPI this will also reduce TE by a factor of two and decrease the total time for the acquisition of all slices, reducing the post labelling delay for the later slices, both of which will boost the SNR. However, since only half of the data are acquired, there is an associated loss of SNR which varies over space, depending on the coil sensitivity profiles. Thus, an experimental comparison is required to determine whether such a technique is beneficial.

One additional advantage of parallel imaging is that through skipping lines of k -space, the effective bandwidth in the phase encode direction is increased, thereby reducing distortion artefacts. For cases where regions of the brain heavily affected by distortion are of primary concern, such a technique may be valuable.

5.3.4 Comparison of Optimisation Techniques

In order to compare these potential improvements to the EPI readout with the original sequence, five healthy volunteers with no known neurological deficit were scanned in the same manner described in Section 5.2.2, except four VEP-CASL scans were performed with various modifications to the EPI readout:

1. "Standard" readout;
2. "Partial Fourier": 6/8 partial Fourier in the phase encode direction;
3. "Parallel imaging": GRAPPA with $R = 2$ in the phase encode direction;
4. "Partial Fourier with parallel imaging": 7/8 partial Fourier in the phase encode direction combined with GRAPPA using $R = 2$.




Figure 5.8: Comparison of the use of an EPI readout with and without background suppression (in a different subject). Note the significant reduction of artefacts arising from movement of the sharp boundary at the scalp (arrows) and similar apparent SNR.

Other than these modifications, the scans were performed using the parameters shown in Table 5.1, except the modification of TE to the minimum allowed by the readout: 23 ms for the standard readout, 14 ms with partial Fourier or parallel imaging alone, and 13 ms for partial Fourier with parallel imaging. All scans were performed with background suppression (with $T_{1,\text{opt}} = 400$ ms) since early experiments showed that artefacts due to motion, particularly around the scalp, were significantly reduced, at no apparent loss of SNR (see Fig. 5.8).

In order to quantitatively compare these four methods, the temporal SNR (tSNR) was calculated. The tSNR allows the assessment of both signal and noise within the brain (rather than a background region, as commonly used to calculate spatial SNR), which is important when using parallel imaging since the noise amplification can vary over the field of view [128]. The tSNR was calculated using ASL difference images obtained from pairs of tag all and control all cycles. This helps to account for the fact that the differences between images are what contribute to the final vascular territory maps. If the raw images were used, the temporal standard deviation would be increased by slow drifts in the signal over time which are effectively subtracted out when the ASL perfusion images are calculated. The tSNR is thus calculated in each voxel as follows:

$$\text{tSNR} = \frac{\langle \Delta S \rangle_t}{\sigma_{\Delta S,t}} \quad (5.4)$$

where $\langle \Delta S \rangle_t$ is the mean ASL difference signal over time and $\sigma_{\Delta S,t}$ is the standard deviation in this difference signal over time. A brain mask was calculated using the Brain Extraction Tool (BET, [129]) and the mean tSNR within this mask calculated for each subject.

A comparison of the four imaging strategies is given in Fig. 5.9. The vascular territories are consistent between the methods, showing that the VEPCASL encoding and Bayesian MAP method for separating the vascular components are performing well. The use of parallel imaging introduces localised noise in the images which makes their interpretation a little more difficult. However, as expected when accelerating by a factor of two, the distortion artefacts appear reduced. All the methods for reducing the TE (both partial Fourier and parallel imaging) reduce signal dropout artefacts near the frontal sinuses as expected. The blurring caused by the partial Fourier method is evident in the fine structure of the image but does not degrade image quality significantly.

It is clear from the bar chart in Fig. 5.9 that the use of partial Fourier increases the measured tSNR, both because of the reduced TE and somewhat because of the blurring which results from it causing an averaging effect on the noise. The use of parallel imaging with $R = 2$ reduces the overall tSNR despite the reduction in TE.

Based on these data, future experiments will be conducted using 6/8 partial Fourier to give a boost in SNR and a reduction in signal dropout without significant loss of image quality. Parallel imaging will not be used, since the artefacts produced are significant and the SNR is not increased over the standard protocol. In cases where distortion is highly problematic, parallel imaging may be of use, although the improvement in these artefacts for the protocol used here is




Figure 5.9: Comparison of VEP-CASL vascular territory maps with the addition of partial Fourier and/or parallel imaging (with $R = 2$) to the EPI readout. In each case transverse (TRA), coronal (COR) and sagittal (SAG) slices are shown, with the location of each plane shown as a thin white line on the other two images (leftmost images only). Distortion artefacts near the frontal sinuses are reduced when parallel imaging is used (yellow arrows), although localised noise is also introduced (purple arrows). Use of partial Fourier leads to some minor blurring in the phase encode (anterior-posterior) direction (orange arrow). The three non-standard methods with reduced TE all reduce signal dropout near the frontal sinus (pink arrows). The mean temporal SNR (tSNR) within a whole brain mask for each method is shown in the bar chart on the right, with the error bars representing the standard deviation across subjects. All the differences between methods are significant at the 1% level, based on paired t -tests conducted between pairs of methods, with the exception of methods 1 and 4.

relatively minor.

5.3.5 Consistency Across Subjects

During the development of these optimisation methods, a total of 10 healthy volunteers (8 men and 2 women, mean age 31.4, age range 23-45) with no known neurological deficit were scanned as described in Section 5.3.4 above. All of these subjects were scanned with the “standard” EPI sequence with background suppression to allow comparison with the optimised methods. Over the course of sequence testing one of these subjects was scanned three times, resulting in a total of 12 scans for all subjects. In order to demonstrate that this method consistently produces similar quality vascular territory maps across subjects and sessions, a single transverse slice near the centre of the brain is shown for each scan in Fig. 5.10.

Note that the image quality is similar across subjects and all images are free from significant artefacts other than distortion near the frontal sinus which is found in all EPI based images. The vascular territory maps from the same subject scanned three times on different occasions are qualitatively very similar, giving some confidence that this method is reproducible. There is considerable variation of the vascular territories across subjects, particularly in the posterior of the brain. A more detailed discussion of such variants of the vasculature can be found in Section 5.5.

5.4 Moving to a 3D-GRASE Readout

Although the 2D multi-slice EPI methods described above yield robust vascular territory maps, they suffer from a number of problems: since each slice is excited and readout in turn, the post labelling delay varies between slices, making




Figure 5.10: Comparison of vascular territory maps across subjects using the standard EPI readout with background suppression. A single transverse slice near the centre of the brain is shown for each scan. All scans are from different subjects, with the exception of those marked with an asterisk (*), which all come from the same subject scanned on different occasions.

optimisation and quantification more cumbersome. In addition, background suppression is poor in later slices since it is only optimised for a single time point. 2D multi-slice readouts tend to have lower SNR than full 3D sequences, since in 3D reconstructions the noise is averaged over the entire 3D k -space rather than only within each slice ([114, 130]). Finally, 2D multi-slice readouts require gaps to be left between the slices to prevent the imperfect slice excitation profiles overlapping, which further reduces the SNR and potentially misses important features in the data.

3D readouts have been used previously in ASL, including a “stack of spirals” approach [122] and 3D-gradient and spin echo (3D-GRASE) [131, 132, 114]. In this section, the 3D-GRASE readout scheme was implemented with the VEPCASL preparation and background suppression, as described above.

5.4.1 Sequence Design

The readout consists of a single 90° excitation pulse followed by a series of EPI echo trains separated by 180° refocussing pulses plus associated gradients and “partition” encoding blips (see Fig. 5.11). Each EPI echo train acquires a plane of k -space at fixed k_z and the partition encoding blips move from one plane of k_z to another. In this manner the whole of 3D k -space can be acquired in a single shot. The refocussing pulses ensure that B_0 inhomogeneity induced signal dephasing is kept to a minimum, resulting in images with T_2 contrast. A centric reordering scheme is used such that the plane corresponding to $k_z = 0$ is acquired first, followed by planes at increasing distances from the centre of k -space, with the outer edges of k -space being acquired last. This ensures a minimum TE is achieved and has good properties relating to blurring in the slice direction (see below). Outer volume suppression (OVS) pulses were added, as per Günther *et al.* [114], to saturate the static magnetisation outside the imaging volume, reducing the need for

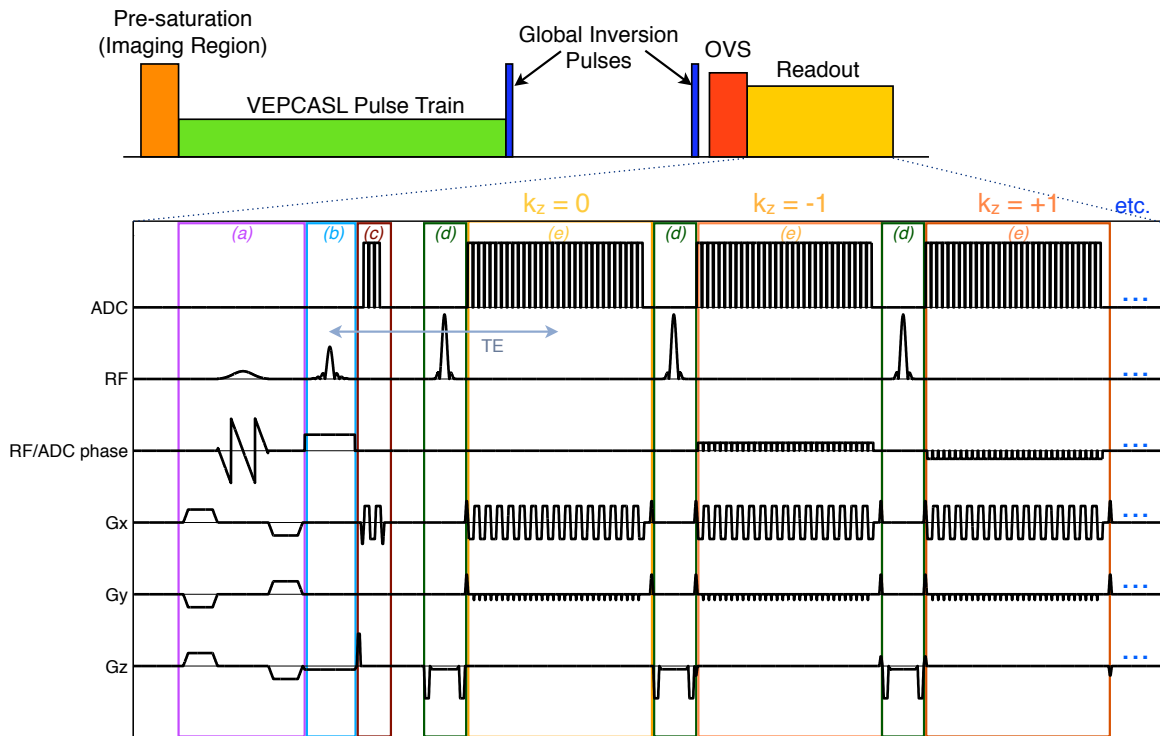


Figure 5.11: Sequence diagram for VEPCASL with a 3D-GRASE readout. Top: Schematic of the preparation scheme, which is identical to that for the EPI sequence, shown in Fig. 5.6, with pre-saturation of the imaging region and two global inversion pulses for background suppression. In addition, an outer volume suppression (OVS) scheme is used to saturate magnetisation on either side of the imaging region, reducing the need for slice oversampling. Bottom: The 3D-GRASE readout module shown in greater detail, consisting of: (a) fat saturation; (b) 90° slab selective excitation pulse; (c) phase correction lines (as for the EPI readout); (d) 180° refocussing pulse; (e) EPI module to acquire a single plane of k -space for a given k_z value. Each EPI echo train is separated by a refocussing pulse and partition encoding blips (here on G_z). The central plane of k -space ($k_z = 0$) is acquired first, followed by planes on either side ($k_z = \pm 1$ in arbitrary units), increasing in distance from the centre until all of 3D k -space has been acquired (a centric ordering scheme).

oversampling in the partition encoding (slice) direction. The sequence parameters used for experiments with the 3D-GRASE sequence are identical to those used with the EPI readout in the previous section to aid comparison (see Table 5.1).

It should be noted that 3D-GRASE is a spin echo technique, so should eliminate static dephasing which causes signal dropout in EPI images. However, since each plane of k -space is acquired using an EPI readout, the phase evolution of

off-resonant spins during the acquisition of each plane is identical to that in EPI images and thus distortion is expected to be the same for both readouts.

5.4.2 Blurring in the Slice Direction

Since the whole of k -space is acquired in a single shot, the total 3D-GRASE echo train is quite long: each plane of k -space takes approximately 41 ms to acquire, giving a readout duration of 738 ms for 18 slices. This is a considerable length of time relative to the T_2 of gray matter, which is approximately 99 ms at 3T [102] (note that the blood magnetisation has exchanged into tissue prior to imaging, so the tissue rather than the blood T_2 is relevant here). Therefore, there will be considerable decay of the signal during the echo train, leading to modulation of the acquired signal, particularly across the k_z direction. This modulation can be estimated for a given T_2 , and described in image space as a perfect image convolved with a kernel, calculated from the Fourier transform of the k -space modulation function (see Fig. 5.12). These calculations show that even slices some distance away from each other will contaminate the other's signal. This "cross-talk" between slices is considerably reduced by using parallel imaging in the first phase encode direction, where using $R = 2$ approximately halves the time to acquire each plane of k -space and thus significantly reduces the total echo train duration.

It is noted that the use of partial Fourier to reduce the total echo train duration may also be beneficial. However, performing partial Fourier in the slice direction will only cause additional blurring (as discussed in Section 5.3.2) and partial Fourier in the first phase encode direction was not compatible with the sequence code available.

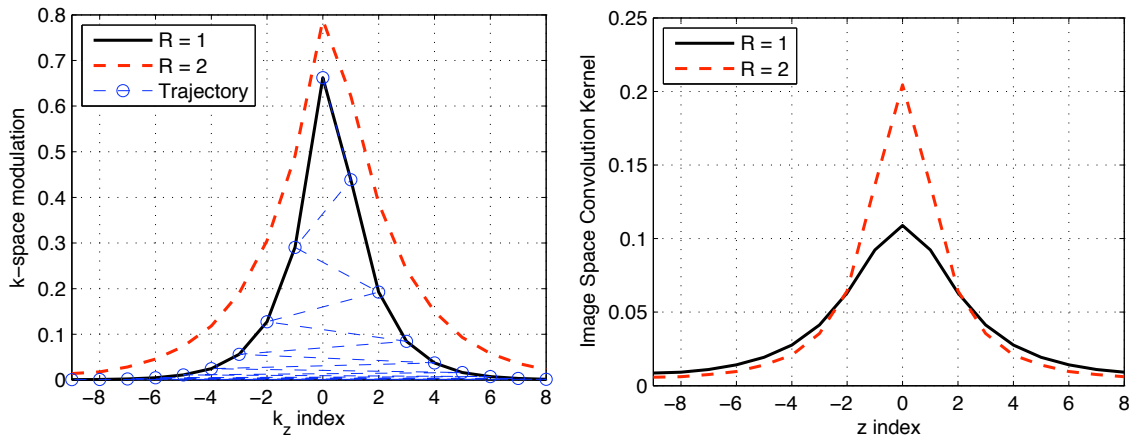


Figure 5.12: Blurring in the slice direction with the 3D-GRASE readout: the modulation of the signal in k -space during the long echo train leads to considerable signal variation along the k_z axis (left), with smaller signals occurring away from the centre of k -space due to the centric k -space trajectory used (shown as a blue dashed line). The use of parallel imaging with $R = 2$ approximately halves the time to acquire each partition, leading to reduced signal attenuation across k_z . The effect of this k -space modulation is the convolution with a kernel in image space (right), which is very broad, but reduced significantly by accelerating with $R = 2$.

5.4.3 Comparison with 2D Multi-Slice EPI

The testing of EPI optimisation methods described in Section 5.3.4 was carried out alongside a comparison with the 3D-GRASE sequence, with and without parallel imaging, forming the basis of a conference abstract [121]. Identical slice prescriptions and imaging parameters were used where possible, with the exception of minimising the TE for each acquisition. Example results in a healthy volunteer and a quantitative tSNR comparison are given in Fig. 5.13. The severe blurring in the z direction expected from the calculations shown above is evident in these 3D-GRASE images, but this is significantly reduced by using parallel acceleration. Unexpectedly, the tSNR is comparable for the 3D-GRASE method with and without parallel imaging. From inspection of the images this appears to be firstly due to the reduced blurring which concentrates the ASL signal inside the brain, rather than having a significant proportion of it blurring into the background. Secondly, the reduced blurring prevents pulsatile artefacts from the large arteries




Figure 5.13: Comparison of VEPCASL vascular territory maps acquired using the 3D-GRASE readout, with and without parallel imaging (with $R = 2$), and the optimised EPI readout with partial Fourier. In each case transverse (TRA), coronal (COR) and sagittal (SAG) slices are shown, with the location of each plane shown as a thin white line on the other two images (leftmost images only). Blurring artefacts in the z (slice) direction are very significant in the 3D-GRASE readout (purple arrows), but a significant improvement is obtained when using parallel imaging (pink arrows). As for EPI images, distortion artefacts near the frontal sinuses are reduced when parallel imaging is used (yellow arrows), although significant localised noise is also introduced (orange arrow, hard to visualise within this dynamic range). The mean tSNR within a whole brain mask for each method is shown in the bar chart on the right, with the error bars representing the standard deviation across subjects. All the differences between methods are significant at the 1% level, based on paired t -tests conducted between pairs of methods, with the exception of methods 2 and 3.

in lower slices appearing higher in the brain, which reduces the average tSNR in the standard 3D-GRASE method.

Despite the increased tSNR, the delineation of boundaries between vascular territories is not as good as the optimised EPI sequence. For this reason, the optimised EPI sequence will be used in future work until considerable improvements can be made to the 3D-GRASE sequence to reduce the blurring artefacts without significant loss of SNR.

5.5 Normal Variants of the Cerebral Vasculature

During these experiments the VEPCASL perfusion technique demonstrated the ability to identify normal variants of the cerebral vasculature through non-standard flow patterns at the tissue level. Two examples, acquired with the optimised EPI sequence, are shown in Fig. 5.14. Both have one recessive vertebral artery, leading to very little signal from this artery being detected downstream. In the first case the contralateral vertebral artery compensates for this, feeding almost the entire cerebellum and posterior of the brain on its own. In the second case, the subject has fetal type circle of Willis (meaning that there is significant collateral flow from the internal carotid arteries through the posterior communicating arteries). This collateral flow makes up for the lack of vertebral blood, providing the majority of the flow to the posterior cerebral artery territories. These conclusions are corroborated by TOF images in these subjects. However, the functional status of the arteries and subsequent downstream perfusion could not be inferred from the TOF data alone, showing the utility of this VEPCASL technique. Indeed, the potential to identify abnormal flow patterns shows promise for application of these methods in patient groups.

These findings are consistent with previously published data [133], where it was shown that the posterior cerebral artery territories were supplied by the ipsilateral internal carotid arteries in patients with a bilateral fetal type circle of Willis. However, in that study both vertebral arteries were labelled simultaneously, and thus it was not possible to identify variants such as the first case described here where one vertebral artery is recessive but the flow patterns around the circle of Willis are otherwise normal.

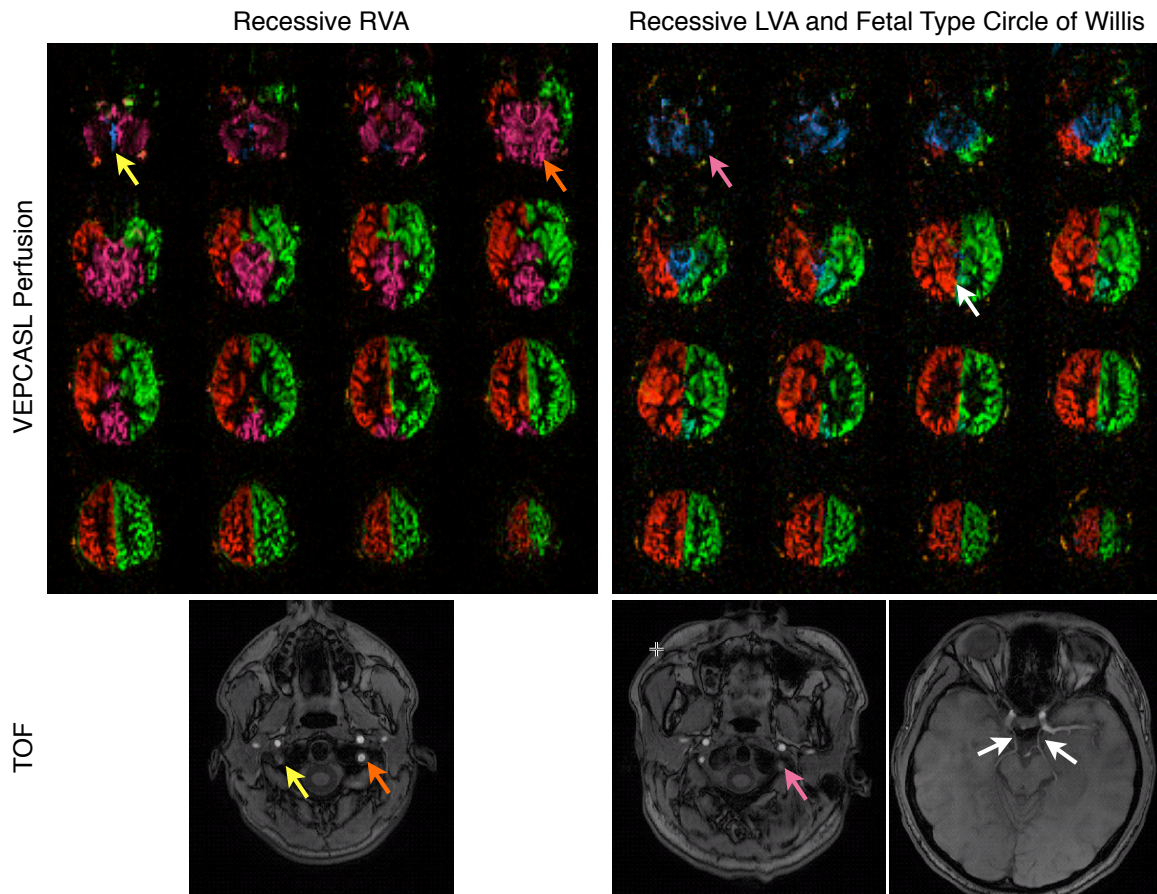


Figure 5.14: Variants of the cerebral vasculature identified in two healthy volunteers using VEPCASL perfusion imaging (top row) with an optimised 2D multi-slice EPI readout. The central 16 transverse slices are displayed running inferiorly to superiorly. Selected slices from a 3D TOF acquisition are shown for comparison (bottom row). The first subject (left) has a recessive RVA, resulting in very little blood signal from this vessel appearing in the VEPCASL perfusion images (yellow arrow). A TOF slice in the neck corroborates this conclusion. The posterior cerebral artery territories and most of the cerebellum are supplied by the LVA (orange arrow) which is considerably wider in the TOF image than most subjects, perhaps to compensate for the weak RVA. The second subject (right) also has a recessive VA, this time on the left side (pink arrow), leading to no LVA signal in the VEPCASL images. This subject also has a fetal type circle of Willis, meaning that there is significant collateral flow from the ICAs across the posterior communicating arteries into the posterior cerebral arteries (white arrows).

5.6 Conclusions

The implementation of the VEPCASL sequence presented here produces well defined vascular territory maps with relatively little artefact in healthy volunteers. A post labelling delay of 1000 ms was found to be sufficient to accurately represent the regional cerebral blood flow in these healthy volunteers, although this value may need to be extended for patient studies or a multi-PLD approach could be used to calibrate the arrival time. A labelling plane location where the internal carotid and vertebral arteries form an approximately rectangular arrangement allows near-optimal vessel encoding and prevents static tissue artefacts from appearing in the brain. The addition of background suppression and 6/8 partial Fourier yielded images with minimal artefacts and the highest SNR for the 2D multi-slice EPI readout. The 3D-GRASE readout has the potential to increase the SNR further, reduce signal dropout artefacts, improve background suppression and facilitate quantification, but was hampered by severe blurring artefacts in the slice direction. Further optimisation of this readout may lead to advantages over the EPI technique.

Vascular territory maps generated using the optimised EPI readout allowed for the inference of abnormal flow patterns, including collateral flow, resulting from normal variation in the cerebral vasculature in some healthy volunteers. This technique thus shows promise for future patient studies. A method for assessing unusual flow patterns in the arteries directly using a dynamic angiographic sequence is discussed in the next chapter.

Chapter 6

Vessel-Encoded Dynamic Angiography

6.1 Introduction

The vessel-encoded tissue perfusion technique described in the previous chapter provides useful information about which arteries are feeding each region of the brain. The presence of collateral flow can be inferred from abnormal territory map patterns, such as signal derived from an internal carotid artery appearing in the posterior regions of the brain. However, in cases with complex flow patterns or abnormal vasculature, looking at the vascular territory maps alone may not provide a clear picture of the flow patterns through the brain feeding arteries. In addition, the anatomical and functional status of the feeding arteries which lead to the abnormal flow patterns are not directly revealed.

In this chapter novel sequences are introduced which combine the vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL) preparation [48] with dynamic ('cine') angiographic readouts, enabling visualisation of the blood flow through the major cerebral arteries in a time-resolved manner. It is hoped

that such methods will be able to visualise directly collateral flow around the circle of Willis or other routes, giving complementary information to the vessel-encoded perfusion method already discussed. Such techniques may also be of use in other clinical areas where x-ray angiography is conventionally used, such as arteriovenous malformation (AVM) and tumours. In addition, it is likely that there is a SNR advantage to such an approach, firstly because the labelled blood is visualised earlier, leading to reduced T_1 decay of the ASL signal, and secondly because the blood is more concentrated whilst still in the arteries.

The first section below (6.2), including many of the figures, is based mostly on work arising from this thesis that has been published [134]. The second section (6.3) describes preliminary results using a more efficient readout method.

6.2 Spoiled Gradient Echo Readout

6.2.1 Sequence Design

The VEPCASL preparation module was combined with a spoiled gradient-echo (SPGR) readout, shown schematically in Fig. 6.1. The VEPCASL pulse train is applied for a specified time (the “tag duration”), creating a bolus of tagged blood water. This is followed by a short inflow delay and then an angiographic image readout module (as described by Günther *et al.* [70] and Sallustio *et al.* [55]) which consists of a 2D thick-slab flow-compensated Fast Low Angle Shot (FLASH) readout combined with a segmented Look-Locker sampling strategy ([135, 92]). The readout is split into a series of readout blocks, within which the same small number of phase encode lines (also called the number of segments) are acquired. The data recorded within each block are allocated to a single time frame. All of the phase encode lines are acquired over multiple VEPCASL preparation periods, allowing the generation of dynamic images with high spatial and temporal reso-

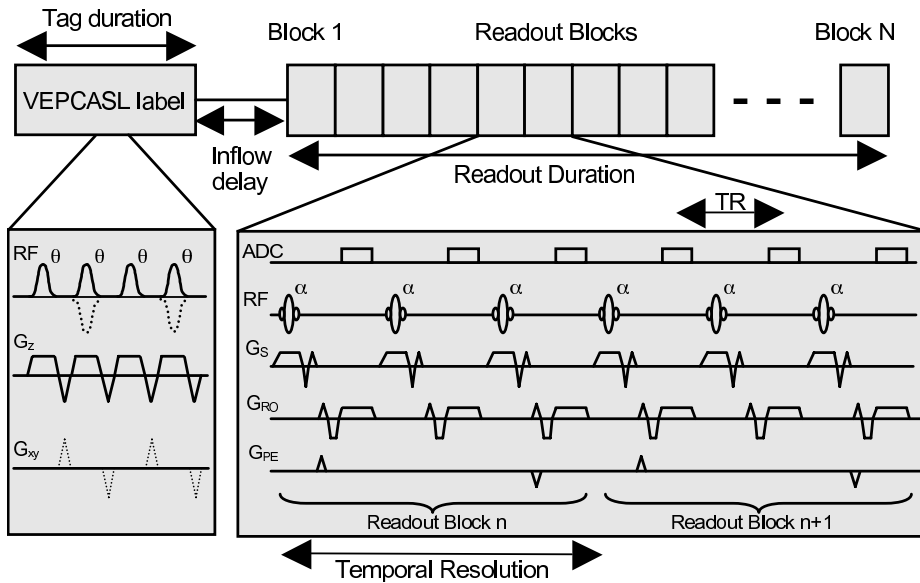


Figure 6.1: Schematic of the vessel-encoded dynamic angiographic pulse sequence. The VEPCASL labelling pulse train consists of a series of radio-frequency (RF) pulses and associated z gradient (G_z) pulses to invert spins passing through an axial labelling plane. In the control condition, consecutive RF pulses have a phase difference of π relative to each other (thick dotted lines). Transverse gradients (G_{xy} , shown as thin dotted lines) are applied along the desired vessel encoding direction only when selective labelling is required. Two readout blocks are shown, each consisting of three single line flow-compensated readouts with different levels of phase encoding (ADC = analogue to digital converter; G_S = slice selection gradient; G_{RO} = readout gradient; G_{PE} = phase encoding gradient; TR = repetition time). Spoiler gradients following each readout are not shown for clarity.

lution, since the temporal resolution is determined by the time between readout blocks, not by the time to acquire an entire image. A 5/8 partial Fourier technique (asymmetric echo) is used in the readout direction to minimise the echo time and gradient moments. This boosts the SNR and reduces signal dephasing due to accelerating blood [29]. Other pulse sequence parameters are given in Table 6.1.

The choice of RF flip angle (α), TR and the number of segments for the readout is a tradeoff between SNR, readout efficiency, signal degradation in later time frames and temporal resolution. With a spoiled readout such as this, use of a high flip angle will give high SNR in early time frames but a poor SNR in later time frames since the ASL contrast is reduced by a factor $\cos \alpha$ after each RF pulse. Reducing the TR will degrade the contrast over time in a similar manner, but will

6.2. SPOILED GRADIENT ECHO READOUT

<i>Group</i>	<i>Parameter</i>	<i>Value</i>
VEPCASL Pulse Train	VEPCASL pulse train duration	1000 ms (long bolus) or 300 ms (short bolus)
	RF Pulse Type	Gaussian
	RF Flip Angle	20°
	RF Duration	600 μ s
	RF Separation	960 μ s
	Mean Tagging Gradient	0.8 mT/m
	Tagging Gradient Amplitude*	6 mT/m
	Inflow delay	1 ms
Readout	Imaging Plane	Any
	Field of View	Approx 205 x 170 mm
	Matrix size	192 x 144 interpolated to 384 x 288
	Slab Thickness	5 - 10 cm
	Readout Bandwidth	78 Hz per pixel
	TR	18 ms
	Segments per Readout Block	3
	Temporal Resolution	55 ms
	TE	3.74 ms
	Excitation Flip Angle	10°
Number of Readout Blocks	20	
Other	Transmit Coil	Body
	Receive Coil	Single or 12 Channel Head
	Time Between VEPCASL Preparations	2000 ms
	Total Imaging Time	Approx 2 mins per VEPCASL cycle

Table 6.1: Default scanning parameters used for vessel selective dynamic angiographic experiments in normal volunteers. *Refers to gradient amplitude during RF pulses.

also increase the temporal resolution and/or readout efficiency (if the number of segments is also increased). The choice of TR = 18 ms, number of segments = 3, and $\alpha = 10^\circ$ had been found to be an appropriate compromise between these competing factors in previous non-selective implementations of this sequence and yielded good results in initial experiments.

Two methods for looping through the phase encode lines and VEPCASL cycles were implemented. The first, termed “cycles outer”, acquires all of the phase encode lines for each VEPCASL cycle before moving on to the next cycle. The second, termed “cycles inner”, acquires the same phase encode lines for each VEPCASL cycle before moving on to the next phase encode lines. The “cycles outer” strategy minimises the time to acquire each full VEPCASL cycle image and thus the effect of motion artefacts or scanner drift within each image is minimised. The “cycles inner” strategy minimises the time between acquisition of the same

phase-encode line for each VEPCASL cycle, which might minimise artefacts upon combining these images.

As for previous experiments with tissue perfusion imaging (see Chapter 5), a transverse labelling plane was positioned approximately 8 cm below the circle of Willis where the two internal carotid and two vertebral arteries run approximately perpendicular to the transverse plane in a roughly rectangular arrangement (see Fig. 6.2). Six cycles of the sequence were performed:

1. Tag all vessels;
2. Control all vessels;
3. Tag left circulation (internal carotid and vertebral) whilst controlling right;
4. Tag right circulation whilst controlling left;
5. Tag anterior circulation (internal carotids) whilst controlling posterior (vertebrals);
6. Tag posterior circulation whilst controlling anterior.

These six cycles were chosen for a number of reasons. Firstly, as discussed in Section 4.11, these cycles form complementary pairs, which allows them to be subtracted from each other without losing vessel specific information. This is particularly important for this angiographic sequence which has a relatively long acquisition time and thus greater potential for artefacts due to subject movement or scanner drift. Pre-subtraction of pairs of cycles was found to significantly reduce these artefacts. Secondly, six was the minimum number of paired cycles that would allow separation of all four vessels in the subsequent analysis (see below). The acquisition of further encoding cycles would better condition the task of separating the vascular components by increasing the rank of the encoding




Figure 6.2: An axial slice from a TOF acquisition in a typical subject that was selected as the labelling plane. In Cycle 3 of the experiment vessels are encoded in the left-right direction such that vessels aligned with Line A are labelled and those aligned with Line B are controlled. In Cycle 4, vessels aligned with B are labelled and those aligned with A are controlled. The labelling in the anterior-posterior direction is performed in a similar manner: in Cycle 5 vessels aligned with Line C are labelled and those aligned with D are controlled. In Cycle 6, vessels aligned with D are labelled and those aligned with C are controlled.

matrix, but would also increase the total acquisition time and therefore artefacts due to subject motion. Finally, this encoding scheme is signal-to-noise ratio (SNR) optimal since each vessel is labelled or controlled an equal number of times [48]. The total acquisition time for all cycles, representing an entire measurement, was approximately 10 minutes.

6.2.2 Experiments

For this study, nine healthy volunteers with no known neurological deficit were recruited (seven men, two women; age range, 23-33; mean age 26.8) each of whom was scanned under a technical development protocol agreed with local ethics and institutional committees. All data were acquired using a 3 Tesla Siemens MRI scanner (both TIM Verio and TIM Trio, Siemens Healthcare, Erlangen, Germany). Comparison was made between a single channel transmit/receive coil and a 12-

channel array receive coil, with the body coil being used for transmission in the latter case. The scanning protocol began with a conventional three-dimensional (3D) multi-slab time-of-flight (TOF) angiographic acquisition of the head and neck. This was used for labelling plane and vessel localisation as well as for comparison with the VEPCASL angiography images which were subsequently acquired in transverse, coronal and/or sagittal oblique imaging planes. In order to encompass the major cerebral arteries around the circle of Willis, a 5 cm thick imaging slab was used in the transverse view and 10 cm slabs for coronal or sagittal oblique views, with positioning as shown in Fig. 6.3. Larger slab thicknesses could have been used to encompass the entire cerebral vasculature, but this would have also increased the amount of static tissue signal which contributes to image noise.

Both the “cycles inner” and “cycles outer” version of the sequence (see above) were tested in the same subject during the same session to determine which method minimised image artefacts. In addition, VEPCASL angiograms were acquired with both a short (300 ms) and long (1000 ms) tag duration. The short duration ensures that the first image frame catches in the inflow of the labelled bolus into the circle of Willis, so the full passage of the blood is clearly visualised. Using a long tag duration will mean that the first acquired image corresponds to the vessels filled with labelled blood, so only outflow can be observed. However, labelling the blood for longer means there is a greater volume of labelled blood, which is likely to increase the SNR. In addition, distal vessels with long blood arrival times may not be seen using the shorter labelling duration and timing information can still be obtained by observing the outflow of blood (see below).




Figure 6.3: Maximum intensity projections of the planning TOF image in sagittal (top right), coronal (top left) and transverse (bottom left) views, upon which typical imaging slab positions are shown for transverse (yellow), coronal (green) and sagittal oblique (cyan) views along with the labelling plane location (dashed red line).

6.2.3 Image Analysis

These dynamic angiographic data sets present some additional challenges compared to the tissue perfusion images described in the previous chapter. In particular, it is vital for this method that complex data be used, since the phase of the blood and static tissue signals can be significantly different from each other. This effect can be caused by the acceleration of blood (i.e. changing speed and/or direction), such that the flow-compensation does not perfectly realign the flowing spins. More importantly, for thick slab readouts such as this, a wide range of off-resonance frequencies will be present in the tissue signal due to imperfect shimming. It is therefore unlikely that the blood and tissue phases will be exactly matched, so using magnitude-only data could lead to poor estimation of the blood signal, as shown schematically in Fig. 6.4. However, the use of complex data makes this technique somewhat more sensitive to certain other effects, such as drift in B_0 , which leads to a drift in the average phase of the magnetisation. In addition, generating robust phase maps from multi-channel coil data can be problematic. The processing steps taken to overcome some of these problems are described below, along with descriptions of how the vascular components were separated and timing information extracted from the data.

Coil Combination with Complex Data

Initial experiments with both single and 12-channel receive coils showed the latter had a considerably higher SNR, by a factor of approximately 1.4 in the centre of the brain and significantly more in cortical regions. In a small proportion of acquisitions the default reconstruction algorithm used with the 12-channel coil produced phase images containing one or more singularities (where the phase varies from zero to 2π around a single point) that arise from the imperfect combination of phase images from separate coils (see Fig. 6.5, left). This leads to a




Figure 6.4: Diagram showing the potential errors which could arise if magnitude-only data are used in ASL angiography. The arrows represent magnetisation vectors in the transverse rotating frame of reference for a single voxel which contains both blood (red lines) and static tissue (blue line). The static tissue magnetisation is approximately the same for each VEPCASL cycle. The blood magnetisation has a different phase from the static tissue, with the solid line showing the control condition and the dashed line the tag condition where the magnetisation has been inverted (T_1 decay is ignored here). The green arrows show the net magnetisation within the voxel for the control (solid line) and tag (dashed line) conditions. If a magnitude only subtraction were performed here, the resulting signal would be close to zero since the net magnetisation has similar magnitude for both tag and control conditions. Performing a complex subtraction accounts for the potential phase differences in blood and tissue magnetisation and yields a signal proportional to the blood magnetisation.

phase instability at the singularities but also in neighbouring areas, resulting in non-negligible signal in the background of the angiogram which obscures smaller vessels (Fig. 6.5, bottom left). The wide range of phases in the image also makes it difficult to perform the median phase correction step described in the following section.

This problem can be overcome by reconstructing the vessel-selective angiograms from the individual coils before they are combined or by using a robust algorithm for combining complex images from multiple coils, such as the adaptive combine method of Walsh *et al.* [136]. This latter method aims to determine a complex filter vector which, when applied to the individual coil images before they are summed, optimises the SNR. This procedure is performed for each small region of the image to account for variation in the complex coil sensitivities over the FOV. For this study the complex filter vector was estimated for each block of 16x16 voxels using

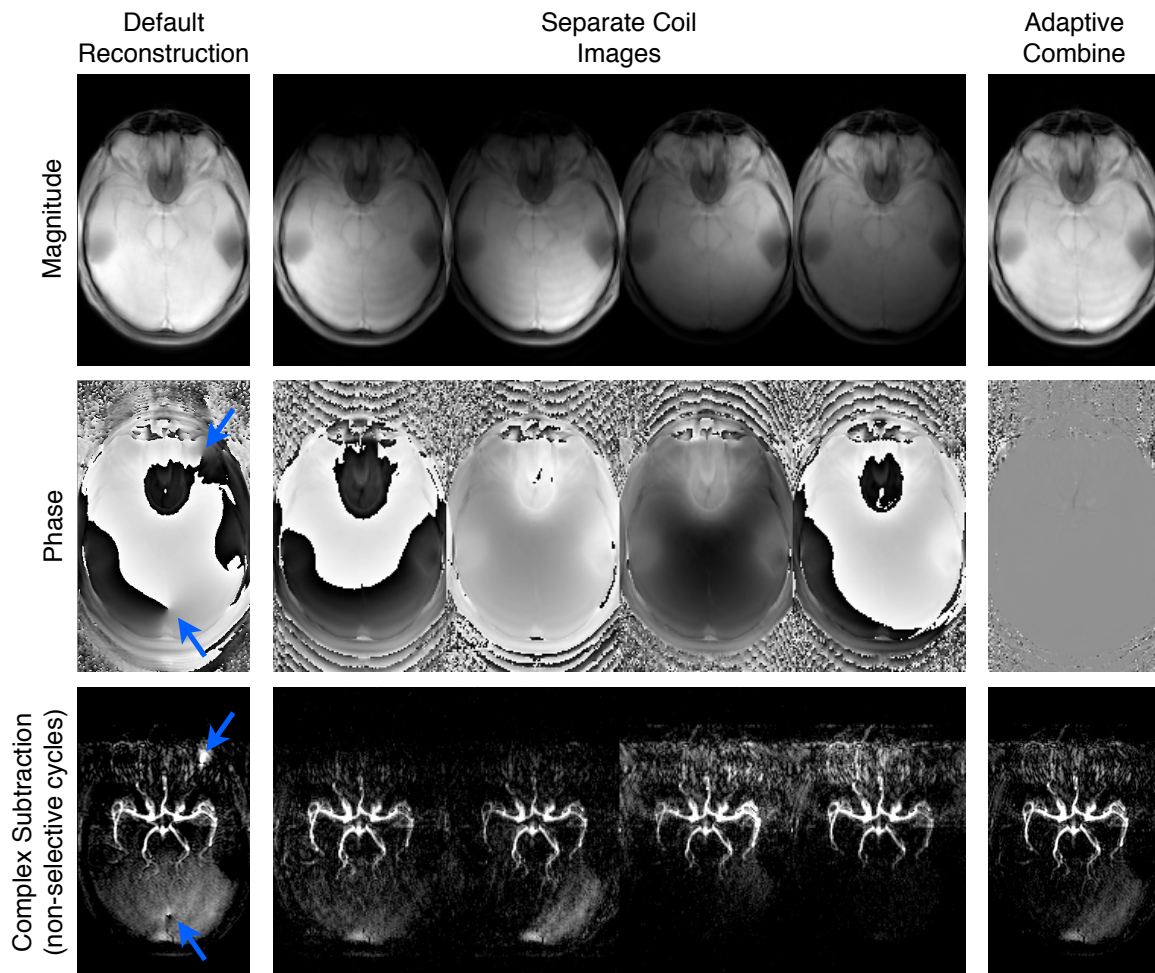


Figure 6.5: Adaptive combination of multi-channel coil data: images from the default image reconstruction algorithm (left) exhibiting phase singularity artefacts (arrows); the 12-channel head coil used has each group of three elements combined in hardware, yielding four separate coil images recorded in the raw data file (middle), which show how the measured phase from each coil differs considerably from the others; the images produced by the adaptive combination algorithm (right) after phase correction (see following section) are free from phase artefacts. Magnitude images are shown on the top row, phase images in the middle row and the complex subtraction of the two non-selective VEP-CASL cycles on the bottom row (only the first time frame after a long tag duration is shown).

a 32x32 voxel estimation region, incorporating all time frames and VEPCASL cycles, and was implemented using Matlab[®] (Mathworks[™], Natick, Massachusetts, USA). The noise correlation matrix was assumed to be equal to the identity matrix to avoid the need for pre-scan noise calibration. The resulting images are free from the phase singularity artefacts (see Fig. 6.5, right).

Phase Correction

When complex data are used, any change in the phase of the images between VEPCASL cycles is interpreted as being due to blood flow. However, there are other factors which may cause the phase to change over time, such as B_0 drift. Since the majority of the raw signal in these angiographic images originates from static tissue, it is possible to estimate such drifts from the data and correct for them. First, a mask is created by thresholding the magnitude images at 10% of their maximum value to exclude low signal regions where the phase is dominated by noise. The median phase within this mask is calculated and subtracted from the phase at every point within the image. This ensures that the median phase of each image becomes zero, correcting for global phase changes.

In addition, the adaptive combine algorithm described above can lead to sharp changes in phase at boundaries between the complex filter estimation regions. Therefore, the phase of each VEPCASL cycle is aligned to that of the control image (i.e. the phase of the control image is subtracted from the phase of all the VEPCASL cycles) to ensure that there are no abrupt changes in phase over the FOV. Since it is only the differences in phase between VEPCASL cycles which are important, this has no impact on the subsequent analysis.

Separation of Vascular Components

In the same manner as described in Section 4.11, individual vessel angiograms can be extracted from the raw data by combining complex images from all six cycles. Unlike the tissue perfusion experiments described in Chapter 5, the acquisition of each VEPCASL cycle takes a significant amount of time, so the number of acquired cycles was minimised in these experiments. However, the data from these six cycles often only provide four unique pieces of information about the vessels supplying blood to each voxel (depending somewhat on the vessel geometry at the labelling plane). In other words, the rank of the encoding matrix is only four. The problem is therefore ill conditioned when solving for the contributions from all four vessels and the static tissue, so simple techniques like the matrix inversion method of Wong [48] would fail. However, the Bayesian inference model developed by Chappell *et al.* [98], outlined in Section 4.11.4, circumvents this problem by determining which combination of vessels is most appropriate to describe the data in each voxel in a probabilistic manner. This allows the individual contributions from all four vessels to be determined and optimises the SNR. This method is particularly useful in cases such as this where the encoding matrix is rank deficient. It can also account for subject motion between the planning TOF acquisition and the VEPCASL acquisition.

In this study we make the assumption that any given voxel is likely to be supplied by only one or two feeding arteries (i.e. two vessels per class, as defined in Section 4.11.4). Using three vessels per class makes the problem more poorly conditioned in this rank-deficient scenario, leading to less stable results in initial tests. Note that it is still possible to represent more than two feeding arteries contributing to a single voxel since the classification is probabilistic.

Timing

The timing information present in each acquisition was condensed into a single 2D map of the estimated blood arrival time at each voxel. This was calculated by first summing all the vascular components then temporally smoothing this signal in each voxel (Savitzky-Golay filter with polynomial order 2 and frame size 19, implemented in the Matlab Signal Processing Toolbox). To exclude low signal regions a mask was produced by thresholding the frame of the angiogram corresponding to the time that the vessels are fully filled with blood, and retaining only the large clusters of voxels to help remove some of the background noise. The time between the end of the VEPCASL pulse train and that at which the blood signal is at 50% of its maximum value after its peak is then determined using linear interpolation. This equates to an approximate time taken for the blood to travel from the labelling plane to each voxel. Simple numerical simulations showed that calculating the time between ceasing the VEPCASL pulse train and detecting the washout of the bolus is equivalent to determining the time between commencing the pulse train and detecting the wash-in of the bolus, with one caveat: dispersion effects must be small relative to the tag duration such that the ASL difference signal in any voxel reaches a plateau before dropping as the bolus washes out. A more sophisticated method for obtaining timing information from these data sets without the caveat above is described in the next chapter.

Inflow Subtraction

The first frame of data sets acquired with a long tag duration corresponds to the vessels filled with labelled blood. Subsequent time frames show the washout of the bolus. Although the same timing information can still be extracted from such data sets, this outflow visualisation is less intuitive than an inflow technique such as x-ray DSA or ASL angiography using a FAIR methodology [70]. However,




Figure 6.6: Example of how inflow subtraction modifies the time series in a single voxel following a long tag duration. At early time points the inflow subtraction signal increases slightly due to the decrease in the original signal caused by the RF excitation pulses. The time at which the original data shows the bolus wash out is identical to the time at which the inflow subtraction shows the bolus wash in. The dynamics of the blood flow are therefore conserved with this method. The time axis is relative to the end of the VEPCASL pulse train.

inflow of labelled blood can be visualised by simply subtracting each time frame from the first and taking the absolute value. A demonstration of this effect is shown in Fig. 6.6 for the time series in a single voxel.

This method has some disadvantages: firstly, the SNR is decreased by a factor of $\sqrt{2}$ due to the subtraction step; secondly, the inflow subtraction is imperfect due to the effect of the RF pulses which attenuate the ASL signal while the bolus is still within the vessels, giving rise to a slowly increasing signal in the inflow subtraction at early time points; finally, it assumes that the first frame corresponds to the vessels completely filled with labelled blood. In voxels where the arrival time of the blood is longer than the tag duration, this inflow subtraction will not produce the expected results. Despite these drawbacks, the resulting movies are generally clearer, particularly for identifying early or late arriving blood in particular vessels.

6.2.4 Results

As for previous chapters, the vessel selective angiograms are generally presented here in colour in order to show the combined information represented by the individual angiograms. The colour indicates the labelled origin of the signal visualised in each voxel (red = right carotid, green = left carotid, blue = right vertebral and magenta = left vertebral). Overlap of vessels or mixing of blood results in intermediate colours (e.g. equal mixing the right internal carotid (red) with the left internal carotid (green) yields yellow voxels). Standard radiological convention is used, such that the right of the image is the subject's left and the top of the image is anterior or superior (depending on the view).

Separation of Vascular Components

The MAP implementation of the Bayesian analysis method used to separate the vascular components from these dynamic angiographic data sets proved highly effective. Each vascular component from a healthy volunteer in the transverse view is displayed both as a combined colour map and separately in Fig. 6.7, demonstrating the high degree of separation. The level of contamination between regions (e.g. RICA signal in the LMCA) is close to the noise level. Note that there is more background noise signal in the RVA and to some extent the LVA components due to the sources of artefact, such as subject motion or scanner drift, correlating more closely with the prescribed encoding of these vessels.

Comparison of Looping Strategies

As mentioned above, the sequence can be run with different looping strategies, with the VEPCASL cycles in the outer loop ("cycles outer"), or the inner loop ("cycles inner"). Fig. 6.8 compares the two strategies in the same healthy volunteer and the same scan session, to ensure identical imaging slab prescription.




Figure 6.7: Demonstration of the high degree with which the four vascular components are separated using the Bayesian MAP method. The first time frame after a long tag duration in a healthy volunteer (transverse view) is shown as a colour map (left) with all components overlaid, and as separate vascular components (right).

These results are typical of those seen in other subjects. Greater suppression of artefacts due to motion or scanner drift is achieved using the “cycles inner” version of the sequence. Therefore, all results presented from this point onwards use the “cycles inner” looping strategy.

Choice of Labelling Duration

For comparison, angiograms were acquired with both a short (300 ms) and long (1000 ms) tag duration, with example results shown in Fig. 6.9. The arrival of the bolus in the imaging region can be observed with a short tag duration, but due to dispersion of the bolus, the distal vessels are less well visualised. When a long tag duration is used, the first frame of the acquired images occurs when all the vessels are filled with inverted blood, so the inflow is not visualised. However, a similar inflow effect can be created using the inflow subtraction technique. Timing information can be obtained from either method by observing the wash-in (short tag duration) or washout (long tag duration) of the bolus from each voxel. The long tag duration allowed the ASL difference signal to plateau before dropping as the bolus washes out (see Fig. 6.6), which did not occur for the short tag duration. As mentioned above, this implies that dispersion effects would confound the blood arrival time estimation for the short but not the long tag duration. As a re-




Figure 6.8: Comparison of sequence looping strategies in a healthy volunteer in the transverse view within the same scan session. The “cycles outer” strategy (left) produces images with more significant artefacts in the background, particularly around the eyes, than the “cycles inner” strategy (right). Only the first frame following a long tag duration is shown here.

sult, all further results presented here were produced using the long tag duration technique.

Example Results Using the Optimal Sequence

Angiograms in transverse, coronal and sagittal (oblique) planes are shown in the same subject in Fig. 6.10, along with maximum intensity projections (MIPs) from the time-of-flight (TOF) data for comparison. In this subject there is little mixing of the blood from the two vertebral arteries when they fuse to form the basilar artery, despite the blood moving to the contralateral side, and only moderate mixing downstream in the posterior cerebral arteries. Some collateral flow from the left internal carotid to the left posterior cerebral artery is visible in the transverse view. In the coronal view the external carotid arteries are also visible. These vessels are not explicitly modelled in the analysis and are therefore assigned to




Figure 6.9: Selected frames from the vessel-encoded dynamic angiography sequence in the transverse view. The top row shows images produced using a short labelling duration of 300 ms, allowing the first frame to catch inflow of the bolus. The middle row shows equivalent images for a longer labelling duration of 1000 ms (in the same subject using an identical imaging slab), which show the distal vessels more clearly (arrows), but the inflow of blood is not visualised. However, a similar effect can be created by performing inflow subtraction using the long bolus data (bottom row). Note that the windowing of each image is different in order to show the features more clearly, giving more pronounced background signal in the short bolus case since the blood signal is weaker. The times shown are at the middle of each readout block relative to the end of the VEPCASL pulse train.

the artery which is encoded in a similar manner: the ipsilateral internal carotid artery. There is a small amount of residual background signal, mainly due to pulsatility effects in the large arteries causing ghosting in the phase encoding direction.

Fig. 6.11 shows that the degree of mixing of vertebral blood in the basilar artery is subject specific since almost complete mixing occurs by the time the blood reaches the posterior cerebral arteries in this case. Timing information extracted from these data sets is also demonstrated in Fig. 6.11, which shows the approximate washout time of the bolus after a long tag duration. The expected pattern of delayed transit to more distal vessels is apparent, along with earlier arrival in central portions of the larger vessels, suggesting faster velocities in the centre of the vessels (e.g. due to laminar flow). In this case the sum of all vascular components was analysed, but analysis could be performed on individual vessel angiograms if desired.

Normal Variants of the Cerebral Vasculature

During this study a number of normal variants of the standard cerebrovasculature were identified in these healthy volunteers. Where appropriate a radiologist was consulted to ensure there was no significant risk to the subject from these findings. Fig. 6.12 shows two examples of such variants. The first has fetal type posterior cerebral arteries, meaning that the proximal segments are hypoplastic (underdeveloped) [14]. The ICAs compensate for this by supplying the majority of the blood to the PCA territories via flow through the posterior communicating arteries. In the second case, the subject has one recessive vertebral artery which does not contribute blood to regions above the circle of Willis, although a small amount of flow to the cerebellum may still be present.

Unfortunately, the first case was acquired during sequence development using




Figure 6.10: Vessel-selective dynamic angiography images (top row) in (left) transverse, (middle) coronal, and (right) sagittal oblique views in the same subject. The sagittal oblique view was rotated by 13.5° about the superior-inferior axis (i.e., from a standard sagittal view toward a coronal view) to reduce the amount of overlap the vessels have in this projection. Only the first frame following a long tag duration (1000 ms) is shown. TOF maximum intensity projections are also shown (bottom row). To aid comparison, these were created by projecting across the same region of space used as the imaging region in the VEPCASL acquisition. Features of interest include little mixing of vertebral blood in the basilar artery despite the blood moving to the contralateral side (white arrows), even downstream in the posterior cerebral arteries (yellow arrows); collateral flow from the left internal carotid to the left posterior cerebral artery (purple arrow); external carotid arteries (blue arrows); the left posterior inferior cerebellar artery (green arrow); and a small amount of residual background signal in some areas (red arrows).




Figure 6.11: The first frame of a long bolus vessel selective dynamic angiography data set in a coronal view (left). In this subject, the blood from the two vertebral arteries is relatively mixed by the time it reaches the posterior cerebral arteries (yellow arrows). Timing information extracted by observing the washout of the long bolus is also shown (right) for the combined signal from all vessels. The color bar indicates the approximate time (in milliseconds) taken for the blood to travel from the labelling plane to the voxel. A mask generated from the vessel-selective dynamic angiographic data was used to exclude low-signal regions. Earlier arrival can be seen in more proximal regions and nearer the centre of the larger vessels where flow is faster (white arrows).

the “cycles outer” strategy, and thus suffers more from background noise. It is worth noting that in such cases where there is significant background noise as well as multiple vascular components (e.g. the right PCA in Fig. 6.12, left), the Bayesian analysis method can struggle to properly separate out all of the components, giving a somewhat blotchy appearance to the vessel. This is probably due to the assumption that there are only two vessels per class, making it hard for the algorithm to properly represent two or three feeding arteries as well as background signal.

6.2.5 Discussion

In this chapter a new vessel-selective dynamic angiographic technique is demonstrated that can produce individual angiograms, non-invasively, by labelling the principal arterial vessels proximal to the circle of Willis. Clear vessel delineation is




Figure 6.12: Examples of normal variants of the cerebral vasculature: fetal type circle of Willis (left) in which the internal carotid arteries are the major feeders to the posterior circulation (arrows) via flow through the posterior communicating arteries; a recessive right vertebral artery (right) means that the left vertebral artery supplies both PCA territories (arrow) in this subject.

achieved and the separation of the four feeding vessels is evident in these healthy volunteers (Fig. 6.7). Static tissue generally subtracted out well and there was only minimal ghosting due to pulsatility effects, despite the lack of cardiac gating.

Fig. 6.8 shows that using the “cycles inner” looping strategy minimises background noise, particularly in areas with significant motion such as the eyes. This might have been expected since the same phase encode lines for each pair of VEPCASL cycles are acquired on consecutive readouts and are therefore only separated by about 2 s. The subsequent subtraction of the complementary pairs of VEPCASL cycles means that only motion or scanner drift which occurs during this short time scale will produce erroneous background signal. The global corruption of the images due to motion (i.e. inconsistency between phase encode lines for a given VEPCASL cycle acquired some time apart) appears to have a small effect by comparison in these healthy volunteers.

Using a long bolus duration yields significantly better images of the distal vessels and improved SNR, which can be seen by the relative size of the background

signal in Fig. 6.9. The inflow subtraction technique applied to the long bolus data yields similar visualisation of the inflow as the short bolus technique, but with better SNR and distal vessel visibility.

The vessels visible with this technique match well those shown in the TOF MIPs, even sometimes showing vessels not visible with the latter technique (particularly the small distal vessels in Fig. 6.9 and 6.10), due to poorer suppression of static tissue in the TOF acquisition. The collateral flow visible in Fig. 6.10 in a healthy volunteer demonstrates the potential of this technique to study this phenomenon in patient groups.

The ability to acquire angiograms in any desired orientation is an advantage over x-ray techniques and also over the original implementation of this sequence which used a FAIR preparation [70]. The former is physically restricted by the subject's body, preventing transverse views which show the circle of Willis so clearly. The latter relies on the inflow of blood from outside the imaging slab. In coronal or saggital views this is likely to encompass very large sections of the feeding arteries, giving a long delay before labelled blood flows into the circle of Willis, leading to greater T_1 decay and reduced SNR. With the VEPCASL method, the labelling plane can be placed anywhere, even within the imaging region, although artefacts occur near the labelling plane where the static magnetisation is perturbed by the preparation, as shown in Section 4.10.

In general the vessels appear wider in the VEPCASL angiographic images than in the TOF MIPs. Slower flow at the edges of vessels reduces the signal enhancement relied upon by the TOF technique. Combined with partial voluming, the signal at the vessel edges is attenuated and therefore often obscured by overlying tissue in the MIP, leading to a reduced apparent vessel diameter.

The lack of mixing of vertebral artery blood within the basilar artery in some volunteers (as shown in Fig. 6.10) has also been reported in vessel encoded per-

fusion studies where cases of both ipsilateral and contralateral supply to the posterior cerebral artery territories have been observed [100].

The timing information shown in Fig. 6.11, generated by observing the washout of the long bolus from each voxel, demonstrates that this technique could be used to quantify blood arrival times in the major cerebral arteries. In patient populations significantly delayed arrival would be suggestive of an underlying stenosis in the feeding vessel or would corroborate the presence of collateral flow. It is expected that relative timing information (e.g. across hemispheres) could be of greater use than absolute arrival times, since the latter depend on the exact location of the labelling plane relative to the vessels of interest. For example, Warmuth *et al.* [137] showed that arrival time differences in the carotid siphon and middle cerebral arteries between hemispheres correlated with the degree of carotid artery stenosis. In addition, it should be possible to quantify flow rates in vessel segments by fitting to an appropriate model (see Chapter 7). In patient populations where significant dispersion of the bolus could occur, the simple timing analysis presented in this chapter might produce biased estimates of blood arrival time (as mentioned previously), motivating further the application of such a kinetic model.

The blotchy appearance of some arteries in Fig. 6.12 could be mediated by the further development and application of the Bayesian analysis technique. In particular, inferring the flow speed in each vessel as well as increasing the number of vessels per class could eliminate this effect. In addition, with a faster acquisition technique, a greater number of VEPCASL cycles could be obtained to increase the rank of the encoding matrix, making the problem of separating vascular components better conditioned.

Vessel selective ASL angiograms have previously been obtained either by selective inversion of a slab oriented to contain a single vessel [71] or by using a

single-artery selective labelling strategy [93, 138]. The former method may suffer from contamination of vascular components due to difficulties in excluding other arteries from the selective inversion slab. In the latter method, only a single artery is labelled per acquisition. This may be suitable for certain applications, but is not time efficient where multiple vessel angiograms are desirable. In addition, there is significant labelling of arteries close to the targeted vessel, leading to signal contamination. This would make it difficult to interpret lower levels of collateral flow, whereas the VEPCASL method can effectively separate out the contributions of multiple feeding vessels, even if they are closely spaced [101]. Phase-contrast magnetic resonance angiography [99] and transcranial colour-coded duplex ultrasonography [139] have also previously been used to study collateral flow in the circle of Willis but these techniques show only the direction of flow within each vessel and thus cannot be used to visualise the amount of collateral flow or mixed supply directly.

It is expected that in some cases the information provided by VEPCASL angiography will be similar to that provided by VEPCASL perfusion imaging (chapter 5) where the blood reaches the tissue or capillary level before image readout. For example, collateral flow observed around the circle of Willis with this angiographic method could be inferred by abnormal or mixed blood supply at the tissue level. However, on theoretical grounds the angiographic method confers a number of advantages. In particular, it allows assessment of vessel stenosis or occlusion (if distal to the labelling plane) and collateral flow directly rather than by inferring it from tissue perfusion. In addition, there is an SNR advantage through imaging the blood sooner after labelling, leading to reduced T_1 decay, as well as the greater concentration of labelled blood in the vessels. This may be particularly important in patient populations where significantly delayed transit to the tissue would lead to very poor SNR at the tissue level. The omission of a

long post-labelling delay makes the angiographic acquisition more time efficient, but differences in the readout technique may be significant. There are also cases where the perfusion technique might be preferred, such as occlusion or stenosis in a small distal vessel that might not be visible in an angiographic image but may be apparent in perfusion images. Experimental comparison of these methods in patient groups is discussed further in Chapter 8.

The VEPCASL angiography sequence described here does have some limitations, however. Firstly, the static tissue is not always perfectly subtracted out, especially if significant subject motion occurs. Secondly, subject motion may be more of a problem in patient groups since images acquired over the entire scan (currently approximately 10 minutes) are combined to form the individual angiograms. Thirdly, the acquired images are two-dimensional so repeated acquisitions are required to obtain multiple projections. Fourthly, as for all ASL techniques, in patient groups where blood arrival is often delayed the T_1 decay of the signal will be more significant, leading to a reduced SNR. Fifthly, there are some minimal artefacts that occur due to the pulsatile nature of the blood flow which can cause variation in the blood signal and changes in the vessel positions and/or diameters. These periodic variations can lead to ghosting in the final image unless cardiac gating is used. Although not very significant in the cerebral arteries, this may be more important in larger arteries if this technique were to be applied elsewhere in the body. Finally, the temporal and spatial resolution does not yet match that of x-ray DSA techniques, limiting application of this method to larger vessels. Some of these issues are addressed in Section 6.3 where the acquisition is greatly speeded up by using a balanced steady state free precession (bSSFP) readout and parallel imaging, allowing a three-dimensional readout with higher resolution, along with cardiac gating to reduce pulsatility effects.

6.3 Readout Optimisation Using Balanced Steady-State Free Precession

To overcome some of the disadvantages of the VEPCASL angiography with a spoiled gradient echo readout described in the previous section, the VEPCASL preparation was combined with a balanced steady-state free precession (bSSFP) readout technique designed by Xiaoming Bi, Peter Schmitt and colleagues [72, 73]. This work led to an abstract submitted to the ISMRM 2011 annual meeting [140] and the contribution of a short section on flow and angiography to a review of SSFP methods for neuroimaging [141]. Some of the advantages and disadvantages of this readout scheme are described here, along with Bloch simulations of signal instabilities in the transient bSSFP regime, a description of the implementation and some preliminary results obtained with this method.

6.3.1 Potential Benefits and Pitfalls of a bSSFP Readout

In spoiled gradient echo (SPGR) readouts, any remaining transverse magnetisation present at the end of each TR period is crushed using gradient and/or RF spoiling. Much more efficient use of the magnetisation can be made by omitting the spoiling step and refocussing the magnetisation at the end of each TR period. When the flip angle and TR are also kept constant, the magnetisation enters a steady-state, following the same path during each TR period (or alternate TR periods if the RF phase is cycled between 0 and π). Such a sequence is known as balanced steady-state free precession (bSSFP) [49]. These sequences are highly SNR efficient and are particularly useful for ASL angiography applications because the attenuation of the ASL contrast by the readout is significantly less than that of a SPGR readout. Therefore, RF pulses can be applied much more rapidly, greatly increasing the readout efficiency. Fig. 6.13 demonstrates how the ASL




Figure 6.13: Bloch simulations of the effect of the readout method on the magnetisation and resulting ASL signal. The top row shows the effect on the z component of the magnetisation (in units of M_0) for both SPGR and bSSFP readouts, along with the unperturbed T_1 decay which would occur in the absence of RF pulses for comparison. Results are shown for magnetisation which is initially at equilibrium (control) and which has just been inverted (tag). The bottom row shows the ASL signal, defined as the absolute difference between the complex transverse magnetisation of tag and control conditions at the echo time (equal to half the TR here), also in units of M_0 . Two regimes are shown: on the left, the same parameters used for the SPGR readout described in Section 6.2, with a flip angle of 10° and a TR of 18 ms; on the right, a more efficient readout with a flip angle of 20° and a TR of 4.5 ms. For fair comparison the bSSFP readout is shown with double the flip angle quoted above to give similar magnitude transverse magnetisation in the initial TR periods. Here we assume no off-resonance and the steady-state is catalysed using the half-angle method of Deimling *et al.* [142].

contrast is maintained much more effectively for a bSSFP readout than a SPGR readout, even when a higher flip angle and shorter TR are used, allowing a greatly accelerated readout with higher SNR.

bSSFP also offers some additional advantages for angiography in general. Firstly, the signal depends approximately on the ratio of T_2/T_1 , which is high for blood [65]. Secondly, bSSFP is somewhat flow compensated: magnetisation moving with constant velocity will experience the same phase accrual in each TR

period due to the gradient pulses. Such constant sources of phase accrual are generally refocussed in bSSFP as long as the phase accrual does not approach $\pm\pi$ within one TR interval.

However, there are additional complications that arise in bSSFP which need not be considered for SPGR. Since the transverse magnetisation is retained for subsequent TR periods, bSSFP is much more sensitive to off-resonance effects. Typically the RF phase is alternated between 0 and π such that magnetisation which is perfectly on resonance is flipped back and forth about the RF axis. However, for magnetisation which is off-resonant and precesses by π radians during the TR period, the effect of the RF pulses is to saturate the magnetisation, giving very low signal and the characteristic banding artefacts of SSFP acquisitions. This effect is described in more detail in Scheffler's review article [49]. Thus the signal strength depends on the local off-resonance frequency as well as other sources of phase accrual, such as those induced by flowing blood during gradient pulses. For the same reason, these sequences are more susceptible to B_0 drift since any change in the precession frequency of the magnetisation will make the banding artefacts move, leading to erroneous signal in ASL subtraction images. The inherent flow-compensation of bSSFP mentioned above only works for spins with constant or slowly changing flow velocity and certain phase-encode ordering schemes in which pairs of TR periods are acquired with similar phase encoding gradients [143]. Out-of-slice effects can also be significant for 2D imaging [144].

In addition, since we want to image the magnetisation soon after the VEPCASL preparation, a steady-state is never truly achieved. Therefore, imaging must be performed in the so-called "transient steady-state" regime. In this period before the magnetisation reaches the true steady-state, signal oscillations can occur which lead to significant image artefacts [49]. A number of methods to catalyse the steady-state and reduce the severity of these oscillations have been proposed,

the simplest of which involves using a half-angle RF pulse in the first TR period which is also half the duration of subsequent TR periods [142]. Another popular method is the use of 5-15 RF pulses with linearly increasing flip angles [145]. Other, more complex methods, such as that by Hargreaves *et al.* [146], are not discussed further here.

6.3.2 Simulation of Signal Instabilities in the Transient Regime

In order to understand the importance of the bSSFP signal oscillations in the transient regime and potential methods for reducing them, Bloch equation simulations were performed as described in Section 4.2. Three steady-state catalysis methods were compared:

1. No catalysis;
2. “Half angle”: the first TR period is half the duration and uses an RF pulse with half the flip angle of all subsequent pulses [142];
3. “Linearly increasing”: each TR period is the same duration but the first 5-15 pulses have flip angles which increase linearly from zero to the final desired value [145];

To simplify the discussion, the blood velocity is assumed to be zero here. The three catalysis schemes are compared for on-resonance magnetisation in Fig. 6.14. It is clear that using no catalysis leads to enormous oscillations in the ASL signal. This is particularly true for long T_1 and T_2 species such as blood, where the oscillations do not die away for a long period of time. Acquiring data with such a scheme would clearly lead to very significant artefacts, so it is not considered further here. Both the “half angle” and “linearly increasing” methods significantly dampen the oscillations, leading to a relatively smooth signal profile over time, although the half angle method works better in this case.




Figure 6.14: Comparison of catalysis schemes for on-resonance magnetisation. T_1 and T_2 values for blood were assumed, with a TR of 4.5 ms and flip angle of 40° . The ASL signal is calculated as described in Fig. 6.13. For the linearly increasing method, 20 catalysis pulses were used.

However, the same does not hold true for all off-resonance frequencies. Fig. 6.15 demonstrates the effect of off-resonance on the “half angle” and “linearly increasing” methods. Although the “half angle” method gives a smoother signal evolution for small off-resonance frequencies, the oscillations are much more severe further from resonance. The “linearly increasing” method gives a relatively stable signal over a wide range of frequencies and is therefore the method implemented in the remainder of this section.

The number of catalysis pulses used with the “linearly increasing” method can be varied, so a range of values were simulated for a variety of off-resonance frequencies to help determine the optimum value for this parameter (see Fig. 6.16). Clearly, increasing the number of catalysis pulses reduces the unwanted signal oscillations. However, no data are acquired during the catalysis period so using too many pulses will prevent imaging of the blood soon after the VEPCASL preparation when the signal is the strongest. A choice of 20 catalysis pulses is a reasonable compromise, giving relatively low signal oscillations and only requiring about 84 ms of the readout period. This value is used for all further




Figure 6.15: Comparison of catalysis schemes off-resonance. Each line represents magnetisation with a different off-resonance frequency (in Hz). As before, a TR of 4.5 ms, flip angle of 40° and 20 catalysis pulses were assumed.

experiments described in this chapter.

6.3.3 Choice of Flip Angle

The simulations described above can also help make decisions on the best parameter choice without the requirement for using significant amounts of scan time. One critical parameter for these experiments is the flip angle. Higher flip angles give higher signal strengths at early time points, but the ASL signal is attenuated more quickly, leading to poorer SNR in later time points (see Fig. 6.17). Lower flip angles give a more consistent signal over time, but with lower signal strength initially. The flip angle is also restricted by SAR considerations. Thus, a flip angle of 40° was chosen for the experiments described here as a suitable compromise between these factors.

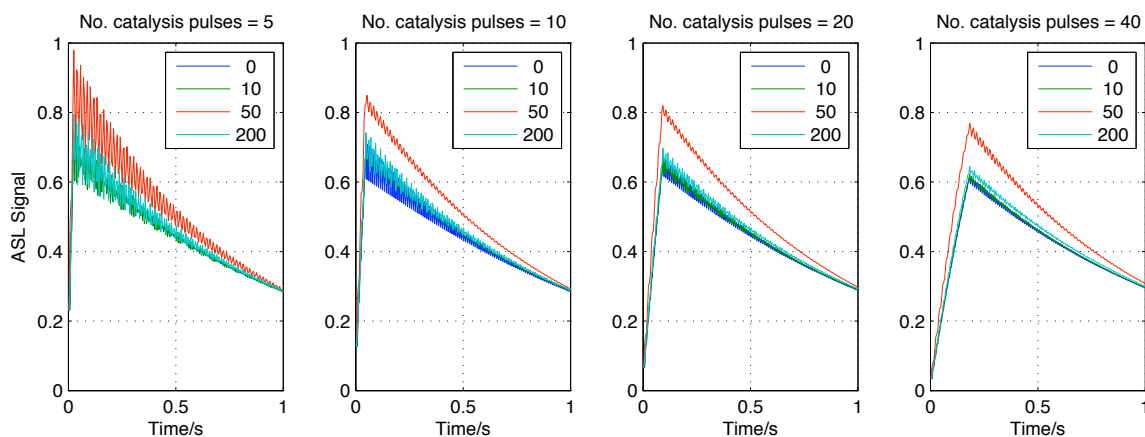


Figure 6.16: Reduction in the ASL signal oscillations with an increasing number of catalysis pulses. A variety of off-resonant frequencies are simulated (see legends, in Hz). A TR of 4.5 ms and flip angle of 40° are assumed here.




Figure 6.17: Variation in the ASL signal over time for various RF flip angles (see legend, in degrees). Lower flip angles give lower signal strength but are more consistent over time. The “linearly increasing” method with 20 catalysis pulses was used with a TR of 4.5 ms here.

6.3.4 Implementation

Sequence Design

The VEPCASL preparation module described previously was combined with a water suppression enhanced through T_1 effects (WET) pre-saturation scheme to suppress the static tissue signal (as used in tissue perfusion experiments in Chapter 5 and similar to that of Golay *et al.* [90]), and the cardiac gated 4D bSSFP readout of Xiaoming Bi, Peter Schmitt and colleagues [72, 73] as part of a collaboration with Siemens Healthcare. The sequence is shown in Fig. 6.18. The presaturation scheme provides some background suppression, reducing the sensitivity of the sequence to subject motion or scanner drift, although there will be considerable T_1 recovery during the VEPCASL pulse train. Unlike perfusion experiments (see Chapter 5), there is no post-labelling delay in which to play out global inversion pulses to further suppress the static tissue. Potential improvements to this simple background suppression scheme are discussed in Chapter 9. Cardiac gating was performed using a pulse oximeter attached to the subject's finger. This reduced sensitivity to pulsatility effects, but also resulted in increased scan time while the scanner waits for the next cardiac trigger. This triggering approach also leads to a variable time between VEPCASL preparations and thus the static tissue magnetisation just prior to the readout. However, the presaturation module ensures that consistent signals are obtained by 'resetting' the static magnetisation shortly before the readout. This sequence uses the "cycles outer" looping strategy by default, such that all phase encode lines for a given VEPCASL cycle are acquired before moving on to the next cycle.

The sequence can either be run in full 4D mode, in which an entire 3D volume is acquired at each time point (readout block), or in dynamic 2D mode in which a 2D image is acquired at each time point, in a similar manner to the SPGR sequence




Figure 6.18: Sequence design for VEPCASL dynamic angiography with a bSSFP readout. The sequence consists of a number of units which commence following an ECG or pulse oximeter trigger: WET pre-saturation of the imaging region for background suppression [90], the VEPCASL pulse train for labelling blood in the neck (only a small section shown here), readout catalysis using 20 RF pulses with linearly increasing flip angles, followed by a number of readout blocks (only two shown here), within which 14 lines of k-space are acquired and binned into a single temporal step. Three blue dots represent continuation of the same repeating pattern, omitted for clarity (ADC = analogue to digital converter; G_x = readout gradient; G_y = first phase encoding gradient; G_z = slice selection/second phase encoding gradient).

described in Section 6.2. The 4D mode allows reformatting of the data in any desired orientation and boosts the SNR due to the additional data acquired, but adds considerably to the scan time: approximately 18 mins in 4D mode compared to about 1.5 mins in dynamic 2D mode, depending on the subject's cardiac cycle.

In order to achieve these scan times at high resolution, considerable acceleration was required. Due to the drastically reduced attenuation of the ASL signal in bSSFP compared to SPGR (see above), the TR can be shortened considerably to 4.2 ms even with a high flip angle of 40° . Maintaining a similar temporal resolution to the original SPGR sequence (59 ms), the number of segments (i.e. k-space lines acquired per readout block) can be increased from 3 to 14. Partial Fourier was used in all three directions (with a zero-padded reconstruction) to reduce the number of k-space lines which need be acquired, along with minimising the TE and TR. Acceleration by a further factor of two was achieved using parallel imaging in the first phase encode direction using the generalized autocalibrating partially parallel acquisition (GRAPPA) technique [59]. Other sequence parameters are given in Table 6.2.

Vessel-Encoding

Due to the increased efficiency of the sequence, eight VEPCASL cycles were performed to increase the rank of the encoding matrix, better conditioning the separation of vascular components. As for the VEPCASL tissue perfusion experiments, discussed in Chapter 5, these cycles were:

1. Tag all vessels;
2. Control all vessels;
3. Tag left circulation (internal carotid and vertebral) whilst controlling right;
4. Tag right circulation whilst controlling left;

6.3. READOUT OPTIMISATION USING BALANCED STEADY-STATE FREE PRECESSION

<i>Group</i>	<i>Parameter</i>	<i>Value (Full 4D mode)</i>	<i>Value (Dynamic 2D Mode)</i>
VEPCASL	VEPCASL pulse train duration		1000 ms
	RF Pulse Type		Gaussian
	RF Flip Angle		20°
	RF Duration		600 μ s
	RF Separation		960 μ s
	Mean Tagging Gradient		0.8 mT/m
	Tagging Gradient Amplitude*		6 mT/m
Readout	Imaging Plane	Transverse slab	Any orientation
	In Plane Field of View	Approx 220 x 177 mm	
	Slab Thickness	64 mm	50 - 100 mm
	Slice Oversampling	12.5%	N/A
	Matrix Size	224 x 180 x 32	224 x 180 x 1
	Parallel Acceleration	GRAPPA factor 2 (1st phase encode direction)	
	Phase Resolution [†]		80%
	Slice Resolution [†]	50%	N/A
	Nominal Voxel Size [†]	1.0 x 1.0 x 2.0 mm	1.0 x 1.0 x (50 - 100) mm
	"True" Resolution [†]	1.0 x 1.2 x 4.0 mm	1.0 x 1.2 x (50 - 100) mm
	Readout Partial Fourier		6.25/8
	Phase Partial Fourier		6/8
	Slice Partial Fourier	6/8	N/A
	Readout Bandwidth		496 Hz per pixel
	TR		4.2 ms
	Segments per Readout Block		14
	Temporal Resolution		59 ms
	TE		1.8 ms
Excitation Flip Angle		40°	
Number of Readout Blocks		12	
Other	Transmit Coil	Body	
	Receive Coil	12 Channel Head	
	Time Between VEPCASL Preps	Approx two cardiac cycles	
	Total Imaging Time	\approx 18 mins	\approx 1.5 mins

Table 6.2: Sequence parameters for the 4D and dynamic 2D VEPCASL dynamic angiography with bSSFP readout. Parameters which are common to both acquisition strategies are placed between the two columns. *Refers to gradient amplitude during RF pulses. [†]This percentage of lines are acquired in the specified direction with the rest being zero padded. The nominal voxel size excludes this effect but the "true" resolution accounts for it (e.g. in 4D mode the nominal slice thickness is 2 mm, but the slice resolution is only 50%, meaning the "true" slice resolution is 4 mm).



Figure 6.19: Vessel-encoding using eight cycles overlaid on a TOF axial slice chosen as the labelling plane in a healthy volunteer. Cycles 1 and 2 are non-selective. Cycle 3: tag at Line A control at Line B. Cycle 4: label B and control A. Cycle 5: label C and control D. Cycle 6: label D and control C. Cycle 7: label E and control F. Cycle 8: label F and control E. Note that there are two lines labelled F since the spatial modulation of inversion efficiency is periodic.

5. Tag anterior circulation (internal carotids) whilst controlling posterior (vertebrals);
6. Tag posterior circulation whilst controlling anterior;
7. Tag diagonally, labelling the right internal carotid and left vertebral whilst controlling the left internal carotid and right vertebral;
8. Tag diagonally, labelling the left internal carotid and right vertebral whilst controlling the right internal carotid and left vertebral.

These tagging locations are shown graphically in Fig. 6.19.

Image Processing

These bSSFP VEPCASL angiographic data were processed as described for the SPGR method (see Section 6.2.3), with the following exceptions:

Firstly, the data were processed using the default reconstruction algorithms on the scanner before exporting off-line. Reconstructing images from the raw k-space data would have required an off-line parallel imaging reconstruction to be implemented, which was beyond the scope of these initial experiments. As mentioned previously, using the default reconstruction leads to phase singularity artefacts in some cases. For this reason, magnitude-only data were used for image processing. As discussed in Section 6.2.3, this can lead to underestimation of the blood signal. However, in these bSSFP data sets the phase appears relatively uniform due to the refocussing effect of the bSSFP readout, meaning that such errors are less significant than for SPGR. Qualitative comparison of initial results using complex or magnitude-only data revealed very similar image features, but using the magnitude-only method suppressed the artefacts due to the phase singularities. Thus for the remainder of this section, magnitude-only data were used for the processing, so no coil combination or phase correction steps were necessary.

Secondly, the brain extraction tool (BET) [129] was applied to the 4D data to segment out non-brain structures which can obscure interesting features in maximum intensity projections (described below), particularly where there are significant motion artefacts (e.g. around the eyes).

The remainder of the processing (separation of vascular components, extraction of timing information and inflow subtraction) was performed in exactly the same manner as described in Section 6.2.3. One additional step was taken in order to display the 4D data sets: maximum intensity projection (MIP) was performed in all three spatial directions in order to obtain transverse, coronal and sagittal time series from the same data set. Note that the MIP was performed for each vascular component separately, so the mixing of blood in the vessels is still represented clearly.

6.3.5 Experiments

This novel sequence was tested in five healthy volunteers with no known neurological deficit (all male, mean age 30.8 years, age range 24-38 years) who were recruited and scanned under a technical development protocol agreed with local ethics and institutional committees. Subjects were scanned on a Siemens 3T TIM Trio system (Siemens Healthcare, Erlangen, Germany) using a 12-channel head coil for receive and body coil for transmission. A multi-slab 3D TOF acquisition was performed for vessel localisation and labelling plane selection. VEPCASL dynamic angiography with bSSFP readout was performed both in full 4D mode and in dynamic 2D mode using the transverse view in the latter case. Exactly the same 64 mm imaging slab was used in both cases to aid comparison.

6.3.6 Results

The bSSFP based sequence produced high resolution angiograms with clear vessel delineation and good suppression of static tissue in all subjects. Example frames from a 4D data set in a healthy volunteer projected across transverse, coronal and sagittal directions are shown in Fig. 6.20. As shown previously, the VEP-CASL preparation and Bayesian analysis method produce a high level of vessel specificity with very little contamination of vascular components. The ability to project across any desired direction is clearly an advantage over 2D angiography, although the resolution and coverage in the superior-inferior direction is somewhat poorer in this case. Normal filling of arteries distal to the circle of Willis is observed in this subject, although a low level of collateral flow from both ICAs to the ipsilateral PCAs can be seen, demonstrating the ability of this sequence to visualise this phenomenon in patient groups. The use of inflow subtraction provides clear visualisation of the blood arriving at the circle of Willis, particularly




Figure 6.20: Selected frames from transverse (TRA), coronal (COR) and sagittal (SAG) MIPs of a single 4D data set collected in a healthy volunteer after separation of vascular components and inflow subtraction. In this case each time frame was subtracted from the second since the first frame contained some residual artefact (see Fig. 6.21). The times shown are those in the centre of each displayed readout block relative to the time at which the VEPCASL train ceased. For clarity, the subject's orientation is marked on each image: R = right, L = left, A = anterior, P = posterior, S = superior, I = inferior.

the delayed arrival from the VAs.

A comparison of the 4D and dynamic 2D modes in the same subject is shown in Fig. 6.21. In order to show some of the image features more clearly, inflow subtraction was not used here. In general, the angiograms are comparable between the two modes, although the 4D mode has higher apparent SNR, showing some of the smaller vessels more clearly. The timing information derived from both acquisitions is also very similar, both showing the delayed blood transit to the posterior circulation, as can be observed in Fig. 6.20.

A number of artefacts are present in these acquisitions: as for the SPGR sequence, any motion, such as that around the eyes, leads to spurious signal in the ASL subtraction. There are some additional artefacts present here which are unique to bSSFP. The first frame of the 4D data set contains signal smeared across the first phase encode direction (left-right) in line with the internal carotid arter-




Figure 6.21: Comparison of 4D mode (top row, shown as a transverse MIP) and dynamic 2D mode (bottom row) acquired in the same subject and session using an identical slice prescription. Inflow subtraction was not used in this case to help highlight some of the image features. The times shown are those in the centre of each displayed readout block relative to the time at which the VEPCASL train ceased. The estimated bolus washout time for both data sets is displayed on the right, showing the later arrival in distal vessels and the posterior circulation (white arrows). A number of artefacts are present in these data sets: residual motion artefacts near the eyes (light blue arrows) and the auditory canal (orange arrow), first frame artefact (yellow arrow, see text) and residual signal in large arteries (purple arrows).

ies, but this is missing from subsequent frames. For this reason the second frame was used for the inflow subtraction in this subject (Fig. 6.20). Interestingly, this artefact was not seen so clearly in other subjects. There is also some residual background signal near the auditory canal in the 4D acquisition. Finally, both sets of images contain residual signal in the large feeding arteries after the bolus of labelled blood has passed through.

6.3.7 Discussion

It is clear from these results that the bSSFP readout has some considerable advantages over the SPGR technique described in Section 6.2. The higher resolution and SNR allow significantly better visualisation of the smaller distal vessels, even

though the comparable dynamic 2D acquisition is performed in approximately one sixth of the scan time. This enormous increase in readout efficiency allows the collection of 4D data sets, which have the great advantage of allowing visualisation of the arteries in any desired plane and increased SNR over the dynamic 2D method. Currently the 4D mode scan time may be too long to fit into a clinical protocol and is also likely to be more susceptible to motion corruption in patient groups. However, the dynamic 2D mode is significantly faster, making this sequence a much more attractive option for clinicians. The 4D mode scan time could potentially be cut by a further 25% by using six VEPCASL cycles rather than eight, as shown previously for the SPGR technique, with a sacrifice in SNR and more difficult separation of vascular components. Further methods for making the 4D sequence viable for a clinical protocol are discussed in Chapter 9.

These images contain some residual artefacts which should be addressed. As for the SPGR sequence, eye motion causes errors in the resulting ASL images, although these local disturbances do not greatly affect the appearance of the main vessels of interest.

In the subject shown in Fig. 6.20, the 4D mode data set had a significant smear of signal across the phase encode (left-right) direction in the first frame. As this artefact was not seen in the 2D data set or to such an extent in other subjects, perhaps the most likely explanation is the timing of the readout relative to the cardiac cycle: after the cardiac trigger, the VEPCASL preparation takes approximately one second, which is close to the cardiac period in most subjects. Thus, the first frame of the angiogram is acquired during systole when the blood flow is changing most quickly. Slight variations in heart rate would lead to some parts of k-space being acquired just before systole, at systole, or just after systole. These considerable variations across k-space could lead to such an artefact which disappears in subsequent frames. Variations in heart rate are perhaps less significant

over the course of the much shorter dynamic 2D acquisition, making this artefact less significant. In future work the labelling duration could be made a little longer to ensure the majority of the readout period occurs during diastole.

A residual artefact is also seen near the auditory canal in the 4D data set. This region usually contains bSSFP banding artefacts due to the B_0 inhomogeneity created by the air/tissue interface. Any subject motion during this relatively long scan will lead to significant artefacts, particularly in regions where there is a sudden change in signal intensity such as near these banding artefacts. Alternatively, a small drift in B_0 during the acquisition would lead to movement of the bands and errors upon subtraction. Experiments on a different 3T scanner which was more prone to B_0 drift led to very significant band movement and artefacts in the resulting angiograms. For a robust implementation on any scanner, it is therefore desirable to repeatedly measure B_0 during the acquisition and update the RF and ADC frequencies accordingly. Such a method is planned in future work.

The residual signal in large vessels after the labelled bolus has washed out is likely to be due to greater pulsatility in these vessels which causes signal fluctuations unrelated to the tagged blood. This may be more significant for bSSFP than SPGR since it is sensitive to effects that influence the steady-state, such as changing blood velocity which will affect the phase accrual of blood moving quickly during gradient pulses.

Initial experiments in a single subject using dynamic 2D mode in coronal or sagittal views (not shown here) revealed unexpected loss of signal in proximal vessels, particularly the ICAs, before the labelled bolus had washed out. This could be due to dephasing caused by gradient pulses where the blood is accelerating quickly. Further work is required to explore this effect further, but potentially more effective flow compensation needs to be implemented before this method could be used clinically, particularly in coronal or sagittal 2D mode.

In addition to reducing these artefacts, a number of improvements could be made to this sequence. Firstly, the sensitivity to motion could be reduced by using the “cycles inner” looping strategy described in Section 6.2. Secondly, implementation of the adaptive combination algorithm for off-line combination of complex multi-coil data, as described in Section 6.2.3, could improve image quality and prevent underestimation of the blood signal, particularly in proximal vessel segments where T_1 decay is less severe. Thirdly, a variable flip angle approach could be used to give a more uniform blood signal over the duration of the measurement. Schmitt *et al.* [73] developed such an approach optimised for a FAIR labelling scheme in which all the blood is labelled within a short period of time. However, blood is continuously labelled by the VEPCASL pulse train, so this particular optimisation would not be appropriate. An optimisation strategy, such as that described by Smith *et al.* [147], could potentially be used to generate a flip angle schedule optimised for VEPCASL which does lead to a more uniform blood signal over time.

6.4 Conclusions

The novel vessel selective dynamic angiography sequences described here provide highly selective angiograms non-invasively that could potentially be used in place of x-ray DSA techniques for larger vessels. The bSSFP readout is considerably more efficient than the initial SPGR implementation, but has some associated artefacts which need to be better resolved before it could be applied clinically.

Applications of this technique include the direct visualisation of collateral flow around the circle of Willis, including an analysis of the relative contributions to outgoing vessels with mixed blood supply. This may be useful in applications such as planning for carotid endarterectomy in patients with bilateral stenoses.

With further improvements in spatial resolution it would also be possible to label the extra-cranial circulation, label vessels above the circle of Willis, or apply this method to other areas of the body. It is hoped that the information provided could help with the diagnosis, prognosis or therapeutic decision making in cerebrovascular disease, brain tumours (e.g. identifying tumour blush and vascular supply) and arteriovenous malformations.

There is considerable scope for improving upon the methods described here, including many possibilities for readout acceleration, some of which are discussed in Chapter 9. In the next chapter a method is described in which VEPCASL dynamic angiographic data can be used to derive quantitative blood volume flow rates in various vessel segments around the circle of Willis.

Chapter 7

Quantification of Dynamic Angiography

7.1 Introduction

Although the images produced by the vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL) dynamic angiography sequences described in the previous chapter provide useful qualitative information on collateral flow and the functional status of each artery, they are somewhat hampered by the lack of quantitative information. In particular, late arriving blood will appear less intense due to the greater T_1 decay which may somewhat bias the observers' interpretation of the data. Comparisons between healthy volunteers and patients or in longitudinal studies would also be difficult. It is therefore desirable to develop a method for quantification of blood volume flow rate in specific vessel segments due to each feeding artery.

A simple modification of the general ASL kinetic model [33] is not appropriate since there is no exchange between the blood water and tissue in large vessels. Indeed, the signal strength measured in these experiments depends mainly on

blood volume, not blood flow. It is the total volume of labelled blood created by labelling for a specific time which is crucial here. The modulation of the signal amplitude due to T_1 decay and the effect of RF pulses must also be accounted for. Quantification in ASL angiography has previously been reported [71], but required an additional calibration scan, is valid only for pulsed ASL experiments and assumed plug flow through the vessel of interest.

The basic idea behind this method is that for a very short labelling duration, all the labelled blood is likely to be contained within a vessel segment of interest at one point in time. If the total volume of labelled blood within this vessel segment can be estimated, then with knowledge of the labelling duration, the volume flow rate can be calculated simply as the volume divided by the time. However, use of very short labelling durations leads to poor SNR. To avoid this problem, a model to describe the time evolution of the ASL signal in a single voxel is derived. Once the arrival time, blood volume and dispersion effects have been measured, a very short bolus duration can be simulated in the absence of T_1 decay and RF effects, allowing quantification to be performed.

In this chapter, first the theoretical model for signal evolution in a single voxel is derived, followed by discussion of a calibration method for relating this signal to blood volume using the same data set. Then the method for flow rate quantification is described along with some preliminary results in a healthy volunteer and a patient with cerebrovascular disease. For simplicity we consider only the case of the spoiled gradient echo dynamic angiography sequence here, but with minor modifications this method should also be applicable to the bSSFP sequence, although some of the artefacts described in the previous chapter would need to be addressed before this is viable. This work led to an a conference abstract (under review at the time of writing) [148].

7.2 Theory

7.2.1 Model of Signal Evolution in a Single Voxel

Let $c(x, y, z, t)$ be the normalised concentration of labelled blood in a small volume, $dx dy dz$, at a given position and time (i.e. equal to one when this small volume is completely filled with labelled blood and zero when outside arteries or in unlabelled blood). The ASL difference signal we expect to measure, $S(x_i, y_j, z_k, t_m)$, at voxel position (x_i, y_j, z_k) and sample time t_m is:

$$S(x_i, y_j, z_k, t_m) = S_0 \iiint_{\text{voxel}} dx dy dz c(x, y, z, t_m) T(t_m, \delta_t(x, y, z)) R(t_m, \delta_t) C_S(x, y, z) \quad (7.1)$$

where S_0 is a calibration factor describing the ASL signal per unit volume of pure blood in the absence of T_1 decay and RF effects, $T(t_m, \delta_t)$ describes the T_1 decay experienced by the blood after labelling where $\delta_t(x, y, z)$ is the blood transit time from the labelling plane to the position of interest, $R(t_m, \delta_t)$ describes the modulation of the signal due to the effect of previous RF pulses and $C_S(x, y, z)$ is the coil sensitivity. From this point onwards it is assumed that the coil sensitivity has been removed as a pre-processing stage and can therefore be neglected.

S_0 will depend on the equilibrium magnetisation of blood, M_{0B} , the excitation flip angle used, α , a calibration factor for the signal strength per unit magnetisation, C , and the inversion efficiency β :

$$S_0 = 2M_{0B}C\beta \sin \alpha \quad (7.2)$$

For the case of continuous or pseudo-continuous labelling, $T(t_m, \delta_t)$ will be only a function of $\delta_t(x, y, z)$ since the time between labelling and arrival at a given location will depend only on the vessel geometry and flow speed (neglecting

pulsatility effects). Thus:

$$T(\delta_t) = \exp(-\delta_t(x, y, z)/T_{1b}) \quad (7.3)$$

where T_{1b} is the longitudinal relaxation time of blood. Each RF pulse reduces the longitudinal magnetisation by a factor of $\cos \alpha$. This applies to both the inverted and non-inverted blood and therefore the ASL difference signal is attenuated by the same factor. Thus at time t_m :

$$R(t_m, \delta_t) = (\cos \alpha)^{N(x, y, z, t_m)} \quad (7.4)$$

where $N(x, y, z, t)$ is the *previous* number of RF excitation pulses experienced by the blood at position (x, y, z) and time t (i.e. zero for the first RF pulse). If we make the further assumption that the number of RF pulses experienced by the magnetisation has no spatial dependence, such that all the magnetisation sees all of the RF pulses, then N can be written using a simple expression:

$$N(x, y, z, t_m) = N(t_m) = mn_{seg} + \text{floor}(n_{seg}/2) \quad (7.5)$$

where the first acquired image corresponds to $m = 0$, the number of segments is n_{seg} and “floor” rounds down to the nearest integer to identify the number of RF pulses before the centre of k-space is acquired in each time frame. In practice this assumption will not hold for labelled blood flowing into the imaging slab after the start of imaging.

Neglecting dispersion and velocity distributions across the elemental voxel, the concentration of labelled spins is a simple “rect” function:

$$c_{\text{ideal}}(x, y, z, t) = \begin{cases} 0 & t < \delta_t \\ \kappa(x, y, z) & \delta_t \leq t < \delta_t + \tau \\ 0 & t \geq \delta_t + \tau \end{cases} \quad (7.6)$$

where τ is the labelling duration and $\kappa(x, y, z) = 1$ inside an artery and is zero elsewhere. In practice dispersion of the bolus will cause a smearing out of this “ideal” profile, which can be described as a convolution:

$$c(x, y, z, t_m) = c_{\text{ideal}}(x, y, z, t_m) \otimes D(x, y, z, t') \quad (7.7)$$

where $D(x, y, z, t')$ is the convolution kernel describing the signal evolution at point (x, y, z) resulting from an infinitely short labelling pulse. Here we use a gamma variate function which has been used in contrast-enhanced MRI studies to model the inflow of a bolus of contrast agent [149]. Although this is not identical to the model required here, it has the desirable properties of being causal and the ability to model a long tail, suitable for describing blood flowing in a laminar fashion. Here we use the parameterisation similar to that described by Rausch [149]:

$$D(x, y, z, t') = \begin{cases} \frac{s}{\Gamma(1+ps)} e^{-st'} (st')^{ps} & st' > 0, ps > -1 \\ 0 & \text{otherwise} \end{cases} \quad (7.8)$$

where $p(x, y, z)$ describes the time to peak of the distribution, $s(x, y, z)$ is the sharpness of the distribution and Γ is the gamma function which normalises the kernel to unit area:

$$\Gamma(1 + ps) = \int_0^\infty e^{-u} u^{ps} du = \int_0^\infty s e^{-st'} (st')^{ps} dt' \quad (7.9)$$

It is noted that this formulation neglects the difference in T_1 decay between different parts of the bolus which have different arrival times. However, this

discrepancy can be absorbed into the shape of the dispersion kernel and should not have a significant impact on attempts to fit this model to the data.

The convolution of the ideal bolus shape with the dispersion kernel can be rewritten as follows:

$$c(x, y, z, t_m) = c_{\text{ideal}}(x, y, z, t_m) \otimes D(x, y, z, t') \quad (7.10)$$

$$= \int_{-\infty}^{\infty} dt' c_{\text{ideal}}(x, y, z, t_m - t') D(x, y, z, t') \quad (7.11)$$

$$= \kappa(x, y, z) \int_{t_m - \delta_t - \tau}^{t_m - \delta_t} dt' D(x, y, z, t') \quad (7.12)$$

$$= \kappa(x, y, z) \left[\int_0^{t_m - \delta_t} dt' D(x, y, z, t') - \int_0^{t_m - \delta_t - \tau} dt' D(x, y, z, t') \right] \quad (7.13)$$

$$= \kappa(x, y, z) [\Gamma_I(s(t_m - \delta_t), 1 + ps) - \Gamma_I(s(t_m - \delta_t - \tau), 1 + ps)] \quad (7.14)$$

where Γ_I is the incomplete gamma function.

Fig. 7.1 shows schematics of the ideal bolus shape, the dispersion kernel and the RF modulation function and how they combine to give the expected signal variation with time.

Putting all this together we have:

$$S(x_i, y_j, z_k, t_m) = S_0 \iiint_{\text{voxel}} [\Gamma_I(s(t_m - \delta_t), 1 + ps) - \Gamma_I(s(t_m - \delta_t - \tau), 1 + ps)] \kappa(x, y, z) e^{-\delta_t/T_{1b}} (\cos \alpha)^N dx dy dz \quad (7.15)$$

If we assume that δ_t , s and p are approximately constant over the dimensions of the voxel in all directions then we arrive at our final model:

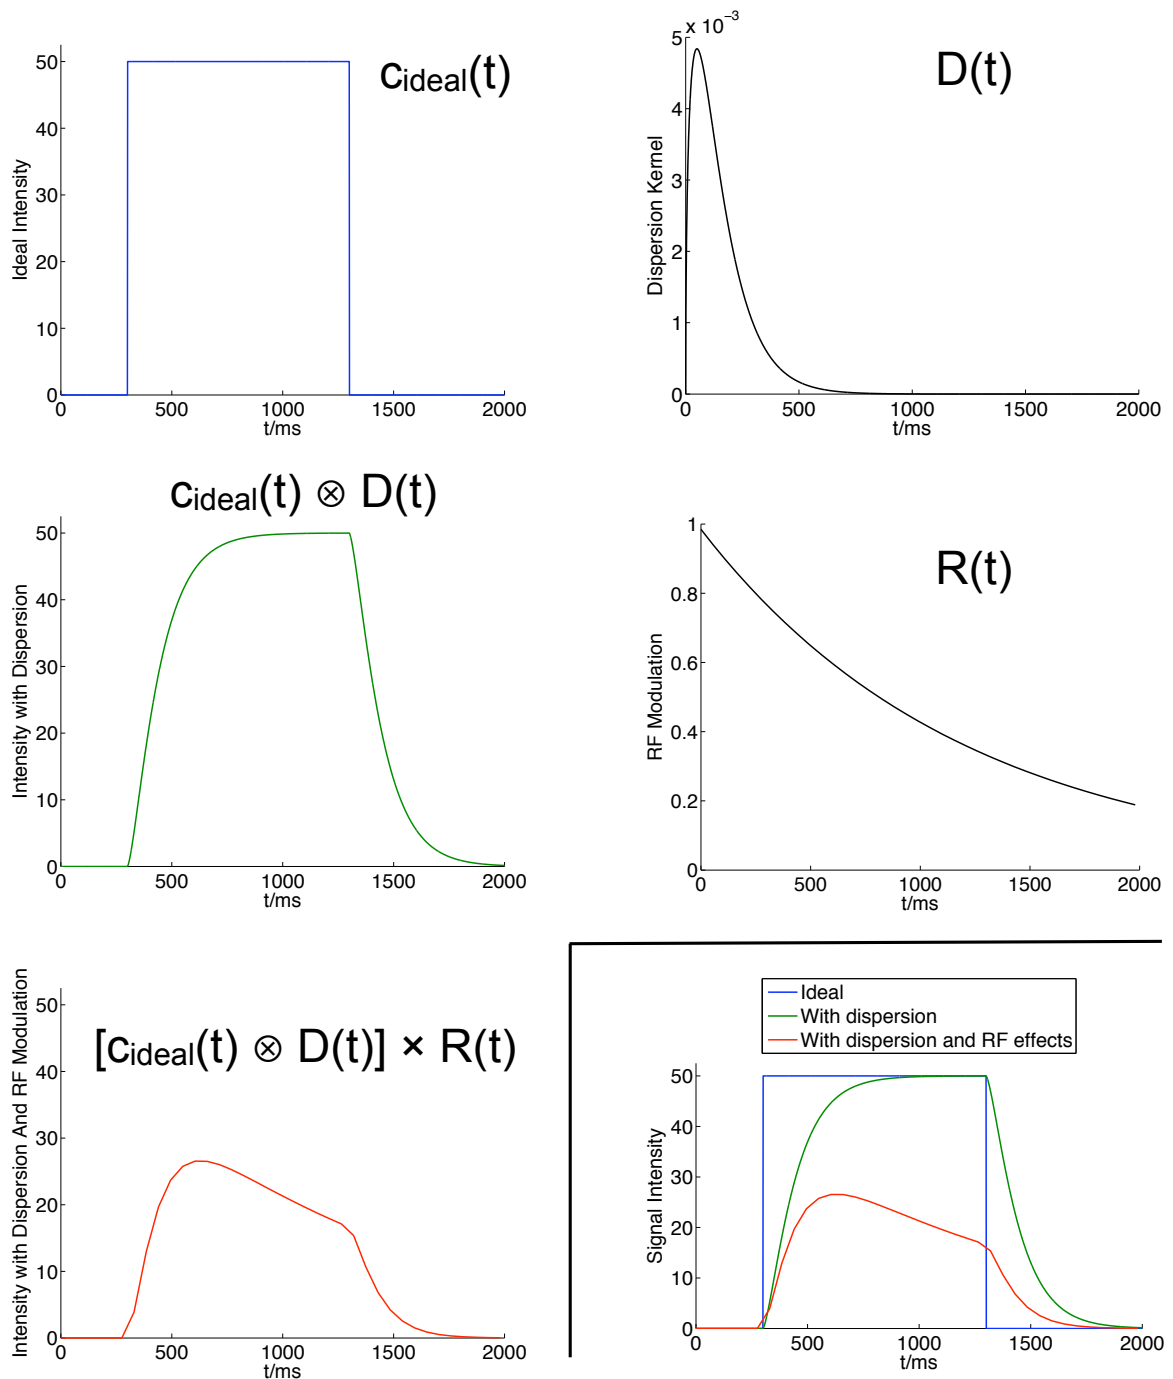


Figure 7.1: Demonstration of how the model curve is modified by the addition of various effects. The ideal rect function (top left), scaled by blood volume, S_0 and the effect of T_1 decay, is convolved with the gamma variate dispersion kernel (top right) and multiplied by the RF modulation function (middle right) to give the final model (bottom left). All three curves are plotted together for comparison (bottom right). In practice, for long labelling durations ($\tau > \delta_t$) only the latter portion of the curve will be observed since imaging cannot commence until the VEPCASL pulse train has finished.

$$S(x_i, y_j, z_k, t_m) = S_0 [\Gamma_I(s(t_m - \delta_t), 1 + ps) - \Gamma_I(s(t_m - \delta_t - \tau), 1 + ps)] e^{-\delta_t/T_{1b}} (\cos \alpha)^N \iiint_{\text{voxel}} \kappa(x, y, z) dx dy dz \quad (7.16)$$

$$= S_0 [\Gamma_I(s(t_m - \delta_t), 1 + ps) - \Gamma_I(s(t_m - \delta_t - \tau), 1 + ps)] e^{-\delta_t/T_{1b}} (\cos \alpha)^N v(x_i, y_j, z_k) \quad (7.17)$$

where $v(x_i, y_j, z_k)$ is the blood volume within vessels in voxel (x_i, y_j, z_k) . For 2D angiography this assumption must hold over the imaging slab used.

In practice it is simpler to combine the factors S_0 and $v(x_i, y_j, z_k)$ into a single scaling factor $A(x_i, y_j, z_k) = S_0 v(x_i, y_j, z_k)$ so our final model to fit to the data is:

$$S(x_i, y_j, z_k, t_m) = A [\Gamma_I(s(t_m - \delta_t), 1 + ps) - \Gamma_I(s(t_m - \delta_t - \tau), 1 + ps)] e^{-\delta_t/T_{1b}} (\cos \alpha)^N \quad (7.18)$$

7.2.2 Calibration

The calibration factor S_0 describes the ASL difference signal per unit volume of labelled blood in the absence of T_1 decay and RF effects. If a 3D angiography acquisition were performed and one or more voxels small enough to fit entirely within an artery could be found then this should be sufficient for the calibration. Fitting Eq. 7.18 to such a voxel would yield an equation where only S_0 is unknown since $v(x_i, y_j, z_k)$ is equal to the voxel volume $\Delta x \Delta y \Delta z$:

$$\hat{S}_0 \approx \left\langle \frac{A}{\Delta x \Delta y \Delta z} \right\rangle_{\text{voxels within arteries}} \quad (7.19)$$

In a 2D acquisition, the volume of blood leading to the measured signal is unknown. However, if it is assumed that the cross section of each vessel is approximately circular and the vessel is running parallel to the imaging plane, then measurements of vessel diameter, \bar{d} , within the imaging plane can be used to estimate the volume of blood producing the measured signal and therefore S_0 . Plotting a profile of A against distance, d , perpendicular to the vessel direction (see Fig. 7.2) allows a fit to both \bar{d} and S_0 simultaneously. The volume of blood contributing to A is $\Delta x \Delta y \Delta z$ where Δz is no longer the voxel dimension, but now the thickness of the blood vessel in the z direction, which can be found from the equation of a circle:

$$(d - d_0)^2 + (\Delta z/2)^2 = (\bar{d}/2)^2 \quad \text{for} \quad |d - d_0| \leq \bar{d}/2 \quad (7.20)$$

where d_0 is the location of the centre of the vessel. Rearranging gives:

$$\Delta z = \begin{cases} \sqrt{\bar{d}^2 - 4(d - d_0)^2} & |d - d_0| \leq \bar{d}/2 \\ 0 & |d - d_0| > \bar{d}/2 \end{cases} \quad (7.21)$$

Therefore the expected variation of A with d , which can be fitted to the data to determine S_0 , is:

$$A(d) = \begin{cases} S_0 \Delta x \Delta y \sqrt{\bar{d}^2 - 4(d - d_0)^2} & |d - d_0| \leq \bar{d}/2 \\ 0 & |d - d_0| > \bar{d}/2 \end{cases} \quad (7.22)$$

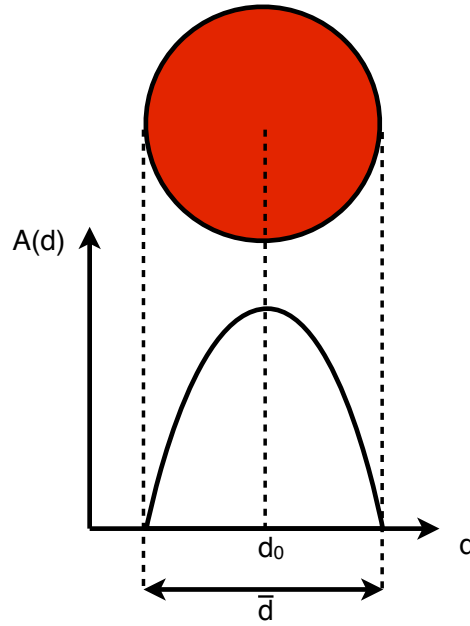


Figure 7.2: Vessel geometry used to estimate the calibration factor, S_0 . Assuming a circular vessel cross-section, A varies parabolically as a function of distance across the vessel, d , with the shape of the curve depending only on S_0 , the position of the vessel centre, d_0 , and the vessel diameter, \bar{d} .

7.2.3 Flow Rate Quantification Within a Vessel Segment

Now, the volume flow rate of blood, F_g , in a given vessel, g , is given by the volume of blood passing through it per unit period of time. Thus, if all of the labelled blood has entered the vessel of interest and is still within the FOV at the time of imaging then we can estimate the volume of blood present, V_g . We also know the time during which this passes into the brain (τ), so can estimate the volume flow rate:

$$F_g = \frac{V_g(\tau)}{\tau} \quad (7.23)$$

$$= \frac{1}{\tau} \iiint_g dx dy dz c(x, y, z, t) \quad (7.24)$$

$$\approx \frac{1}{\tau} \sum_{i,j,k \in g} \frac{S(x_i, y_j, z_k, t_m)}{S_0 e^{-\delta t / T_{1b}} (\cos \alpha)^N} \quad (7.25)$$

Note that the left side of this equation should be independent of the time at which it is measured. Therefore, for any t_m where the labelled bolus is within the vessel of interest but has not moved out of the FOV or started exchanging with tissue, the same volume flow rate should be calculated. Measuring F_g at multiple time points thus provides an opportunity to test the model and calculate a more robust estimate of F_g .

However, to satisfy these assumptions only a very short labelling duration could be used and, even then, if dispersion is significant and the FOV restricted then it may be impossible to catch the whole labelled bolus within the vessel of interest. In addition, as mentioned in Section 6.2.4, using a short labelling duration reduces the SNR and provides poorer visualisation of the vessels. It is therefore desirable to be able to quantify volume blood flow without these restrictions.

If an accurate fit to Eq. 7.18 can be made then it should be possible to simulate the signal, S' , which would arise from a very short labelling duration, τ' , in the absence of T_1 decay and RF effects:

$$S'(x_i, y_j, z_k, t_m) = A [\Gamma_I(s(t_m - \delta_t), 1 + ps) - \Gamma_I(s(t_m - \delta_t - \tau'), 1 + ps)] \quad (7.26)$$

This would effectively allow the apparent use of a very short labelling duration but without the associated SNR penalty. The quantification could then be performed in a straightforward manner:

$$F_g \approx \frac{1}{\tau'} \sum_{i,j,k \in g} \frac{S'(x_i, y_j, z_k, t_m)}{S_0} \quad (7.27)$$

However, in cases where dispersion is very significant or where the vessel segment of interest is small, this may not be sufficient to ensure the entire labelled

bolus is within the vessel and FOV at some point in time. In this case it may be more appropriate to attempt to simulate the bolus in the absence of dispersion (i.e. $s \rightarrow \infty, p = 0$). However, it should be noted that this may introduce some additional errors if the dispersion model is not highly accurate. For example, in 2D angiography there may be voxels where two vessels cross and the dispersion model may provide a reasonable fit to the data where there are in fact two boluses present. When the dispersion is removed from the model, the simulated signal, S' , may no longer well represent the real blood flow.

7.2.4 Bayesian Inference

The fit to Eq. 7.18 in each voxel was performed within a Bayesian framework. This helps to prevent spurious fits to the data where the SNR is poor and allows estimation of the errors on the fitted parameters. Bayes' theory tells us that we can calculate the probability distribution of a set of parameters, Θ , given a set of data, \mathbf{D} [150]:

$$P(\Theta|\mathbf{D}) = \frac{P(\mathbf{D}|\Theta)P(\Theta)}{P(\mathbf{D})} \quad (7.28)$$

where $P(\Theta|\mathbf{D})$ is the posterior distribution of the parameters given the data, $P(\mathbf{D}|\Theta)$ is the likelihood function, describing how likely the data were to be generated by a set of parameters, $P(\Theta)$ is the prior distribution on the parameters, incorporating any previous knowledge, and $P(\mathbf{D})$ is the evidence, which can be treated as a normalising factor since it does not depend on the parameter values.

It is assumed that the measurement of the signal in each voxel at each time point is independent and normally distributed with standard deviation σ . Thus, for a single data point:

$$P(S_m|\Theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\left(\frac{S_m - \hat{S}_m(\Theta)}{\sigma}\right)^2\right) \quad (7.29)$$

where S_m is the measured ASL signal and $\hat{S}_m(\Theta)$ is the model fit to the data using the set of parameters Θ . Since each data point is an independent measurement, the probability of obtaining a set of data, D , with M data points, given some parameter set, Θ , is just the multiplication of the probabilities of obtaining each data point on its own:

$$P(D|\Theta) = \prod_m^M \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\left(\frac{S_m - \hat{S}_m(\Theta)}{\sigma}\right)^2\right) \quad (7.30)$$

$$= \frac{1}{(2\pi)^{M/2}\sigma^M} \exp\left(-\frac{1}{2}\sum_m^M \left(\frac{S_m - \hat{S}_m(\Theta)}{\sigma}\right)^2\right) \quad (7.31)$$

Since we do not know the noise level, σ , beforehand and it may vary over the FOV (e.g. due to pulsatility effects), we also try to estimate it from the data so $\Theta = \{A, \delta_t, s, p, \sigma\}$. For simplicity we define the prior on each parameter in the model to be normally distributed and independent, i.e.:

$$P(\Theta) = P(A)P(\delta_t)P(s)P(p)P(\sigma) \quad (7.32)$$

$$= \frac{1}{\sqrt{2\pi\sigma_A^2}} e^{-\frac{1}{2}\left(\frac{A-\mu_A}{\sigma_A}\right)^2} \frac{1}{\sqrt{2\pi\sigma_{\delta_t}^2}} e^{-\frac{1}{2}\left(\frac{\delta_t-\mu_{\delta_t}}{\sigma_{\delta_t}}\right)^2} \dots \quad (7.33)$$

where μ_X and σ_X represent the mean and standard deviation on parameter X , respectively. The priors used for the fitting of the model to 2D GRE VEP-CASL dynamic angiography data are given in Table 7.1. It should be noted that these values were determined empirically from initial experience with the fitting

<i>Parameter</i>	<i>Mean</i>	<i>Standard Deviation</i>	<i>Units</i>
A	0	30	Arbitrary
δ_t	800	300	ms
s	0.01	0.002	ms ⁻¹
p	10	10	ms
σ	2	0.5	Arbitrary

Table 7.1: Chosen values for the prior distributions on each parameter based on some initial experience fitting to 2D GRE VEPCASL angiography data. Note that the prior on A has a mean of zero to prevent large values being estimated in noisy voxels which could corrupt later attempts at blood flow rate quantification. s is also tightly constrained to discourage very small sharpness fits which approximate a flat curve in noisy voxels. δ_t is given a large mean value since voxels with high SNR are generally those with lower δ_t (due to reduced T_1 decay) and so the data can drive the parameter estimate down to the correct value. Voxels with poorer SNR often have longer arrival times and therefore the prior drives the fitting more strongly so needs to be closer to the expected value. Note that the priors on parameters A and σ depend on the scaling of the images and thus will vary between scanners.

procedure. In addition, Gaussian priors allow for potentially negative parameter values which are non-physical (although the fitting procedure is constrained to positive parameter values). More realistic and robust priors which do not rely on empirical assessments would benefit this method in future work.

Since we are only interested in the shape of the posterior distribution and location of its maximum (i.e. the best fit to the data and priors), and not the absolute value of the posterior probability, we can discard the terms which do not depend on the parameters (since these are just scaling factors) to give:

$$P(\Theta|\mathbf{D}) \propto P(\mathbf{D}|\Theta)P(\Theta) \quad (7.34)$$

$$\propto \frac{1}{\sigma^M} e^{-\frac{1}{2} \left[\left(\sum_m^M \left(\frac{S_m - \hat{S}_m(\Theta)}{\sigma} \right)^2 \right) + \left(\frac{A - \mu_A}{\sigma_A} \right)^2 + \left(\frac{\delta_t - \mu_{\delta_t}}{\sigma_{\delta_t}} \right)^2 + \left(\frac{s - \mu_s}{\sigma_s} \right)^2 + \left(\frac{p - \mu_p}{\sigma_p} \right)^2 + \left(\frac{\sigma - \mu_\sigma}{\sigma_\sigma} \right)^2 \right]} \quad (7.35)$$

To find the optimal solution we want to maximise the posterior probability of the parameters given the data which is equivalent to minimising the negative logarithm of the posterior probability distribution, J :

$$J = -\ln(P(\Theta|\mathbf{D})) = M \ln \sigma + \frac{1}{2} \left[\left(\sum_m \left(\frac{S_m - \hat{S}_m(\Theta)}{\sigma} \right)^2 \right) + \left(\frac{A - \mu_A}{\sigma_A} \right)^2 + \left(\frac{\delta_t - \mu_{\delta_t}}{\sigma_{\delta_t}} \right)^2 + \left(\frac{s - \mu_s}{\sigma_s} \right)^2 + \left(\frac{p - \mu_p}{\sigma_p} \right)^2 + \left(\frac{\sigma - \mu_\sigma}{\sigma_\sigma} \right)^2 \right] \quad (7.36)$$

Using this formulation it is also possible to extract estimates of the errors on the parameters. We make the Laplace approximation [150] that the posterior distribution on the parameters can be approximately described by a multi-variate normal (MVN) distribution, i.e.:

$$P(\Theta|\mathbf{D}) \approx \frac{1}{\sqrt{(2\pi)^k |\Sigma|}} \exp \left(-\frac{1}{2} (\Theta - \Theta^*)^T \Sigma^{-1} (\Theta - \Theta^*) \right) \quad (7.37)$$

where k is the dimensionality of the parameter space, Σ is the covariance matrix which describes the confidence in the fitted parameters and Θ^* is the best estimate of the parameters. It has been shown that the Hessian matrix, \mathbf{H} , evaluated at the best parameter fit, Θ^* , is equal to the inverse of the covariance matrix [151]:

$$\mathbf{H}_{l,l'}(\Theta^*) = \left. \frac{\partial^2 J}{\partial \Theta_l \partial \Theta_{l'}} \right|_{\Theta^*} = (\Sigma^{-1})_{l,l'} \quad (7.38)$$

Therefore, if a numerical estimate of the Hessian matrix at the solution can be obtained, then the uncertainty in each parameter can be determined. In addition, various quantities which rely on the parameter estimates, such as the mean and variance of the simulated signal, can be calculated by numerically integrating over the MVN distribution in Eq. 7.37, weighted by the quantity of interest.

7.3 Methods

The proposed method was applied in a normal volunteer and in a patient with significant unilateral vertebral artery stenosis (part of the study described in Section 8.2) to obtain estimated volume flow rates in the major cerebral arteries. Six cycle 2D GRE VEPCASL angiography images in both transverse and coronal planes were processed as described in Section 4.11.4 and the maximum *a posteriori* (MAP) method used to separate the vascular components (see Section 4.11.4).

The output of the MAP processing is a complex signal value in each voxel at each time point for each vascular component. However, the model described above assumes purely real data. Simply taking the magnitude of the signal could skew the model fit, particularly when the signal is low and noise dominates, since the signal will then always be positive. Therefore a phase correction step was performed in pre-processing: the mean phase over time in each voxel for each vascular component was calculated after weighting by the magnitude of the signal, since the phase of a low magnitude signal is less reliable. This mean phase was subtracted from the phase at each time point before taking the real part of the signal. It is assumed that the phase of the blood signal does not vary significantly over time and any phase variation is due to noise alone. This should be a reasonable assumption since each time frame is acquired a few tens of milliseconds after the previous frame so the phase variation due to external sources such as B_0 drift should be minimal. This was confirmed by comparisons of phase corrected and magnitude images which show negligible differences when not in the noise floor.

A mask was used to exclude background tissue by calculating the mean signal intensity over time in each voxel, determining the 99th percentile of this value across the whole image and then excluding voxels which have mean signal of less than 30% of this value. This speeds up the fitting procedure and prevents fitting the model to pure noise which may, on occasion, produce spurious parameter

values which could corrupt subsequent attempts at quantification.

The derived theoretical model (Eq. 7.18) was fitted to the time series in each voxel within the mask by minimisation of the negative log posterior probability (Eq. 7.36) in Matlab[®] (Mathworks[™], Natick, Massachusetts, USA). The Hessian matrix at the optimal solution was also calculated to allow estimation of the covariance matrix of the parameters, Σ . The calibration factor, S_0 , was determined by fitting Eq. 7.22 to a profile through ten vessel segments. These segments were chosen using a time-of-flight angiogram taken in the same session to ensure that the vessels were running approximately parallel to the imaging plane. S_0 was then taken as the average of ten such measurements.

Volume flow rate quantification was performed by simulating a very short bolus of duration ($\tau' = 1$ ms) in the absence of T_1 decay and RF effects (using Eq. 7.26). The mean and variance values of this signal at a number of simulated time points were calculated by numerical integration over the estimated MVN distribution (Eq. 7.37). For comparison, the quantification was also performed without dispersion by setting $s \rightarrow \infty$ and $p = 0$ everywhere. When simulating a very short bolus length and no dispersion, only a few voxels will contribute to the estimated volume flow rate at a given time, making this process very noisy and prone to error. Increasing the simulated bolus length mediates this problem to some degree so τ' was set to 50 ms when dispersion was ignored. It is also noted that when ignoring the effects of dispersion the simulated bolus occupies much shorter sections of the arteries, so a larger τ' can be used without danger of the simulated bolus becoming larger than the vessel segment of interest.

The estimated volume flow rate within a given mask is expected to plateau during the time that the bolus is entirely within the mask. To estimate this plateau region the approximate start and end of the whole bolus were determined by thresholding the estimated volume flow rate at 20% of its maximum. Assum-

ing the plateau occupies a significant proportion of this broad peak, the median value within it should be approximately equal to the volume flow rate within the plateau region and relatively insensitive to noisy peaks and troughs in the curve. The plateau region was defined as the interval bounded by the times at which the estimated volume flow rate first and last exceeds this median value. The final estimated flow rate and its error were calculated as the mean and standard deviation within this plateau region.

Quantification was performed within a number of masks which were drawn manually for each data set (examples shown in Fig. 7.3). Note that the anterior cerebral arteries (transverse view) and posterior cerebral arteries (coronal view) were not analysed as they quickly move out of the imaging plane, rendering attempts at quantification inaccurate. In addition, the vertebral and basilar arteries were analysed together to ensure the mask was large enough to encompass the simulated bolus. In order to demonstrate the potential error due to violations of the assumption that all the blood has experienced all of the previous RF pulses, two masks were used for each vessel in the transverse view. The first was a larger mask that encompassed the proximal segments of the artery and the second was a smaller mask which included only the distal segments.

7.4 Preliminary Results

The theoretical model derived above appears to fit well to the acquired data (see examples in Fig. 7.4), except a slight deviation at early time points in proximal vessel segments. Fitting this model to the entire 2D transverse image from a healthy volunteer produces the parameter maps shown in Fig. 7.5. The expected patterns can be seen here: high A values in larger vessels which contain more blood volume, smaller δ_t in proximal vessel segments and near the centre of large




Figure 7.3: Manually drawn masks used for flow rate quantification in the major cerebral arteries overlaid on a map of parameter A in a healthy volunteer in both transverse and coronal views: right and left middle cerebral arteries (yellow and purple), right and left posterior cerebral arteries (cyan and orange), anterior cerebral arteries (brown) combined due to their proximity, right and left internal carotid arteries (red and green) and combined vertebral and basilar arteries (blue). In the transverse view, each mask is split into a larger mask (darker colour) and a smaller mask (brighter colour), with the latter excluding the proximal artery segment.

vessels where the blood arrives sooner, and lower s and higher p values in distal vessels where the bolus has had more time to disperse.

The calibration method described above appeared to work well for these data, although the fit to Eq. 7.22 was imperfect at the edges of the vessel (Fig. 7.6). The derived calibration factor, S_0 , was consistent both between transverse and coronal views, as well as between subjects. Note that the data have already been corrected for coil sensitivity profiles at this point, otherwise S_0 would vary with position.

The quantification performed in a healthy volunteer in the transverse view is shown in Fig. 7.7. To clarify here, these graphs show the estimated volume flow rate as a function of simulated time. This flow rate is calculated by integrating the total simulated signal (which excludes T_1 decay and RF effects) over the specified vessel mask for each simulated time point and dividing this value by S_0 to give the estimated blood volume within the mask at the specified time. This value is then divided by the simulated bolus duration, τ' , to give the estimated volume flow rate at each time point. It is expected that if the quantification is performing




Figure 7.4: Examples of the model fit to the data in a healthy volunteer using the transverse view. Two locations in the proximal (circle) and distal (square) right MCA are overlaid on the colour coded map of A (left). The corresponding time series and model fit at these locations to the right ICA component are also shown (right). The time axis is relative to the start of the VEPCASL pulse train.

well, then for the time during which the simulated bolus lies entirely within the mask the estimated volume flow rate should be constant (i.e. plateau) at its true value. At earlier or later time points when only a fraction of the simulated bolus lies within the vessel mask, the estimated volume flow rate is expected to be smaller than its true value.

However, from Fig. 7.7, it can be seen that the estimated volume flow rate does not always plateau as a function of time in the larger masks which include proximal vessel segments, particularly in the MCAs where the flow is faster. The plateau region is clearer in the smaller, distal masks, although the estimated volume flow rate is reduced. Results without dispersion are significantly less stable and have much greater errors associated with them, although the mean volume flow rate within the plateau regions is generally comparable. Due to this instability only results with dispersion are considered further here. Contamination from

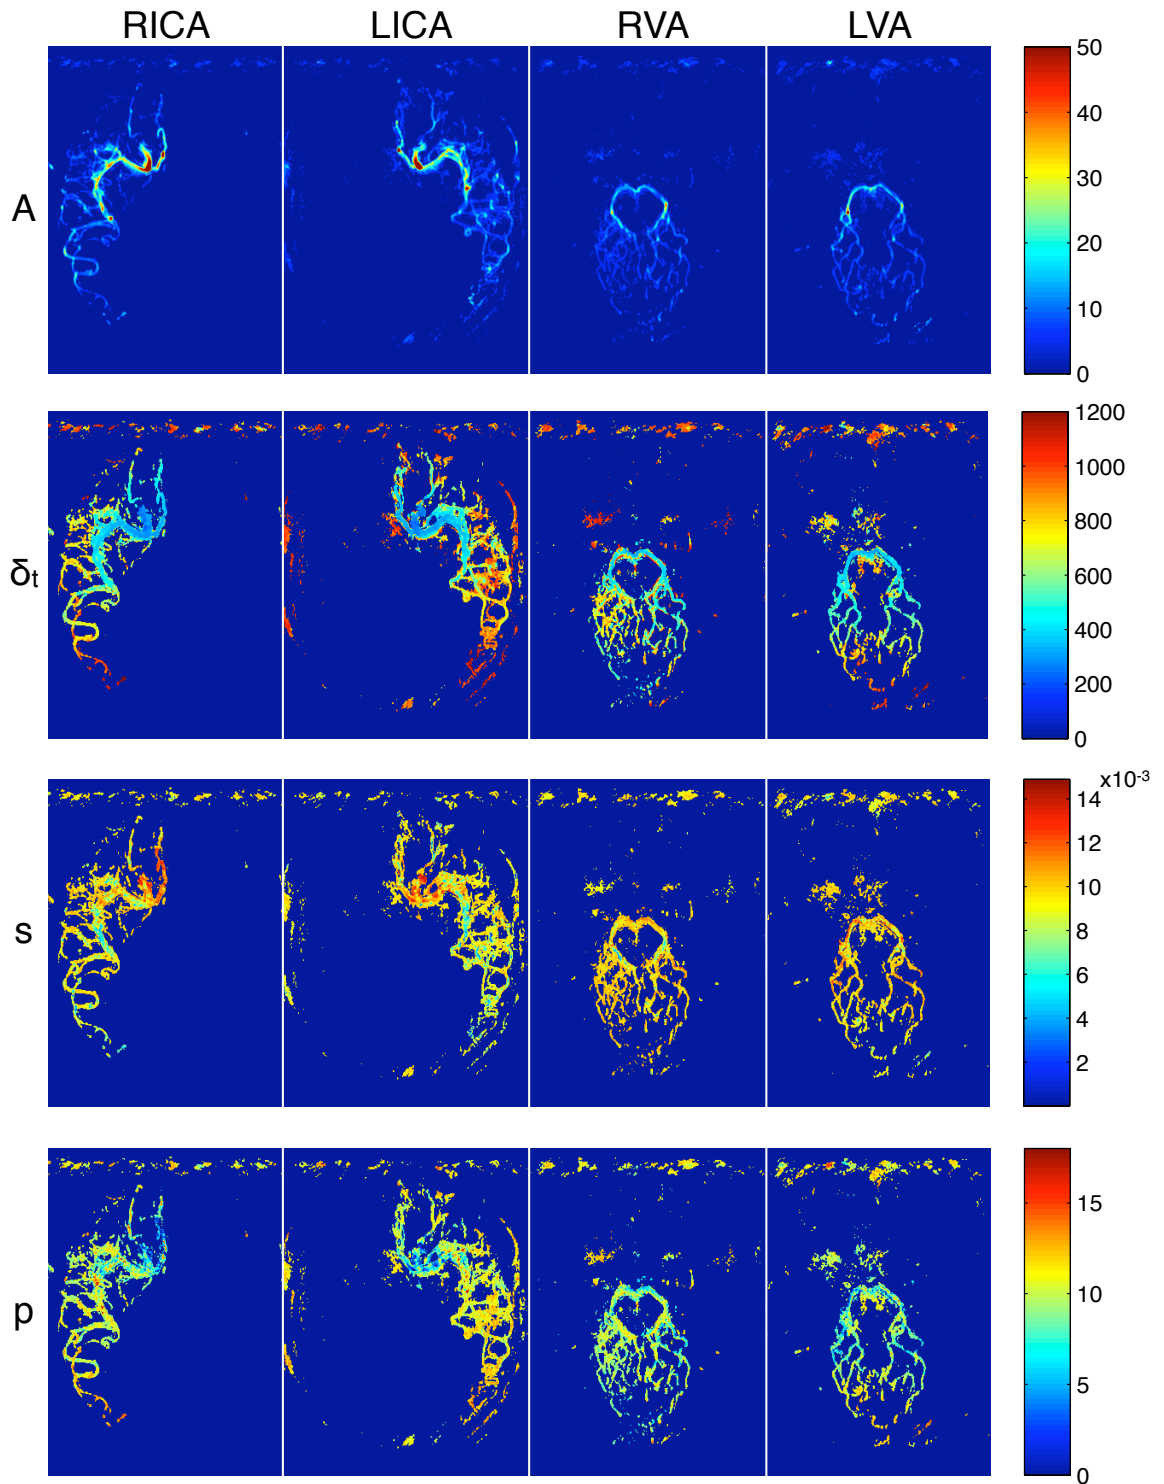


Figure 7.5: Example parameter maps generated by fitting the model in Eq. 7.18 to 2D GRE VEP-CASL dynamic angiography data from a healthy volunteer after separation of the vascular components. Units are arbitrary (A), ms (δ_t and p) or ms^{-1} (s).




Figure 7.6: Calibration of the dynamic angiography signal: profile positions used for the calibration overlaid on a map of parameter A in a healthy volunteer, transverse view (left); plot of A against distance across the profile, d , for the green line in the overlay alongside the fit to Eq. 7.22 (middle) which yielded the following parameter estimates: $S_0 = 30 \pm 4 \text{ mm}^{-3}$, $d_0 = 2.78 \pm 0.08 \text{ mm}$, $\bar{d} = 2.9 \pm 0.2 \text{ mm}$; box and whisker plot of the estimated S_0 values in transverse (TRA) and coronal (COR) views in both the healthy volunteer (H) and patient (P) data sets (right).

vascular components other than the expected feeding artery is minimal, showing that the separation of vascular components using the MAP method (see Section 4.11.4) is highly effective.

The quantification from the coronal view in the same subject in the same session (Fig. 7.8) appears to give a more stable quantification over time in the output vessels from the circle of Willis (MCAs and ACAs). The input vessels (ICAs and VAs) can also be quantified in this view, but the fast passage of blood through these vessels makes it harder to estimate a plateau region. Measurements in the VAs are also complicated by the outflow of blood to the cerebellar arteries, meaning that a percentage of the labelled bolus leaves the ROI soon after entering it.

A summary of the estimated volume flow rates within each input and output artery for the healthy volunteer are shown in Fig. 7.9. To check for consistency within the data, Gaussian noise was assumed and two-tailed z -tests were per-

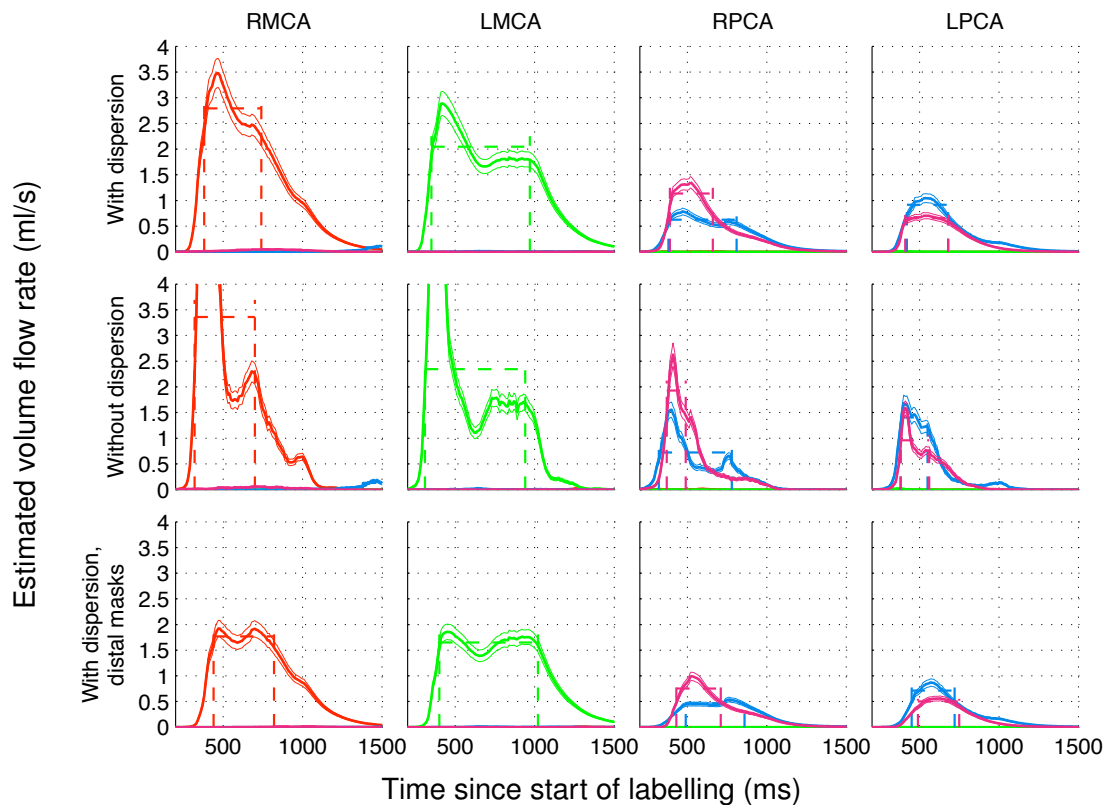


Figure 7.7: Volume flow rate quantification in a healthy volunteer using the transverse view, with (top and bottom rows) and without (middle row) inclusion of dispersion effects. The same quantification performed within the smaller distal masks is also shown (bottom row). For each vessel mask the estimated volume flow rate (thick lines) and associated error (thin lines) at each simulated time point are shown. Dashed lines represent the start and end of the estimated plateau region and the mean volume flow rate within it. Line colour represents the vascular component being analysed, in keeping with previous figures (e.g. red = right ICA).

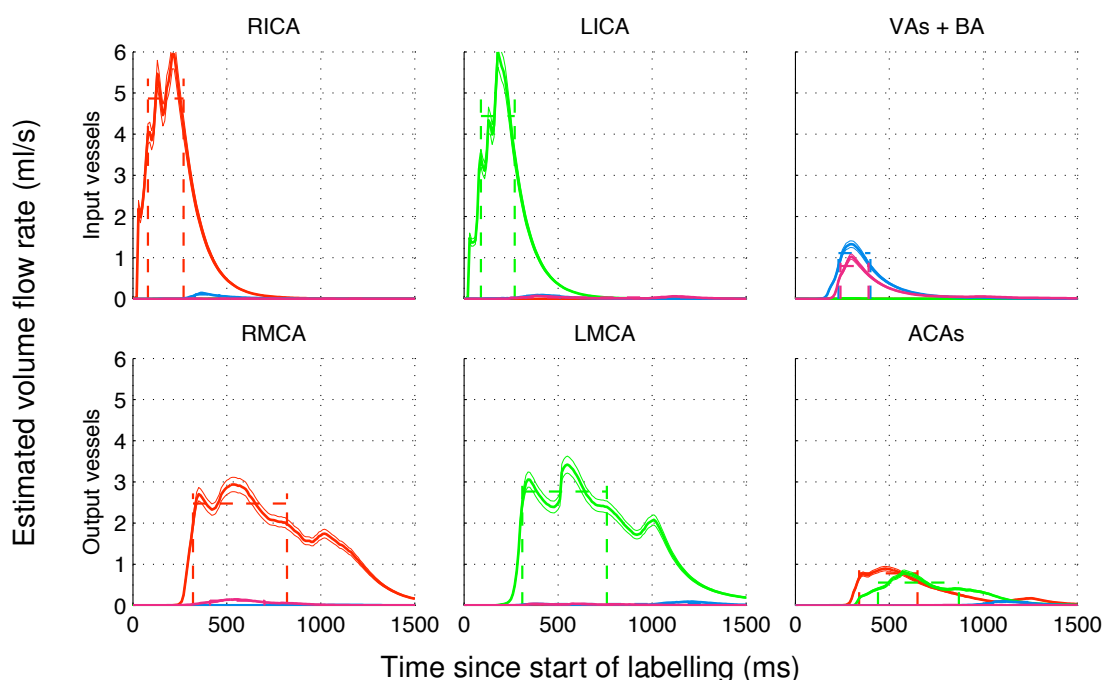


Figure 7.8: Volume flow rate quantification in a healthy volunteer using the coronal view with dispersion included. Input vessels to the circle of Willis can be quantified in this view (upper row) along with the output vessels (lower row).

formed at the 1% significance level. Using this approach it was found that flow rates from the main feeding arteries are consistent between transverse and coronal views and the total flow rate in output vessels is consistent with the flow rate for the relevant input vessel, with the exception of the left vertebral artery. Here the flow was more dispersed so the simulated bolus was not entirely captured by the mask. Ignoring the effects of dispersion allowed the simulated bolus to be captured entirely within the vertebral and basilar artery mask, but gave less stable results, as mentioned previously. Estimated flow rates were generally consistent with a previous study: van Osch *et al.* [71] used an ASL angiographic method, validated against phase-contrast measurements, to calculate volume flow rates in the major cerebral arteries. They estimated (in ml/s) 2.56 ± 0.48 in each MCA, 1.20 ± 0.14 in each PCA, 3.16 ± 0.36 for both ACAs combined, 8.3 ± 1.4 for both ICAs combined (i.e. approximately 4.2 ± 0.7 for each ICA) and 2.40 ± 0.34 in




Figure 7.9: Summary of quantification results in a healthy volunteer. All results were generated using dispersion except the input vessels in the coronal view (right), for comparison.

the BA. Using the z -test described above, only their ACA result is significantly different from the values presented here.

Similar results were obtained from the patient data set (Fig. 7.10). The model fit appears better in proximal vessels than comparable fits in the healthy volunteer data set, as does the quantification in the MCAs using the transverse view. Very limited flow in the LVA and clear collateral flow from ICAs to PCAs are evident from these results (Fig. 7.11), although there is a little more contamination of the signals (e.g. contributions to the LMCA from the RICA and both VAs in the coronal view) due to background noise. Also, reduced flow in the LPCA, where the degree of collateral flow is much lower, could account for the patient's symptoms which indicated TIA/minor stroke in the posterior of the brain.

7.5 Discussion

In general, the fit of the theoretical model to the time course in each voxel was good, with the resulting parameter maps having the expected appearance over the cerebral vasculature. The deviation at early time points in proximal vessel

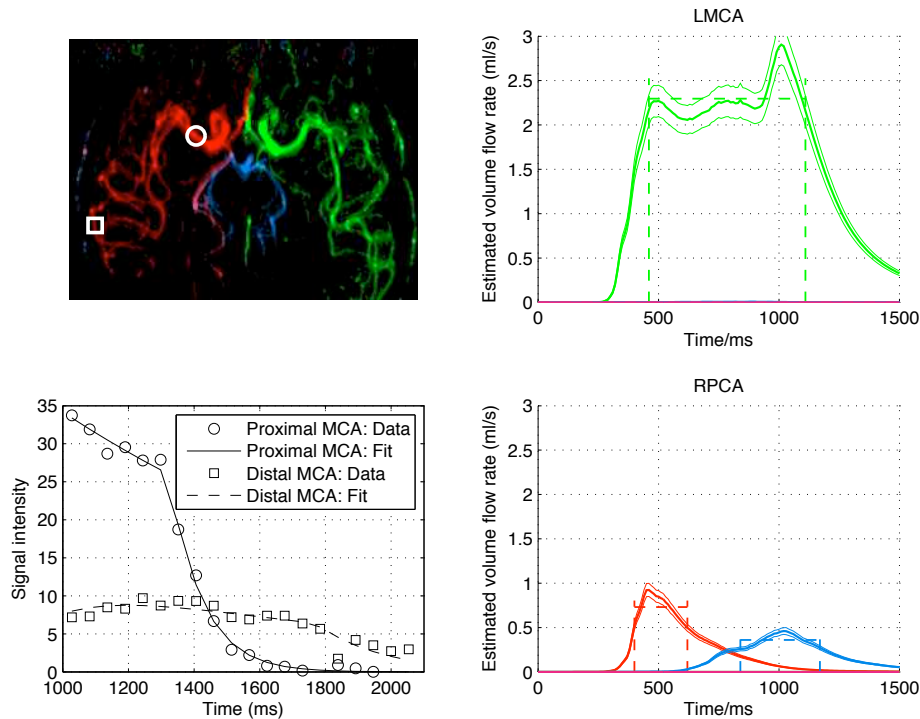


Figure 7.10: Quantification in a patient with vertebral stenosis (transverse view). Locations in the proximal (circle) and distal (square) right MCA are overlaid on the map of A (left,top). The corresponding time series and model fit at these locations to the right ICA component are also shown (left,bottom). The time axis is relative to the start of the VEPCASL pulse train. Estimated volume flow rates over time are shown for the LMCA (right,top) and RPCA (right,bottom).




Figure 7.11: Summary of quantification results in a patient with vertebral stenosis. All results were generated using dispersion except the input vessels in the coronal view (right), for comparison.

segments in the transverse view is likely to be due to violation of the assumption that all of the labelled blood has experienced all the previous RF pulses. For example, the plot from the proximal MCA location shown in Fig. 7.4 is approximately flat after the first time point before falling off as the bolus washes out. This implies that the blood arriving at this point has only experienced a finite number of RF pulses, so while the bolus of labelled blood remains within the voxel the signal is approximately constant. This effect is much less significant in distal vessel segments, coronal views, and also in the patient data set, where the flow is slower and thus the blood is within the imaging volume for a longer time period before reaching the proximal MCA segment.

This is likely to be the cause of the non-plateau like behaviour of the MCAs in the transverse view shown in Fig. 7.7. The false assumption that the magnetisation has experienced all the previous RF pulses will artificially inflate the scaling factor, A , in proximal vessel segments, as well as potentially biasing the arrival time, δ_t , leading to an overestimate of the volume flow rate while the bolus is still within the proximal vessels. This is seen as the initial peak before the plateau in the estimated flow rate. When the smaller distal mask is used, this initial peak disappears, giving the expected plateau behaviour. It is likely that this effect is also the cause of the cleaner plateau region in the patient data set (Fig. 7.10) where the blood flow is slower, as mentioned above. However, excluding proximal vessel segments is not ideal, as blood leaving the mask earlier (e.g. to the ophthalmic arteries) will be ignored, potentially leading to an underestimated volume flow rate. Modelling the finite RF history effect in future work should provide a more robust estimation of flow rates in proximal vessels where the flow is faster.

The calibration process (exemplified in Fig. 7.6) appears to work well in the two data sets presented here. The fit to the expected variation (Eq. 7.22) is poor at the edges of the profile, which can be attributed to partial voluming and blur-

ring in the readout direction due to the partial Fourier (asymmetric echo) readout method used. However, when averaging over ten such profiles the derived calibration factor is encouragingly consistent between subjects and image orientations, allowing some confidence that this procedure is reasonably robust, despite the imperfect fit to the model. Studies using much larger numbers of data sets are required to confirm this.

Estimated volume flow rates when dispersion effects are ignored were highly variable. This is likely to be due to the dispersion kernel absorbing some uncertainties in the fitting process. For example, a voxel which contains blood flowing at different velocities, or blood flowing perpendicular to the imaging plane, will have a range of arrival times associated with it. The dispersion kernel will model this range of arrival times as a blurring of the resulting signal. If dispersion is ignored, a single arrival time is imposed on all of these different signals, leading to sharp peaks and troughs in the estimated flow rate when the simulation time falls within or outside the single value for blood arrival time in high signal voxels. However, it is anticipated that in certain situations where it is difficult to draw a mask which encompasses the whole dispersed bolus, this will provide at least some estimate of volume flow rate, albeit with a large error associated with it. With more sophisticated modelling or extra processing stages this approach may become more reliable.

Low estimated flow rates in the vertebral and basilar arteries are due to the dispersed simulated bolus exceeding the size of the mask which can be drawn for these arteries. Blood running to the cerebellum from the VAs also makes the quantification difficult, although the broad estimate achieved should be approximately representative of the average signal in the VAs. This highlights one of the limitations of this approach, in which the mask must be larger than the dispersed bolus for accurate quantification. Potential solutions include ignoring dispersion

effects, at the cost of accuracy, or labelling in a more proximal region, allowing a larger mask to be drawn without impinging upon the labelling plane. Alternatively, a further step in the modelling process could be undertaken, in a similar manner to [71], in which the summed signal within a smaller vessel mask is modelled to estimate the total transit time of blood within the mask and flow rate simultaneously. However, this extra level of complexity could introduce further errors and sensitivity to violation of the assumptions made by the model.

The lower than expected anterior cerebral artery flow rates could also be caused by the dispersed bolus not fitting completely within the vessel mask, or because these narrow vessels tend to be more heavily eroded by the masking procedure than the wider MCAs and PCAs. Alternatively, portions of the ACAs may be outside the field of view used. In future work, a field of view which covers the entire brain would provide reduced susceptibility to such an effect, but at the cost of SNR since there is more static tissue to provide background noise. It would also be advantageous to better model the background noise so less restrictive vessel masks could be used.

It should also be noted that pulsatility effects have been ignored in this work. As noted in Chapter 6, the lack of pulsatility artefacts in the data suggests that the pulsatility is minimal in the cerebral vessels. The data acquisition occurs over a relatively long timescale (about 10 minutes), which is likely to have an averaging effect over the cardiac cycle, further reducing any remaining pulsatility effects. However, in cases where the vessels are more pulsatile and cardiac gating is required, modification of the described approach will be necessary to ensure accurate flow rates can be calculated.

Finally, the method used to identify the plateau region in this chapter is empirical and appears to work reasonably well in the cases shown here. However, it is likely that this simple method may fail in more complicated situations. It would

therefore be desirable to find a more robust way of determining this plateau region in further work, such as fitting a function to the estimated volume flow rate curve against simulated time.

7.6 Conclusions

The quantification approach taken here appears to give some promising results, consistent with previous publications and between scans in different orientations. The ability to quantify volume flow rates provides a significant advantage over qualitative x-ray techniques and allows comparison between subjects or centres. However, there are a number of difficulties associated with this methodology, including the requirement for a relatively large manually drawn mask in which the quantification is performed and inaccuracies in the modelling of the signal time-course in proximal vessel segments. It is hoped these difficulties can be overcome in future work to provide a more robust process for flow rate quantification.

Chapter 8

Patient Studies

The vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL) techniques developed in previous chapters show potential to visualise the degree and source of collateral flow in patient groups. However, it is important to establish that these techniques firstly can perform adequately in patients where subject motion tends to be more severe, and secondly that useful clinical information is obtained from the resulting images.

Previous studies have shown that vessel-selective ASL techniques provide useful information in patient groups [152]. For example, vessel-specific approaches using PASL have been shown to allow visualisation of collateral flow around the circle of Willis [153], have shown extra-cranial contributions to brain perfusion in patients with internal carotid artery occlusion [154], have been applied in patients who have had extra-intracranial bypass surgery or severe stenoses [155], and been shown to provide comparable information to x-ray DSA in the assessment of collateral flow when combined with time-of-flight (TOF) magnetic resonance angiography (MRA) [91].

Vessel-encoded pseudo-continuous ASL (VEPCASL) has also been used in patient studies, showing the reduction in collateral flow following endarterectomy

[100], giving reasonable agreement with x-ray DSA and demonstrating the ability to detect the presence of leptomeningeal collateral pathways [99].

In this chapter preliminary results are presented from a number of on-going patient studies where the VEPCASL techniques developed for this thesis have been applied. General methods for these studies are given in Section 8.1, followed by some example data from studies of patients with vertebro-basilar disease (Section 8.2), carotid disease (Section 8.3) and arteriovenous malformation (Section 8.4).

8.1 General Methods

For vessel-selective tissue perfusion imaging, the optimised 2D multi-slice echo planar imaging (EPI) readout with partial Fourier described in Chapter 5 was used. The only differences to the protocol given in Table 5.1 are the reduction of the slice thickness to 4 mm, with a slice gap of 1 mm, and use of 24 slices to achieve better slice resolution and to match previous ASL protocols in patient groups. As discussed in Section 5.2.3 the optimum post labelling delay (PLD), which must allow enough time for the labelled blood water to reach the brain and exchange into tissue, is around 1000 ms in healthy volunteers. However, in patient groups where there is likely to be delayed transit of blood to the brain, the optimum PLD may be higher. Initial studies using a PLD of 1500 ms yielded vascular territory maps with poor SNR. This could potentially be counterbalanced by acquiring a greater number of averages, but at the cost of increasing the total scan time, which is not desirable in already busy clinical protocols.

Therefore, in these initial studies the PLD was maintained at 1000 ms, meaning that the relative intensities of voxels in the resulting images may not truly represent the relative difference in regional cerebral blood flow (CBF) in cases

where blood transit is delayed. However, provided the transit time to the tissue is less than the VEPCASL tag duration plus the PLD (i.e. 2400 ms for the standard parameters given in Chapter 5), then some ASL signal should still be detected. Thus, except in cases of extremely delayed transit, the identification of the arteries responsible for supplying different regions of the brain should be possible, even if the relative signals are not truly representative of their contributions to the regional CBF.

At the time of the onset of these studies, the balanced steady-state free precession (bSSFP) method for VEPCASL angiography (see Section 6.3) was still under development. Therefore for angiographic studies the implementation using a spoiled gradient echo readout (see Section 6.2) was used. Early experiments found that the transverse and coronal views provided sufficient information on collateral flow in these subjects; the overlap of vessels in the sagittal view makes their interpretation more difficult, so this slab orientation was only used in a few specific cases.

8.2 Vertebro-basilar Disease

This ongoing project involves the study of patients with a history of transient ischaemic attack (TIA) or minor stroke, particularly those with symptoms suggestive of disease of the vertebral or basilar arteries. Arterial spin labelling (ASL) techniques are used to assess abnormal regional perfusion levels in these patients and any changes which result from endovascular procedures (such as angioplasty or stenting) where possible. This study is being undertaken in collaboration with Dr. Ursula Schulz (patient recruitment and clinical input) and Dr. Jingyi Xie (physics support for non-selective ASL techniques).

The study protocol consists of T_1 weighted structural scans to aid registra-

tion, non-selective pulsed ASL using the flow-sensitive alternating inversion recovery (FAIR) technique [43] with a 3D-GRASE readout (to maintain consistency with previous studies) at multiple inversion times plus calibration images for CBF quantification, 3D multi-slab time-of-flight (TOF) angiography for labelling plane selection and vessel localisation, eight-cycle VEPCASL tissue perfusion imaging and six-cycle VEPCASL dynamic angiography in transverse and coronal views, centred on the circle of Willis. The total scan time for this protocol is approximately 45 minutes.

Fig. 8.1 shows example VEPCASL perfusion and dynamic angiography images in a patient scanned for this study. The TOF data show a stenosed LVA (abbreviations are defined in Fig. 2.1) and significant signal in the right posterior communicating artery, but it is not clear from this alone the extent to which each vessel contributes to blood flow into the brain. The reduced flow in the LVA is evident in both the VEPCASL perfusion and angiography images. The vascular territory maps reveal that the LVA supplies only a small region of the cerebellum, but very little above the circle of Willis. This is confirmed by the angiography data, with no detectable LVA signal at the level of the circle of Willis shown in the transverse view, and a highly stenosed LVA evident in the coronal view. Note that the PLD used here may be insufficient to show the cerebellar perfusion clearly with these highly stenosed vessels.

Collateral flow around the circle of Willis is clear in the VEPCASL images of this patient, providing unambiguous information on flow directionality and arterial source. A high level of collateral flow from the RICA through the right posterior communicating artery into the RPCA territory is visible in both the perfusion and angiography images. In addition, the dynamic nature of the angiography data reveals significantly delayed arrival of blood from the VAs, suggestive of upstream disease. A lower level of collateral flow from the LICA to the LPCA is

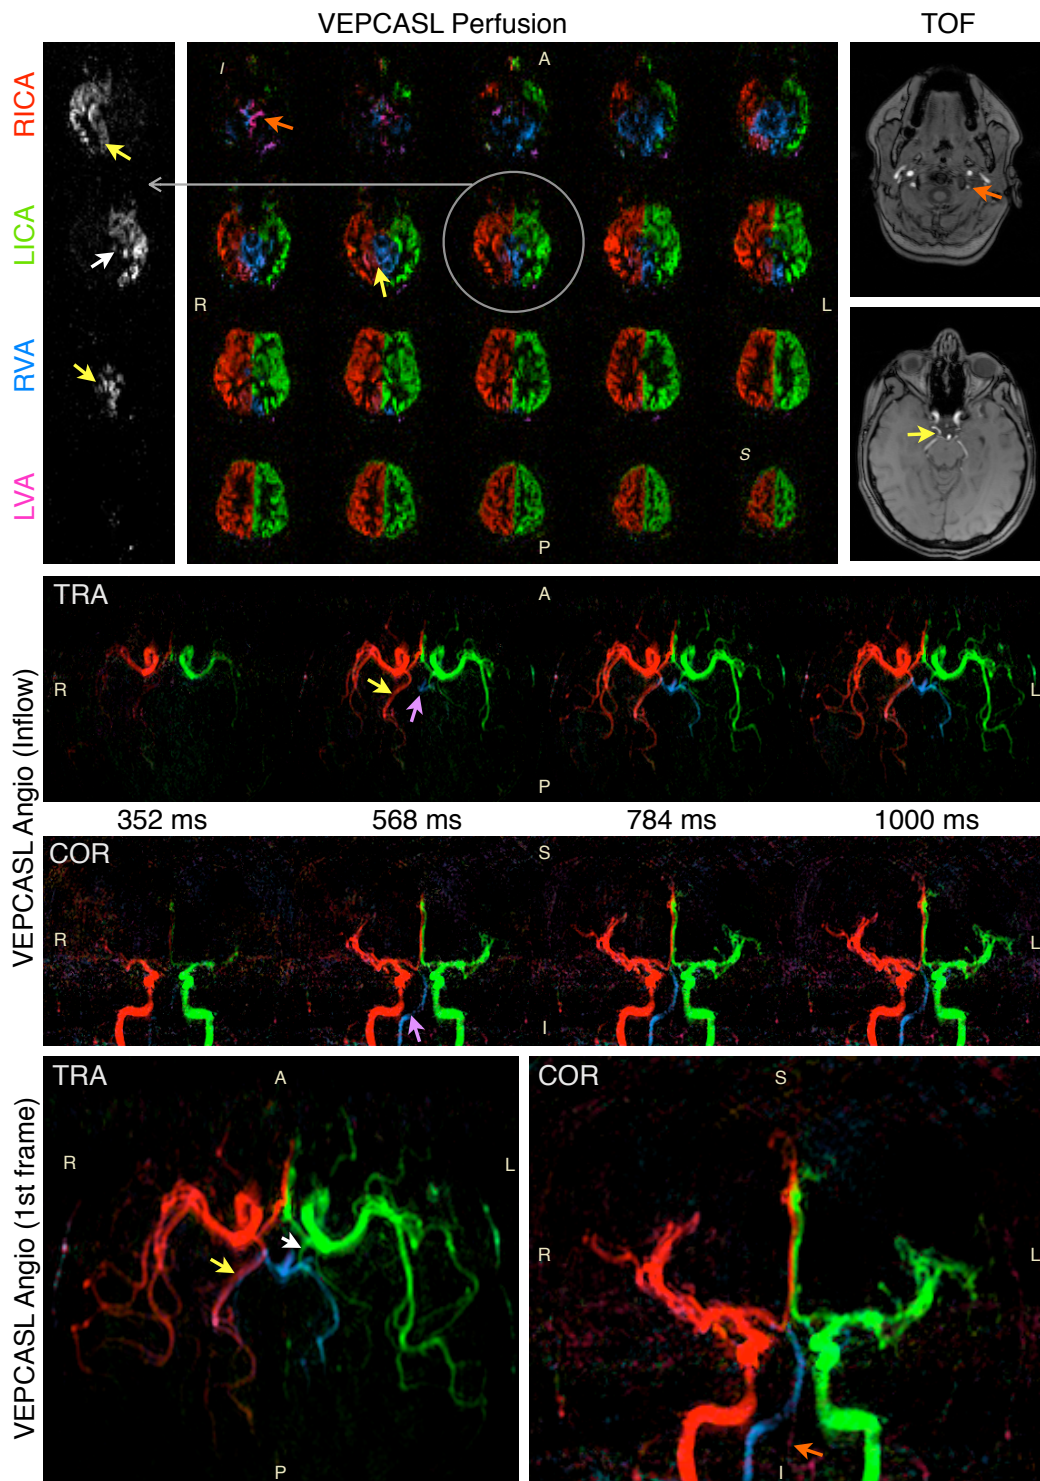


Figure 8.1: Images from VEPCASL perfusion and dynamic angiography acquisitions of a patient with a history of posterior circulation TIA and minor stroke, alongside TOF data for comparison. Note reduced flow in the LVA (orange arrows), delayed arrival in the RVA (purple arrows), significant collateral flow from the RICA through the posterior communicating artery to the RPCA (yellow arrows) and to a lesser extent on the left side (white arrows).

also evident in the VEPCASL angiography and to a lesser extent in the perfusion images, which is harder to assess from the TOF data, particularly in the MIPs (not shown here) where the overlying static tissue obscures the flow in smaller vessels.

Another patient with significant stenoses or occlusion in both VAs is shown in Fig. 8.2. The VEPCASL techniques clearly show collateral flow from the ICAs to the PCAs bilaterally, and even retrograde flow down the BA to supply the cerebellum. Again, the PLD used in the perfusion images may have been insufficient to see the full extent of perfusion to the cerebellum via this collateral route, but the presence of ICA blood there is still clear. As with the previous patient, the TOF images corroborate the conclusions drawn from the VEPCASL data, but the presence, direction, arterial source and relative contribution to downstream perfusion of the collateral flow are all much clearer with VEPCASL techniques.

The final example described here demonstrates the ability to identify changes in blood flow patterns before and after surgery with VEPCASL angiography and perfusion imaging (see Fig. 8.3). In this patient, who had an occluded LVA and stenosed RVA, collateral flow from the ICAs was required to ensure enough blood reached the PCA territories. After angioplasty to widen the RVA, the blood flow increased enough for the collateral flow to no longer be necessary.

It is worth noting that the VEPCASL dynamic angiography images contain some significant motion artefact in some cases. Useful information about vessel-specific flow patterns was still available in all subjects scanned thus far, but the visualisation of smaller vessels is somewhat inhibited where motion artefacts are significant. This highlights the requirement for an increase in angiographic acquisition speed, perhaps through the use of a bSSFP readout (Section 6.3) or other acceleration techniques (see Chapter 9). The perfusion measurements using an EPI readout are much less affected by motion since each volume is acquired in less than a second and post-processing can eliminate the majority of motion arte-




Figure 8.2: Example of a patient with severe stenosis or occlusion of both VAs. The VEPCASL perfusion images show no detectable blood signal from either VA in the brain. This is confirmed by very little signal enhancement in the VAs on the TOF image and no detectable VA signal in the VEPCASL angiography (orange arrows). Note that the vessels shown in blue (purple arrows) are the external carotid arteries which are not properly accounted for in the analysis and thus assigned to the RVA here. The lack of VA perfusion results in significant bilateral collateral flow from the ICAs to the PCAs via the posterior communicating arteries (yellow arrows), evident in all the imaging modalities, including the transverse TOF MIP. To compensate for the lack of VA flow, retrograde flow down the basilar artery from both ICAs is seen in the VEPCASL angiography images, leading to ICA signal in the cerebellum (white arrows).




Figure 8.3: Modified flow patterns before and after RVA angioplasty. This patient has an occluded LVA (yellow arrows) and stenosis in the RVA. Pre-surgery there is significant collateral flow from the ICAs, through the posterior communicating arteries (purple arrows) to compensate for the lack of blood to supply the PCA territories (orange arrows), particularly on the left side. After angioplasty the RVA is capable of supplying sufficient blood to the PCA territories that the collateral flow from the ICAs is no longer required. A subtle enhancement in the TOF image following surgery is also seen (green arrows). There are significant motion artefacts present in these images (white arrows).

facts.

8.3 Carotid Stenosis

A second study involving scanning patients with carotid stenosis before and after endarterectomy is being piloted in collaboration with Dr. Saran Shantikumar. The study protocol also includes high resolution vessel wall imaging at the level of the carotid bifurcation. To keep to a reasonable clinical scan time, VEPCASL angiography was only performed in the transverse view along with the EPI perfusion method.

Only a single patient has been scanned for this study thus far, yielding the images shown in Fig. 8.4. Stenosis of the LICA leads to extremely delayed arrival of the blood into the circle of Willis, visualised clearly with the dynamic angiographic technique. Neither the PLD of the perfusion sequence nor the acquisition window of the angiography protocol allow sufficient time to visualise the flow of blood into distal LMCA and LPCA territories and exchange into tissue.

This example shows that, again, the VEPCASL methods are able to visualise collateral flow and delayed blood transit. However, in future studies of these patients, multiple PLDs should be used to ensure late perfusion of the tissue can be observed and a greater number of angiography frames acquired to see the delayed blood reach the distal vessels.

8.4 Arteriovenous Malformation

The final patient group to be described in this chapter are those with arteriovenous malformation (AVM). These patients were scanned as part of a study on epilepsy, in collaboration with Dr. Natalie Voets. The main aim of that study is to




Figure 8.4: Patient with carotid stenosis. The severe LICA stenosis results in extremely delayed arrival of LICA blood into the circle of Willis and downstream vessels (yellow arrows). Areas of the brain presumably fed by this late arriving blood have no ASL signal in the VEPCASL perfusion images at the PLD used here. This patient also has LVA stenosis/occlusion (orange arrows), leading to no detectable LVA signal in the VEPCASL images. The cerebellum is supplied entirely by the RVA and the PCA territories by a combination of the RVA and collateral flow from the ICAs. The VEPCASL angiography data are shown *without* inflow subtraction, with the times shown being relative to the start of the VEPCASL pulse train. Note that the proximal RPCA is not well depicted here, perhaps due to imperfect imaging slab positioning, but collateral flow from the RICA into the distal RPCA is shown clearly.

map critical brain areas (such as those involved with language or motor function) close to the AVM (or other lesion) using blood oxygen level dependent (BOLD) based fMRI to help guide surgical resection of the lesion.

At the time of writing four patients with cerebral AVM had been scanned using the VEPCASL angiography sequence. The ultimate goal of such a technique would be to selectively label potential AVM feeding arteries. Identification of the major vessels supplying the lesion could help guide embolisation therapy. However, for these pilot scans, standard labelling at the level of the neck was performed to allow observation of blood flow to the AVM from the major brain feeding arteries without the complication of attempting to label many small arteries higher in the brain.

An example data set is shown in Fig. 8.5. The AVM is highly conspicuous in the VEPCASL angiography images and appears to be fed both by the RICA (via the RMCA) and the LICA (via the ACAs). The high signal intensity within the large vessels indicates the considerable quantity of blood which is not perfusing the tissue but instead draining directly into the venous system. This draining process is not directly visualised here, suggesting that a larger number of frames should be acquired in subsequent scans of these patients.

The abnormally large vessels feeding the AVM are also visualised in the TOF MIP, but based on this image alone it might be assumed that the AVM is fed largely via the ACAs. However, in the later frames of the VEPCASL angiography sequence it is apparent that the relative contributions from the ACAs and RMCA are approximately the same.

These images demonstrate the potential of VEPCASL angiography to provide clinically useful information in AVM patients. However, considerable optimisation of the technique to label the smaller feeding arteries is required, perhaps also with an increase in spatial resolution and acquisition speed to allow a greater

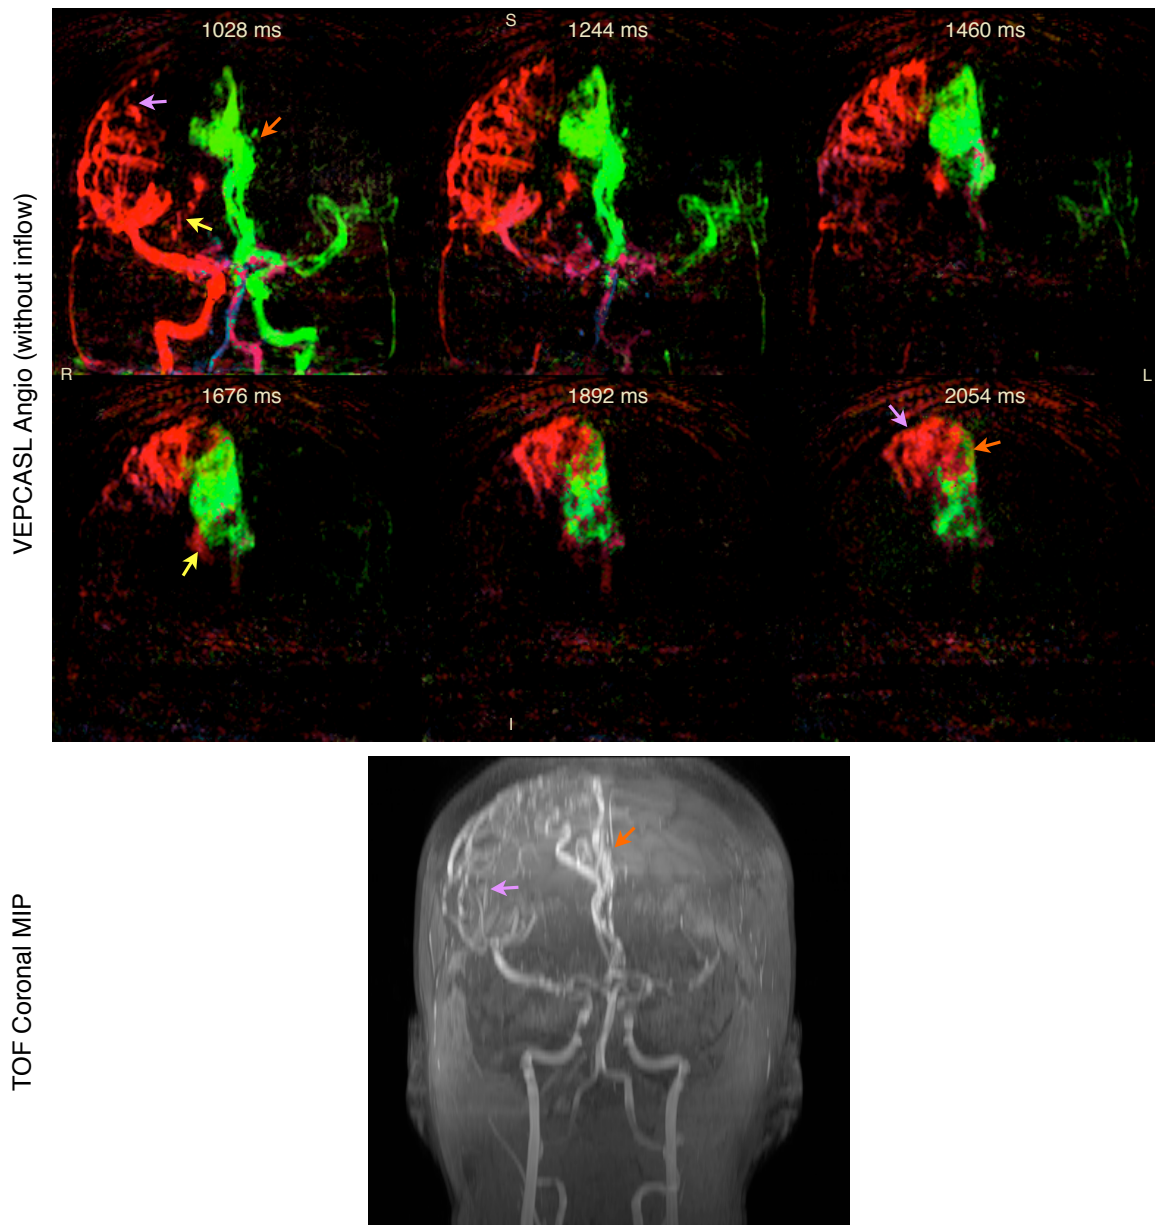


Figure 8.5: Imaging AVM with VEPCASL angiography. Selected frames are shown (without inflow subtraction), showing major feeding vessels to the AVM from both the RICA via lateral (purple arrows) and medial (yellow arrows) branches of the RMCA and the LICA via the ACAs (orange arrows). The times shown are relative to the start of the VEPCASL pulse train. Similar features can be observed in the TOF MIP (bottom).

number of encodings to be performed.

8.5 Conclusions

These preliminary patient studies demonstrate that the vessel-encoded ASL techniques developed in this thesis to observe both vessel-selective tissue perfusion and angiography are capable of providing useful clinical information in patients. The source and extent of collateral flow around the circle of Willis can be clearly visualised and the effect this has on downstream tissue perfusion observed. The angiographic technique (particularly in the coronal view) also enables the assessment of the brain feeding arteries directly, with stenosis or occlusion clearly visible. It also provides dynamic information, allowing the assessment of delayed arrival from compromised vessels.

However, there are still a number of potential improvements which could be made to these techniques to make their application in patient groups more useful and robust, particularly regarding sensitivity to motion and flow quantification. A summary of this thesis and ideas for future research are discussed in the next chapter.

Chapter 9

Summary and Future Work

In this final chapter, a summary of the work undertaken for this thesis is given, including a discussion on the relative merits of the vessel-encoded perfusion and angiography techniques described in previous chapters. There follows a discussion of ideas for future research based on the work completed here, with the aim of developing methods for the better understanding and assessment of collateral blood flow in the brain.

9.1 Thesis Summary

The aim of this work was to develop non-invasive methods to study the phenomenon of collateral blood flow in the brain using magnetic resonance imaging (MRI). A study of previous literature on the topic revealed many methods which are capable of providing some information regarding collateral flow, each with its own benefits and problems. However, vessel-encoded pseudo-continuous arterial spin labelling (VEPCASL) [48] possessed many of the desired qualities: it is completely non-invasive, provides vessel-selective potentially dynamic information for each of the brain-feeding arteries, allows the generation of perfusion weighted maps showing not only the vascular territories but also the relative perfusion of

the tissue at each point, is highly signal-to-noise ratio (SNR) efficient, produces minimal artefacts beyond a few centimetres from the labelling plane, is insensitive to magnetisation transfer effects, allows the relatively simple prescription of the target vessels within a single plane and has been shown to provide useful information on collateral flow in patient groups [99].

Simulations of the Bloch equations and experiments using a flowing water phantom allowed exploration of the sensitivities of the VEPCASL scheme to a variety of effects. The phantom data matched the simulation results very closely, allowing a greater range of parameters to be simulated than it would be possible to test experimentally within a reasonable time frame. It was found that VEPCASL results in highly efficient inversion of magnetisation flowing through a defined labelling plane (up to approximately 90% of the theoretical maximum). The control condition, which matches MT effects and eddy currents almost exactly, has very little effect on the magnetisation. There is a significant influence of blood flow velocity on the labelling efficiency of VEPCASL, although within the range expected in the brain-feeding arteries (5-50 cm/s) the inversion efficiency remains above 80%.

Optimal pulse sequence parameters were confirmed to be close to those used by Wong [48] and these were used for the remainder of the thesis. The spatial modulation of the inversion efficiency across the labelling plane was shown to allow the selective labelling of target vessels with high efficiency, although the form of this curve has some velocity dependence. Off-resonance effects were shown to be significant only above about 100 Hz and vessel angulation only above 60 ° away from the direction perpendicular to the labelling plane. Some differences in the effect of the VEPCASL pulse train on static magnetisation near the labelling plane between tag and control conditions resulted in the recommendation to place the imaging region at least 3 cm away from the site of labelling for the standard

parameters used here.

A number of methods for the separation of vascular components from vessel-encoded data were tested, but the Bayesian method of Chappell *et al.* [98] has many advantages over the other methods, including the ability to cope with a rank deficient encoding matrix and to compensate for subject motion between the vessel-localisation scan and the VEPCASL acquisition. Highly vessel-specific images are produced using this method, with almost no detectable contamination of vascular components, a clear advantage over some of the vessel-specific ASL techniques where contamination can be significant [85, 93].

The VEPCASL module was attached to a two-dimensional (2D) multi-slice echo planar imaging (EPI) sequence, allowing the generation of vascular territory maps weighted by relative regional cerebral blood flow (CBF) in healthy volunteers. Experiments were conducted to determine that a post labelling delay of 1000 ms is sufficient to give high SNR perfusion weighted images that well represent CBF in healthy volunteers. In addition, the optimum labelling plane location was found to be approximately level with the lower jaw at a point where the two internal carotid arteries (ICAs) and two vertebral arteries (VAs) lie in an approximately rectangular arrangement in nearly all subjects. This allows the efficient encoding of all four vessels using eight vessel-encoding cycles. The resulting vascular territory images were highly vessel selective and of good image quality.

However, the EPI readout suffers somewhat from signal dropout and distortion artefacts and has sub-optimal SNR. Optimisation experiments resulted in the conclusion that the use of background suppression (pre-saturation of the imaging region plus two global inversion pulses) reduces motion artefacts significantly. In addition, using a 6/8 partial Fourier technique in the phase-encode direction yielded higher SNR images and reduced signal dropout without significant loss of image quality. Acceleration using parallel imaging yielded reduced distortion

artefacts but at a significant cost in SNR and was thus not used in future EPI experiments.

The optimised EPI sequence was shown to be able to clearly identify normal variants of the cerebral vasculature in healthy volunteers, including the presence of collateral flow in some cases. Comparison was also made with a 3D gradient and spin echo (3D-GRASE) sequence, which has theoretical benefits in terms of SNR and improved background suppression. However, this long echo train technique suffers from severe blurring in the slice direction. This was improved somewhat through the use of parallel imaging. However, the blurring of the borders between vascular territories was still significant, reducing its effectiveness for this application. With future pulse sequence developments it is likely that 3D-GRASE can surpass EPI in terms of SNR, improved background suppression and ease of CBF quantification.

These experiments showed the utility of VEPCASL to produce tissue perfusion weighted images, but the direct visualisation of collateral flow around the circle of Willis (or via other pathways) and assessment of stenosis or occlusion in brain-feeding arteries was lacking. To remedy this, a novel sequence combining the VEPCASL preparation with a 2D flow-compensated spoiled gradient echo dynamic (“cine”) readout was implemented. Proximal vessel status and collateral flow around the circle of Willis were clearly visualised. Excellent separation of vascular components was again achieved and the dynamic information allowed the clear depiction of delayed arrival in certain vessels. It was found that the use of a relatively long labelling duration (1000 ms) produced images with higher SNR and better distal vessel visibility, although the inflow of the labelled bolus is not visualised in this case. However, timing information can still be extracted by observing outflow of the bolus and the effect of inflow recreated using a simple post-processing subtraction method. The resulting angiograms matched well

time-of-flight magnetic resonance angiography (TOF MRA) images acquired in the same session, sometimes providing better distal vessel visibility due to improved static tissue suppression in the VEPCASL images. Normal variants of the cerebral vasculature were again identified with this method, showing its potential for use in patient groups.

However, these VEPCASL dynamic angiographic data sets took a relatively long time to acquire (about 10 minutes for a dynamic 2D data set) and were therefore susceptible to motion artefacts. Simply increasing the acquisition speed would result in greatly reduced ASL contrast in later time frames. However, a readout method based on balanced steady-state free precession (bSSFP), in which the transverse magnetisation is not spoiled at the end of each repetition time (TR) period, was shown to result in very little perturbation of the ASL contrast in simulations. The cardiac gated readout method of Bi, Schmitt and others [72, 73] allows considerable acceleration of the image readout, to the extent that 4D (i.e. dynamic 3D) imaging is possible. When combined with the VEPCASL preparation, high quality vessel-selective angiograms were produced.

The ability to reformat in any given orientation is a significant advantage, although the full 4D acquisition time was still too long for a standard clinical protocol at around 18 minutes. However, when run in dynamic 2D mode, image quality comparable to the initial implementation with a spoiled gradient echo readout was obtained in a fraction of the scan time (approximately 1.5 mins). A number of minor artefacts were present in these images but future sequence developments are likely to help in their reduction.

A novel method for the quantification of VEPCASL dynamic angiographic data was developed, allowing the estimation of blood volume flow rates in various vessel segments proximal and distal to the circle of Willis without the need for additional calibration scans. This involved the fitting of a model to the timecourse

in each voxel for each vascular component using a Bayesian inference procedure. Once the flow dynamics in each voxel are known, it is possible to simulate the expected signal for a very short bolus duration. Summation of this signal over the vessel of interest allows calculation of the volume of blood which flows through the artery during this short simulated duration, and thus calculation of volume flow rate. Calibration is achieved by observing the signal in a profile across larger vessels and assuming circular cross-sections to derive a factor which relates the measured signal to blood volume.

This method was shown to produce volume flow rates consistent with the literature for the major vessels around the circle of Willis in both a healthy volunteer and a patient with cerebrovascular disease. Such a method gives VEPCASL dynamic angiography a considerable advantage over many other qualitative techniques, including the gold-standard for collateral flow measurement, x-ray digital subtraction angiography (DSA).

Improvements to this quantification approach are required, including better modelling of magnetisation which has not experienced all the previous RF pulses, and the avoidance of the requirement for a relatively large mask in which the quantification is performed.

Finally, the optimised EPI and spoiled gradient echo dynamic angiographic VEPCASL methods were applied in a number of patient groups. The longest running of these studies is in a group of patients with disease of the vertebral or basilar arteries. Both perfusion and angiographic methods were shown to clearly visualise collateral flow in these patients and provided consistent results. Dynamic angiography in the coronal view also allowed direct visualisation of stenosed or occluded arteries, consistent with other imaging previously undertaken in these patients. When x-ray DSA data become available in this patient group it is hoped that the VEPCASL methods can be validated against this current gold-standard.

Preliminary results in patients with carotid stenosis and arteriovenous malformation were also presented, showing promise for these techniques in a variety of diseases. However, subject motion was significant in all of these patient groups, causing significant artefacts in the dynamic angiographic images, although useful diagnostic information was still obtained in all cases. In cases of severely delayed blood transit, tissue perfusion and even arrival in distal vessels with angiography was not always visualised clearly, highlighting the importance of moving to multiple post labelling delays for the perfusion technique and longer data readout times for angiography.

These sequences provided information similar to that which can be achieved using x-ray DSA, although the spatial and temporal resolution is still some way behind. It is hoped that with further improvements, these techniques will provide useful clinical information in these and other patient groups non-invasively and relatively quickly compared to conventional methods.

9.2 VEPCASL Tissue Perfusion or Angiography?

The vessel-encoded tissue perfusion and dynamic angiographic sequences described in this thesis provide consistent information (e.g. see Section 8.2). In a busy clinical protocol it may not always be possible to include both sequences, so which one should be chosen?

The perfusion technique has the advantage that it is more robust to motion and is faster to acquire. It should also be noted that ultimately it is the perfusion of tissue which is critical in cerebrovascular disease patients, and assessment of increased flow in large vessels after surgery does not necessarily mean that the downstream tissue is successfully reperfused [4].

However, the angiographic method has a theoretical SNR advantage in that

the blood is imaged sooner after labelling when it is both more concentrated and less attenuated by T_1 decay. In addition, collateral flow can be visualised directly, rather than inferred from tissue perfusion measurements, sometimes allowing more subtle effects to be observed (e.g. the low level of collateral flow from LICA to LPCA in Fig. 8.1, which is less apparent in the perfusion images). Stenosis or occlusion in large arteries can also be assessed and more information is available on the dynamics of the blood flow (e.g. the presence of significantly delayed blood transit in diseased vessels).

For a general clinical protocol ideally both acquisitions would be performed, but if scan time is very limited the faster perfusion technique would probably provide sufficient information in most cases.

9.3 Future Research

Although the VEPCASL methods described in this thesis are able to demonstrate the presence of collateral flow, there are a number of improvements which could be made to overcome some of their limitations. Some ideas for future research are given in this section, along with preliminary results showing the proof of principle in some cases.

9.3.1 Improved Sequence Design

With the continuing rapid development of MRI pulse sequence research, many potential improvements to both perfusion and angiographic sequences are possible. Some of these are outlined here.

Background Suppression

The background suppression scheme used in the tissue perfusion chapter was sub-optimal due to the restriction that the two global inversion pulses had to be played out after the VEPCASL pulse train had finished. This restriction means that such pulses cannot be used in the angiographic readout since there is very little time between the end of the labelling period and the start of the readout.

Recently a method by which such global inversion pulses can be placed within a PCASL pulse train has been described [94], allowing more freedom in the choice of the T_1 species which are optimally nulled. This would also allow a greater number of inversion pulses to be used, if desired, although it is noted that for imperfect inversion pulses the ASL contrast is diminished with each additional pulse. This scheme should allow much better background suppression for both the tissue perfusion and angiographic sequences, reducing sensitivity to motion and scanner drift.

Tissue Perfusion Imaging

Without much modification it should be possible to run the tissue perfusion sequence at multiple post labelling delays, which would allow fitting to a theoretical ASL kinetic model [33], and quantification of CBF in absolute units (with suitable calibration scans). Quantification of VEPCASL data has previously been performed in patients [99], although only a single post labelling delay was used here, which may result in inaccuracies for areas with significantly delayed arrival. In particular, it will be interesting to note if the quantification is significantly improved by considering each vascular component separately. It seems likely that poor fitting to the standard model will be obtained in non-selective ASL acquisitions in cases where there is combined flow from two vessels, with one significantly delayed relative to the other.

To aid the quantification process and improve background suppression a 3D readout is desirable. Improvement to the 3D-GRASE sequence through the use of partial Fourier, more extensive parallel acceleration and potentially segmentation of the readout are all likely to reduce the through-slice blurring significantly, making this sequence a strong candidate for VEPCASL tissue perfusion imaging. However, the impact of these acceleration techniques on SNR relative to the optimised EPI readout must be assessed.

Experiments above the circle of Willis are also planned. Design of a robust and flexible encoding strategy is likely to be challenging, although perhaps a randomised encoding method would be beneficial in such a case [156], particularly when combined with the powerful analysis method of Chappell *et al.* [98].

Dynamic Angiography

The VEPCASL dynamic angiography methods show promise for the visualisation of collateral flow directly in the brain-feeding arteries. There are a number of potential improvements which could be made, however. Since complex data are required for the analysis, there is considerable sensitivity to phase drift. This can be somewhat compensated for by phase-correction steps in post-processing, but for the bSSFP readout B_0 field drift results in movement of the banding artefacts which cannot be corrected for. Therefore, a method for monitoring and correcting for changes to B_0 is desirable.

Both the spoiled gradient echo and bSSFP readouts could benefit from variable flip angle approaches like that of Schmitt *et al.* [73] to create a more even signal profile over time. For the bSSFP readout, catalysis of the steady-state could be improved by using a smoother variation in the flip angle of the preparation pulses [157], rather than a simple linear increase. This might help reduce artefacts in the initial time frames. Reduced sensitivity to motion is also likely to be achieved by

changing the looping structure to “cycles inner” (see Section 6.2), so each phase encode line is acquired for all VEPCASL cycles as close together as possible.

However, perhaps the biggest drawback of the angiographic method is its long acquisition time which leads to increased motion sensitivity. This could be addressed using acceleration techniques such as compressed sensing, described in the next section.

Acquisition Acceleration Using Compressed Sensing

Recently, it has been shown that pseudo-random undersampling of k -space with relatively high acceleration factors can still yield high quality images without a significant loss of SNR using a technique known as compressed sensing (CS) [158]. Other techniques, making use of temporal undersampling strategies in dynamic imaging have also been proposed [60]. A requirement for the CS approach is that the resulting images have a sparse representation in some mathematical domain, which allows regularisation of the image reconstruction. Dynamic angiographic data sets are ideal for this application, since they are sparse both in the spatial domain (vessels occupy only a small fraction of the image) and the temporal domain (each time frame is closely related to adjacent frames). This redundancy potentially allows considerable acceleration which would be highly beneficial for the VEPCASL angiography method.

To show the potential of CS in VEPCASL angiography, fully sampled data sets were retrospectively undersampled using a pseudo-random pattern in both k -space and time (although denser sampling was used at the centre of k -space to maintain reasonable SNR). Sparsity was enforced in the spatial and temporal frequency domains simultaneously. Fig. 9.1 shows examples of the first frame of such a data set undersampled by a factor of 2, 3 and 5 and compares it to a naive reconstruction with zero-filling of “unacquired” data, compensated for the

density of k -space sampling which was higher in the centre than at the edges.

These results show the potential of CS to dramatically undersample the data, even in this quick implementation. Non-cartesian trajectories such as radial and spiral are likely to improve these reconstructions, and may be beneficial for such acquisitions [158].

9.3.2 Dynamic Angiography Quantification

Although the quantification of dynamic angiographic data described in Chapter 7 yields reasonable volume flow rates in output vessel segments, there are a number of improvements which could be made. Firstly, the model does not accurately represent blood which has not experienced all the previous RF pulses. A relatively straightforward modification of the model to include an extra fitting parameter which describes the maximum number of previous RF pulses experienced by the magnetisation in each voxel yields considerable improvement in the single-voxel fit in proximal vessel segments (see Fig. 9.2). This modification is likely to remove the early peaks in the estimated volume flow rate seen in Chapter 7, giving more robust quantification in proximal vessel segments. This may improve robustness in the absence of dispersion effects, allowing quantification in smaller vessel segments.

This quantification approach would also benefit from validation in a flow phantom where the volume flow rate is known and can be varied. Comparison with phase-contrast MRA would also be useful, both in phantom and *in vivo*. Adaptation of the model for bSSFP acquisitions would also be beneficial. It is expected that this might present some challenges due to the more complex signal dependence of bSSFP on factors such as off-resonance frequency. However, if a robust method can be developed, this would be of significant value.




Figure 9.1: Preliminary results using the compressed sensing (CS) method of Lustig *et al.* [158] on non-selective VEPCASL angiography data. The first frame is shown with reconstructions using fully sampled k -space (left), undersampled k -space with zero filling and density compensation (DC) (middle), and a CS reconstruction enforcing sparsity in spatial and temporal domains (right). Acceleration factors (R) of 2, 3 and 5 are shown. Note the significant artefacts in the zero filled reconstruction, even at low undersampling factors (red arrows). The CS reconstruction maintains good image quality, even at relatively high acceleration factors, although there is some loss of distal vessel visibility (yellow arrows).




Figure 9.2: Improved modelling of VEPCASL dynamic angiography data in a proximal MCA voxel from a healthy volunteer. The “new” fit allows variation of the maximum number of RF pulses the magnetisation has previously experienced, whereas the “old” fit assumes the magnetisation has experienced all previous pulses. This modification gives a much better fit to the data in this case.

9.3.3 Applications

It is hoped that the methods developed within this thesis will continue to be applied in patient groups. This will allow validation of the method in patients who also undergo examination with x-ray DSA or other established imaging modalities. The information on collateral flow obtained with these VEPCASL techniques will be analysed against patient outcome measures to assess the usefulness of these data for patient prognosis.

The projects in patients with vertebro-basilar disease, carotid stenosis and arteriovenous malformation are on-going. It is hoped that sequence improvements will benefit these studies. Considerable optimisation is required to allow the assessment of significantly delayed blood arrival and efficient encoding of blood vessels supplying an AVM. In addition, collaboration with a group interested in collateral flow in sub-arachnoid haemorrhage patients has been initiated. Variation in the collateral flow patterns over the course of this disease, particularly in the case of vasospasm, is likely to be clinically significant.

References

- [1] A. Srinivasan, M. Goyal, F. Al Azri, and C. Lum, "State-of-the-art imaging of acute stroke," *Radiographics*, vol. 26 Suppl 1, pp. S75–95, 2006.
- [2] A. C. Guyton and J. E. Hall, *Textbook of medical physiology*. Philadelphia, Pa. ; London: W. B. Saunders, 10th ed. ed., 2000.
- [3] D. S. Liebeskind, "Collateral circulation," *Stroke*, vol. 34, no. 9, pp. 2279–84, 2003.
- [4] D. S. Liebeskind, "Reperfusion for acute ischemic stroke: arterial revascularization and collateral therapeutics," *Curr Opin Neurol*, vol. 23, pp. 36–45, Feb 2010.
- [5] S. P. Kloska, M. Wintermark, T. Engelhorn, and J. B. Fiebach, "Acute stroke magnetic resonance imaging: current status and future perspective," *Neuroradiology*, vol. 52, pp. 189–201, Mar 2010.
- [6] O. Y. Bang, J. L. Saver, B. H. Buck, J. R. Alger, S. Starkman, B. Ovbiagele, D. Kim, R. Jahan, G. R. Duckwiler, S. R. Yoon, F. Vinuela, and D. S. Liebeskind, "Impact of collateral flow on tissue fate in acute ischemic stroke," *J Neurol Neurosurg Psychiatry*, 2007.
- [7] G. A. Christoforidis, Y. Mohammad, D. Kehagias, B. Avutu, and A. P. Slivka, "Angiographic assessment of pial collaterals as a prognostic indicator following intra-arterial thrombolysis for acute ischemic stroke," *AJNR Am J Neuroradiol*, vol. 26, no. 7, pp. 1789–97, 2005.
- [8] A. I. Qureshi, "New grading system for angiographic evaluation of arterial occlusions and recanalization response to intra-arterial thrombolysis in acute ischemic stroke," *Neurosurgery*, vol. 50, no. 6, pp. 1405–14; discussion 1414–5, 2002.
- [9] C. E. Jones, R. L. Wolf, J. A. Detre, B. Das, P. K. Saha, J. Wang, Y. Zhang, H. K. Song, A. C. Wright, E. M. Mohler, 3rd, R. M. Fairman, E. L. Zager, O. C. Velazquez, M. A. Golden, J. P. Carpenter, and F. W. Wehrli, "Structural MRI of carotid artery atherosclerotic lesion burden and characterization of hemispheric cerebral blood flow before and after carotid endarterectomy," *NMR Biomed*, vol. 19, pp. 198–208, Apr 2006.

- [10] J. R. Romero, A. Pikula, T. N. Nguyen, Y. L. Nien, A. Norbash, and V. L. Babikian, "Cerebral collateral circulation in carotid artery disease," *Curr Cardiol Rev*, vol. 5, pp. 279–88, Nov 2009.
- [11] W. D. Jeans, "The development and use of digital subtraction angiography," *Br J Radiol*, vol. 63, pp. 161–8, Mar 1990.
- [12] M. Bendszus, M. Koltzenburg, R. Burger, M. Warmuth-Metz, E. Hofmann, and L. Solymosi, "Silent embolism in diagnostic cerebral angiography and neurointerventional procedures: a prospective study," *Lancet*, vol. 354, pp. 1594–7, Nov 1999.
- [13] T. J. Kaufmann, J. Huston, 3rd, J. N. Mandrekar, C. D. Schleck, K. R. Thielens, and D. F. Kallmes, "Complications of diagnostic cerebral angiography: evaluation of 19,826 consecutive patients," *Radiology*, vol. 243, pp. 812–9, Jun 2007.
- [14] A. Krishnaswamy, J. P. Klein, and S. R. Kapadia, "Clinical cerebrovascular anatomy," *Catheter Cardiovasc Interv*, vol. 75, pp. 530–9, Mar 2010.
- [15] E. N. Marieb, *Human anatomy & physiology*. San Francisco: Benjamin Cummings, 5th ed ed., 2001.
- [16] B. R. Radzik, H. A. Puttgen, R. F. Gottesman, and E. M. Aldrich, "The role of interventional radiology in cerebrovascular disease: A review of cerebrovascular anatomy and disease," *Journal of Radiology Nursing*, vol. 26, no. 2, pp. 39–43, 2007.
- [17] K. W. Muir, A. Buchan, R. von Kummer, J. Rother, and J.-C. Baron, "Imaging of acute stroke," *Lancet Neurol*, vol. 5, pp. 755–68, Sep 2006.
- [18] J. Cahill, W. J. Cahill, J. W. Calvert, J. H. Calvert, and J. H. Zhang, "Mechanisms of early brain injury after subarachnoid hemorrhage," *J Cereb Blood Flow Metab*, vol. 26, pp. 1341–53, Nov 2006.
- [19] T. Thom, N. Haase, W. Rosamond, V. J. Howard, J. Rumsfeld, T. Manolio, Z.-J. Zheng, K. Flegal, C. O'Donnell, S. Kittner, D. Lloyd-Jones, D. C. Goff, Jr, Y. Hong, R. Adams, G. Friday, K. Furie, P. Gorelick, B. Kissela, J. Marler, J. Meigs, V. Roger, S. Sidney, P. Sorlie, J. Steinberger, S. Wasserthiel-Smoller, M. Wilson, P. Wolf, and American Heart Association Statistics Committee and Stroke Statistics Subcommittee, "Heart disease and stroke statistics—2006 update: a report from the american heart association statistics committee and stroke statistics subcommittee," *Circulation*, vol. 113, pp. e85–151, Feb 2006.
- [20] R. Al-Shahi and C. Warlow, "A systematic review of the frequency and prognosis of arteriovenous malformations of the brain in adults," *Brain*, vol. 124, pp. 1900–26, Oct 2001.

REFERENCES

- [21] C. L. Taylor, K. Dutton, G. Rappard, G. L. Pride, R. Replogle, P. D. Purdy, J. White, C. Giller, T. A. Kopitnik, Jr, and D. S. Samson, "Complications of preoperative embolization of cerebral arteriovenous malformations," *J Neurosurg*, vol. 100, pp. 810–2, May 2004.
- [22] R. M. Scott and E. R. Smith, "Moyamoya disease and moyamoya syndrome," *N Engl J Med*, vol. 360, pp. 1226–37, Mar 2009.
- [23] F. Bloch, "Nuclear induction," *Physical Review*, vol. 70, no. 7-8, pp. 460–474, 1946.
- [24] F. Bloch, W. W. Hansen, and M. Packard, "The nuclear induction experiment," *Physical Review*, vol. 70, no. 7-8, pp. 474–485, 1946.
- [25] P. Lauterbur, "Image formation by induced local interactions: examples employing nuclear magnetic resonance," *Nature*, vol. 242, no. 5394, pp. 190–191, 1973.
- [26] P. Mansfield and P. K. Grannell, "NMR 'diffraction' in solids?," *Journal of Physics C: Solid State Physics*, vol. 6, no. 22, pp. L422–L426, 1973.
- [27] P. C. Lauterbur, D. M. Kramer, W. V. House Jr., and C. . Chen, "Zeugmatographic high resolution nuclear magnetic resonance spectroscopy. Images of chemical inhomogeneity within macroscopic objects [4]," *Journal of the American Chemical Society*, vol. 97, no. 23, pp. 6866–6868, 1975.
- [28] P. Jezzard, P. M. Matthews, and S. M. Smith, *Functional MRI: an introduction to methods*. Oxford: Oxford University Press, 2001.
- [29] M. A. Bernstein, K. F. King, and Z. J. Zhou, *Handbook of MRI pulse sequences*. Amsterdam: Academic Press, 2004.
- [30] E. M. Haacke, *Magnetic resonance imaging: physical principles and sequence design*. New York: J. Wiley & Sons, 1999.
- [31] J. A. Detre, J. S. Leigh, D. S. Williams, and A. P. Koretsky, "Perfusion imaging," *Magn Reson Med*, vol. 23, no. 1, pp. 37–45, 1992.
- [32] D. S. Williams, J. A. Detre, J. S. Leigh, and A. P. Koretsky, "Magnetic resonance imaging of perfusion using spin inversion of arterial water," *Proc Natl Acad Sci U S A*, vol. 89, no. 1, pp. 212–6, 1992.
- [33] R. B. Buxton, L. R. Frank, E. C. Wong, B. Siewert, S. Warach, and R. R. Edelman, "A general kinetic model for quantitative perfusion imaging with arterial spin labeling," *Magn Reson Med*, vol. 40, no. 3, pp. 383–96, 1998.
- [34] E. L. Barbier, L. Lamalle, and M. Decorps, "Methodology of brain perfusion imaging," *J Magn Reson Imaging*, vol. 13, no. 4, pp. 496–520, 2001.

- [35] E. T. Petersen, I. Zimine, Y. C. Ho, and X. Golay, "Non-invasive measurement of perfusion: a critical review of arterial spin labelling techniques," *Br J Radiol*, vol. 79, no. 944, pp. 688–701, 2006.
- [36] R. L. Wolf and J. A. Detre, "Clinical neuroimaging using arterial spin-labeled perfusion magnetic resonance imaging," *Neurotherapeutics*, vol. 4, no. 3, pp. 346–59, 2007.
- [37] X. Golay, E. T. Petersen, I. Zimine, and T. C. Lim, "Arterial spin labeling: a one-stop-shop for measurement of brain perfusion in the clinical settings," *Conf Proc IEEE Eng Med Biol Soc*, vol. 1, pp. 4320–3, 2007.
- [38] T. T. Liu and G. G. Brown, "Measurement of cerebral perfusion with arterial spin labeling: Part 1. methods," *J Int Neuropsychol Soc*, vol. 13, no. 3, pp. 517–25, 2007.
- [39] A. R. Deibler, J. M. Pollock, R. A. Kraft, H. Tan, J. H. Burdette, and J. A. Maldjian, "Arterial spin-labeling in routine clinical practice, part 1: technique and artifacts," *AJNR Am J Neuroradiol*, vol. 29, pp. 1228–34, Aug 2008.
- [40] W. Dai, D. Garcia, C. de Bazelaire, and D. C. Alsop, "Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields," *Magn Reson Med*, vol. 60, pp. 1488–97, Dec 2008.
- [41] R. R. Edelman, B. Siewert, D. G. Darby, V. Thangaraj, A. C. Nobre, M. M. Mesulam, and S. Warach, "Qualitative mapping of cerebral blood flow and functional localization with echo-planar MR imaging and signal targeting with alternating radio frequency," *Radiology*, vol. 192, pp. 513–20, Aug 1994.
- [42] R. R. Edelman and Q. Chen, "EPISTAR MRI: multislice mapping of cerebral blood flow," *Magn Reson Med*, vol. 40, pp. 800–5, Dec 1998.
- [43] S. G. Kim, "Quantification of relative cerebral blood flow change by flow-sensitive alternating inversion recovery (FAIR) technique: application to functional mapping," *Magn Reson Med*, vol. 34, pp. 293–301, Sep 1995.
- [44] K. K. Kwong, D. A. Chesler, R. M. Weisskoff, K. M. Donahue, T. L. Davis, L. Ostergaard, T. A. Campbell, and B. R. Rosen, "MR perfusion studies with T1-weighted echo planar imaging," *Magn Reson Med*, vol. 34, pp. 878–87, Dec 1995.
- [45] D. C. Alsop and J. A. Detre, "Multisection cerebral blood flow MR imaging with continuous arterial spin labeling," *Radiology*, vol. 208, no. 2, pp. 410–6, 1998.
- [46] D. Garcia, C. de Bazelaire, and D. Alsop, "Pseudo-continuous flow driven adiabatic inversion for arterial spin labeling," in *Proceedings of the 13th Annual Meeting of ISMRM*, (Miami Beach, FL, USA), p. 37, 2005.

REFERENCES

- [47] W.-C. Wu, M. Fernández-Seara, J. A. Detre, F. W. Wehrli, and J. Wang, "A theoretical and experimental investigation of the tagging efficiency of pseudocontinuous arterial spin labeling," *Magn Reson Med*, vol. 58, pp. 1020–7, Nov 2007.
- [48] E. C. Wong, "Vessel-encoded arterial spin-labeling using pseudocontinuous tagging," *Magn Reson Med*, vol. 58, no. 6, pp. 1086–91, 2007.
- [49] K. Scheffler and S. Lehnhardt, "Principles and applications of balanced SSFP techniques," *Eur Radiol*, vol. 13, no. 11, pp. 2409–18, 2003.
- [50] E. T. Petersen, T. Lim, and X. Golay, "Model-free arterial spin labeling quantification approach for perfusion MRI," *Magn Reson Med*, vol. 55, no. 2, pp. 219–32, 2006.
- [51] L. M. Bidaut, C. Laurent, M. Piotin, P. Gailloud, M. Muster, J. H. Fasel, D. A. Rüfenacht, and F. Terrier, "Second-generation three-dimensional reconstruction for rotational three-dimensional angiography," *Acad Radiol*, vol. 5, pp. 836–49, Dec 1998.
- [52] J. J. Barfett, J. Fierstra, P. W. A. Willems, D. J. Mikulis, and T. Krings, "Intravascular functional maps of common neurovascular lesions derived from volumetric 4D CT data," *Investigative radiology*, vol. 45, no. 7, pp. 370–377, 2010.
- [53] J. C. Tan, W. P. Dillon, S. Liu, F. Adler, W. S. Smith, and M. Wintermark, "Systematic comparison of perfusion-CT and CT-angiography in acute stroke patients," *Ann Neurol*, vol. 61, pp. 533–43, Jun 2007.
- [54] E. de Camargo and W. Koroshetz, "Neuroimaging of ischemia and infarction," *NeuroRx*, vol. 2, no. 2, p. 265, 2005.
- [55] F. Sallustio, R. Kern, M. Günther, K. Szabo, M. Griebe, S. Meairs, M. Hennerici, and A. Gass, "Assessment of intracranial collateral flow by using dynamic arterial spin labeling MRA and transcranial color-coded duplex ultrasound," *Stroke*, vol. 39, pp. 1894–7, Jun 2008.
- [56] J. Hendrikse, C. J. M. Klijn, A. C. van Huffelen, L. J. Kappelle, and J. van der Grond, "Diagnosing cerebral collateral flow patterns: accuracy of non-invasive testing," *Cerebrovasc Dis*, vol. 25, no. 5, pp. 430–7, 2008.
- [57] H. Zhang, J. H. Maki, and M. R. Prince, "3D contrast-enhanced MR angiography," *J Magn Reson Imaging*, vol. 25, pp. 13–25, Jan 2007.
- [58] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI," *Magn Reson Med*, vol. 42, pp. 952–62, Nov 1999.

- [59] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magn Reson Med*, vol. 47, pp. 1202–10, Jun 2002.
- [60] J. Tsao, P. Boesiger, and K. P. Pruessmann, "k-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations," *Magn Reson Med*, vol. 50, pp. 1031–42, Nov 2003.
- [61] P. H. Kuo, E. Kanal, A. K. Abu-Alfa, and S. E. Cowper, "Gadolinium-based MR contrast agents and nephrogenic systemic fibrosis," *Radiology*, vol. 242, pp. 647–9, Mar 2007.
- [62] T. J. Masaryk, M. T. Modic, J. S. Ross, P. M. Ruggieri, G. A. Laub, G. W. Lenz, E. M. Haacke, W. R. Selman, M. Wiznitzer, and S. I. Harik, "Intracranial circulation: preliminary clinical results with three-dimensional (volume) MR angiography," *Radiology*, vol. 171, pp. 793–9, Jun 1989.
- [63] P. R. Moran, "A flow velocity zeugmatographic interlace for NMR imaging in humans," *Magn Reson Imaging*, vol. 1, no. 4, pp. 197–203, 1982.
- [64] R. Bammer, T. A. Hope, M. Aksoy, and M. T. Alley, "Time-resolved 3D quantitative flow MRI of the major intracranial vessels: initial experience and comparative evaluation at 1.5T and 3.0T in combination with parallel imaging," *Magn Reson Med*, vol. 57, pp. 127–40, Jan 2007.
- [65] M. Miyazaki and V. S. Lee, "Nonenhanced MR angiography," *Radiology*, vol. 248, pp. 20–43, Jul 2008.
- [66] W. T. Dixon, L. N. Du, D. D. Faul, M. Gado, and S. Rossnick, "Projection angiograms of blood labeled by adiabatic fast passage," *Magn Reson Med*, vol. 3, no. 3, pp. 454–62, 1986.
- [67] D. G. Nishimura, A. Macovski, J. M. Pauly, and S. M. Conolly, "MR angiography by selective inversion recovery," *Magn Reson Med*, vol. 4, pp. 193–202, Feb 1987.
- [68] S. J. Wang, D. G. Nishimura, and A. Macovski, "Multiple-readout selective inversion recovery angiography," *Magn Reson Med*, vol. 17, pp. 244–51, Jan 1991.
- [69] R. R. Edelman, B. Siewert, M. Adamis, J. Gaa, G. Laub, and P. Wielopolski, "Signal targeting with alternating radiofrequency (STAR) sequences: application to MR angiography," *Magn Reson Med*, vol. 31, pp. 233–8, Feb 1994.
- [70] M. Günther, C. Warmuth, and C. Zimmer, "Sub-millimeter dynamic spin labeling cerebral 2D-angiography with 40ms temporal resolution," in *Proceedings of the 10th Annual Meeting of ISMRM*, p. 1100, 2002.

REFERENCES

- [71] M. J. van Osch, J. Hendrikse, X. Golay, C. J. Bakker, and J. van der Grond, "Non-invasive visualization of collateral blood flow patterns of the circle of willis by dynamic MR angiography," *Med Image Anal*, vol. 10, no. 1, pp. 59–70, 2006.
- [72] X. Bi, P. Weale, P. Schmitt, S. Zuehlsdorff, and R. Jerecic, "Non-contrast-enhanced four-dimensional (4D) intracranial MR angiography: a feasibility study," *Magn Reson Med*, vol. 63, pp. 835–41, Mar 2010.
- [73] P. Schmitt, P. Speier, X. Bi, P. Weale, and E. Mueller, "Non-contrast-enhanced 4D intracranial MR angiography: Optimizations using a variable flip angle approach," in *Proceedings 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Stockholm*, p. 402, 2010.
- [74] M. Wintermark, M. Sesay, E. Barbier, K. Borbely, W. P. Dillon, J. D. Eastwood, T. C. Glenn, C. B. Grandin, S. Pedraza, J. F. Soustiel, T. Nariai, G. Zaharchuk, J. M. Caille, V. Dousset, and H. Yonas, "Comparative overview of brain perfusion imaging techniques," *Stroke*, vol. 36, no. 9, pp. e83–99, 2005.
- [75] M. Ebinger, D. A. De Silva, S. Christensen, M. W. Parsons, R. Markus, G. A. Donnan, and S. M. Davis, "Imaging the penumbra - strategies to detect tissue at risk after ischemic stroke," *J Clin Neurosci*, vol. 16, pp. 178–87, Feb 2009.
- [76] P. Schramm, P. D. Schellinger, E. Klotz, K. Kallenberg, J. B. Fiebach, S. Küllkens, S. Heiland, M. Knauth, and K. Sartor, "Comparison of perfusion computed tomography and computed tomography angiography source images with perfusion-weighted imaging and diffusion-weighted imaging in patients with acute stroke of less than 6 hours' duration," *Stroke*, vol. 35, pp. 1652–8, Jul 2004.
- [77] A. E. Baird, M. C. Austin, W. J. McKay, and G. A. Donnan, "Sensitivity and specificity of 99mTc-HMPAO SPECT cerebral perfusion measurements during the first 48 hours for the localization of cerebral infarction," *Stroke*, vol. 28, pp. 976–80, May 1997.
- [78] J. C. Baron, "Clinical use of positron emission tomography in cerebrovascular diseases," *Neurosurg Clin N Am*, vol. 7, pp. 653–64, Oct 1996.
- [79] P. A. Barber, D. G. Darby, P. M. Desmond, Q. Yang, R. P. Gerraty, D. Jolley, G. A. Donnan, B. M. Tress, and S. M. Davis, "Prediction of stroke outcome with echoplanar perfusion- and diffusion-weighted MRI," *Neurology*, vol. 51, pp. 418–26, Aug 1998.
- [80] M. Essig, R. von Kummer, T. Egelhof, R. Winter, and K. Sartor, "Vascular MR contrast enhancement in cerebrovascular disease," *AJNR Am J Neuroradiol*, vol. 17, pp. 887–94, May 1996.

- [81] G. Zaharchuk, R. Bammer, M. Straka, A. Shankaranarayan, D. C. Alsop, N. J. Fischbein, S. W. Atlas, and M. E. Moseley, "Arterial spin-label imaging in patients with normal bolus perfusion-weighted MR imaging findings: pilot identification of the borderzone sign," *Radiology*, vol. 252, pp. 797–807, Sep 2009.
- [82] F. F. Paiva, A. Tannus, and A. C. Silva, "Measurement of cerebral perfusion territories using arterial spin labelling," *NMR Biomed*, vol. 20, no. 7, pp. 633–42, 2007.
- [83] J. A. Detre, W. Zhang, D. A. Roberts, A. C. Silva, D. S. Williams, D. J. Grandis, A. P. Koretsky, and J. S. Leigh, "Tissue specific perfusion imaging using arterial spin labeling," *NMR Biomed*, vol. 7, no. 1-2, pp. 75–82, 1994.
- [84] R. Werner, K. Alfke, T. Schaeffter, A. Nabavi, H. M. Mehdorn, and O. Jansen, "Brain perfusion territory imaging applying oblique-plane arterial spin labeling with a standard send/receive head coil," *Magn Reson Med*, vol. 52, no. 6, pp. 1443–7, 2004.
- [85] R. Werner, D. G. Norris, K. Alfke, H. M. Mehdorn, and O. Jansen, "Continuous artery-selective spin labeling (CASSL)," *Magn Reson Med*, vol. 53, no. 5, pp. 1006–12, 2005.
- [86] M. Helle, S. Rüfer, K. Alfke, O. Jansen, and D. G. Norris, "Perfusion territory imaging of intracranial branching arteries - optimization of continuous artery-selective spin labeling (CASSL)," *NMR Biomed*, Nov 2010.
- [87] N. P. Davies and P. Jezzard, "Selective arterial spin labeling (SASL): perfusion territory mapping of selected feeding arteries tagged using two-dimensional radiofrequency pulses," *Magn Reson Med*, vol. 49, no. 6, pp. 1133–42, 2003.
- [88] G. Furst, H. Steinmetz, H. Fischer, B. Skutta, M. Sitzler, A. Aulich, T. Kahn, and U. Modder, "Selective MR angiography and intracranial collateral blood flow," *J Comput Assist Tomogr*, vol. 17, no. 2, pp. 178–83, 1993.
- [89] J. Hendrikse, J. van der Grond, H. Lu, P. C. van Zijl, and X. Golay, "Flow territory mapping of the cerebral arteries with regional perfusion MRI," *Stroke*, vol. 35, no. 4, pp. 882–7, 2004.
- [90] X. Golay, E. T. Petersen, and F. Hui, "Pulsed star labeling of arterial regions (PULSAR): a robust regional perfusion technique for high field imaging," *Magn Reson Med*, vol. 53, no. 1, pp. 15–21, 2005.
- [91] S. M. Chng, E. T. Petersen, I. Zimine, Y.-Y. Sitoh, C. C. T. Lim, and X. Golay, "Territorial arterial spin labeling in the assessment of collateral circulation: comparison with digital subtraction angiography," *Stroke*, vol. 39, pp. 3248–54, Dec 2008.

REFERENCES

- [92] M. Gunther, M. Bock, and L. R. Schad, "Arterial spin labeling in combination with a look-locker sampling strategy: inflow turbo-sampling EPI-FAIR (ITS-FAIR)," *Magn Reson Med*, vol. 46, no. 5, pp. 974–84, 2001.
- [93] W. Dai, P. M. Robson, A. Shankaranarayanan, and D. C. Alsop, "Modified pulsed continuous arterial spin labeling for labeling of a single artery," *Magn Reson Med*, vol. 64, pp. 975–82, Oct 2010.
- [94] P. M. Robson, W. Dai, A. Shankaranarayanan, N. M. Rofsky, and D. C. Alsop, "Time-resolved vessel-selective digital subtraction MR angiography of the cerebral vasculature with arterial spin labeling," *Radiology*, Sep 2010.
- [95] M. Helle, D. G. Norris, S. Rüfer, K. Alfke, O. Jansen, and M. J. P. van Osch, "Superselective pseudocontinuous arterial spin labeling," *Magn Reson Med*, vol. 64, pp. 777–86, Sep 2010.
- [96] I. Zimine, E. T. Petersen, and X. Golay, "Dual vessel arterial spin labeling scheme for regional perfusion imaging," *Magn Reson Med*, vol. 56, no. 5, pp. 1140–4, 2006.
- [97] M. Gunther, "Efficient visualization of vascular territories in the human brain by cycled arterial spin labeling MRI," *Magn Reson Med*, vol. 56, no. 3, pp. 671–5, 2006.
- [98] M. A. Chappell, T. W. Okell, P. Jezzard, and M. W. Woolrich, "A general framework for the analysis of vessel encoded arterial spin labeling for vascular territory mapping," *Magn Reson Med*, vol. 64, pp. 1529–39, Nov 2010.
- [99] B. Wu, X. Wang, J. Guo, S. Xie, E. C. Wong, J. Zhang, X. Jiang, and J. Fang, "Collateral circulation imaging: MR perfusion territory arterial spin-labeling at 3T," *AJNR Am J Neuroradiol*, vol. 29, pp. 1855–60, Nov 2008.
- [100] A. P. Kansagra and E. C. Wong, "Mapping of vertebral artery perfusion territories using arterial spin labeling MRI," *J Magn Reson Imaging*, vol. 28, no. 3, pp. 762–6, 2008.
- [101] E. Wong and A. Kansagra, "Mapping middle cerebral artery branch territories with vessel encoded pseudo-continuous ASL: Sine/cosine tag modulation and data clustering in tagging efficiency space," in *Proceedings 16th Scientific Meeting of the ISMRM*, p. 182, 2008.
- [102] G. J. Stanisz, E. E. Odobina, J. Pun, M. Escaravage, S. J. Graham, M. J. Bronskill, and R. M. Henkelman, "T1, T2 relaxation and magnetization transfer in tissue at 3T," *Magn Reson Med*, vol. 54, no. 3, pp. 507–512, 2005 Sep.
- [103] B. Sui, P. Gao, Y. Lin, B. Gao, L. Liu, and J. An, "Blood flow pattern and wall shear stress in the internal carotid arteries of healthy subjects," *Acta Radiologica*, vol. 49, no. 7, pp. 806–814, 2008.

- [104] F. M. A. Box, R. J. van der Geest, J. van der Grond, M. J. P. van Osch, A. H. Zwinderman, I. H. Palm-Meinders, J. Doornbos, G.-J. Blauw, M. A. van Buchem, and J. H. C. Reiber, "Reproducibility of wall shear stress assessment with the paraboloid method in the internal carotid artery with velocity encoded MRI in healthy young individuals.," *J Magn Reson Imaging*, vol. 26, no. 3, pp. 598–605, 2007 Sep.
- [105] E. B. Ringelstein, B. Kahlscheuer, E. Niggemeyer, and S. M. Otis, "Transcranial doppler sonography: anatomical landmarks and normal velocity values.," *Ultrasound Med Biol*, vol. 16, no. 8, pp. 745–761, 1990.
- [106] W. Sorteberg, I. A. Langmoen, K. F. Lindegaard, and H. Nornes, "Side-to-side differences and day-to-day variations of transcranial doppler parameters in normal subjects.," *J Ultrasound Med*, vol. 9, no. 7, pp. 403–409, 1990 Jul.
- [107] P. Kalvach, D. Gregova, O. Skoda, T. Peisker, R. Tumova, J. Termerova, and J. Korska, "Cerebral blood supply with aging: normal, stenotic and recanalized.," *J Neurol Sci*, vol. 257, no. 1-2, pp. 143–148, 2007 Jun 15.
- [108] H. H. Hu and K. S. Nayak, "Quantification of absolute fat mass using an adipose tissue reference signal model.," *J Magn Reson Imaging*, vol. 28, pp. 1483–1491, Dec 2008.
- [109] H. E. Möller, H. K. Klocke, G. M. Bongartz, and P. E. Peters, "MR flow quantification using phase-contrast MRI: clinical application to the carotid arteries," *J Magn Reson Imaging*, vol. 6, no. 3, pp. 503–12, 1996.
- [110] A. Kansagra and E. Wong, "Automated segmentation of multiple vascular territories from vessel encoded pseudo-continuous arterial spin labeling MRI data," in *Proceedings 17th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Honolulu*, p. 3652, 2009.
- [111] M. A. Chappell, T. W. Okell, P. Jezzard, and M. W. Woolrich, "Vascular territory image analysis using vessel encoded arterial spin labeling," *Med Image Comput Comput Assist Interv*, vol. 12, no. Pt 2, pp. 514–21, 2009.
- [112] M. A. Chappell, T. W. Okell, P. Jezzard, M. W. Woolrich, and S. J. Payne, "Fast analysis of vessel encoded ASL perfusion and angiographic images for clinical applications," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, p. 366, 2011.
- [113] P. Mansfield, "Multi-planar image formation using NMR spin echoes," *Journal of Physics C: Solid State Physics*, vol. 10, p. L55, 1977.
- [114] M. Gunther, K. Oshio, and D. A. Feinberg, "Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements," *Magnetic Resonance in Medicine*, vol. 54, no. 2, pp. 491–498, 2005.

REFERENCES

- [115] A. Jesmanowicz, E. Wong, and J. Hyde, "Phase correction for EPI using internal reference lines," in *Proceedings of the 12th Annual Meeting of SMRM, New York*, 1993.
- [116] B. J. MacIntosh, N. Filippini, M. A. Chappell, M. W. Woolrich, C. E. Mackay, and P. Jezzard, "Assessment of arterial arrival times derived from multiple inversion time pulsed arterial spin labeling MRI," *Magn Reson Med*, vol. 63, pp. 641–7, Mar 2010.
- [117] A. P. Kansagra and E. C. Wong, "Quantitative assessment of mixed cerebral vascular territory supply with vessel encoded arterial spin labeling MRI," *Stroke*, 2008.
- [118] M. Jenkinson, P. Bannister, M. Brady, and S. Smith, "Improved optimization for the robust and accurate linear registration and motion correction of brain images," *NeuroImage*, vol. 17, no. 2, pp. 825–841, 2002.
- [119] D. C. Alsop and J. A. Detre, "Reduced transit-time sensitivity in noninvasive magnetic resonance imaging of human cerebral blood flow.," *J Cereb Blood Flow Metab*, vol. 16, pp. 1236–1249, Nov 1996.
- [120] D. Gallichan and P. Jezzard, "Variation in the shape of pulsed arterial spin labeling kinetic curves across the healthy human brain and its implications for CBF quantification.," *Magn Reson Med*, vol. 61, pp. 686–695, Mar 2009.
- [121] T. W. Okell, M. A. Chappell, M. Günther, and P. Jezzard, "A comparison of 3D-GRASE and EPI for vessel-encoded arterial spin labeling," in *Proceedings 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Stockholm*, p. 4071, 2010.
- [122] F. Q. Ye, J. A. Frank, D. R. Weinberger, and A. C. McLaughlin, "Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST)," *Magn Reson Med*, vol. 44, no. 1, pp. 92–100, 2000.
- [123] G. Krüger and G. H. Glover, "Physiological noise in oxygenation-sensitive magnetic resonance imaging," *Magn Reson Med*, vol. 46, pp. 631–7, Oct 2001.
- [124] S. Mani, J. Pauly, S. Conolly, C. Meyer, and D. Nishimura, "Background suppression with multiple inversion recovery nulling: applications to projective angiography," *Magn Reson Med*, vol. 37, pp. 898–905, Jun 1997.
- [125] D. M. Garcia, G. Duhamel, and D. C. Alsop, "Efficiency of inversion pulses for background suppressed arterial spin labeling," *Magn Reson Med*, vol. 54, pp. 366–72, Aug 2005.
- [126] M. Silver, R. Joseph, and D. Hoult, "Highly selective $\pi/2$ and π pulse generation," *Journal of magnetic resonance*, vol. 59, no. 2, pp. 347–351, 1984.

- [127] G. McGibney, M. Smith, S. Nichols, and A. Crawley, "Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI," *Magnetic Resonance in Medicine*, vol. 30, no. 1, pp. 51–59, 1993.
- [128] D. J. Larkman and R. G. Nunes, "Parallel magnetic resonance imaging," *Phys Med Biol*, vol. 52, pp. R15–55, Apr 2007.
- [129] S. M. Smith, "Fast robust automated brain extraction," *Hum Brain Mapp*, vol. 17, pp. 143–55, Nov 2002.
- [130] B. A. Poser, P. J. Koopmans, T. Witzel, L. L. Wald, and M. Barth, "Three dimensional echo-planar imaging at 7 Tesla," *Neuroimage*, vol. 51, pp. 261–6, May 2010.
- [131] K. Oshio and D. A. Feinberg, "GRASE (gradient- and spin-echo) imaging: A novel fast MRI technique," *Magnetic Resonance in Medicine*, vol. 20, no. 2, pp. 344–349, 1991.
- [132] D. A. Feinberg and K. Oshio, "GRASE (gradient- and spin-echo) MR imaging: A new fast clinical imaging technique," *Radiology*, vol. 181, no. 2, pp. 597–602, 1991.
- [133] P. J. van Laar, J. Hendrikse, X. Golay, H. Lu, M. J. van Osch, and J. van der Grond, "In vivo flow territory mapping of major brain feeding arteries," *Neuroimage*, vol. 29, no. 1, pp. 136–44, 2006.
- [134] T. W. Okell, M. A. Chappell, M. W. Woolrich, M. Günther, D. A. Feinberg, and P. Jezzard, "Vessel-encoded dynamic magnetic resonance angiography using arterial spin labeling," *Magn Reson Med*, vol. 64, pp. 698–706, Sep 2010.
- [135] D. Look and D. Locker, "Time saving in measurement of NMR and EPR relaxation times," *Review of Scientific Instruments*, vol. 41, no. 2, pp. 250–251, 1970.
- [136] D. O. Walsh, A. F. Gmitro, and M. W. Marcellin, "Adaptive reconstruction of phased array MR imagery," *Magn Reson Med*, vol. 43, pp. 682–90, May 2000.
- [137] C. Warmuth, M. Ruping, A. Forschler, H. C. Koennecke, J. M. Valdueza, A. Kauert, S. J. Schreiber, R. Siekmann, and C. Zimmer, "Dynamic spin labeling angiography in extracranial carotid artery stenosis," *AJNR Am J Neuroradiol*, vol. 26, no. 5, pp. 1035–43, 2005.
- [138] P. M. Robson, W. Dai, A. Shankaranarayanan, N. M. Rofsky, and D. C. Alsop, "Time-resolved vessel-selective digital subtraction MR angiography of the cerebral vasculature with arterial spin labeling," *Radiology*, vol. 257, pp. 507–15, Nov 2010.

REFERENCES

- [139] A. W. Hoksbergen, B. Fülesdi, D. A. Legemate, and L. Csiba, "Collateral configuration of the circle of willis: transcranial color-coded duplex ultrasonography and comparison with postmortem anatomy," *Stroke*, vol. 31, pp. 1346–51, Jun 2000.
- [140] T. W. Okell, P. Schmitt, X. Bi, M. A. Chappell, R. H. Tijssen, K. L. Miller, and P. Jezzard, "4D vessel-encoded arterial spin labeling angiography," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, p. 4034, 2011.
- [141] K. L. Miller, R. H. Tijssen, N. Stikov, and T. W. Okell, "Steady-state mri: methods for neuroimaging," *Imaging in Medicine*, vol. 3, no. 1, pp. 93–105, 2011.
- [142] M. Deimling and O. Heid, "Magnetization prepared true FISP imaging," in *Proceedings 2nd Annual Meeting, Society for Magnetic Resonance, San Francisco*, p. 495, 1994.
- [143] O. Bieri and K. Scheffler, "Flow compensation in balanced SSFP sequences," *Magn Reson Med*, vol. 54, pp. 901–7, Oct 2005.
- [144] M. Markl, M. T. Alley, C. J. Elkins, and N. J. Pelc, "Flow effects in balanced steady state free precession imaging," *Magn Reson Med*, vol. 50, pp. 892–903, Nov 2003.
- [145] D. Nishimura and S. Vasanawala, "Analysis and reduction of the transient response in SSFP imaging," in *Proceedings 8th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Denver*, p. 301, 2000.
- [146] B. A. Hargreaves, S. S. Vasanawala, J. M. Pauly, and D. G. Nishimura, "Characterization and reduction of the transient response in steady-state MR imaging," *Magn Reson Med*, vol. 46, pp. 149–58, Jul 2001.
- [147] T. Smith, Z. Zun, E. C. Wong, and K. S. Nayak, "Design and use of variable flip angle schedules in transient balanced SSFP subtractive imaging," *Magn Reson Med*, vol. 63, pp. 537–42, Feb 2010.
- [148] T. W. Okell, M. A. Chappell, U. Schulz, and P. Jezzard, "Quantification of vessel-encoded arterial spin labeling dynamic angiography with auto-calibration," in *Proceedings 19th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Montreal*, p. 3405, 2011.
- [149] M. Rausch, K. Scheffler, M. Rudin, and E. W. Radü, "Analysis of input functions from different arterial branches with gamma variate functions and cluster analysis for quantitative blood volume measurements," *Magn Reson Imaging*, vol. 18, pp. 1235–43, Dec 2000.
- [150] D. J. C. MacKay, *Information theory, inference, and learning algorithms*. Cambridge, UK: Cambridge University Press, 2003.

-
- [151] K.-V. Yuen, *Bayesian methods for structural dynamics and civil engineering*. Singapore: John Wiley and Sons, 2010.
- [152] P. J. van Laar, J. van der Grond, and J. Hendrikse, "Brain perfusion territory imaging: methods and clinical applications of selective arterial spin-labeling MR imaging," *Radiology*, vol. 246, pp. 354–64, Feb 2008.
- [153] J. Hendrikse, M. J. Hartkamp, B. Hillen, W. P. Mali, and J. van der Grond, "Collateral ability of the circle of willis in patients with unilateral internal carotid artery occlusion: border zone infarcts and clinical symptoms," *Stroke*, vol. 32, no. 12, pp. 2768–73, 2001.
- [154] P. J. van Laar, J. van der Grond, J. P. Bremmer, C. J. M. Klijn, and J. Hendrikse, "Assessment of the contribution of the external carotid artery to brain perfusion in patients with internal carotid artery occlusion," *Stroke*, vol. 39, pp. 3003–8, Nov 2008.
- [155] X. Golay, J. Hendrikse, and J. Van Der Grond, "Application of regional perfusion imaging to extra-intracranial bypass surgery and severe stenoses," *J Neuroradiol*, vol. 32, no. 5, pp. 321–4, 2005.
- [156] J. Guo and E. C. Wong, "From optimized vessel encoded PCASL (opt-VEPCASL) to randomly-encoded VEPCASL (re-VEPCASL)," in *Proceedings 18th Scientific Meeting, International Society for Magnetic Resonance in Medicine, Stockholm*, p. 4069, 2010.
- [157] P. Le Roux, "Simplified model and stabilization of SSFP sequences," *J Magn Reson*, vol. 163, pp. 23–37, Jul 2003.
- [158] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn Reson Med*, vol. 58, pp. 1182–95, Dec 2007.