Development of an optimised non-invasive MRI method to measure renal perfusion in patients with impaired renal function

Fábio Rui Alves Nery

University College London

Developmental Imaging and Biophysics Section

UCL Great Ormond Street Institute of Child Health

London, U.K.

PhD Thesis

Submitted for Doctor of Philosophy Degree

2017

Declaration

I, Fábio Nery, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Fábio Nery

Abstract

Arterial Spin Labelling (ASL) is a unique Magnetic Resonance Imaging (MRI) approach for quantifying tissue perfusion non-invasively. More than two decades of technical developments established ASL as a valuable tool in neuroimaging, having more recently began its translation to the clinic. ASL holds great potential for the assessment of kidney disease given that it does not require contrast agents which are typically contraindicated for patients with impaired renal function. However, renal ASL applications remain limited and the technique has yet to be incorporated into clinical practice. The sensitivity of ASL to patient movement, which severely corrupts the renal perfusion estimates, is arguably one of the greatest factors hindering a wide adoption of this technique.

This thesis begins with an overview of the main concepts addressed in this work (kidney physiology, MRI and ASL) and a thorough literature review of previous renal ASL work. The problem of patient movement is then addressed at all levels of the ASL framework by combining a motion-insensitive ASL acquisition scheme with a specifically tailored image processing pipeline. The feasibility of this technique to provide repeatable renal perfusion measurements is demonstrated in the first paediatric cohort with impaired renal function to undergo renal ASL. Finally, the critical findings of this thesis are summarised and prospective future research directions are outlined.

Acknowledgements

Thanks to Danilo Maziero, my brother from another mother. Danilo has a contagious passion for life and has the ability of bringing people together and connecting with everyone. From the long walks in London to him throwing me into a stage in Hawaii to play music with complete strangers, I will always appreciate the time we spent together – I am looking forward to our future shenanigans.

To Luís Lacerda, from geeking out with me about anything and everything on a daily basis to engaging into deep conversations over endless cups of coffee, he has been a true friend and helped me cope with the frenzied pace of the final year of my PhD. I hope we continue to work together for years to come.

I am lucky to have had a great support system not only in the UK but also abroad. In particular, I would like to thank my good friends André Ferreira, André Gaspar, Luís Martins and Tiago Carvalho for keeping in close touch despite the distance. Our reunions around Christmas time are a precious tradition for me and never fail to bring fond memories of times spent in Coimbra back to life. A special thanks to Miguel Patrício for always being there to listen to my ramblings and sharing valuable lessons on life and research.

I find it impossible to accurately articulate just how lucky I am to have such a wonderful family. Everything I am today, I owe to them. My parents have been a source of unwavering support and love for every second of my existence. I hope I have, and continue to make you proud. My grandparents and my godparents are another great source of strength and motivation, I can only apologize to all of you for rarely going back home. Hugo, my thanks to you as well. I wish you every success in your endeavours.

Finally, to Andreia. You are the reason I moved to the UK, which was the second best decision I have made in my life (the first was the start of our relationship). You have been my love and my best friend for almost half of my life. I cannot thank you enough for your patience, for staying by my side every second of this journey, for encouraging me to live a balanced life and making me a better person. I love and trust you with all my heart. Thanks for sticking around. The best is yet to come.

Table of Contents

Table of Figures

Table of Tables

Abbreviations

ASL	Arterial Spin Labelling
ATN	Acute Tubular Necrosis
BH	Breath-hold
BOLD	Blood Oxygenation Level-Dependent
BS	Background Suppression
CASL	Continuous ASL
CBF	Cerebral Blood Flow
CC	Correlation Coefficient
CKD	Chronic Kidney Disease
CV	Coefficient of Variation
DCE	Dynamic Contrast-Enhanced
ED	Euclidean Distance
eGFR	Estimated Glomerular Filtration Rate
EPI	Echo Planar Imaging
EPISTAR	Echo-Planar Imaging and Signal Targeting with Alternating Radio frequency
ETD	Echo-Train Duration
ETL	Echo Train Length
FAIR	Flow-sensitive Alternating Inversion Recovery
FID	Free Induction Decay
FOV	Field of View
FRP	Functional Renal Parenchyma
FSE	Fast Spin-Echo
FT	Fourier Transform
GFR	Glomerular Filtration Rate
GRASE	Gradient and Spin Echo
HASTE	Half-Fourier Acquisition Single-Shot Turbo Spin Echo
ICC	Intra-class Correlation Coefficient
LOA	Limits of Agreement
LS	Loop Structure
MRI	Magnetic Resonance Imaging
MT	Magnetisation Transfer
NMR	Nuclear Magnetic Resonance
NS	Non Slice-selective
PAH	Para-aminohippurate
PASL	Pulsed ASL
pCASL	Pseudo-continuous ASL
PD	Proton-density
PD-PWI	Proton-density Perfusion-weighted Fusion Image
PE	Phase encoding
PLD	Post-label Delay
PWI	Perfusion-weighted Image
Q2TIPS	QUIPSS II with Thin-Slice Periodic Saturation
QUIPSS	Quantitative Imaging of Perfusion Using a Single Subtraction
RARE	Rapid Acquisition with Relaxation Enhancement
RAS	Renal Artery Stenosis
RBF	Renal Blood Flow

RF	Radio-frequency
RMSE	Root Mean Squared Error
ROI	Region of Interest
RPF	Renal Plasma Flow
SAR	Specific Absorption Rate
SCr	Serum Creatinine
SNR	Signal-to-Noise Ratio
SR	Saturation Recovery
SS	Slice-selective
STD	Standard Deviation
T_1	Longitudinal relaxation time
T_2	Transverse Relaxation Time
T_2^*	Apparent Transverse Relaxation Time
TD	Post-saturation Delay
TE	Echo Time
TI	Inflow Time
TR	Repetition Time
tSNR	Temporal Signal-to-Noise Ratio
tSTD	Temporal Standard Deviation
WK	Whole-kidney
WSCV	Within-subject Coefficient of Variation

1 Context

The aim of this thesis is to develop an optimised non-invasive Magnetic Resonance Imaging (MRI) method to measure renal perfusion in paediatric patients with impaired renal function. In particular, the technique known as Arterial Spin Labelling (ASL), originally proposed by Detre et al. will be employed (Detre et al. 1992; Williams et al. 1992).

This chapter provides context and the motivation for the work carried out. It begins with an introduction to the basics of renal anatomy and physiology, with special focus on describing the blood circulation within the kidneys. The problem of Chronic Kidney Disease (CKD) is stated, together with a discussion on the pressing need for new approaches to obtain *in vivo* functional information of the human kidneys, both in healthy and diseased states*.*

Given the specific focus of this work, I will review existing methods for measuring Renal Blood Flow (RBF) and some of their drawbacks. Finally, ASL is briefly introduced as a new method for the non-invasive *in vivo* measurement of renal perfusion.

1.1 Basic Renal Anatomy and Physiology

The kidneys are a pair of intra-abdominal organs, located near the posterior abdominal wall, behind the peritoneum. Each adult human kidney has a weight ranging from 115-170g, with approximate length, width and thickness ranging from 11-12cm, 5.0-7.5cm and 2.5-3.0cm, respectively (Taal et al. 2012). The main anatomical features of the kidneys are depicted in Figure 1-1. Two main regions can be identified upon dissection of a kidney. The outer region, the cortex, with approximately 1cm thickness (Taal et al. 2012), receives the majority of the RBF and is where ultrafiltration occurs. The darker inner region, the medulla, contains the renal pyramids which extend towards the renal calyces. Urine passes through these, and continues towards the renal pelvis, where the calyces converge. The urine then finally leaves the kidneys through the ureters.

The kidneys are responsible for a range of intricate and crucial tasks which are key to the overall normal function of the human body. These include a major role in the regulation of fluid volume within the internal environment, achieved by tightly controlling the ion and acid-base balance and adjusting the amount of water lost in urine. They also filter the blood, removing waste products as well as potentially toxic foreign chemicals, while avoiding the excretion of useful substances such as glucose or protein. They also have an important role in the production and secretion of hormones (e.g. erythropoietin) and enzymes (e.g. renin). Finally, the kidney plays an important role in gluconeogenesis (Gerich et al. 2001).
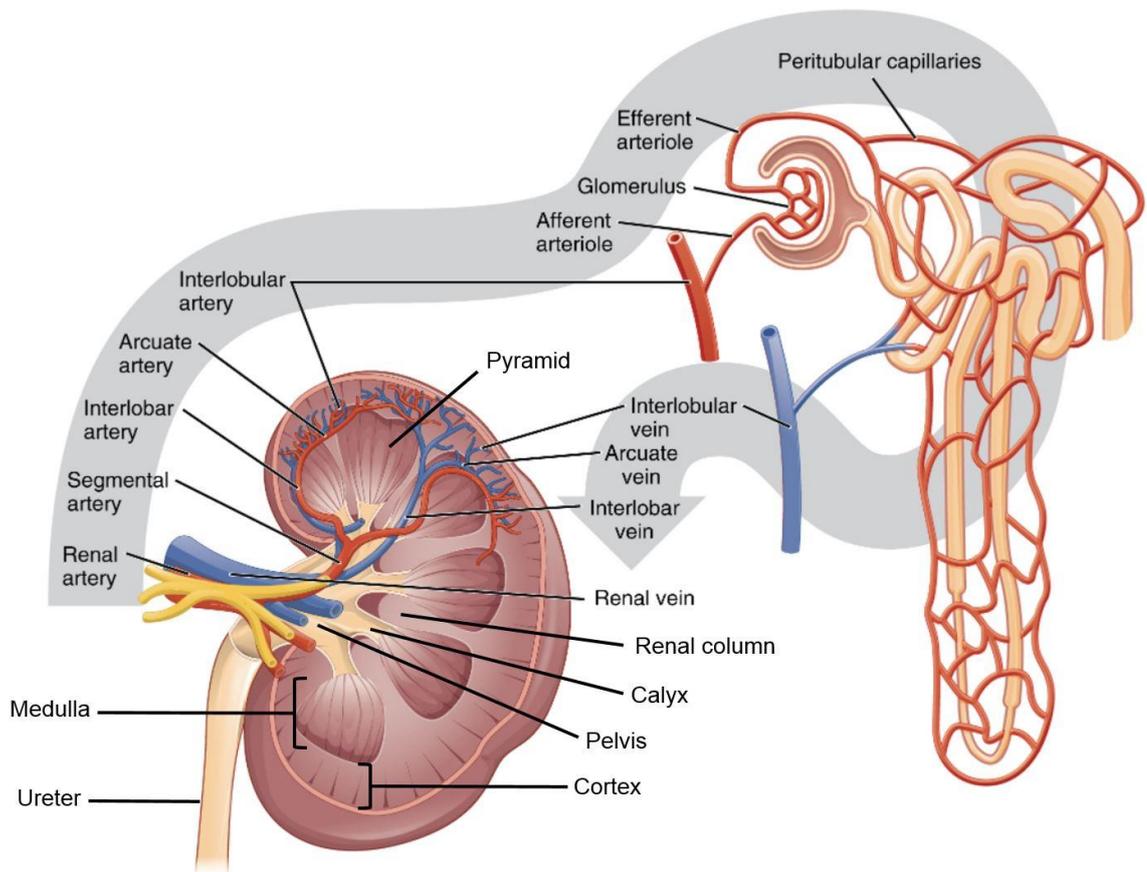
Figure 1-1 Renal Anatomy, with focus on the main arteries and veins. Adapted from (OpenStax 2016) under a Creative Commons Attribution 4.0 International License.

1.1.1 Blood circulation in the kidneys

The kidneys are extremely well perfused organs. Even though they make up less than 1% of the total body mass, they receive 20% to 25% of the total cardiac output (Regan et al. 1995), corresponding to approximately 1–1.2L/min in healthy adults (Taal et al. 2012). This amount far surpasses their metabolic requirements. However, it is key to enable high glomerular filtration rates to ensure appropriate waste removal and a tight control over the blood volume and concentration of solutes in the body.

The renal cortex and medulla have different physiological roles and metabolic demands. Accordingly, there are regional variations in blood flood throughout the kidney to meet the specific requirements of these regions. As it arrives into the kidney, all of the inflowing blood first reaches the renal cortex. Blood is then supplied to the renal medulla by the efferent arterioles of the deep nephrons, closest to the renal medulla (juxtamedullary nephrons) (Taal et al. 2012). Blood flow also varies within the medulla, where per unit tissue it ranges from about 40% to 10% in the outer and inner medulla, respectively, relative to that of the cortex (Evans et al. 2004).

Blood is supplied to the kidneys by one of the branches of the abdominal aorta, the renal artery. Upon entering the kidney, the renal arteries divide into segmental arteries, which in turn branch into interlobar arteries. These traverse the renal columns and as they reach the renal cortex they further divide into arcuate arteries, which in turn branch into interlobular arteries. These finally give rise to the afferent arterioles which supply the nephrons, which will be described later.

Each afferent arteriole leads into an intertwined network of capillaries, known as a glomerulus, which together with the Bowman's capsule forms the renal corpuscle, which acts as a filter for the inflowing blood. The blood then exits the glomerulus through the efferent arteriole. Depending on the location of the nephron (superficial or deep cortex), the efferent arteriole leads into different vessel networks. In superficial cortical nephrons (approximately 85% of the total number of nephrons in the kidneys (Saladin et al. 2016)), the efferent arteriole leads to a plexus of peritubular capillaries. In nephrons located in the deep cortex, closest to the medulla (juxtamedullary nephrons), the efferent arterioles descend to the renal medulla giving rise to long capillary loops, the vasa recta. The peritubular capillaries then reunite forming the interlobular veins, from which the blood flows through to the arcuate veins. The vasa recta drain into the arcuate veins. From these, the blood then proceeds to the interlobar veins and then finally leaves the kidney through a single renal vein and continues towards the vena cava. Figure 1-1 highlights the main arteries and veins involved in the blood circulation in the kidneys.

1.1.2 The Nephron

The nephron is the functional unit of the kidney. The number of nephrons in human kidneys averages approximately 1 million per kidney, and is fixed at birth. There is nevertheless a wide variation in their number among subjects (exceeding a 10-fold range (Bertram et al. 2011)). Interestingly, recent studies have been addressing the feasibility of using MRI for counting the number of glomeruli (and thereby nephrons) in the rat kidney *ex vivo* (Beeman et al. 2011; Heilmann et al. 2012).

Two main components are part of the nephron: the renal corpuscle, where the blood filtration takes place, and the renal tubules through which the resulting filtered fluid flows (also referred to as glomerular filtrate or ultrafiltrate). These and their main roles will be briefly described in the following sections.

1.1.2.1 Renal corpuscle: filtration

The difference in pressures in the afferent and efferent arterioles (at the vascular pole of the renal corpuscle) is largely responsible for driving the ultrafiltration process. The shorter

diameter of the efferent arteriole compared to the afferent arteriole results in a very high hydrostatic pressure in the glomerulus. Both the hydrostatic pressure exerted in the Bowman's capsule by the filtrate and the oncotic pressure caused by proteins in blood are not enough to counteract the glomerular hydrostatic pressure. This forces fluid to move into the capsular space. Relatively large components, such as blood cells and large proteins are unable to traverse the filtration barriers in the glomerulus and therefore remain in the bloodstream. On the other hand, water and small molecules such as glucose, amino acids and urea (among others) readily cross these barriers and move into the glomerular space, making up the ultrafiltrate. This fluid, which already differs significantly from the blood in terms of its composition, then continues to move through the renal tubule, where further changes to its composition will occur. The rate at which ultrafiltrate is generated in all glomeruli in both kidneys is known as glomerular filtration rate (GFR). Normal values for GFR in adults are approximately 125mL/min (Taal et al. 2012) and decrease with age (Delanaye et al. 2012). The GFR is kept at a relative constant level despite changes in blood pressure, owing to several mechanisms both intrinsic (autoregulation) and extrinsic to the kidney (neural and hormonal regulation). These act mainly by regulating the amount of blood entering and leaving the glomerulus and adjusting the surface area of the glomerular capillaries involved in filtration (Tortora & Derrickson 2014).

Another important parameter to characterise renal function is the filtration fraction. This corresponds to the fraction of blood plasma which becomes glomerular filtrate. It is therefore computed as the ratio of GFR to renal plasma flow (RPF) and is hypothesized to reflect glomerular hydraulic pressure (Taal et al. 2012). A typical value for the filtration fraction in humans is 20% (Hall 2016). It is worth highlighting that from the approximately 1500L/day of blood reaching the kidney (of which around 750L/day would correspond to blood plasma), an excess of 150L/day of filtrate is formed. Given the small fraction of this which will correspond to the amount of urine produced, clearly there must exist reabsorption mechanisms in place to ensure that the vast majority of filtrate returns to circulation. These are described in the next section.

1.1.2.2   Renal tubule: reabsorption and secretion

The renal tubule initiates at the urinary pole of the renal corpuscle, opposite to the vascular pole where the afferent and efferent arterioles connect to the glomerulus. Its components are the proximal convoluted tubule, the loop of Henle, and the distal convoluted tubule. The distal convoluted tubules of independent nephrons then drain into the collecting duct, which transports urine to the calyces, which in turn drain into the ureter. Each segment of the renal

tubule has its own specialized roles in the production of urine. Two main basic processes occur: tubular reabsorption and tubular secretion. During reabsorption, filtered substances return to the bloodstream by moving into the peritubular capillaries. Proximal convoluted tubule cells, with their brush border to increase surface area, perform most of the reabsorption. Some of the solutes which are reabsorbed include glucose, amino acids and ions. The increase in osmolarity within the peritubular capillaries causes water to follow by osmosis. This is as opposed to water reabsorption that occurs in the collecting ducts, which is dependent on the presence of antidiuretic hormone. In fact, the reabsorption of both water and ions in later distal parts of the renal tubule (i.e. after the proximal convoluted tubule) is regulated by a variety of hormones. This allows reabsorption to occur in variable amounts which is key for ensuring a fine control over the volume and osmolarity of body fluids. Tubular secretion, essentially the opposite of reabsorption, is the process by which selected substances leave the bloodstream into the renal tubules. It therefore allows the elimination of substances that could not be filtered in the glomerulus (e.g. when such substances are attached to proteins). Unwanted substances that might have been reabsorbed passively, such as nitrogenous wastes are also secreted and make part of urine.

1.2   Chronic Kidney Disease and the need for new biomarkers

CKD is characterised by abnormalities in kidney function and usually structure, present for over 3 months, with implications for health (KDIGO 2013). The loss in renal function may be progressive and irreversible and in its most advanced stage leads to end-stage renal disease. At this point, renal replacement therapy is required for survival, in the form of dialysis or kidney transplantation. CKD is recognised as a global health burden affecting more than 10% of the population in many countries (Eckardt et al. 2013), including the United Kingdom (Archibald et al. 2007). Worldwide, the costs of treating patients with end-stage renal disease have been estimated to exceed US$1 trillion (Stenvinkel 2010). In England, more than 1% of the total National Health Service (NHS) spending in 2009–2010 was attributed to CKD, where renal replacement therapy accounted for over half of these costs (Kerr et al. 2012). It is expected that in the future this financial burden will continue to intensify as the incidence of CKD continues to increase (Hoerger et al. 2015). Contributing factors for this include population ageing, obesity and diabetes (Kainz et al. 2015). It is also a well-known fact that CKD patients are at a much higher risk of developing cardiovascular disease, compared to the general population (Sarnak et al. 2003; Daly 2007).

In clinical practice, CKD is classified according to the GFR and proteinuria levels. Kidney transplant recipients are considered to have CKD, regardless of their GFR levels or presence of

other markers of kidney damage (KDIGO 2013). Serum creatinine concentration is commonly used to estimate the GFR (Levey et al. 2009) (see Section 7.2.4.2 for the approach used to estimate GFR in this work). Recently other markers such as Cystatin C have been proposed to improve the precision in GFR estimation procedures (Inker et al. 2012). Nevertheless, estimated GFR (eGFR) measures obtained with creatinine-based equations remain inaccurate in certain patient groups (Stevens & Levey 2009). In fact, an early diagnosis may be hindered when relying on these methods as alterations in the GFR may only occur when irreversible damage to the kidneys has already taken place (Haufe et al. 2006; Artz 2010). Clearance studies using exogenous filtration markers may be necessary to confirm eGFR values (Stevens & Levey 2009), but many of these are costly, cumbersome to use in the clinic, or involve radiation exposure. As such, additional new biomarkers are necessary to allow for an earlier diagnosis and identification of individuals at higher risk of developing progressive CKD which can lead to kidney failure (James et al. 2010; Levin et al. 2011).

Advances in molecular biology combined with a very active field of research in renal biomarkers has resulted in the introduction of a wide array of new potential biomarkers over recent years (Fassett et al. 2011; Lopez-Giacoman & Madero 2015; Wasung et al. 2015). Medical imaging, in particular through MRI, has also been emerging as a promising complement to these, offering distinctive advantages and a unique wealth of information on renal physiology. In addition to the established capabilities of MRI in providing exquisite anatomical detail and tissue contrast (allowing for precise total kidney volume or independent cortex/medulla volume measurements), MRI is also capable of providing information on blood flow and perfusion (Becker & Rossi 2017), microstructure (Thoeny & De Keyzer 2011), and related to oxygenation (Niendorf et al. 2015), among other physiological parameters. Importantly, all these measures are obtained in a completely non-invasive manner and allow for a separate assessment of both kidneys and even among different regions of the same kidney.

1.3   Assessing Renal Perfusion

Renal perfusion relates to the delivery of oxygen and nutrients to tissue, by means of blood flow occurring on the level of its capillary bed. It is quantified as a volume of blood delivered per unit time and mass of tissue. Common units are ml/100g/min. In the kidneys, it is intimately connected to glomerular filtration and tissue oxygenation. Therefore, it is a crucial parameter in the pathophysiology of renal disease, making its measurement highly desirable in the clinical setting. Biomarkers related to renal haemodynamics have the potential of providing an earlier and more sensitive diagnosis. This is especially important given the drawbacks of current methods used in the clinic (as discussed in the previous section). As such, they would not only

assist in the determination of the best treatment options at an early stage, but also in providing the clinicians with a window into the progression of disease. Furthermore, they would offer a novel follow up method to monitor therapeutic interventions. This additional information might also have the unexpected benefit of providing further understanding on the pathogenesis of specific diseases affecting the kidney.

Investigating the differences in regional perfusion in the kidneys has long been of interest to physiologists (Aukland 1980; Knox et al. 1984). However, despite decades of research, developing methods with the ability to quantify absolute regional changes in renal perfusion is still a challenge (Pallone et al. 2012; Beierwaltes et al. 2013).

Para-aminohippurate (PAH) clearance studies are the classical gold standard approach for the estimation of effective RPF (Sandilands et al. 2013; Schneider et al. 2013). This approach is based in the fact that the vast majority of PAH is extracted from the blood, due to a combination of filtration at the glomerulus and especially excretion by the renal tubules. If one assumes the extraction to be 100% then the clearance of this indicator allows an estimation of RPF, hence the term "effective" renal plasma flow. The disadvantage of this approach, which makes it unsuitable for clinical practice, is its invasive and time-consuming nature. In addition, the estimation of RPF in renal disease may be confounded by damage in the renal tubules, which impairs the secretion process, rendering the assumption of 100% extraction of PAH invalid. Nevertheless, PAH clearance may serve as a benchmark to validate newer methods to measure renal perfusion in healthy volunteers. Indeed, recent work by Ritt et al. has shown a correlation between renal perfusion measurements using ASL MRI and RPF measurements using PAH (Ritt et al. 2010), before and after a pharmacological intervention.

Radio-labelled microspheres have been used to assess organ perfusion (Aukland 1980). However, this method is limited to animal studies as it involves tissue collection. Artz et al. found a good correlation between perfusion estimates in the renal cortex obtained with both microspheres and ASL in a swine model (Nathan S Artz et al. 2011).

Warmuth et al. validated renal ASL in the kidney in extracorporeal porcine kidneys, comparing perfusion estimates obtained with ASL and calibrated vessel flow, using an ultrasound flowmeter (Warmuth et al. 2007), after normalizing the ultrasound flow measurements per 100g of cortical tissue.

Scintigraphy has long been used to assess GFR (Chantler et al. 1969) and to this day continues to be used for the same purpose (Cutajar et al. 2015). Positron emission tomography (PET) can

also be used to assess renal perfusion (Green & Hutchins 2011). The disadvantage with all nuclear medicine techniques is their associated radiation burden, invasiveness and low spatial and temporal resolution.

Computed tomography (Lerman et al. 1999) offers better spatial resolution, although also requires the use of ionizing radiation. Contrast agents are typically required which can be contraindicated, especially when complications in renal function are already present (Beierwaltes et al. 2013).

Doppler ultrasound has been proposed as a method to measure RBF but its accuracy is limited (Wan et al. 2008). The more recently introduced contrast-enhanced ultrasound technique may provide a favourable alternative (Schneider et al. 2011) but there are still concerns regarding its safety (Herget-Rosenthal 2011).

MRI has been increasingly playing a larger role in the understanding and diagnosis of renal disease (Zhang et al. 2014). As with all MR-based methods, exposure to ionizing radiation is not a concern. Information regarding oxygenation and metabolism can be obtained with Blood oxygenation level-dependent (BOLD)-MRI (Prasad et al. 1996). There is data suggesting limitations of BOLD for CKD assessment (Michaely et al. 2012), although the future utility of renal BOLD-MRI is still under debate (Inoue et al. 2012). For a recent review on this topic, refer to (Niendorf et al. 2015). MRI can measure perfusion using exogenous intravascular tracers (based on gadolinium chelates) or endogenous diffusible tracers (with ASL). Over the years, Dynamic contrast-enhanced (DCE)-MRI has been delivering promising results and providing a wealth of information regarding kidney function (e.g. (Lee et al. 2007; Sourbron et al. 2008; Tofts et al. 2012), among others). Nevertheless, there are still strict impediments to the use of contrast agents in certain patient groups (Kaewlai & Abujudeh 2012), which have proved useful in diminishing the incidence of nephrogenic systemic fibrosis due to exposure of gadolinium-based contrast agents (Wang et al. 2011). A disadvantage with this technique is the fact that multiple measurements within the same scanning session are not possible due to the relatively long clearance times of the contrast agent. A review of the DCE-MRI methods used to assess renal function is beyond the scope of this thesis. For a review of this topic, refer to (Jones et al. 2011).

ASL is an MRI technique that allows a quantitative assessment of perfusion without the need for exogenous contrast agents. In fact, the subject's own blood is used for this purpose. At the time of writing, no other method that can be used to assess perfusion in the kidneys does so in a completely non-invasive way, while being able to deliver regional perfusion measurements. This makes ASL extremely interesting both from the academic and clinical point of view. This work

focuses on the optimisation of ASL measurements for the kidneys. As such, both MRI and ASL will be introduced in its own chapter together with a comprehensive overview of previous efforts to assess renal perfusion with this technique.

2   Principles of MRI

MRI is a powerful imaging modality, unique in its ability to non-invasively provide detailed anatomical and physiological information *in vivo*. It emerged in the 1970s, the fruit of key developments which allowed the application of Nuclear Magnetic Resonance (NMR) principles within the context of medical radiology. As such, its basis rests upon the interaction of nuclear spins and externally applied magnetic fields.

This section aims to briefly summarize the fundamentals of MRI. The discussion is limited to a physical description of spin dynamics based on the classical Bloch formalism. This allows a description of the macroscopic behaviour as it pertains to the NMR phenomenon, which is sufficient for the understanding of the work presented in this thesis. Readers interested in a quantum mechanical treatment of the subject are referred to (Brown et al. 2014). In particular, I will address the origin of the NMR signal, MRI physics and the process of image formation. MRI sequences most relevant to this work will also be described. Furthermore, a significant focus will be given to ASL, the functional MRI technique developed and applied within the context of this thesis.

2.1   Classical Description of NMR

2.1.1   Spin and precession

All subatomic particles, such as protons, possess an intrinsic nuclear angular momentum, known as *spin*. The same is true about atoms with an odd atomic number and/or an odd number of neutrons. This fundamental property is in some ways analogous to the classical angular momentum. An important difference is the fact that spin is quantized, taking only discrete integer or half-integer values. Hydrogen nuclei (^{1}H) are composed of a single proton. Therefore, they have non-zero spin and thus exhibit the NMR phenomenon. This fact, combined with their extreme abundance in biological tissue as well as relatively high intrinsic NMR sensitivity make them most studied and widely used nuclei in MRI. Henceforth, I will refer to hydrogen nuclei (^{1}H) nuclei simply as spins. Additionally, unless noted, I will also assume we are dealing with proton (^{1}H) imaging throughout this thesis.

In accordance to the laws of electromagnetic induction, a particle with net spin and an unbalanced charge (e.g. a proton) produces a magnetic field around itself. Thus, protons also possess a magnetic moment. If an external magnetic field (B_0) is absent, the spins assume random orientations, causing their net magnetic moment (resulting from the individual contributions of each spin) to be zero. An external magnetic field acting upon a sample causes

spins to align with it. A typical convention is to refer to the direction along which B_0 is applied as the "z" or longitudinal direction. A number of spin states are possible for a nuclei of spin I, given by $2I + 1$. Therefore, two ($2.\frac{1}{2} + 1$) states are possible for ^{1}H. This energy splitting becomes observable when the sample is under the influence of B_0. As a result, two spin populations emerge: spins aligned with B_0 in a parallel orientation (low-energy state) and spins aligned with B_0 in an anti-parallel orientation (high-energy state). The difference in energy between the two states depends on the strength of the applied magnetic field B_0 (see Figure 2-1). Thermal energy causes a slight preponderance for spins in the lower energy state. The ratio of the numbers of spins in each of the two populations is described by the Boltzmann distribution:

$$\frac{n^-}{n^+} = e^{-\Delta E/kT}$$

Equation 2-1

Where n^- and n^+ are the populations of spins in the anti-parallel and parallel orientations, ΔE is the difference in energy between the states, k is Boltzmann's constant and T is the absolute temperature.

Solving Equation 2-1 for a temperature of 293.15K (i.e. $20°C$) and a field strength of 1.5T yields a population ratio n^-/n^+ equal to 0.999998. Despite this ratio being close to one, typical samples contain an immense number of protons (on the order of the Avogadro number). This translates to an excess of spins in the n^+ compared to those in the n^- population. The contributions of all the individual magnetic moments add up and the result is enough to generate a net macroscopic magnetisation, often denoted by the vector $\boldsymbol{M}$. After further manipulation (as discussed below), this magnetisation becomes measurable, forming the basis for the NMR signal.

As we can see from Equation 2-1, the difference in the number of spins in each energy state is proportional to ΔE, which in turn increases with B_0. Thus, as B_0 increases, so does the available net magnetisation $\boldsymbol{M}$. This is one of the main reasons driving the increase in the magnetic field strengths achievable by the superconducting magnets of MRI scanners over the last decades. A larger amount of available magnetisation is desirable from a signal-to-noise ratio (SNR) perspective, provided technical challenges associated with high field imaging can be overcome (e.g. magnetic field inhomogeneities).
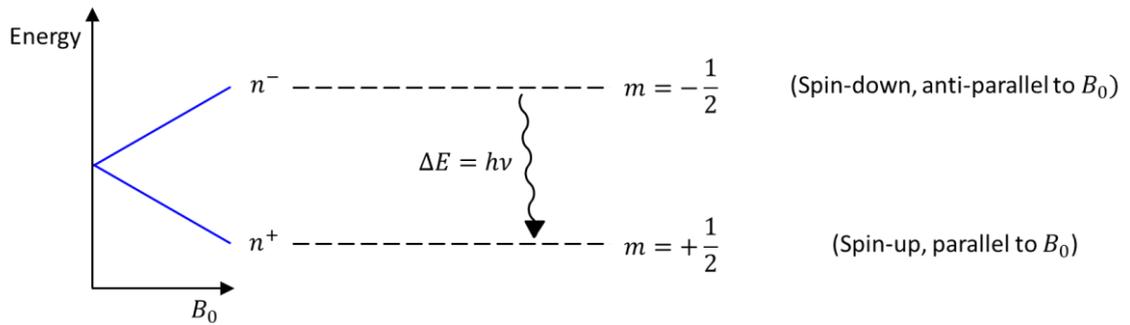
Figure 2-1 Zeeman splitting in a spin ½ system. In a sample of 1H nuclei, two spin populations are observed when an external magnetic field B_0 is applied. The energy gap ΔE between the two corresponding spin states is proportional to the strength of the applied magnetic field B_0. There is a slight preponderance for spins to exist in the lower energy state, i.e. aligned to B_0. The emission of a photon of energy ΔE occurs with a transition from the high to the lower energy state.

Another key phenomenon in NMR/MRI is *precession*. The interaction of spins with B_0 produces a torque which acts perpendicular to the magnetic field and in the direction of the angular momentum. This gives rise to a circular motion of the individual magnetic moments about the field direction, analogous to the precessional movement of a gyroscope when its angular momentum interacts with the gravitational field. The precession frequency is given by the Larmor equation (Equation 2-2):

$$\omega_0 = \gamma B_0$$
<div align="right">Equation 2-2</div>

Where γ is a constant known as the gyromagnetic ratio (γ = 2.675x10^8 rad/s/T, for 1H). Frequently, the gyromagnetic ratio is also referred to as $\bar{\gamma}$ where $\bar{\gamma} = \gamma/2\pi$ = 42.58 MHz/T. It is worth noticing that the gyromagnetic ratio varies among nuclei. For 23Na, $\bar{\gamma}$ = 11.26 MHz/T. This fact can be exploited for performing "nucleus-specific" imaging, because if these two nuclei species are subjected to the same external magnetic field, their different Larmor frequencies allow any one of them to be excited while leaving the remaining nuclei undisturbed. Most of today's clinical MRI systems operate at B_0 fields of strength 1.5 and 3 Tesla. From Equation 2-2 it follows that the corresponding Larmor frequencies for 1H are in the radio wave range. Electromagnetic radiation at this range of the electromagnetic spectrum is able to penetrate the patient with relatively low power deposition. This allows the phenomenon of NMR to be exploited in MRI as a technique with ability to probe inside of the human body non-invasively.

2.1.2 Excitation

Even though spin state population differences result in a net magnetisation, this is not directly measurable. This is because at rest, individual magnetic moments precess at random phases and consequently, no measurable transverse magnetisation exists. This is depicted in Figure 2-2.
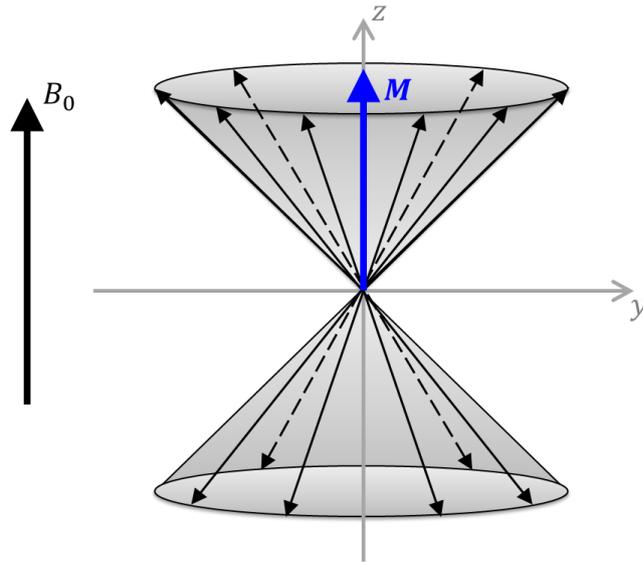
Figure 2-2 Snapshot of a spin ensemble in dynamic equilibrium, under the influence of an external static magnetic field B_0. The spins align with B_0, with two populations arising. A slight preponderance exists (exaggerated in this diagram) for the lowest energy state (spin-up, parallel to B_0). Individual magnetic moments precess at random phases. The lack of phase coherence results in no measurable transverse magnetisation component.

In order to generate a measurable NMR signal, the equilibrium magnetisation must be tipped away from the longitudinal (z) axis into the transverse (xy) plane. This can be achieved through the application of an oscillating magnetic field B_1 at the Larmor frequency, perpendicular to B_0 (i.e. in the xy plane). This is also referred to as a radiofrequency (RF) pulse, owing to the short duration during which it is applied (on the order of a few milliseconds) and the fact that its frequency is tuned to the radio wave range. The matching of the frequency of B_1 and the Larmor frequency of the spin system allows *resonance* to occur. As a result, B_1 causes the magnetisation vector *M* to experience a torque, being "excited" out of equilibrium. Individual spins will now precess synchronously. Consequently, *M* also initiates precession, generating a time-varying transverse magnetisation which is readily detectable as it generates an electromotive force in a nearby receive coil, according to Faraday's law of induction. Different amounts of energy can be supplied to the spin system by adjusting the strength and duration of the B_1 field, resulting in different degrees of "conversion" of longitudinal magnetisation to transverse magnetisation. By carefully tuning these parameters, the equilibrium magnetisation can be rotated by 90°. This brings all of the longitudinal magnetisation into the transverse component, which maximises the amount of measurable signal. The corresponding RF pulse is aptly name 90°- or $\pi/2$-pulse. In general, the angle of rotation is known as the flip angle. This gives rise to the simplest MR signal that can be recorded, and is known as free induction decay (FID), a rather descriptive name as it encapsulates the fact that the measured signal results from a *free* precession of the net

magnetisation vector, which is measured via *induction* of current in a nearby coil. Furthermore it hints that the measured signal will eventually *decay*, through mechanisms that will be described shortly.

A depiction of the excitation process can be seen in Figure 2-3 a). Before the application of the RF pulse, all available coherent magnetisation is aligned with the external magnetic field B_0, along the longitudinal axis. This is frequently referred to as "equilibrium magnetisation" M_0. As soon as B_1 is applied, M starts being "pushed" into the transverse plane (xy) while simultaneously precessing about the z-axis at the Larmor frequency. In the laboratory frame of reference (i.e. defined with respect to the magnet in the MR scanner), this results in a spiralling trajectory (Figure 2-3 a)). A "rotating frame of reference" is frequently considered during the analysis of the behaviour of M, as it simplifies the description of many MR phenomena. This can be thought of as a coordinate system where the xy plane itself rotates at an angular frequency matching the Larmor frequency of the spin system. The excitation process as seen from the rotating frame of reference is depicted in Figure 2-3 b). Note the usage of x', y' and z' in the naming of the orthogonal axes in this coordinate system for distinguishing it from the laboratory frame of reference. In this case, the precession motion of M about the z-axis disappears.

a)   b)
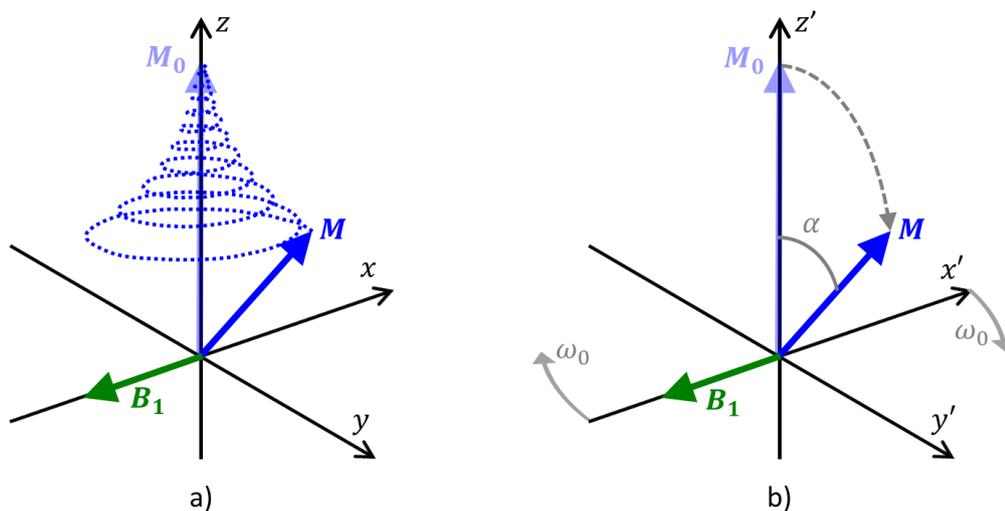
Figure 2-3 Excitation due to application of a RF pulse as seen in the laboratory and rotating frames of reference. a) In the laboratory frame of reference, the tipping of the net magnetisation vector **M** towards the transverse plane is accompanied by precession about the longitudinal axis. b) In the rotating frame of reference, itself rotating at the Larmor frequency ω_0, precession is no longer visible, and the behaviour of **M** is simplified to a monotonic tipping towards the $x'y'$ plane. The angle α between **M** and the longitudinal axis is typically referred to as flip angle.

Even though the behaviour of individual magnetic moments is governed by the laws of quantum mechanics, a classical description of the behaviour of the bulk magnetization vector $\boldsymbol{M}$ is given by Equation 2-3:

$$\frac{d\boldsymbol{M}}{dt} = \gamma \boldsymbol{M} \times \boldsymbol{B} \qquad\qquad \text{Equation 2-3}$$

Where $\boldsymbol{M} = [M_x, M_y, M_z]$, γ is the gyromagnetic ratio and $\boldsymbol{B}$ is the applied magnetic field ($\boldsymbol{B} \equiv B_0$). This equation describes precession of $\boldsymbol{M}$ about the direction of $\boldsymbol{B}$.

2.1.3 Relaxation

As was alluded to before, as soon as the magnetisation vector $\boldsymbol{M}$ is tipped to the transverse (xy) plane, it immediately starts experiencing relaxation. As opposed to the individual magnetic moments, this causes the magnitude of $\boldsymbol{M}$ to vary with time. These relaxation effects are in fact one of the major sources of image contrast in MRI, as if these did not exist we would have to rely on proton densities only which vary over a much smaller range when compared to the relaxation times. Two independent effects are at play. Any perturbation to the equilibrium magnetisation M_0 will be followed by a recovery process whereby the magnetisation will return to its thermal equilibrium state, along the z-axis:

$$\frac{dM_z}{dt} = \frac{1}{T_1}(M_0 - M_z) \qquad\qquad \text{Equation 2-4}$$

Where T_1 is the *spin-lattice* (or longitudinal) relaxation time, which characterises the rate at which excitation energy is lost from the spin system as heat to the surrounding environment (in early NMR experiments this was most frequently the *lattice* of crystalline solids). In practice, it is the time required for the longitudinal component of the magnetisation (M_z) to return to 63% of M_0 following the application of a 90°-pulse. Equation 2-4 shows that the rate at which M_z changes at any point in time during recovery is proportional to the difference $M_0 - M_z$ at that instant (see Figure 2-4).

Unlike longitudinal relaxation, there is no energy loss involved in the second relaxation mechanism. Rather, it hinges on the loss of phase coherence of the individual magnetic moments, which manifests macroscopically as a decay of the transverse component of the magnetisation (M_{xy}), as depicted in Figure 2-5. This is given by Equation 2-5:
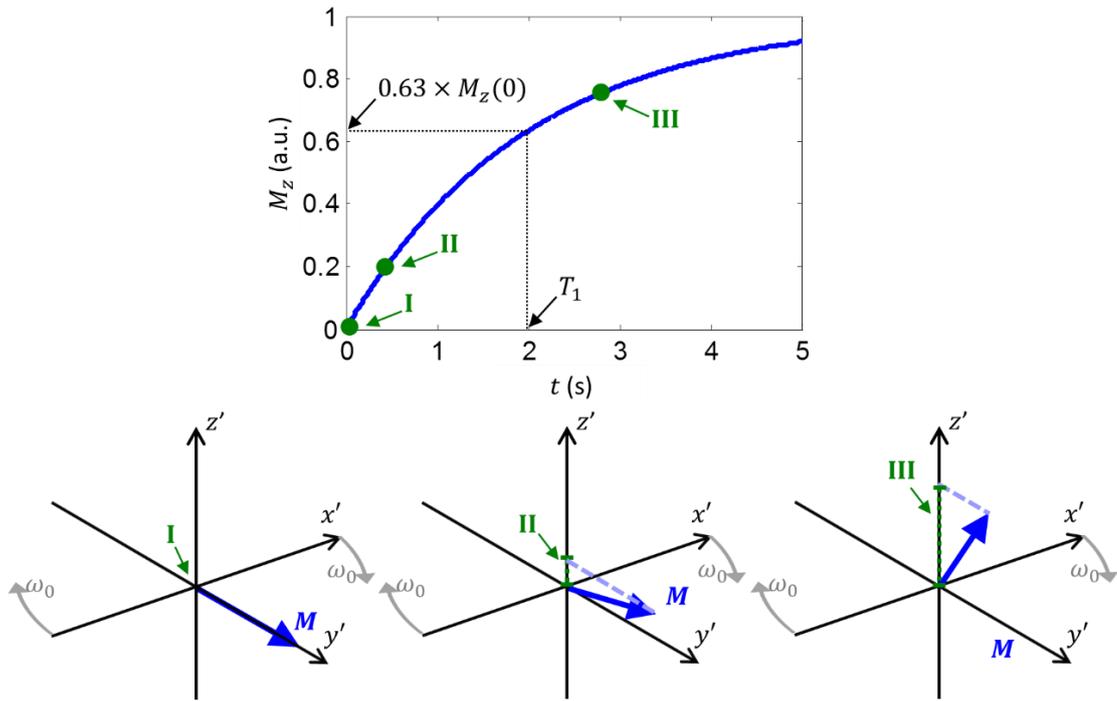
Figure 2-4 T_1 relaxation. Simulation of the effect of T_1 relaxation on the longitudinal magnetisation (top) and corresponding depiction of longitudinal magnetization recovery in the rotating frame of reference (bottom).
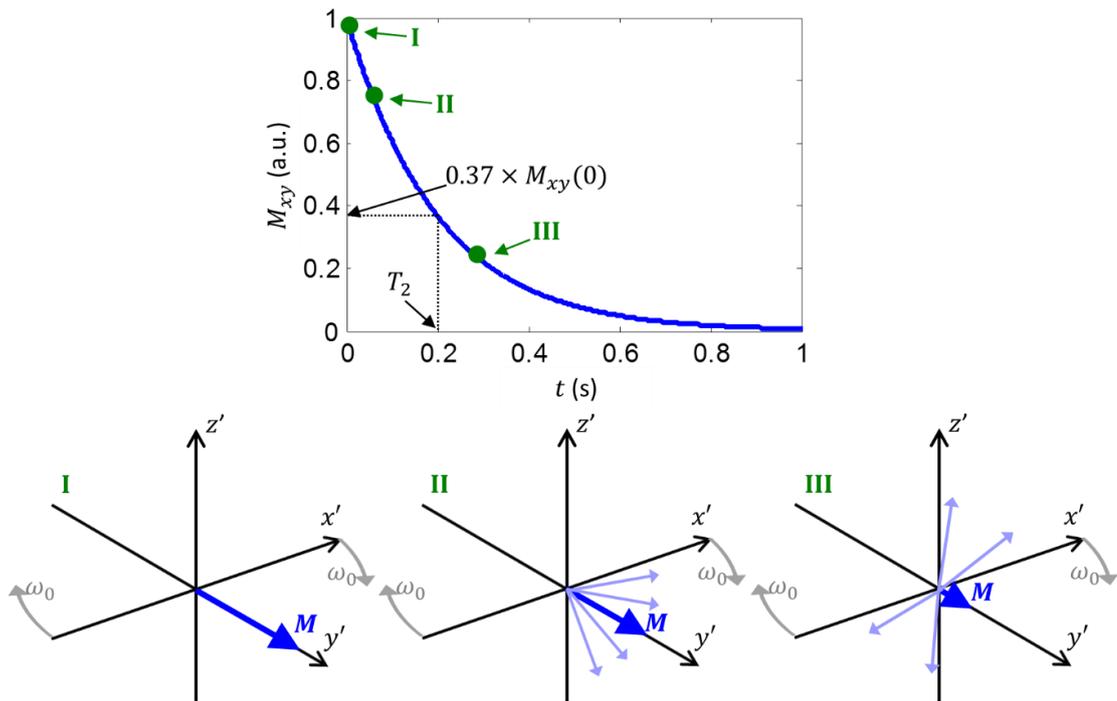
Figure 2-5 T_2 relaxation. Simulation of the effect of T_2 *relaxation on the transverse magnetisation* (top) and corresponding depiction of spin dephasing in the rotating frame of reference (bottom).

$$\frac{dM_{xy}}{dt} = -\frac{M_{xy}}{T_2}$$

Equation 2-5

Where T_2 is the *spin-spin* (or transverse) relaxation time. The physical basis of this effect is directly linked to local magnetic field variations, which cause the precession frequencies of individual magnetic moments to change. The effects of this are easily understood when considered in the rotating frame of reference as depicted in Figure 2-5. In short, individual magnetic moments will deviate from the Larmor frequency, causing them to start fanning out. This phase dispersion causes the measurable signal to be attenuated compared to what it would have been if all moments kept precessing at the Larmor frequency, where the net sum of their vectors would yield the maximum possible signal.

It is worth noticing that all relaxation processes occur at much slower time scales than precession by several orders of magnitude. Whereas relaxation occurs within the realm of fractions to a couple of seconds, precession occurs approximately 10^7 to 10^9 faster. Also, note that T_1 determines the upper limit for T_2 (i.e. $T_2 \leq T_1$), given that in addition to the *spin-spin* effects described above, the transverse magnetization is inherently reduced as T_1 effects take place (i.e. return of $\boldsymbol{M}$ towards the z-axis). The latter increases with the strength of the static field B_0.

Combining Equation 2-3, Equation 2-4 and Equation 2-5 we obtain:

$$\frac{d\boldsymbol{M}}{dt} = \gamma \boldsymbol{M} \times \boldsymbol{B} + \frac{1}{T_1}(M_0 - M_z)\hat{\boldsymbol{z}} - \frac{1}{T_2}(M_x\hat{\boldsymbol{x}} + M_y\hat{\boldsymbol{y}})$$

Equation 2-6

This is known as the Bloch equation, introduced by Felix Bloch in 1946 (Bloch 1946), where $\hat{\boldsymbol{x}}$, $\hat{\boldsymbol{y}}$ and $\hat{\boldsymbol{z}}$ are unit vectors in the laboratory frame of reference. This provides a phenomenological description of the precession and relaxation behaviours of the net magnetisation vector, and accurately describes the signal measured by the receiver coils.

When B_1 is no longer being applied, the only remaining field is the constant field B_0. In this case, $\boldsymbol{B} = B_0\hat{\boldsymbol{z}}$, and the three component equations can be solved to yield:

$$M_x(t) = e^{-t/T_2}(M_x(0)\cos(\omega_0 t) + M_y(0)\sin(\omega_0 t))$$

Equation 2-7

$$M_y(t) = e^{-t/T_2}(M_y(0)\cos(\omega_0 t) + M_x(0)\sin(\omega_0 t))$$

Equation 2-8

$$M_z(t) = M_z(0)e^{-t/T_1} + M_0(1 - e^{-t/T_1})$$

Equation 2-9

The local magnetic field variations related to the T_2 effect are intrinsic to the spin system, as they relate to the interactions between the dipole moments of neighbouring spins. In practice, extrinsic factors can also contribute to an enhanced transverse relaxation, such as hardware-dependent B_0 inhomogeneities and those caused by the object under study itself (e.g. magnetic susceptibility differences due to air-tissue interfaces). The extrinsic field inhomogeneities are frequently captured by a separate decay term T_2'. This combines with T_2 to give the *apparent transverse relaxation time* T_2^*:

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}$$

Equation 2-10

Unlike pure T_2 signal losses, those caused by T_2' can be recovered. This leads us into another key concept in NMR, that of spin-echo, described in the next section.

2.1.4 Spin Echo

In the spin echo technique, the first step is to apply a 90° pulse tuned to the Larmor frequency to generate transverse magnetisation. Natural spin dephasing occurs as spins in different positions experience slightly different magnetic fields due to local inhomogeneities, causing their precession frequencies to vary. The corresponding signal decays (free induction decay) under a T_2^* envelope as shown in Figure 2-5. After waiting for a given amount of time ($TE/2$), an additional 180° pulse (also referred as *refocussing pulse*, depending on the context) is played out. This does not have an effect on the rate of phase accrual by different spins, assuming they remain in the same position and therefore continue to experience the same field inhomogeneities. However, their phase angles are reversed. They continue to accrue phase as before, but because the phases had been reversed, this time they start rephasing.

After another $TE/2$ time interval following the 180° pulse their phases converge, yielding a total signal which is no longer limited by T_2^* but rather T_2 (see Figure 2-5). Note that if we keep playing out refocussing pulses at multiples of the echo time TE following the first $TE/2$ interval we are able to acquire data corresponding to multiple echoes, even though their signal continues to decrease as governed by the T_2 of the sample. Several imaging sequences make use of this approach, such as rapid acquisition with relaxation enhancement (RARE) (Hennig et al. 1986) and 3D Gradient and Spin Echo (3D-GRASE) (Feinberg & Oshio 1991; Oshio & Feinberg 1991). The latter was the basis for all imaging experiments in this thesis and therefore will be described in its own section (see Section 2.2.4.5).
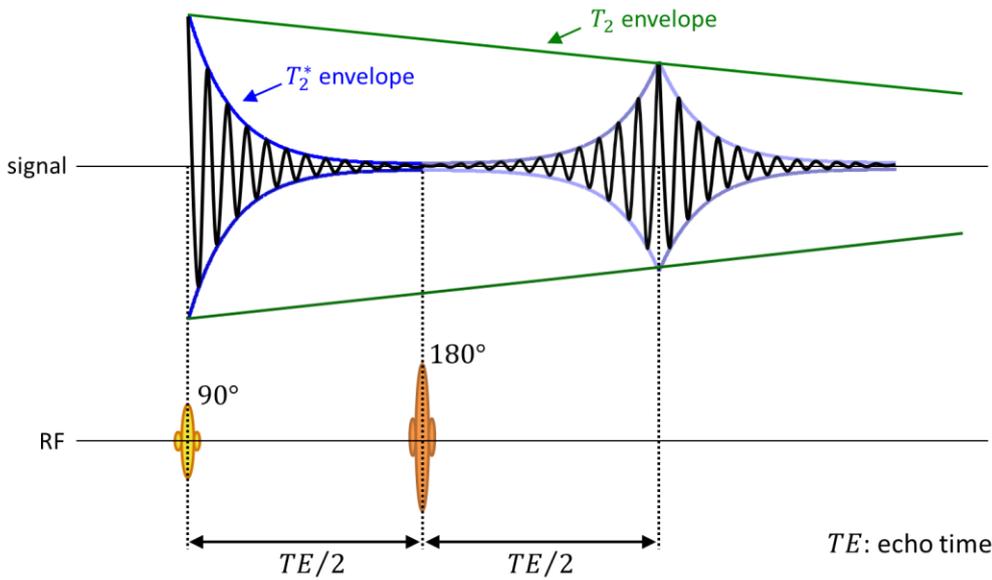
Figure 2-6 Spin echo pulse sequence and corresponding signal evolution. The horizontal axis represents time.

2.1.5 Inversion and saturation recovery

By carefully tuning the parameters of an NMR acquisition, one can set up experiments that are sensitive to different tissue characteristics, such as magnetic relaxation properties. Several approaches have been developed for the *in vivo* measurement of the longitudinal relaxation times of tissue. A simple, yet powerful way of measuring the longitudinal relaxation time (T_1) of tissue is the inversion recovery (IR) method (Vold et al. 1968). Closely related is the saturation recovery technique (SR) (McDonald & Leigh Jr. 1973). The latter was used within this thesis for mapping the T_1 relaxation times of the kidneys (see Chapter 6).

If, for simplicity, we consider an NMR experiment, the IR/SR methods are conceptually very similar to a simple FID experiment, differing only in two main ways. First, prior to the excitation pulse, an additional pulse is played out. This causes the longitudinal magnetisation to differ from M_0 at the time of the excitation pulse. Second, this is repeated multiple times, each with a different time delay from the "magnetisation preparation" pulse and the excitation pulse. During this delay, different components of the tissue will recover exponentially at a rate given by their particular T_1 and thus, a T_1 contrast will be generated among them. This is then exploited by converting the available longitudinal magnetisation to transverse magnetisation, through the excitation pulse and subsequent data acquisition. Repeating this at multiple delay times allows us to sample the T_1 recovery curves of the different tissues. The obtained signal intensities are then fit to an equation describing longitudinal relaxation (see Figure 2-7).
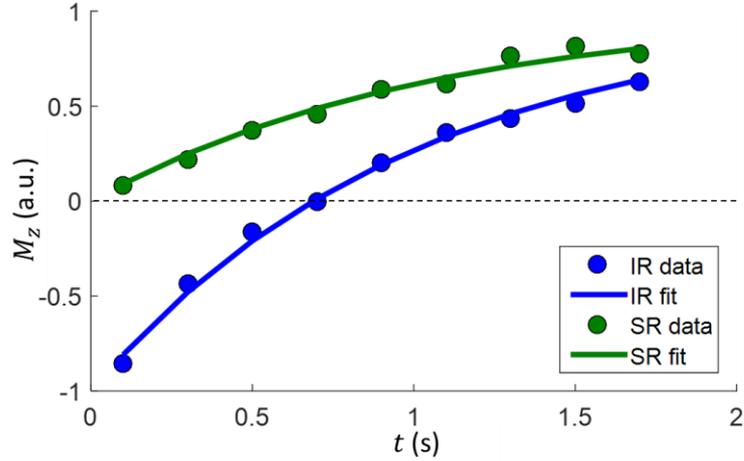
Figure 2-7 Recovery of the longitudinal magnetisation in IR and SR for a tissue species with T_1 = 1s and assuming M_0 = 1. The curves correspond to the theoretical IR and SR fits to simulated data points, obtained from Equation 2-12 and Equation 2-13, respectively for IR and SR data.

As the name hints, IR experiments start with a 180° inversion pulse, which inverts the longitudinal magnetisation from the z to the $-z$ axis (i.e. M_0 becomes $-M_0$). Therefore, it is necessary to allow for a full recovery of the longitudinal magnetisation between repetitions, and so the time needed for a full inversion recovery experiment for T_1 mapping can be quite long. SR methods address this by setting the longitudinal magnetisation to zero in the initial step of the experiment, replacing the 180° pulse by a 90° pulse. Therefore, the range of M_z is reduced from $[-M_0, M_0]$ in IR to $[0, M_0]$ in SR (see Figure 2-7). For SR, the state of the longitudinal magnetisation immediately prior to the saturation pulse is not important, since a 90° pulse will always set M_z to zero, irrespective of its current value. As such, SR is a more time-efficient method for T_1 mapping.

The Bloch equation for the longitudinal magnetisation (Equation 2-4) can be solved to yield:

$$M_z(t) = M_0[1 - (1 - \cos(\theta))e^{-t/T_1}] \qquad\qquad \text{Equation 2-11}$$

If θ = 180° (IR), we have:

$$M_z(t) = M_0(1 - 2e^{-t/T_1}) \qquad\qquad \text{Equation 2-12}$$

If θ = 90° (SR), we have:

$$M_z(t) = M_0(1 - e^{-t/T_1}) \qquad\qquad \text{Equation 2-13}$$

35

To obtain T_1 maps in a MRI experiment, we combine the methods described to measure T_1 with an imaging pulse sequence for data acquisition (see Section 2.2.4 for a description of several pulse sequences).

2.2 MRI Concepts

So far, we have considered the signal coming from a large sample, where all spins are precessing at the Larmor frequency. To obtain images, we need to limit the resonance condition (i.e. when the frequency content of the RF pulse matches the precessional frequency of a spin isochromat) to a well-defined region in space. For standard (2D) MRI methods, spatial localisation can be thought of as a stepwise approach, where initially a single slice with a given thickness is "selected" and then further in-plane localisation is performed to restrict the MRI signal generation to small volume elements (voxels) which are the basic components of an image. Magnetic field gradients are at the heart of MRI. Their usage for encoding spatial information into NMR signals was initially proposed in 1973 by Paul Lauterbur (Lauterbur 1973). Comprehensive descriptions on the topics of signal localisation and pulse sequences in MRI can be found in the outstanding books by Bernstein *et. al* (Bernstein et al. 2004) and Brown, Cheng and Haacke *et. al* (Brown et al. 2014)*.*

2.2.1 Slice selection

The process of slice selection is depicted in Figure 2-8. The application of linear gradient fields causes the spin precession frequency to linearly vary in space. If we play an excitation pulse simultaneously while the gradient is being applied, only the narrow band of spins whose precession frequencies match the frequency content of the excitation pulse B_1 will meet the resonance condition. Thus, only the spins in this region will be excited, delineating a well-defined 2D plane within the patient. The physical shape of the excited region depends on the shape of the spectrum of the RF pulse. In practice, to obtain approximately rectangular slice profiles regions it is typical to apply apodized sinc RF pulses (see Figure 2-8). Changing the carrier frequency of the RF pulse allows one to change the position of a slice, whereas the slice thickness can be specified by adjusting the bandwidth of the RF pulse or the amplitude of the slice-selection gradient. It is from these regions that all signals which will be manipulated during an MRI pulse sequence will arise.
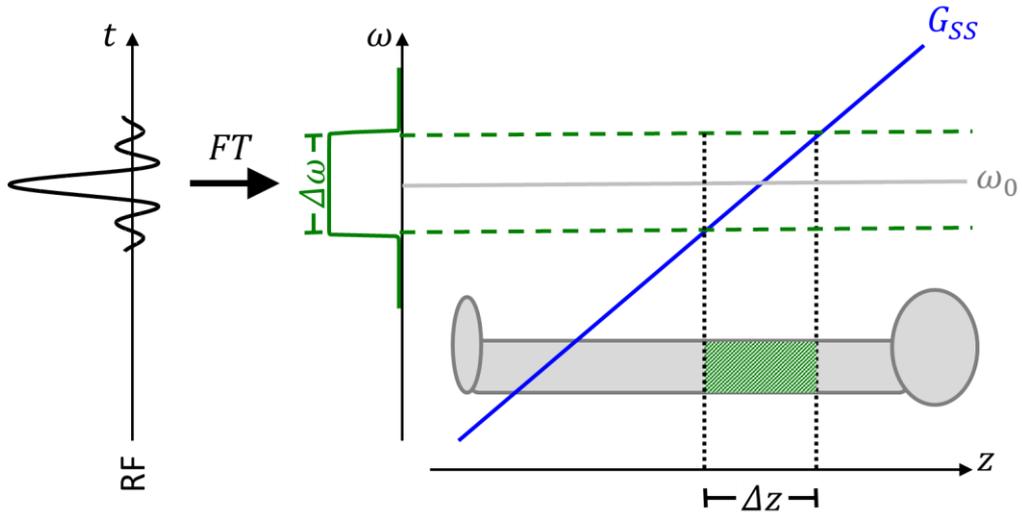
Figure 2-8 Slice selection using a sinc RF pulse. While a slice selection gradient G_{SS} is being applied, the excitation of spins is restricted to a location with thickness Δz. The range of frequencies Δw is centered on the Larmor frequency ω_0, typically at the isocentre of the magnet. Note that the sinc and the rect functions are a Fourier pair. This RF pulse in practice would correspond to a function truncated in time.

2.2.2 Frequency encoding and K-space

Having completed slice selection, the task of spatial localisation is effectively reduced to two dimensions. The next step is to do in-plane localisation within the selected region.

After switching off the slice selection gradient (G_{SS} in Figure 2-8) all spins are back to precessing at the Larmor frequency ω_0. To obtain information along one of the in-plane dimensions, another magnetic field gradient is superimposed on the main static field B_0, this time orthogonal to the slice selection gradient.

There are independently controlled gradient fields along each of the scanner's axes G_x, G_y and G_z, which unlike B_0, can be switched on and off. Mathematically, they can be defined as:

$$G_x(t) = \frac{\partial B_z}{\partial x}$$

Equation 2-14

$$G_y(t) = \frac{\partial B_z}{\partial y}$$

Equation 2-15

$$G_z(t) = \frac{\partial B_z}{\partial z}$$

Equation 2-16

We can express the total gradient field as a vector, which is the 3D gradient of the z component of the magnetic field, that is: $\boldsymbol{G}(t) = \nabla G_Z$.

The z component of the magnetic field due to the linear gradients at a point in space $\boldsymbol{r}$ is:

37

$$B_z(\boldsymbol{r}, t) = \boldsymbol{G}(t) \cdot \boldsymbol{r}$$

Equation 2-17

Therefore the Larmor frequency depends on the superposition of static and gradient magnetic fields:

$$\omega(\boldsymbol{r}, t) = \gamma(B_0 + \boldsymbol{G}(t) \cdot \boldsymbol{r}) = \omega_0 + \omega_G(\boldsymbol{r}, t)$$

Equation 2-18

In the rotating frame, the term ω_0 disappears. Therefore, during the application of a gradient of duration t, spins isochromats at position $\boldsymbol{r}$ have accrued phase given by:

$$\phi_G\,(\boldsymbol{r}, t) = -\int_0^t \omega_G(\boldsymbol{r}, t')\,dt' = -\int_0^t \gamma \boldsymbol{G}(t') \cdot \boldsymbol{r}\,dt'$$

$$= -\gamma \boldsymbol{r} \int_0^t \boldsymbol{G}(t')\,dt' = -\boldsymbol{k}(t)\boldsymbol{r}$$

Equation 2-19

Where the vector $\boldsymbol{k}$ corresponds to a point defining a specific spatial frequency in the so-called k-space:

$$\boldsymbol{k}(t) = \gamma \int_0^t \boldsymbol{G}(t')\,dt'$$

Equation 2-20

It can be shown (see (Brown et al. 2014) for an extended discussion) that the demodulated measured signal $s(t)$ can be described by Equation 2-21:

$$s(t) = \int \rho(x)\,e^{i\phi_G\,(x,t)}dx$$

Equation 2-21

Where $\rho(x)$ is the spin density of the object at location x.

In spatial frequency terms, for the x dimension:

$$s(k_x) = \int \rho(x)\,e^{-i2\pi k_x x}dx$$

Equation 2-22

This shows that for the case of linear magnetic field gradients the signal corresponds to the Fourier transform (FT) of the spin density. Therefore, the inverse FT allows one to recover the underlying spin density:

$$\rho(x) = \int s(k_x)\,e^{i2\pi k_x x}dk_x$$

Equation 2-23

As we continuously sample data as the gradient is being applied (therefore referred to as readout gradient, G_{RO}), we are creating a mapping between frequency and location - *frequency encoding* (FE). The recorded data will consist of the superimposition of the signals at different frequencies. A 1-D FT on the acquired signal then allows to obtain its amplitude (related to the proton density) for each of the signal's component frequencies, which correspond to each location in the first in-plane dimension. The time between the slice-selective excitation RF pulse and the maximum signal being recorded during the application of G_{RO} is referred to as echo time (TE).

The goal of MRI is to sample the whole of *k*-space corresponding to the object (or part of it) for which magnetisation is excited. The combination of gradients and RF pulses define the *k*-space trajectory. The longer a gradient with a given polarity is on, the further away from the origin of *k*-space ($k_x = k_y = k_z = 0$) a data point will be. The location of a data point in *k*-space can also be affected by RF pulses. A typical example is the change of phase caused by a 180° pulse (as seen in Section 2.2.4.1) which effectively reflects the coordinates of the *k*-space point across both k_x and k_y axis. There are countless possible *k*-space trajectories, particularly if non-Cartesian imaging is considered. In practice, as discrete sampling of a continuous-time signal is involved, the Nyquist sampling criteria must be fulfilled to avoid aliasing artefacts. This states that the sampling rate must be twice the maximum frequency component of the signal being measured.

The concept of *k*-space and *k*-space trajectories proves very helpful in the understanding of many MRI concepts. This the case for pulse sequences, image artefacts and other advanced techniques, particularly those related to partial sampling, such as partial Fourier (Feinberg et al. 1986) or parallel imaging (Pruessmann et al. 1999; Griswold et al. 2002).

After acquisition, the process known as *image* reconstruction allows one to convert the signal acquired in the frequency domain to the spin density in the image (real) space, where the FT plays a critical role as seen above.

2.2.3   Phase Encoding

The frequency encoding step, on its own, is not enough to match the acquired signal to the location of its source. Localisation along the second in plane dimension is achieved by the process of *phase-encoding* (PE). The PE concept is based in the fact the application of a gradient causes dephasing of the components of the transverse magnetisation. Therefore, after a gradient is applied for a given amount of time, phase is accrued or lost relative to some reference. Moreover, if we disregard transverse relaxation, this effect persists indefinitely after

the gradients are switched off, hence the term phase-encoding. This is contrarily to the spin precession frequencies, which revert to their original values. As we have seen, the amount (angle) of dephasing by a spin isochromat in the presence of a gradient is related to the product of the gradient amplitude and its duration. In practice, phase-encoding relies on the application of a separate gradient field G_{PE}, applied along the second in plane direction, before G_{RO} and signal acquisition. For complete localisation of the signal, different measurements are performed, where the only variable during the acquisition is the strength of the PE gradient magnetic field.

In the end, frequency and phase, are direct correlates of locations in the image. The mapping between the spatial frequency domain (k-space) and the image domain is achieved by 2D FT. If during slice selection a thick volume of spins is excited (also referred to as slab), spatial localisation along the third dimension (following the above discussion) can be achieved by extending the PE principles along this dimension (3D-FT imaging).

2.2.4   Pulse Sequences

In an actual imaging experiment, all the concepts addressed thus far are combined into MR pulse sequences. Here I will outline the two elemental MR pulse sequences – the spin echo and gradient echo sequences. The concept of echo-train pulse sequences is briefly introduced as relating to the Echo-planar Imaging (EPI) and Rapid Acquisition with Relaxation Enhancement (RARE) techniques. Furthermore, a more advanced echo-train pulse sequence, which derives from both the SE and GE, will be described. It is aptly known as 3D Gradient and Spin Echo (3D-GRASE) and was the foundation for the all imaging experiments in this thesis.

2.2.4.1   Spin Echo

The physical basis of the NMR spin echo effect was described in Section 2.1.4. For imaging purposes (following the principles of slice-selection outlined in Section 2.2.1), the RF excitation pulse is played during the application of a slice-selection gradient (G_{SS}). Along the same axis, an additional gradient of opposite polarity, known as *rephasing* gradient, is used to compensate for the phase dispersion caused by the main slice-selection gradient. If not accounted, this would result in signal loss throughout the slice. The first step in the process of in-plane signal localisation is the application of a *pre-phasing* gradient which prepares the excited magnetisation for an echo to occur later. In terms of k-space trajectory the phase accrual caused by the pre-phasing gradient moves the magnetisation into high spatial frequency areas at the edge of k-space (see Figure 2-9 right location 2). Next, a slice-selective 180° refocussing pulse is played to form a spin-echo. This occurs when the area under the readout gradient (G_{RO}) matches

the area of the prephasing gradient lobe. This framework is repeated a number of times (*repetitions*) equal to number of required lines along the phase-encoding direction. The only parameter that varies within each repetition is the gradient moment of the phase-encoding gradient. The time delay between excitation pulses is known as *repetition time* (TR).
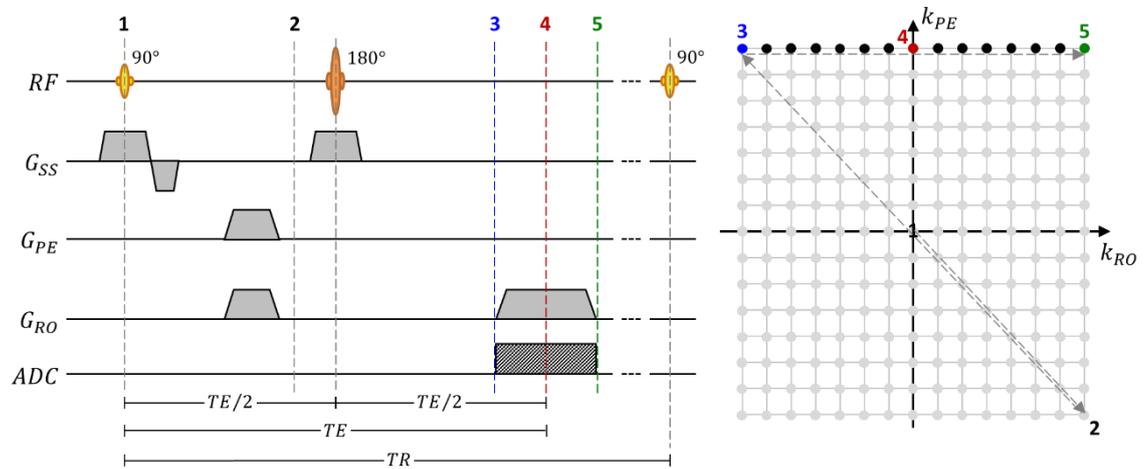
Figure 2-9 Spin echo pulse sequence diagram and k-space trajectory corresponding to the acquisition of the outermost positive phase-encoding k-space line. TE = Echo time. ADC = Analog-to-Digital Converter

2.2.4.2   Gradient Echo

The main difference between the gradient echo sequence (Figure 2-10) and the spin echo sequence is that the echo is achieved by inverting the polarity of the readout gradient. The inversion of polarity is necessary because there is no refocussing pulse, whose effect is to reflect a k-space point with respect to the k-space centre. Furthermore, this time spin rephasing is limited by T_2^*. The echo time (TE) is defined as the time from the excitation pulse until the echo is formed. Gradient-echo based sequences can be made extremely fast, for example if low flip angles are used, which removes the need to wait for an extended T_1 recovery period. Notable sequences that make extensive use of the gradient-echo principles include the fast low-angle shot (FLASH) (Haase et al. 1986) and echo planar imaging (EPI) (Mansfield 1977) techniques.
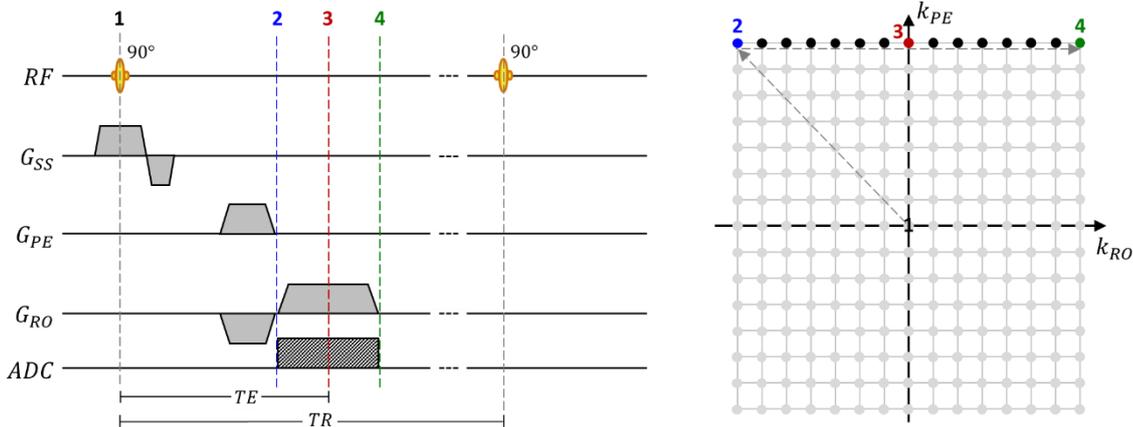
Figure 2-10 Gradient echo pulse sequence diagram and k-space trajectory corresponding to the acquisition of the outermost positive phase-encoding k-space line.

2.2.4.3   Echo-planar Imaging (EPI)

The main disadvantage of the basic spin echo and gradient echo techniques is that they are inefficient in their use of transverse magnetisation: only one k-space line is acquired per RF excitation making these sequences prohibitively long in practice. Echo-train sequences, such as Echo-planar Imaging (EPI), Rapid Acquisition with Relaxation Enhancement (RARE) (section 2.2.4.4) and Gradient and Spin-Echo (GRASE) (section 2.2.4.5) address this issue by acquiring a fraction or the entirety of k-space after a single excitation.

In EPI, originally proposed by Sir Peter Mansfield (Mansfield 1977), a train of gradient echoes is generated following the excitation pulse through the application of a series of bipolar alternating gradients along the frequency encoding direction. Small phase-encoding gradients (typically referred to as *blips*) allow each echo to be differently phase-encoded to cover different regions of k-space. This results in a rectilinear k-space trajectory. The amount of time saved in EPI relative to the standard gradient echo sequence is directly related to the number of echoes sampled following a single excitation, known as the echo-train length (ETL). The maximum achievable ETL depends on the time window during which there is usable transverse magnetisation which in turn is determined by the T_2^* relaxation properties of the tissue. Extremely high temporal resolution and motion-robust data can be obtained from EPI which has allowed this technique to find wide use in functional and real-time imaging. An in depth treatment of the family of EPI acquisition schemes, reconstruction approaches and strategies to deal with common EPI image artefacts is given in (Schmitt et al. 1998).

2.2.4.4   Rapid Acquisition with Relaxation Enhancement

RARE (Hennig et al. 1986), also referred to as Fast Spin-Echo (FSE) or Turbo Spin-Echo, is another key echo-train sequence which in turn uses a series of RF refocussing pulses to generate a train

of spin-echoes during which the transverse magnetisation is sampled to fill multiple k-space lines. RARE is relatively insensitive to off-resonance artefacts given that its primary means for echo generation if RF refocussing. Minimum scan time can be achieved by acquiring the entire k-space after a single excitation (*single-shot* RARE), particularly when coupled with k-space undersampling methods such as partial Fourier (Feinberg et al. 1986). This is known as the Half-Fourier Acquisition Single-shot Turbo spin Echo imaging (HASTE) technique (Semelka et al. 1996). It allows high-quality "snapshot" T_2-weighted scans to be performed, which are able to freeze patient movement and as such have found wide use in abdominal imaging. In the context of RARE, ETL is defined as the number of RF refocussing pulses. Even though in RARE the signal decays according to T_2 (rather than T_2^* as in EPI), achievable ETLs with RARE are typically lower than with EPI. This is due to the longer echo-spacing in RARE resulting from the fact that the rate of production of echoes per unit time is lower when relying on RF refocussing pulses than in the case of an echo train induced by gradient switching. An important limitation of RARE is the increased RF power deposition owing to its large number of RF pulses. This is often addressed by reducing the number of slices acquired (particularly as B_0 increases), which limits patient coverage.

2.2.4.5   3D-GRASE

Gradient and Spin Echo (GRASE), a pulse sequence introduced by Oshio and Feinberg in 1991 (Oshio & Feinberg 1991; Feinberg & Oshio 1991) incorporates the principles of gradient and spin-echo. Each time the magnetisation is excited (i.e for each repetition of the sequence), a train of RF refocusing pulses is combined with an alternating polarity readout gradient to enable the acquisition of multiple k-space lines. Different excitations of the magnetisation are frequently referred to as "shots" or "segments". The time between these is the TR of the sequence. Günther *et. al* (Günther et al. 2005) extended the original 2D implementation for 3D imaging, first applying it for brain perfusion mapping with ASL. This sequence was used for all ASL data acquisitions within this thesis. A schematic of this pulse sequence is shown in Figure 2-11, as well as the corresponding k-space trajectory for the case of a multi-shot acquisition with three segments.

The sequence starts with a slab-selective excitation using a 90° pulse. This is followed by a spin-echo train, where each refocussing pulse is followed by a phase-encoding step along the second phase-encoding direction (also called the partition direction). Flip angle error accumulation over the spin-echo train which causes loss of transverse magnetisation is avoided by employing the Carr-Purcell-Meiboom-Gill (CPMG) phase cycling technique (Meiboom & Gill 1958). In this approach, there is a 90° offset between the phases of the excitation and refocusing pulses (e.g.

applying the excitation along the x axis and the refocusing pulses along the y axis). Within the envelope of each spin echo, an EPI-like echo-train is used to acquire a complete partition. This phase-encoding strategy effectively allows separation of off-resonance and T_2 effects along each of the phase-encoded directions (Mugler 1999).

Furthermore, a centric ordering is used in the partition encoding direction. Here, the acquisition of the k-space centre is performed within the first spin echo. This ensures that the lower spatial frequencies in k-space (which specify the contrast of the image) are sampled with the largest amount of transverse magnetisation, and thus signal, available. Subsequently, partitions with positive values of k (i.e. $+k_{PAR}$) are acquired by even-numbered spin-echoes. Partitions in the $-k_{PAR}$ region of k-space are acquired by odd-numbered spin echoes (see Figure 2-11).

The number of spin echoes following the initial excitation is called ETL, and corresponds to the number of measured partitions. This may or may not correspond to the final number of reconstructed slices. A situation where the number of reconstructed slices is larger than the number of measured partitions is when k-space undersampling techniques are used. Partial Fourier (D. A. Feinberg, Hale, Watts, Kaufman, & Mark, 1986), is one of such techniques. It harnesses the complex conjugate symmetry property of the Fourier Transform (Bracewell 2000), and can be used to reduce the overall echo train duration (ETD) in a 3D-GRASE readout. Using this technique, the number of partitions that need to be explicitly measured can be reduced, while still being able to reconstruct a full image dataset.

As previously mentioned, T_2 effects are most important along the partition encoding direction. The longer the total ETD, the more transverse magnetisation will decay throughout the acquisition of a given shot. This translates into blurring along the partition direction, which limits the effective image resolution. Segmentation of the acquisition into multiple excitations per 3D k-space (also referred to as *shots* within this context), can be used to reduce the ETL (and therefore ETD). Figure 2-11 shows one possibility for such a strategy, in this particular case showing segmentation along the phase-encoding direction, as implemented in the experiments in Chapter 4.
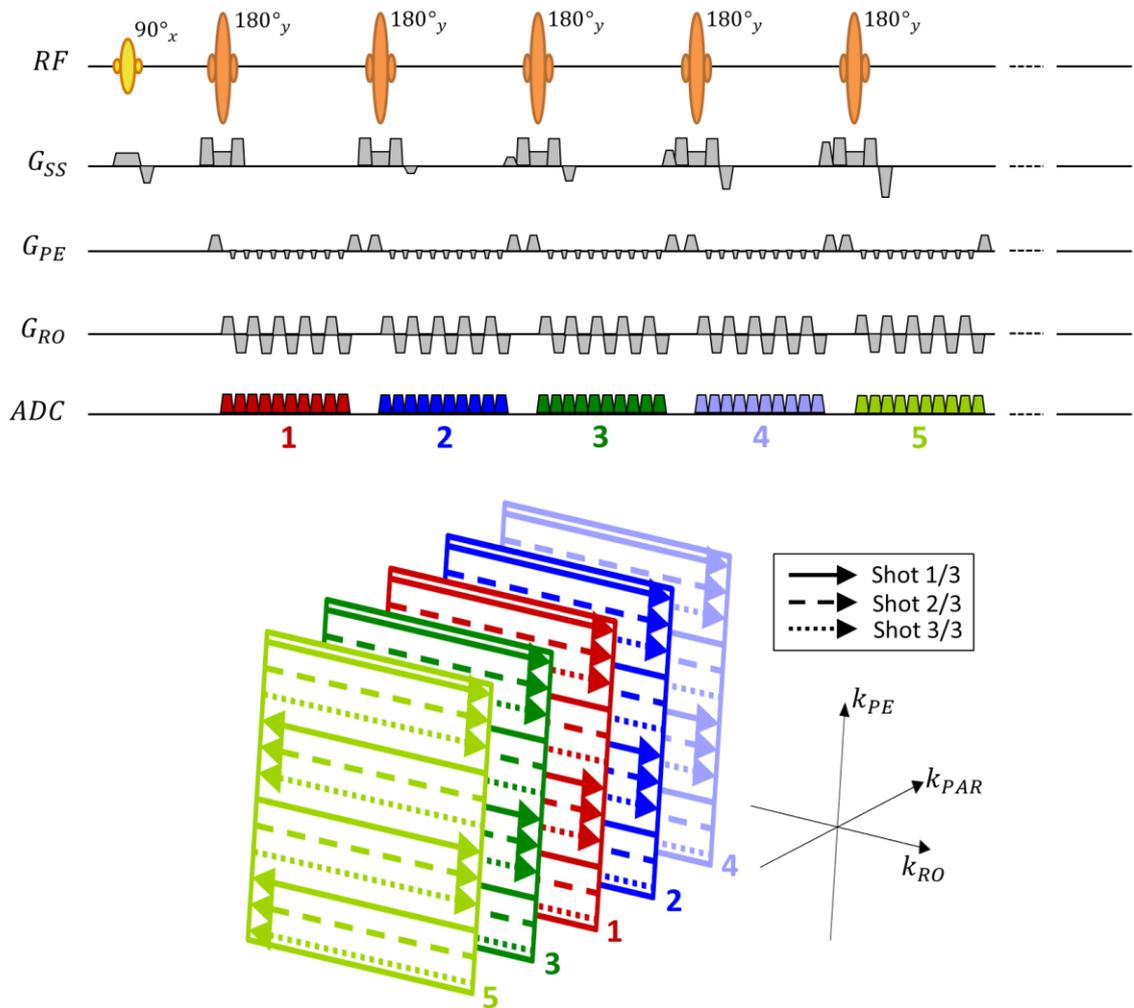
Figure 2-11 3D-GRASE pulse sequence diagram and k-space trajectory. This particular k-space trajectory corresponds to an acquisition with a centric reordering strategy for the different partitions and segmentation along the phase-encoding direction (number of shots = 3) – not shown in the pulse sequence diagram.

2.3   Arterial Spin Labelling

ASL is a non-invasive MRI technique which uses blood water as a freely diffusible tracer to quantify tissue perfusion (Detre et al. 1992; Williams et al. 1992). Perfusion relates to the delivery of blood and thus oxygen and nutrients to tissue at the capillary bed level. This technique is based on the acquisition of a pair of images. One of these images is known as a *tag* (or *label*) image. The label image is flow-sensitized through the use of RF pulses that alter the longitudinal magnetization of arterial blood before it flows into the imaging plane, within the organ of interest.   The remaining image, known as "control" image, is acquired without perturbing the magnetisation of the inflowing blood. Ideally, provided the two images were acquired in the exact same way (except for the labelling module) and the organ of interest did not move, a simple subtraction should cancel the signal from the static tissue. This would leave

us with a purely perfusion-weighted image (PWI), also called ΔM (or difference) image. In practice, this is rarely the case, with several factors introducing errors into the PWIs. This thesis focuses on one of these – that of subject motion. Finally, the obtained PWI is then fed into an appropriate mathematical model describing the relationship between the ΔM signal and the actual blood perfusion. In the end, we obtain a quantitative perfusion map, in relevant physiological units (typically ml/min/100g of tissue). The major steps in a typical ASL experiment are depicted in Figure 2-12.
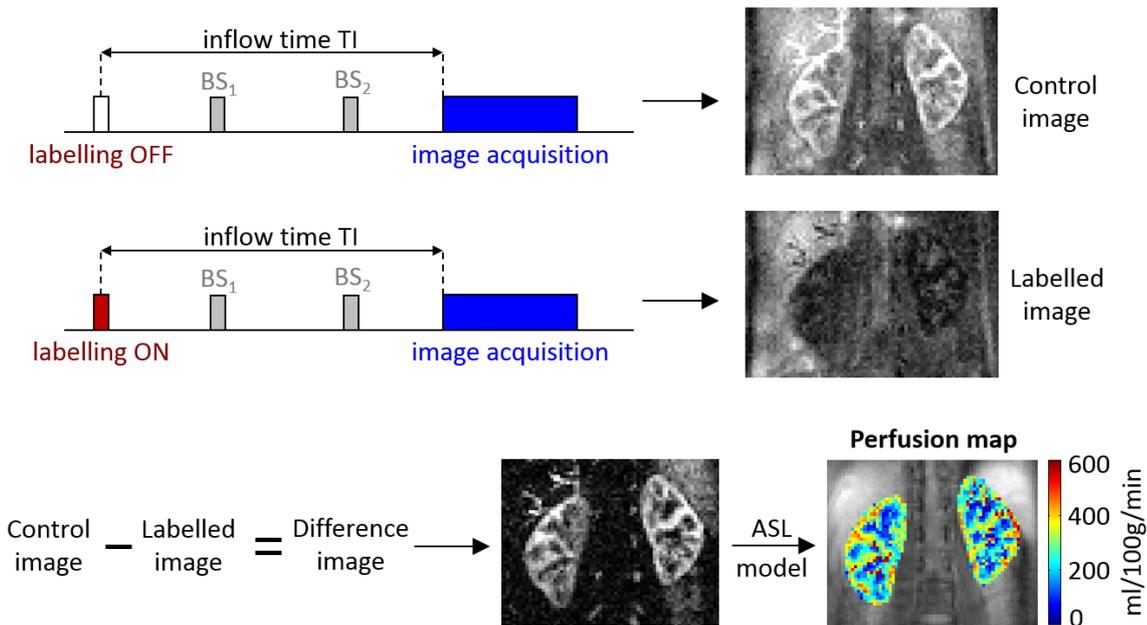
Figure 2-12 ASL overview. In ASL, one PWI is obtained by subtraction of two separate scans: one acquired after inverting the magnetisation of the inflowing blood and the other without perturbing it. The result is then fed to an appropriate model to compute quantitative perfusion maps. Background suppression (BS) pulses (see Section 2.3.4), which are optional, were used to obtain the data in this Figure (using protocols described in Section 4.3.3). This explains the large signal intensity difference between the control and labelled images.

Two important features characterise the ASL tracer and therefore the types of measurements that can be derived from it. First, the magnetic labelling mechanism is based on the excitation of the longitudinal magnetisation away from equilibrium. The longitudinal magnetisation naturally tends to recover which means that the label will decay with T_1, therefore lasting only for a couple of seconds. This gives rise to an important trade-off in ASL. With the exception of velocity-selective labelling (Wong et al. 2006) (see Section 2.3.2) one has to wait a certain amount of time to allow the magnetic tag to move from the labelling site to the tissue of interest. Clearly if the measurement is performed too soon after labelling, no perfusion-weighted signal will be obtained. On the other hand, the longer one waits (potentially allowing a greater amount of labelled blood to arrive), the greater will be the amount of T_1 relaxation occurring, yielding a

lower SNR measurement. Second, the label can freely exchange between the intravascular and extravascular compartments. This makes the transit time of a given labelled water molecule from the arterial to the venous side of the circulation longer that the time it takes for its magnetic label to vanish. Effectively, this means that once the ASL tracer reaches the tissue it will no longer leave, making the ASL measurement intrinsically proportional to the local perfusion.

Two important shortcomings of ASL can be identified. First, ASL is that it is an inherently SNR-limited technique. This is due to the fact that the amount of signal that can be attributed purely to blood perfusion is a small fraction (1-5%) of the tissue signal. As such, any fluctuation to this signal caused by factors extrinsic to blood perfusion will confound the quantitative ASL measurements. To address this, in addition to being performed at relatively low spatial resolution, ASL experiments typically include multiple repeated acquisitions (often referred to as averages) for the purposes of signal averaging. The second shortcoming, particularly important when imaging regions of the body and patients cohorts prone to moving, is ASL's sensitivity to motion. This arises from the need to perform an image subtraction to obtain a PWI. Any subtraction errors resulting from an inconsistent position of the tissue of interest will introduce severe artefacts in the perfusion maps.

Compared to more traditional methods that only allow a global assessment of Renal Blood Flow (RBF), ASL has the advantage of yielding parametric maps which enable an evaluation of focal/regional perfusion defects. Furthermore, unlike imaging methods based in exogenous contrast agents, ASL measurements can be performed multiple times in a single scanning session. In addition, nephrogenic systemic fibrosis is not a concern with ASL, making it ideal for the perfusion assessment of patients with impaired kidney function and paediatric populations (see Section 1.3).

After more than 20 years of development, ASL is now regarded as a well-established technique, in particular in neuroimaging for cerebral blood flow quantification (CBF) (Detre et al. 2012). Significant effort is being put forth by the research community for widespread adoption of ASL into routine clinical practice (Alsop et al. 2015; Mutsaerts et al. 2015), specifically in the standardization of acquisition and analysis strategies.

2.3.1 Magnetisation transfer

Different pools of hydrogen nuclei (^{1}H) contribute to the generation of the MRI signal in varying proportions. The T_2 relaxation times of "bound" protons such as those associated with membranes or macromolecules are too short for their signal to be detected in conventional MRI

(i.e. excluding Ultrashort TE (UTE) pulse sequences (Robson et al. 2003; Tyler et al. 2007). On the other hand, mobile protons in "free" water have long enough relaxation times making them the largest source of MRI signal in most tissues (excluding fat). Nevertheless, it is possible for an exchange of magnetisation to occur between the bound and the free ^{1}H pools. This is a process known as Magnetisation Transfer (MT) (Wolff & Balaban 1989; Henkelman et al. 1993; Henkelman et al. 2001) and can be useful for several applications as an additional source of image contrast (Filippi & Rocca 2007). In ASL, though, it acts mainly as a confounding factor which if ignored may result in an overestimation of perfusion. The RF pulses used to create the ASL tracer are off-resonance with respect to the water frequency and consequently saturate spins associated with macromolecules at the imaging location. Transfer of saturation to the spins in the free water pool then takes place due to the MT effect, reducing the available longitudinal magnetization. This reduces the signal intensity of the resulting tagged image which is a similar effect to the reduction of magnetisation caused by the arrival of tagged spins to the imaging region. Many different approaches have been devised to mitigate this problem, as discussed in the next section.

2.3.2 Labelling Schemes

Throughout the development of ASL, a wide range of approaches were implemented to magnetically label the inflowing blood (Petersen et al. 2006). These can be broadly divided into three main categories: Continuous ASL (CASL), Pulsed ASL (PASL) and Velocity-Selective ASL. Both CASL and PASL have been used extensively for the assessment of CBF for a wide range of applications (Detre et al. 2012; Telischak et al. 2015). This work, as well as the majority of renal ASL implementations (see Section 2.3.6) uses the Flow-sensitive Alternating Inversion Recovery (FAIR) PASL technique (Kim 1995; Kwong et al. 1995).

CASL was the labelling approach employed in the first demonstrations of ASL (Detre et al. 1992; Williams et al. 1992). Typically, a long-duration RF pulse is applied proximal to the region of interest while a gradient is applied in the direction of flow. Blood spins are then inverted as they pass through the resulting inversion plane by the process of adiabatic fast passage (originally used for MR angiography (Dixon et al. 1986)). The temporal bolus width is defined by the duration of the RF preparation pulse. Therefore, there is no uncertainty regarding the duration of the bolus except for dispersion effects en route to the imaging slice(s). By the time any of the labelled spins reach the tissue of interest they have experienced the same amount of T_1 decay. Since in theory all the spins are inverted closer to the tissue of interest, there will be less decay of the magnetic tag and consequently the SNR ratio of the resulting PWI will be higher than in PASL techniques (Eric C Wong et al. 1998).

The initial implementation of CASL was limited to a single slice, centered between the tag and control planes. To allow for multi-slice acquisitions, Alsop *et. al* (Alsop & Detre 1998) proposed a modification to the control acquisition whereby the MT effects are controlled for using a double adiabatic inversion technique. Two inversion planes are employed at the same offset frequency to replicate the MT effects of the label acquisition and their combined effect leaves the net magnetisation of the blood undisturbed.

A different approach which also permits multi-slice studies and is robust to MT effects is to use an additional local coil for labelling. This was originally implemented in animal studies (Detre et al. 1994; Silva et al. 1995; Zhang et al. 1995) and later in human studies (Zaharchuk et al. 1999). The robustness to MT effects exists as long as the RF field from the labelling coil does not interfere with the region to be imaged. Another advantage of this approach compared to single coil CASL methods is the global lower power deposition (Specific Absorption Rate, SAR) which is particularly beneficial at higher magnetic field strengths. The requirement of specialist hardware has limited the widespread adoption of this technique.

A relatively new technique, based on the same flow-driven adiabatic inversion principle as CASL is pulsed-continuous (or pseudo-continuous) ASL (pCASL) (Dai et al. 2008). In this approach, the long RF pulse used in CASL is broken down into multiple short RF pulses, both reducing MT effects and the SAR. This technique also presents the advantage of being easier to implement in modern hardware, which is optimized for short, high-powered RF pulses. This is currently the recommended labelling scheme for the assessment of CBF with ASL in the clinic (Alsop et al. 2015) and it is being increasingly used in renal ASL in recent studies (see Table 2-1).

In PASL, a large volume of blood is inverted almost instantaneously (on the order of a few milliseconds). The inversion of the blood spins is achieved by a spatially-selective frequency-modulated adiabatic inversion slab, rather than flow-induced adiabatic inversion (as in CASL). The first PASL technique was EPISTAR (Echo-Planar Imaging and Signal Targeting with Alternating Radio frequency) proposed by Edelman *et. al* (Edelman et al. 1994) which controlled for MT effects by playing out identical inversion RF pulses in symmetrically opposite locations with reference to the imaging slice. However, this technique had limited multi-slice capabilities as the MT effects could only be completely eliminated in the tag-control subtraction of the slice in the central location. Improvements to this technique to enable multi-slice acquisitions consisted in the application of a double inversion pulse for the control image at the same location as the inversion pulse for the tag image (Edelman & Chen 1998). This implementation also solves the

issue of negative blood signal which may flow into the imaging region after the control pulse in the original EPISTAR method.

Shortly after the introduction of EPISTAR, the Flow-sensitive Alternating Inversion Recovery (FAIR) PASL technique was proposed (Kwong et al. 1995; Kim 1995). The tag image is acquired after a non-spatially-selective RF pulse. In order for the static tissue signal in the imaging region to be the same as in the tag condition, the control image is acquired after an inversion pulse which is slice-selective for the imaging region (see Figure 2-13). An immediate advantage of this approach is the fact that no off-resonance RF pulses are applied with respect to the imaging region making this technique insensitive to MT effects. FAIR is the most widely utilised PASL technique, especially in renal applications (see Section 2.3.6).
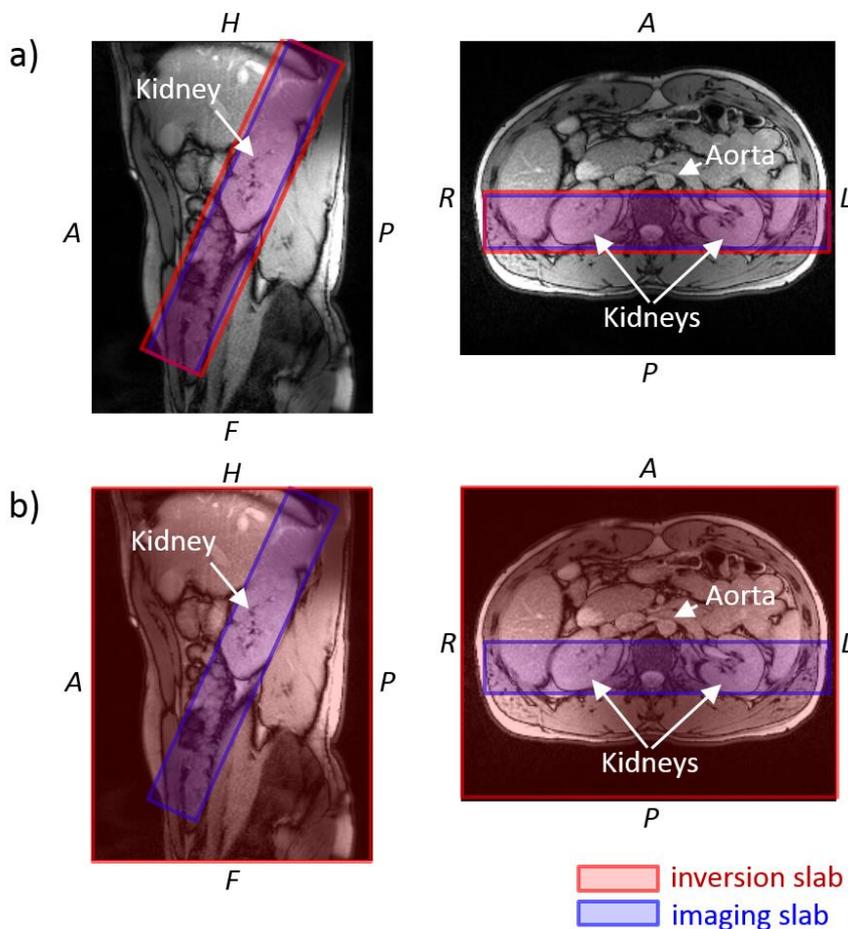
Figure 2-13 FAIR labelling module. Prescription of a) slab-selective and b) non-selective labelling and imaging slabs.

Unlike in CASL, where the bolus length is determined by the duration of the labelling RF pulse, in single inflow time (TI) PASL studies, the accuracy of the perfusion quantification can be improved by making the bolus duration a known parameter. Several techniques were developed for this purpose: initially QUIPSS (Quantitative Imaging of Perfusion Using a Single Subtraction)

and QUIPSS II (Eric C. Wong et al. 1998) and shortly after Q2TIPS (QUIPSS II with thin-slice TI_1 Periodic Saturation (Q2TIPS) (Luh et al. 1999). In QUIPSS, a saturation pulse is applied at given time TI_1 after the labelling pulse, removing the contribution of the labelled spins which arrived to the imaging region before this time. In the second version of this technique, QUIPSS II, the saturation pulse is instead applied to the labelling slab at time TI_1 which cuts off the trailing edge of the bolus. The resulting bolus has then known duration equal to TI_1. Another advantage of QUIPSS II is that it allows for multi-slice acquisitions. Q2TIPS further improves perfusion quantification by providing a greater saturation efficiency using a periodic train of thin-slice saturation pulses whose slice profile is also a better match to the tagging pulse profile. Q2TIPS was used in all acquisitions throughout the thesis where renal perfusion was quantified (i.e. Chapters 5, 6 and 7).

Several other PASL techniques were introduced which differ mainly in the procedure for the mitigating MT effects on the signal intensity of the ASL difference image and resulting perfusion quantification, as summarized in a number of review articles (Golay et al. 2004; Günther 2014; Wong 2014; Petersen et al. 2006).

However, many of these approaches are no longer as relevant nowadays since MT-effects can be controlled by saturating the imaging volume either before and/or after the labelling/control inversion pulses.

Overall the main advantages of PASL compared to CASL are a lower RF power deposition and lower MT effects due to the shorter duration RF pulses and high inversion efficiency. Disadvantages are the lower SNR and the dependency on the RF coil coverage.

Velocity-selective ASL (Wong et al. 2006) is a relatively new approach, which makes use of the decelerating property of arterial blood. Its main advantage is an insensitivity to transit delay effects since there is no spatial gap that the blood label must travel before it reaches the tissue of interest. This is achieved by applying non spatially-selective, but rather velocity-selective pulses which will label blood spins closer and within the imaging region, as long as they are travelling above a certain cut-off velocity. After a given time delay allowing the labelled blood to reach the capillary bed, an image is acquired whose signal originates from spins with velocity lower than the cut-off velocity. This condition is enforced by additional velocity-selective pulses during the imaging module.

2.3.3   Pre and post-labelling saturations

Pre and post-labelling saturation pulses increase the accuracy of ASL measurements by reducing systematic errors which introduce signal differences between the label and control conditions which are not due to tissue perfusion (i.e. result in artificial perfusion-weighted signal). Such errors may be caused by residual magnetisation transfer effects, differences in inversion efficiencies between label and control conditions or imperfect slice profiles. Furthermore, saturating the imaging volume also ensures that at the start of each TR the state of the longitudinal magnetisation is always known (nulled), regardless of the duration of the TR of the sequence, which may vary due to respiratory-triggering. This allows background-suppression pulses to be used to null the static tissue magnetisation at the time of the excitation pulse.

2.3.4   Background Suppression

Background suppression (BS) is a technique that aims to improve the sensitivity of ASL measurements by attenuating signal from static tissue at the imaging region, and thereby reducing the noise contribution of this tissue (Ye et al. 2000). It builds upon the principles of multiple inversion recovery suppression introduced by Dixon *et. al* (Dixon et al. 1991). In this work, the BS scheme proposed by Günther *et. al* (Günther et al. 2005) was used. Two non-selective inversion pulses are applied at times τ_1 and τ_2 to null tissue components with longitudinal relaxation rates $R_{1,opt}$ and $0.5 \times R_{1,opt}$. For each TI, τ_1 and τ_2 are computed using the following equations:

$$\tau_1 = TI + 2\ln\left(0.25 + 0.75e^{-0.5 \cdot TI \cdot R_{1,opt}}\right)/R_{1,opt} \qquad \text{Equation 2-24}$$

$$\tau_2 = TI + 2\ln\left(0.75 + 0.25e^{-0.5 \cdot TI \cdot R_{1,opt}}\right)/R_{1,opt} \qquad \text{Equation 2-25}$$

Where $R_{1,opt}$ equals $1/T_{1,opt}$. For all our measurements, $T_{1,opt}$ was assigned the value of 700ms. An example of the evolution of the magnetisation for two tissues with different longitudinal relaxation times before data acquisition is shown in Figure 2-14.
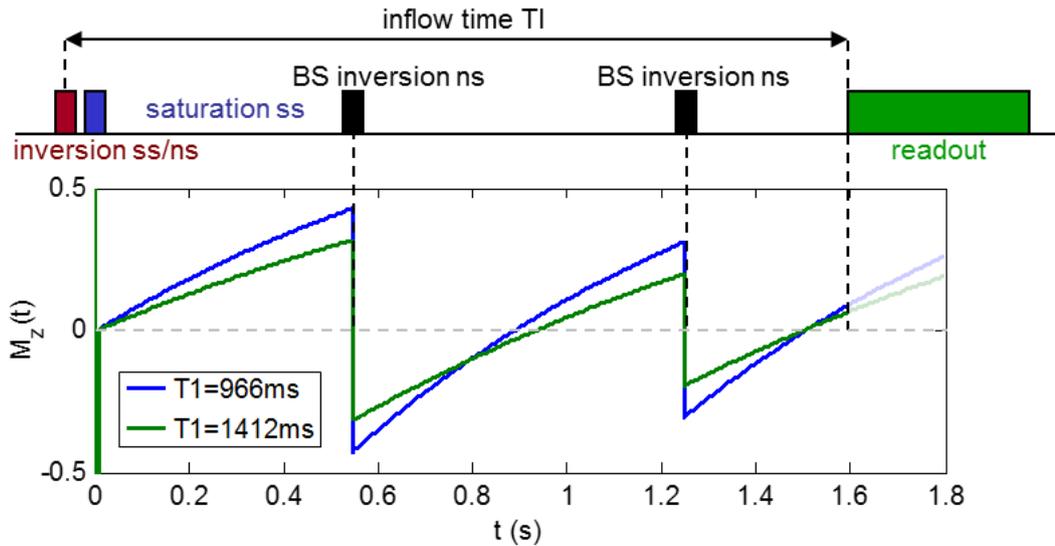
Figure 2-14 Incorporation of BS inversion pulses during the ASL M_z preparation in the imaging region prior to the readout module. In this example, TI=1600ms, with nulling of the magnetisation occurring at t=1500ms, by applying BS pulses at 549ms and 1249ms. At the time of the acquisition, the longitudinal magnetisation values in the renal cortex (T_1 = 966ms) and medulla (T_1 = 1412ms) (de Bazelaire et al. 2004) are approximately 9% and 6% of the equilibrium magnetisation, respectively. The sequence blocks are not to scale with respect to their temporal width.

2.3.5 Quantification in ASL

As previously discussed, the most common application of ASL is CBF quantification. Single-compartment models commonly applied for CBF quantification are typically employed in renal studies. One might argue that this is a limitation since the unique anatomy and physiology of the kidneys may not be accurately represented by these models. On the other hand, the application of complex quantification models is not straightforward and may entail loss of precision in renal ASL due to current technical limitations, in particular the low SNR of the measurements.

Perfusion quantification in the original ASL studies was achieved by modifying the Bloch equation for the longitudinal magnetization by adding a flow term accounting for arterial inflow and venous outflow (Williams et al. 1992).

A more general approach was proposed by Buxton *et. al* (Buxton et al. 1998) which describes the evolution of the magnetisation as a sum over the history of the delivery of the magnetic tag by arterial blood and clearance by venous outflow and longitudinal relaxation:

$$\Delta M(t) = 2M_{0b}f \int_0^t c(t')r(t - t')m(t - t')dt'$$

$$= 2M_{0b}fc(t) * (r(t)m(t))$$

Equation 2-26

53

In this equation, $\Delta M(t)$ is the signal intensity in the difference image at time t, M_{0b} is the equilibrium magnetisation of arterial blood, f is the flow, $c(t')$ is the delivery function, representing the normalised concentration of magnetization arriving at the voxel at time t, $r(t, t')$ is the residue function, corresponding to the fraction of tagged water molecules that arrived at time t' and remain at the voxel at time t. Finally, $m(t, t')$ is the relaxation function, corresponding to the fraction of original longitudinal magnetisation in the tag that arrived at the voxel at time t' and remains there at time t and $*$ denotes the convolution operator.

Assuming plug flow conditions, single-compartment kinetics and that as soon as the magnetic label reaches the tissue voxel it starts decaying with the T_1 of the tissue, the following expressions can be used for $c(t)$ (PASL case), $r(t)$ and $m(t)$ (PASL/CASL case):

$$c(t) = \begin{cases} 0 \, , & 0 < t < \Delta t \\ \alpha e^{-t/T_{1b}} \, , & \Delta t < t < \Delta t + \tau \\ 0 \, , & \Delta t + \tau < t \end{cases}$$

Equation 2-27

$$r(t) = e^{-ft/\lambda}$$

Equation 2-28

$$m(t) = e^{-t/T_1}$$

Equation 2-29

Where Δt is the arterial transit time, τ is the temporal bolus width, α is the labelling efficiency, T_{1b} is the longitudinal relaxation time of blood, λ is the tissue-blood partition coefficient for water (describing the relative concentration of water in tissue and blood), and T_1 is the longitudinal relaxation time of the tissue.

The signal in each voxel of the difference image given in Equation 2-26 under the assumptions represented by Equation 2-27, Equation 2-28 and Equation 2-29 then becomes:

$$c(t) = \begin{cases} 0 \, , & 0 < t < \Delta t \\ 2M_{0b}f(t - \Delta t)\alpha e^{-t/T_{1b}}q_p(t) \, , & \Delta t < t < \Delta t + \tau \\ 2M_{0b}f\tau\alpha e^{-t/T_{1b}}q_p(t) \, , & \Delta t + \tau < t \end{cases}$$

Equation 2-30

Where:

$$q_p(t) = \begin{cases} \dfrac{e^{k \cdot t}(e^{-k \cdot \Delta t} - e^{-k \cdot t})}{k(t - \Delta t)} \, , & \Delta t < t < \Delta t + \tau \\ \dfrac{e^{k \cdot t}(e^{-k \cdot \Delta t} - e^{-k(\tau + \Delta t)})}{k \cdot \tau} \, , & \Delta t + \tau < t \end{cases}$$

Equation 2-31

$$k = \frac{1}{T_{1b}} - \frac{1}{T_1} - \frac{f}{\lambda}$$

Equation 2-32

Quantification can be achieved by performing one or multiple measurements at different TIs after the labelling. In the former case (i.e. single-TI measurements), the appropriate sub-function in Equation 2-30 can be solved for f. In multi-TI studies, the signal intensities in the difference images at the acquired TIs are fit to Equation 2-30, from which we obtain values for flow, arterial transit time and, if not pre-determined by the experiment settings, bolus width (see Figure 2-15).
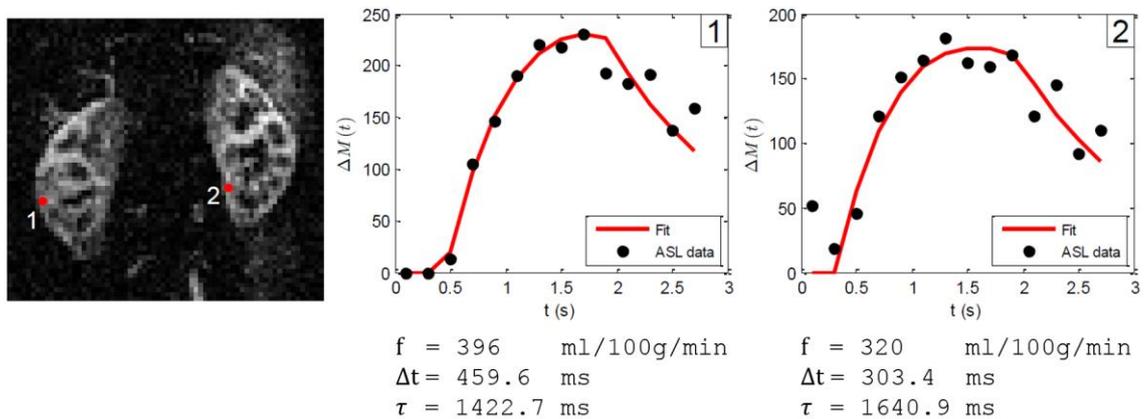
Figure 2-15 Example of voxelwise multi-TI ASL fitting. Left – ASL PWI. Middle/right – Buxton model fit to ASL TI-series for voxels 1 and 2, respectively.

2.3.6 Renal ASL

The majority of ASL studies focus on brain applications (Detre et al. 2012). One can argue that the smaller amount of motion involved when imaging the brain plays a role in this. This is especially true in ASL, since the PWIs result from the subtraction of two images acquired at different points in time. Susceptibility artefacts can also be problematic, especially when imaging the abdomen as shown by Chen *et. al* (Chen et al. 1997). Fortunately, ASL is simply a preparation of the magnetization for the subsequent imaging module, which can be optimised according to the intended application. Despite these difficulties, the use of ASL to assess renal function has been steadily increasing, with more than 20 studies in humans published in the last 5 years (see Table 2-1). Furthermore, comparisons with more established methods to assess blood flow have been performed. Examples include comparisons to Para-aminohippurate (PAH) clearance: (Ritt et al. 2010), microspheres: (Nathan S Artz et al. 2011), ultrasound (Warmuth et al. 2007) and Dynamic contrast-enhanced (DCE)-MRI: (Winter et al. 2011; Wu et al. 2011; Zimmer et al. 2013; Cutajar et al. 2014; Conlin et al. 2017), generally supporting the premise that ASL can deliver reasonable estimates of blood flow. A considerable amount of published renal ASL studies

address the quantification of perfusion in healthy volunteers and reproducibility of the technique. More recently, ASL has been used to compare RBF measurements in healthy controls and patients. Table 2-1 summarises RBF quantification results in the majority of these studies, in humans.

Roberts *et. al* conducted the first renal ASL experiments (Roberts et al. 1995). They modified a 2D gradient echo sequence by adding a CASL module for labelling and acquired one axial slice in 7 healthy volunteers during breath hold (BH). Intravascular signal was reduced using a bipolar gradient pulse after excitation. Chen *et. al* (Chen et al. 1997) compared two PASL approaches with echo planar imaging (EPI) and half-Fourier single-shot turbo spin-echo (HASTE) readout modules in the brain and kidneys. This study alluded to the potential issues of using EPI in the abdomen, namely susceptibility effects. Wang *et. al* (Wang et al. 1998) demonstrated the value of ASL in the early detection of transplant rejection by monitoring renal perfusion changes following renal transplantation in an animal model.

Berr *et. al* (Berr et al. 1999) utilised a partial Fourier turbo spin-echo sequence with a custom labelling scheme (Berr & Mai 1999). To reduce artefacts, they segmented the acquisition (number of shots = 2) and acquired multiple averages, resulting in an approximately 30-second acquisition for a single slice. Their semi-quantitative approach showed a signal decrease in the PWIs that correlated with the degree of stenosis in diseased kidneys. Karger *et. al* (Karger et al. 2000) provided additional data showing realistic values for renal perfusion in 10 healthy volunteers. They combined a FAIR labelling scheme with a UFLARE (Ultra Fast Low Angle Rare) (Norris 1991) imaging readout. Michaely *et. al* (Michaely et al. 2004) used a FAIR HASTE sequence at 1.5T to acquire a single slice at a TI of 1.2s in healthy volunteers and patients with different types of renal disease. Three measurements were acquired under a 15s BH, and measurements with motion artefacts were manually discarded (retrospective sorting). They performed only a semi-quantitative analysis and report significant differences in the SNR between healthy and diseased kidneys but no significant differences between the different groups of patients.

Martirosian *et. al* (Martirosian et al. 2004) presented the first study using a TrueFISP (Oppelt et al. 1986) imaging readout in renal ASL. They report an overall better image quality than when using EPI. The advantage of this method when compared to the previously referred non-EPI studies is mainly the shorter scan time. They scanned 12 volunteers and 1 renal transplant patient in BH and free breathing conditions, where subjects were asked to synchronize breathing with the acquisition. However, they only report quantitative perfusion values for one volunteer,

which agree with previous studies. Of the 36 acquired control-tag pairs (free-breathing), only 26 were used for quantification due to differences in position of the kidneys.

As the previous study alluded to, motion is probably the greatest challenge in renal ASL. Two conference abstracts have been presented employing navigators for motion-correction (Warmuth et al. 2005; Gach et al. 2006).

De Bazelaire *et. al* assessed the feasibility of ASL for the study of metastatic renal cell carcinoma (de Bazelaire et al. 2005). They used a background-suppressed FAIR labelling module in combination with a single-shot fast spin echo readout with multiple averages at 3T. As a partial validation of the technique, they scanned 5 healthy volunteers and obtained quantitative values of perfusion consistent with existing literature. Interestingly, their data suggested the potential value of ASL for tumour blood flow monitoring after antiangiogenic therapy. Fenchel *et. al* (Fenchel et al. 2006) utilised the previously described FAIR TrueFISP scheme (Martirosian et al. 2004) to scan patients with (n=12) and without (n=6) history of renal artery stenosis (RAS). In patients without RAS, the perfusion values obtained under BH and non-BH conditions were not statistically different. This time they used rigid registration on the non-BH scans, independently for right and left kidneys. They were able to differentiate patients with no or low-grade RAS from those with high-grade RAS.

The first study addressing specifically the problem of motion in renal ASL was published by Robson *et. al* (Robson et al. 2009). Using a single-shot sequence (FSE), they were able to focus specifically on reducing motion-related subtraction errors. For this purpose, different combinations of BS and breathing strategies were evaluated. Both the variance of the measurements and image quality scoring by radiologists demonstrated the advantage of using BS. From an SNR point-of-view, controlled breathing (synchronizing the respiratory cycle with the image acquisition) was superior to BH acquisitions. However, this requires subjects to be compliant, which may be a problem in the clinic, especially with younger patients. They also showed that the image quality of free-breathing acquisitions can be improved by performing retrospective image sorting, allowing only measurements close to end-expiration to be considered in the averaging procedure. Kiefer *et. al* (Kiefer et al. 2009) scanned 11 healthy volunteers with a single-slice FAIR TrueFISP sequence. Forty measurements were performed under multiple BHs (total scan time of 08m24s), of which some were rejected to obtain a consistent dataset before averaging. Perfusion quantification was performed with both the classical single-compartment model and a two-compartment model which accounts for

restricted capillary permeability (Parkes & Tofts 2002), which reduced the overall variability of the perfusion measurements among subjects.

In the work of Gardener *et. al* (Gardener & Francis 2010), the use of EPI was revisited. However, this time parallel imaging was used allowing them to overcome many of the drawbacks of EPI when applied to the abdomen. Specifically, PI revealed usefulness in decreasing the TE and the EPI echo-train length (ETL) which both contributed to a better image quality and to the possibility of covering the entire kidney with 7 slices. Similarly to a previous study (Robson et al. 2009), they compared different breathing strategies and their combination with a BS scheme. Interestingly, their data suggested that BS reduced the SNR of the perfusion weighted image and caused an underestimation of perfusion. They also utilised rigid image registration which improved image quality in all acquisition methods. After realignment, they found no significant difference between perfusion rates measured from free-breathing and BH data.

Song *et. al* (Song et al. 2010) introduced a free-breathing dual-navigated FAIR TrueFISP approach. The tagged images (i.e. acquired after slice-selective inversion) were discarded if the lung-liver interface position was not within the acceptance windows of both navigators. They acquired 20 single-slice control/tag pairs with acceptance rates of 40-60%, resulting in scan times of 6-9 min. Ritt *et. al* (Ritt et al. 2010) a FAIR TrueFISP sequence to scan 24 patients with metabolic syndrome (3 scans were non diagnostic due to poor image quality), before and after pharmacological intervention. Importantly, this is the first study comparing the ASL perfusion measurements with renal plasma flow (RPF) measurements obtained by a clearance study with PAH, where a moderate correlation was found. This study also suggests the capability of ASL to approximate changes in kidney perfusion. Lanzman *et. al* (Lanzman et al. 2010) scanned 20 renal allograft recipients (FAIR TrueFISP at 1.5T, 30 averages, single-slice, TI of 1.2s, scan time 4m30s). They have shown a significant reduction of cortical perfusion in the group of patients with an acute deterioration of renal function. Interestingly, they reported the smaller effect of respiratory motion on transplanted kidneys (located in the iliac fossa).

Artz *et. al* (Nathan S. Artz et al. 2011) assessed inter and intra-visit reproducibility in adult healthy subjects and transplant recipients. The scanning protocol consisted of a single slice, single TI (TI=1.2s) FAIR TrueFISP acquisition with 32 measurements. They used respiratory triggering and coached the subjects to avoid breathing during the data readout. They report reproducible measurements for the cortex, but not for the medulla. Cutajar *et. al* (Cutajar et al. 2012) demonstrated good intra- and inter-session repeatability of renal ASL measurements in healthy volunteers using a FAIR segmented 3D-GRASE acquisition scheme. Using this efficient

readout module, they were able to achieve whole kidney coverage at multiple inversion times. Prospective respiratory triggering and BS were used to maximise the accuracy of the measurements. Rossi *et. al* (Rossi et al. 2012) also employed a FAIR TrueFISP sequence to assess the spatial heterogeneity of perfusion in healthy volunteers and patients with mild CKD (CKD stages 1-3). They acquired 30 single-slice, single-TI ASL pairs during 30 BHs of 14s. Image registration (rigid) was used to correct for kidney shifts between the measurements. They found significant differences both in cortical and parenchymal (i.e. cortex + medulla) perfusion of the healthy volunteers compared to the patients, although the differences in the latter case were smaller.

Schneider *et. al* (M. P. Schneider et al. 2012) have shown, in a group of 34 hypertensive patients, an increase in renal perfusion during aliskiren therapy. Post-aliskiren, RBF returned to original levels, even though the effect of blood pressure reduction was maintained. Imaging was performed at 1.5T using a standard FAIR TrueFISP under multiple BHs. Wang *et. al* (Wang et al. 2012) hypothesized ASL could be sensitive to changes in renal haemodynamics following the intravenous administration of furosemide. They scanned 11 healthy subjects with a FAIR single-shot FSE sequence at 3T. They measured a single slice which required 2 BHs of 20-30 seconds. A region-of-interest (ROI) analysis of cortical RBF. Three out of the 11 subjects failed the BH and image registration had to be used for correcting for the resulting motion artefacts. Both cortical and medullary perfusion was shown to decrease following the administration of furosemide.

Dong *et. al* (Dong et al. 2013) scanned 18 healthy volunteers and 13 acute kidney injury patients with a FAIR-FSE sequence at 3T. Three ASL pairs of a single slice at a fixed TI of 1.4s were acquired during a BH with duration of 20-30s. Despite respiratory coaching two patients had to be excluded from the analysis due to BH failure. Data processing consisted of non-linear registration for realignment of the different measurements, followed by a standard single-compartment ASL analysis. They report significantly lower renal perfusion values in the patient group when compared to healthy volunteers.

Cutajar *et. al* (Cutajar et al. 2014) assessed the reproducibility of ASL and DCE-MRI and compared the perfusion estimates obtained with both techniques. For ASL, a similar protocol was used as described previously (Cutajar et al. 2012). Both techniques were shown to be reproducible, although interestingly their data suggests ASL is more reproducible than DCE-MRI. No significant differences were found between the RBF values obtained with both techniques. Another reproducibility study was performed by Gillis *et. al* (Gillis et al. 2014). They also demonstrate reproducibility in the ASL measurements, this time at 3T. A single slice was acquired with a FAIR

TrueFISP approach during BH. Image registration algorithms were then used to re-align the different measurements. They report no significant differences between two different ASL MRI measurements performed on different occasions.

Tan *et. al* (Tan et al. 2014) used a FAIR TrueFISP sequence incorporating a low resolution 2D FLASH navigator to track the position of the lung-liver interface (acquired at the end of the main readout), which allowed for selective averaging according to the phase of the respiratory cycle. This was tested in 10 healthy volunteers and 5 patients with CKD (CKD stage 2-4). They have shown significantly higher SNR in free-breathing acquisitions compared to acquisitions during BH due to the higher number of measurements, even after data rejection. The acquisition efficiency was 35% in patients. They obtained significantly lower cortical and medullary perfusion rates in patients. In the work of Heusch *et. al* (Heusch et al. 2014)*,* 98 transplant patients underwent renal ASL using a FAIR TrueFISP sequence at 1.5T (n = 65) and 3T (n = 33). They were divided in two groups according to their estimated glomerular filtration rate (eGFR) (>30 or <30 ml/min/1.73m^2). Significantly higher perfusion was found in the group with the higher eGFR. There were no significant differences between the perfusion estimates obtained at 1.5T or 3T for both groups. In addition, the perfusion values correlated with eGFR suggesting the potential value of renal ASL to provide a marker of renal function.

Cutajar *et. al* (Cutajar et al. 2015) conducted the first study assessing renal perfusion in living kidney donors (n = 8) pre and 1 year post-nephrectomy. Data acquisition was performed with a similar FAIR multi-shot 3D-GRASE scheme as previously reported (Cutajar et al. 2012). Formal GFR measurements were also performed to allow the calculation of the filtration fraction. They reported an increase in GFR in all donors which was accompanied by a comparatively smaller, but significant, increase in RPF per kidney, suggesting an increase in blood flow in every glomerulus (the number of glomeruli is constant). Hueper *et. al* (Hueper et al. 2015) scanned 46 patients 4–11 days after transplantation. Twenty-six of the patients had delayed graft function and were shown to have significantly reduced perfusion (using a FAIR TrueFISP ASL method). ASL perfusion correlated with renal function assessed with eGFR. They reported higher perfusion in the living donor grafts compared to deceased donor grafts. Additionally, after the 1 year follow up their data suggests the potential of ASL measurements to help determine whether dialysis will be required after transplantation. This study provides further evidence of the potential value of ASL for monitoring renal allograft rejection, potentially reducing the need for invasive biopsies in the early post-transplantation period, as Wang *et. al* (Wang et al. 1998) had previously suggested.

A recent study by Gillis *et.* al (Gillis 2016) has also shown significantly decreased renal perfusion in CKD patients compared to healthy volunteers. Imaging was performed using a single-slice, single-TI (TI = 900ms) FAIR TrueFISP scheme during multiple BHs. Affine image registration with a correlation coefficient (CC) type cost function was used to ensure alignment of M_0, T_1 and perfusion data. Work by Ren *et. al* (Ren et al. 2016) has shown a reduction in renal cortical perfusion comparing healthy controls to renal transplant recipients with normal allograft function. A further reduction was found when comparing the latter group to recipients with impaired allograft function. A positive correlation between renal perfusion and eGFR in the transplant recipients was also reported. Scans were performed in free-breathing without respiratory triggering owing to the location of the transplanted kidney. Niles *et. al* published the first longitudinal study with matched pairs renal allograft donors and recipients (all clinically stable) (Niles et al. 2016). Scans were performed with a standard FAIR TrueFISP approach with a coached breathing strategy. They found a progressive decrease in renal perfusion in donors compared to the baseline scan (i.e. before transplant surgery) over a period of two years. Half of the transplant recipients received losartan therapy at 3 months after transplant. Recipients who did not undergo losartan therapy showed a more marked decrease in cortical perfusion as compared to the donors. At the 2 year point, cortical perfusion of the losartan group had decreased about 50% less than the losartan therapy group.

A recent study by Robson *et. al* presents an innovative approach for acquiring volumetric renal ASL data at near-isotropic resolution (Robson et al. 2016). This is based in a background-suppressed multi-shot 3D FSE pulse sequence with a radial k-space trajectory, which is preceded by a PCASL labelling module. This was compared to a more standard ASL approach using a single-shot 2D FSE sequence with the same labelling technique, both in terms of perfusion quantification and SNR of the measurements. 3D acquisitions were performed both during synchronized and free breathing, whereas the 2D acquisition was performed with synchronized breathing only. Similar cortical perfusion values and SNR in healthy volunteers were obtained with the two techniques. Particularly interesting is that this result was found in both types of breathing strategies, including in the segmented multi-shot 3D acquisitions with a 23 shots, even though it was acknowledged that image artefacts and blurring were seen in the free-breathing acquisitions. A semi quantitative evaluation by radiologists of both acquisitions in patients with renal masses has shown clinical advantages of the 3D sequence due to higher spatial coverage and thinner slices which reduced partial volume effects.

Shimizu *et. al* (Shimizu et al. 2017) investigated how the choice of post-labelling delay time in pCASL affects perfusion estimates in the renal cortex. ASL data was acquired at 3T with a

multislice 2D spin-echo EPI at 3 different post-labelling delay times (0.5, 1, and 1.5s) and was fitted to a typical single compartment model, to yield arterial transit time corrected RBF (ATC-RBF). More perfusion estimates were also obtained using each of the 3 PLDs independently. All except the PLD of 1.5s exhibited a moderate correlation to eRBF obtained with 99mTc-MAG3 scintigraphy. A recent study by Kim et al (Kim et al. 2017) not only assessed the reproducibility of RBF but also of arterial transit time in both the renal cortex and medulla. For this, they scanned 25 adult healthy volunteers at 3T, twice (scans 1 hour apart). The acquisition consisted of PCASL labelling with 4 post labelling delays (0.5, 1, 1.5 and 2s), a BS module, and a 2D spin-echo EPI readout. Each PLD was repeated 12 times for signal averaging, requiring multiple BHs. This resulted in a scan time of ~12 min per subject. They found RBF to be repeatable in the cortex but not in the medulla (see Table 2-1).

Li *et. al* (Li et al. 2017) used a FAIR TrueFISP sequence at 3T with retrospective motion correction (2D navigator + rigid registration) to evaluate differences in RBF between 30 healthy controls and 33 patients with diabetes and stage 3 CKD (median eGFR 46.7 ml/min/1.73m^2). A single coronal-oblique slice was acquired at an inflow time of 1.5s (50 control/label pairs) and 2s (100 control/label pairs), for controls and patients, respectively. A constant T_1 value of 1.15 seconds was used. Cortical and medullary ROIs were manually selected in the final RBF maps. RBF in CKD patients was found to be approximately half to that of controls. Cortical and medullary RBF correlated with eGFR. Interestingly, they determined optimal thresholds for separating controls and CKD patients of 142.9 and 24.1 ml/100g/min, with corresponding Youden indexes of 0.85 and 0.64, respectively for cortical and medullary RBF.

A number of studies also shown the usefulness of ASL for the assessement of perfusion in renal cell carcinoma (Cedric de Bazelaire et al. 2008; C de Bazelaire et al. 2008; Lanzman et al. 2012; Zhang et al. 2016; Yuan et al. 2016) (not included in Table 2-1).

In summary, the majority of studies limit the acquisition to a single slice at a fixed TI (typically ranging from 1-1.5s). One recent study used a different fixed TI for healthy volunteers (TI=1.5s) and patients (TI=2.0s) to account for possible delays in the arrival of the label in the latter case (Tan et al. 2014). Multiple measurements (up to 40) are performed for SNR benefits and in certain cases to obtain sufficient data to cope with retrospective sorting algorithms (e.g. to reject data acquired during unwanted parts of the respiratory cycle). Breath-hold acquisitions have been tested, but even in compliant adult subjects this does not completely solve the problem of motion. Reasons for this include the fact that the kidneys change position between different BHs and residual motion exists within BHs (Holland et al. 1998). In addition, the SNR per unit time of

these approaches is lower than in free-breathing scans. A common approach is to train subjects so that they match their respiratory cycle to the specific repetition time of the sequence (e.g.: (Martirosian et al. 2004; Nathan S Artz et al. 2011; Robson et al. 2016), among others). If they are able to comply with these instructions, this essentially results in a series of quick BHs during which the data acquisition is performed. The applicability of this approach may be limited, as compliance is likely to decrease in patient and/or young/elderly populations.

Although BS is recommended as the main method to mitigate motion artefacts in brain ASL (Alsop et al. 2015), further work is required to verify if this holds true for renal imaging. BS has been used sparingly by certain groups, with different conclusions regarding the effectiveness of the method (Robson et al. 2009; Gardener & Francis 2010). It is also important to note that as the degree of BS increases, the effectiveness of image registration approaches is likely to decrease. This is the case given that BS reduces the amount of anatomical information to drive the registration algorithms, particularly if more advanced non-rigid registration transformations are employed. Gardener and Francis (Gardener & Francis 2010) used a combination of BS and image registration, but one might argue that given their 2D acquisition scheme, BS is not as effective as what could be achieved using 3D acquisitions, such as in the work presented by Cutajar *et. al* (Cutajar et al. 2012). Overall, image registration algorithms are commonly used, most frequently employing rigid transformations to re-align the different measurements (averages). Image-based navigators have also been gaining attention in the last few years.

ASL has been used for the investigation of the effects of certain drugs in renal perfusion. Examples include furosemide (Wang et al. 2012), aliskiren (M. P. Schneider et al. 2012), and losartan (Niles et al. 2016). Other therapies have also been researched (e.g. renal denervation (Ott et al. 2013)). Changes in in renal cortical perfusion following intravenous administration of colloid fluids in healthy volunteers were also assess using ASL by Chowdhury *et. al* (Chowdhury et al. 2012).

Table 2-1 Review of Renal ASL studies

Reference	N_{subj}	B_0 (T)	Labelling	Readout	RBF (ml/100g/min)	Highlights
(Roberts et al. 1995)	7 (H)	1.5	CASL	2D-GE	Cortex = 278±55; Medulla = 55±25	First renal ASL study in humans
(Chen et al. 1997)	2 (H)	1.5	PASL (STAR)	HASTE	---	Highlights robustness to susceptibility and chemical shift artefacts of HASTE readouts vs. EPI
(Berr et al. 1999)	6 (H); 7 (P)	1.5	PASL	HASTE	---	Renal Vascular Disease. ASL signal correlated to % stenosis from X-ray angiograms
(Karger et al. 2000)	10 (H)	1.5	PASL (FAIR)	UFLARE	WK = [140-319], 213±55 ([range], mean±SD)	Quantitative assessment; First using FAIR.
(Martirosian et al. 2004)	12 (H); 1 (TrR)	1.5	PASL (FAIR)	TrueFISP	Cortex = 220-280; Medulla = 50-120 (range)	First time using TrueFISP readout
(de Bazelaire et al. 2005)	5 (H)	3	PASL (FAIR)	FSE	Cortex = 275±14 (mean±SD)	First study @ 3T using FSE readout
(Warmuth et al. 2005)	---	1.5	PASL (FAIR)	TrueFISP	Image visual assessement only	First study using navigators for prospective motion correction
(Fenchel et al. 2006)	12+6 (P)	1.5	PASL (FAIR)	TrueFISP	Patients without RAS: WK: 240±57 (149–313); Medulla: 117±25 (80–151) Cortex: 316±77 (204–434)	12 Patients with RAS, 6 without. Correlation between ASL perfusion vs. stenosis grade and ASL vs SPECT perfusion
(Gach et al. 2006)	---	1.5	CASL	Single-shot spiral	---	Free-breathing with real-time motion tracking using a navigator
(Robson et al. 2009)	4 (H)	1.5	PCASL	FSE	WK = 332±140, (231-601) Cortex = 410±150, (287-706) Medulla = 223±76, (165-346) (mean±SD, range)	Found BS, controlled breathing and image sorting helpful to reduce motion artefacts

Reference	N_{subj}	B_0 (T)	Labelling	Readout	RBF (ml/100g/min)	Highlights
(Kiefer et al. 2009)	11 (H)	3	PASL (FAIR)	TrueFISP	Cortex = 245±11; Medulla = 109±5 (mean±SD)	Two-compartment model reduced variance of the ASL data compared to standard single-compartment model
(Gardener & Francis 2010)	9 (H)	1.5	PASL (FAIR)	SE-EPI	Cortex = 341±45; Medulla = 102±32 (mean±SD)	Used parallel imaging to reduce EPI ETL, achieving full kidney coverage (7 slices). BS reduced SNR of PWIs. Image registration improved image quality.
(Lanzman et al. 2010)	20 (P)	1.5	PASL (FAIR)	TrueFISP	Stable >3 weeks post-transp: 305±34, (258-347) Stable <3 weeks post-transp: 297±44, (215-346) Acute decrease in function: (182±53, 113-278); (mean±SD, range); [all cortex]	First study to assess the perfusion of transplanted kidneys with ASL. Used fixed renal allograft T_1 (mean of cortical ROIs in the entire patient group).
(Ritt et al. 2010)	24 (P)	1.5	PASL (FAIR)	TrueFISP	253 ± 20 (Before telmisartan) 268 ± 25 (After 2 weeks of treatment with telmisartan) WK measurements only.	Metabolic syndrome patients. Moderate correlation between ASL and PAH clearance, as well as between perfusion changes assessed with both techniques after treatment with telmisartan.
(Schewzow et al. 2010)	---	3	PASL (FAIR)	TrueFISP	SNR analysis of PWIs only.	Non-rigid registration improved PWI quality. Shown 16 averages as a good compromise between SNR and scan time.
(Song et al. 2010)	8 (H)	3	PASL (FAIR)	TrueFISP	Cortex = 376, 191-378 (mean, range)	Proposed dual navigator approach for motion correction.
(Nathan S. Artz et al. 2011)	24	1.5	PASL (FAIR)	TrueFISP	Cortex: Native = 334±123, (81-456) Transplant = 273±82 (125-409) Medulla: Native = 72±29 (33-121) Transplant = 35±15 (11-70)	14 transplant patients; 10 native kidneys broad range of function. Cortical perfusion reproducibility: Intra-visit CV: 4.8%-6.0% Inter-visit CV: 7.6%-13.1%

Reference	N_{subj}	B_0 (T)	Labelling	Readout	RBF (ml/100g/min)	Highlights
(Wu et al. 2011)	19 (H)	3	PCASL	GE-EPI	Cortex = 227±30; Medulla = 101±21 (mean±SD)	Comparison of RBF estimates with both ASL and DCE-MRI. Moderate correlation found. ICC = 0.41. Low SNR hindered medulla measurements with ASL.
(Cutajar et al. 2012)	20 (H)	1.5	PASL (FAIR)	3D-GRASE	Cortex = 196 and 204 (mean, two scans 7–55 days apart)	First time 3D readout used in renal ASL. Multiple TI perfusion quantification. T_1s determined on a per-subject basis. Shown good short and long term repeatability of RBF.
(Rossi et al. 2012)	8 (H); 9 (P)	1.5	PASL (FAIR)	TrueFISP	Healthy subjects: Cortex = 329±53; Parenchyma = 301±51; CKD patients - GFR 69±12 (Inulin clearance): Cortex = 263±81; Parenchyma = 244±77 (mean±SD)	Significant difference was found between RBF in healthy subjects and patients with mild CKD
(A. G. Schneider et al. 2012)	34 (P) HT	1.5	PASL (FAIR)	TrueFISP	Pre-aliskiren = 272±25ml/100g/min; During-aliskiren = 287±29 ml/100g/min; Post-aliskiren = 272±26 ml/100g/min	Renal perfusion increased during aliskiren therapy. After withdrawal of aliskiren, renal perfusion returned to original levels.
(Wang et al. 2012)	11 (H)	3	PASL (FAIR)	SS-FSE	Pre-furosemide: Cortex = 367±41; Medulla = 119±25 Post-furosemide: Cortex = 314±49; Medulla = 97±18 (mean±SD)	Shown ASL to be sensitive to changes in renal haemodynamics following administration of furosemide. Used non-linear registration for correcting for failed BHs (3 out of 11 subjects)
(Dong et al. 2013)	18 (H); 13 (P)	3	PASL (FAIR)	SS-FSE	Healthy volunteers: Cortex = 357-426, 399±20; Medulla = 107-126, 118±20; AKI patients: Cortex = 223-392, 292±47; Medulla = 89-116, 105±17 (range, mean±SD)	Shown significantly lower renal perfusion values in an acute kidney injury patient group when compared to healthy volunteers.

Reference	N_{subj}	B_0 (T)	Labelling	Readout	RBF (ml/100g/min)	Highlights
(Ott et al. 2013)	19 (P) HT	1.5	PASL (FAIR)	TrueFISP	1-day before RDN = 256.8, 241–278 1-day post-RDN = 263.4, 252–277 3-mon post-RDN = 261.2, 240–285 (mean, interquartile range) (all WK perfusion values)	RBF assessed in patients with treatment-resistant hypertension before and after renal denervation (RDN). Renal perfusion and function did not change after RDN.
(Park et al. 2013)	1 (H)	3	PCASL	TrueFISP	Cortex = ≈320	Part of a study assessing feasibility of PCASL TrueFISP ASL in brain, retina and kidney
(Sugimori et al. 2013)	5 (H)	3	PASL (EPISTAR)	EPI	Cortex = 287±49	Acquired data at multiple TIs, but performed single TI quantification using the TI yielding the highest signal intensity
(Cutajar et al. 2014)	16 (H)	1.5	PASL (FAIR)	3D-GRASE	Cortex = 263±41 (mean±SD)	Shown good agreement between RBF estimates with ASL and DCE. Also found ASL more repeatable than DCE MRI.
(Gillis et al. 2014)	12 (H)	3	PASL (FAIR)	TrueFISP	Cortex = 327±63; WK = 229±41 (mean±SD)	Shown ASL to be repeatable over a period of up to 28 days (ICC=0.85; CV=9.2% for cortical perfusion)
(Heusch et al. 2014)	98 (TrR)	1.5 (n=65); 3 (n=33)	PASL (FAIR)	TrueFISP	Cortex: 283±61 (eGFR>30 ml/min/1.73m^2; n=60); 178±63 (eGFR≤30 ml/min/1.73m^2; n=38)	No significant differences between perfusion estimates obtained at 1.5T or 3T (both groups). RBF correlated with eGFR.
(Tan et al. 2014)	10 (H); 5 (P)	3	PASL (FAIR)	TrueFISP	Cortex: Healthy volunteers = 269±27; CKD patients = 155±25; (mean±SD)	Employed a navigator technique to enable free breathing scans. BH scans with low number of measurements yielded very low SNR data. Good quality perfusion maps in patients with reduced perfusion.

Reference	N_{subj}	B_0 (T)	Labelling	Readout	RBF (ml/100g/min)	Highlights
(Breidthardt et al. 2015)	20 (H); 20 (P)	1.5	PASL (FAIR)	TrueFISP	Cortical perfusion (mean±SD) Young volunteers = 278±59 Age-matched volunteers = 274±65 HF eGFR > 60 ml/min/1.73m^2 = 171±31 HF eGFR < 60 ml/min/1.73m^2= 146±50	RBF correlated to eGFR in patients with chronic heart failure (HF). Correlation between volume-adjusted global kidney perfusion measured by PC-MRI and RBF measured by ASL. HF patients with and without renal dysfunction shown to have comparable RBF.
(Cutajar et al. 2015)	8 Kidney Donors	1.5	PASL (FAIR)	3D-GRASE	Pre-nephrectomy: 186±36 ml/100g/min; 214±32 ml/kidney/min Post-nephrectomy: 184±37 ml/100g/min; 253±45 ml/kidney/min (mean±SD) (values for right kidney only - left kidney nephrectomy)	First study to assess RPF in healthy living kidney donors, pre and post-donation. RPF per kidney increased in almost all donors, but not as much as GFR. Filtration fraction increased in all but one donor post-nephrectomy
(Hueper et al. 2015)	46 TrR	1.5	PASL (FAIR)	TrueFISP	231±15 vs. 331±15 (DGF vs. patients with initial function) 339±19 vs. 243±14 (Living donor vs. deceased donor grafts) 202±35 vs. 246±14 (Allograft rejection vs. DGF but no rejection) (all cortical perfusion values)	Transplant patients with delayed graft function shown to have significantly reduced perfusion. ASL perfusion correlated with eGFR. Higher perfusion in living donor grafts compared to deceased donor grafts.
(Gillis 2016)	24 (H); 17 (P) CKD	3	PASL (FAIR)	TrueFISP	Healthy volunteers: Cortex = 279±69; WK = 221±38 CKD patients: Cortex = 136±37; WK = 146±24 (mean±SD)	ASL perfusion correlated with eGFR.

Reference	N_{subj}	B_0 (T)	Labelling	Readout	RBF (ml/100g/min)	Highlights
(Hammon et al. 2016)	5 (H); 9 (P) HT	1.5	PASL (FAIR)	TrueFISP	WK = 307±26 Cortex = 337±35 Medulla = 280±27 (mean±SD)	Shown good reproducibility study of RBF using ASL for both WK, cortex and medulla measurements
(Ren et al. 2016)	20 (H); 62 (TrR)	3	PASL (FAIR)	TrueFISP	Healthy volunteers: Right kidney = 390±61 Left kidney = 392±64 Transplant recipients: eGFR > 60 ml/min/1.73m^2 (n = 34) ≈290±75 eGFR < 60 ml/min/1.73m^2 (n = 19) ≈200±90	Cortical RBF lower in allografts with good function (eGFR≥60ml/min/1.73m^2) compared to volunteers. Cortical RBF lower in allografts with impaired function (eGFR<60ml/min/1.73m^2) compared to those with good function. RBF correlated with eGFR in TrR.
(Niles et al. 2016)	15 (TrR); 15 (TrD)	1.5	PASL (FAIR)	TrueFISP	Donors: Baseline = 422±17; Decrease 2 years post-transplant = 50±19; Transplanted kidneys: Baseline = 412±19; Decrease 2 years post-transplant = 141±21; (all cortical perfusion values)	ASL performed 3, 12 and 24 months post-transplant in both TrR and TrD. In TrD also pre-transplantation. Both TrD and TrR: perfusion decreased over time, more so on TrR. Smaller decrease in TrR with losartan therapy.
(Robson et al. 2016)	4 (H) 4 (P)	1.5	PCASL	3D-FSE (radial trajectory)	Healthy volunteers: Cortex = 284±21; Patients (renal masses): semi-quantitative assessment only	First 3D near-isotropic resolution renal ASL study. No difference between RBF measured by 3D FSE and 2D SS-FSE. No difference in RBF obtained in free vs. synchronized breathing with 3D FSE. Better renal mass detection with 3D imaging readout.

Reference	N_{subj}	B_0 (T)	Labelling	Readout	RBF (ml/100g/min)	Highlights
(Shimizu et al. 2017)	14 (H)	3	PCASL	SE-EPI	Whole group = 139±38 Young volunteers (n=8, ages 22-39) = 157±38; Older volunteers (n=6, ages 53-75) = 117±24 (cortical perfusion obtained with multi-PLD modelling)	RBF Quantification performed at 3 PLDs. Moderate correlation of RBF and effective RBF using ^{99m}Tc-MAG3 scintigraphy, except for longest PLD (1.5s), when performing single-TI quantification.
(Kim et al. 2017)	23 (H)	3	PCASL	SE-EPI	Scan 1: Cortex = 215±65; Medulla = 81±21; Scan 2: Cortex = 215±63; Medulla = 79±28.	Multi-slice acquisition (4 post-labelling delays). Two scans (1-hour apart) shown good reproducibility of cortical perfusion and low reproducibility of cortical and medullary arterial transit time.
(Conlin et al. 2017)	7 (H), 8 (P)	3	PASL (FAIR)	TrueFISP	Healthy subjects: Cortex = 151±37; Medulla = 25±22 Patients: Cortex = 158±103; Medulla = 36±31 (units: ml/min)	RBF derived from slope of ASL difference signal (averaged within manually drawn ROIs) using up to 16 different TIs. ASL RBF estimates agree with those obtained with DCE-MRI. Reliability of perfusion measurement dependent on having more than two TIs.
(Li et al. 2017)	30 (H); 33 (P)	3	PASL (FAIR)	TrueFISP	Healthy subjects: Cortex = 207±42; Medulla = 47±16 Patients (median eGFR 46.7 ml/min/1.73m^2) Cortex = 106±36; Medulla = 23±9	RBF approximately half in patients compared to controls. Cortical and medullary RBF correlated with eGFR. RBF thresholds were determined to separate controls from patients.

H: Healthy volunteers; P: Patients; RAGE = Rapid Gradient Echo; WK: whole-kidney; TrR: Transplant recipient; TrD: Transplant Donor. CV: Coefficient of variation; ICC: Intraclass correlation. HF: Chronic Heart Failure; PC-MRI: Phase-contrast MRI; DGF: Delayed graft function; HT: Hypertension.

3 Motion Correction in MRI

Motion has been a problem in MRI for as long as the technique has existed. It has deleterious effects on image quality and often renders scans non-diagnostic. As a result, certain MR sequences need to be repeated during an examination, which in addition to adding to the patient's discomfort, incurs additional financial costs. A recent study estimates yearly costs on the order of $100,000 per scanner due to the added scan-time required to repeat MR scans (Andre et al. 2015).

A major focus in the MR research community is in establishing MR scanners as fully fledged scientific measuring instruments capable of delivering *quantitative* measurements related to Nuclear Magnetic Resonance (NMR) properties of tissue (such as T_1, T_2 and proton density). This represents a key advance compared to the more traditional MR acquisitions which instead provide qualitative or "weighted" measurements related to the fundamental NMR properties. Ultimately, with accurate quantitative measurements, biological tissue (both healthy and in diseased states) could potentially be characterised based on these parameters.

Many of these modern MRI approaches are based on the acquisition of large amounts of data, routinely involving hundreds of images and different types of contrasts. The full potential of these rich datasets can only be harnessed if techniques are in place which ensure that all images are not only free of motion artefacts but also aligned in the same space allowing them to be combined. This is the case for many of the MRI methods able to provide physiologically relevant parameter maps, such as blood flow maps as obtained with ASL. Ensuring a voxel-to-voxel correspondence throughout a time series is a challenge when such acquisitions take several minutes, and even shorter periods of time when the organs of interest are under the influence of involuntary motion, such as that caused by breathing. Furthermore, certain groups of patients are particularly prone to moving during the scans, due to a variety of medical conditions (such as Parkinson's disease), or reduced tolerance to the MRI environment (e.g. children).

Research on motion correction started in the early MRI days (Schultz et al. 1984; Wood & Henkelman 1984; Haacke & Patrick 1986), and over the years, a plethora of techniques have been developed to prevent, reduce and correct the negative influence of motion on image quality (Zaitsev et al. 2015). Nevertheless, there is no universally optimal approach for motion correction. The wealth of MR techniques existing and the variety of applications in the clinical arena often mean that tailored motion correction solutions are necessary for a given task. In fact, researchers frequently combine multiple approaches from the large pool of techniques that have been developed, often acting at different stages of the imaging workflow. These range

from effective scanning standard operating procedures to prevent motion artefacts in the first place (effective patient training and communication and immobilization devices), avoiding and reducing the negative influence of motion on image quality using specialized motion insensitive image acquisition techniques and correcting artefacts through advanced retrospective image processing methods. A significant challenge is to determine the best approach to perform motion correction in any particular application, as noted in a recent review by Zaitsev *et. al* (Zaitsev et al. 2015).

This section aims to introduce the physical basis of motion-related artefacts in MRI and to summarise different approaches that have been developed over the years and are most related to the work in this thesis.

3.1   Effects of motion in MRI

Before delving into the basic principles which describe the effects of motion in MRI, it is pertinent to consider the time scale of physical movement by the subject and how it relates to the different steps in image acquisition. Firstly, motion artefacts are typically avoided along the frequency-encoding direction simply because this stage of the spatial localization process is fast enough (usually <10ms) that any movement by the subject is effectively "frozen" in time. This is not the case for phase-encoding, and therefore this is the direction along which motion artefacts are most common. Nevertheless, in many applications translational motion still occurs on a long time-scale when compared to the time required for acquisition of a complete 2D or 3D k-space. This is particularly true in single-shot acquisitions, where all data is acquired following a single MR excitation. In these cases, the total length of time during which data is being acquired is necessarily small due to relaxation effects, which may allow intra-shot motion to be assumed negligible.

Multi-shot sequences are inherently motion-sensitive owing to the extended time gap between the different excitations required for the complete data acquisition. In this scenario, both involuntary physiological motion (e.g. respiration) and bulk body motion are no longer slower than the timescale of data acquisition and can have a dramatic effect on the quality of a given image. Furthermore, even if multi-shot motion artefacts are avoided or single-shot sequences are used, movement can still play a detrimental role in the consistency of the complete MR dataset, by causing the misalignment of successive images in a time series. This can be addressed by image registration methods (see Section 3.5).

As discussed before (see Section 2.2), MR data is acquired in the frequency domain (k-space) and thereafter reconstructed into images using the FT. Intrinsic to this process is the assumption

that the sampled object remains static during the entirety of the data acquisition process. When this assumption is violated, the reconstructed image will contain artefacts. From the Fourier relationship it follows that each k-space sample is a complex number which specifies the contribution of a given spatial frequency to the entirety of the object in the field of view (FOV). As such, a single point in k-space influences the image intensity of every pixel in the image. Consequently, movement at any point in the scan has the potential to corrupt the entire image.

The simplest types of motion include translations and rotations. Their effects in the frequency domain are well-known and described by the Fourier Shift and Rotation theorems, respectively (Bracewell 2000). Consider a 2D object $f(x, y)$. A translated version of this object can be given by $f(x - x_0, y - y_0)$. Assuming $F(k_x, k_y)$ corresponds to the FT of $f(x, y)$, the Fourier Shift theorem states:

$$f(x - x_0, y - y_0) \overset{FT}{\Leftrightarrow} F(k_x, k_y)e^{-i2\pi(k_x x_0 + k_y y_0)} \qquad \text{Equation 3-1}$$

Therefore, a translation in image space results in phase changes in the frequency domain. In addition, the Fourier rotation theorem states that:

$$\begin{aligned} f(x \cos\theta - y \sin\theta, x \sin\theta + y \cos\theta) \overset{FT}{\Leftrightarrow} F(k_x \cos\theta \\ - k_y \sin\theta, k_x \sin\theta + k_y \cos\theta) \end{aligned} \qquad \text{Equation 3-2}$$

As such, a rotation of the object in the (x, y)-plane causes a rotation of k-space in the (k_x, k_y)-plane by the same amount and direction (Bracewell 2000). These Fourier relationships are widely used in algorithms that retrospectively correct movement-related k-space inconsistencies (see Figure 3-1). If there is a means to determine the amount of translational movement by an object during the acquisition, multiplying the raw data by the complex conjugate of the added phase term due to this motion corrects for the corresponding image artefact. Rotational movement may in addition result in missing k-space sections and interpolation algorithms are required to satisfy the requirement of a regularly sampled grid imposed by Fast Fourier Transform (FFT)-based reconstructions (Atkinson & Hill 2003).

Unlike translational and rotational motion, there is not a direct relation between non-rigid motion in image space and its effect in the frequency domain. This limits the types of techniques that can be used to tackle this problem. The physics of the acquisition process and artefact causes can be elegantly described using matrix notation (Batchelor et al. 2005). This can be solved iteratively with a conjugate gradient scheme to recover the motion-free image, as long as the underlying motion is known. White *et. al.* provided this prior information in the form of a

patient-specific respiratory motion model, obtained by non-rigid registration of a low resolution training dataset (White et al. 2009). External sensors can also be used for this purpose (Odille et al. 2008).
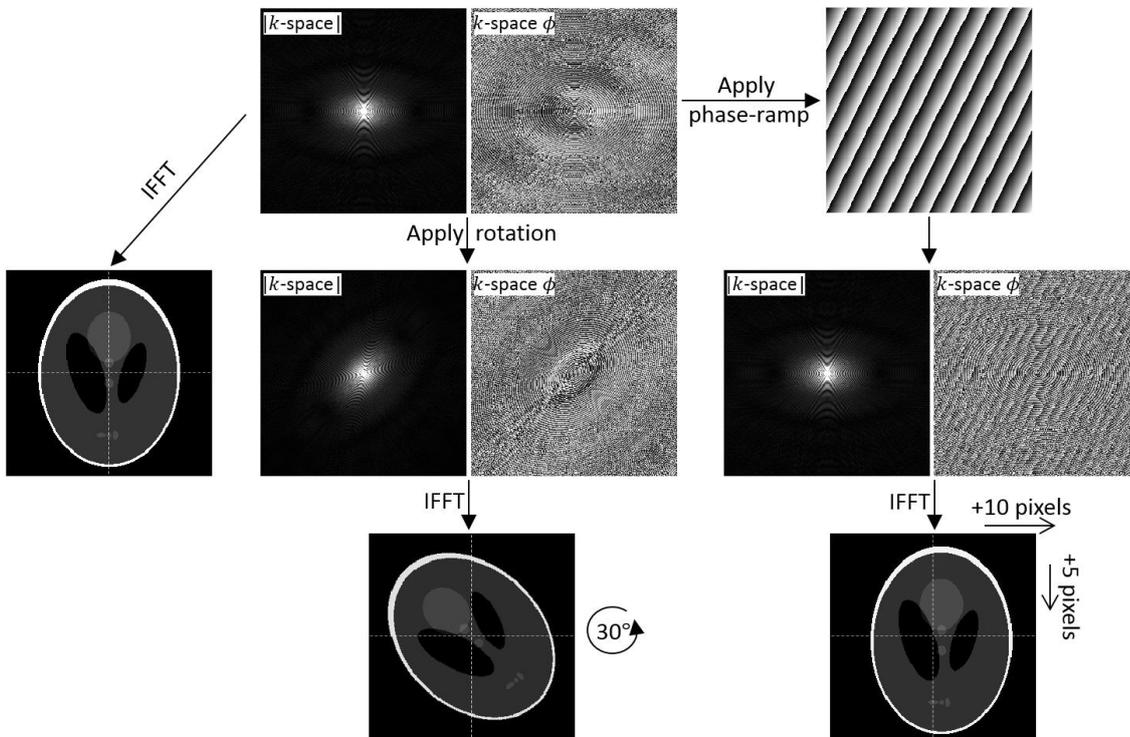
Figure 3-1 k-space operations corresponding to rotations and translations in image space. These are some of the building blocks of k-space-based motion correction techniques, such as autofocus or navigator-based approaches.

A substantial amount of work on prospective motion correction using external tracking systems has been performed (e.g. (Maclaren et al. 2012), (Ooi et al. 2013)) since the initial proof-of-concept paper by Zaitsev *et. al* (Zaitsev et al. 2006). However, at the moment these solutions are targeted exclusively to neuroimaging (Maclaren et al. 2013). Interestingly, recent research addresses the possibility of extracting physiological signals using a markerless external tracking system (Maclaren et al. 2015). Prospective and retrospective motion correction methods have also been used in combination (MacLaren et al. 2011).

3.2   Autofocus

Autofocus (Atkinson et al. 1997; Atkinson et al. 1999) is a retrospective motion-correction approach whereby trial corrections are iteratively performed on the raw data and then evaluated according to a particular cost function (e.g. an image quality metric), typically after reconstruction (see Figure 3-2). Ideally, an artefact-free image yields the minimum value for the cost function. Various image quality metrics have been reviewed previously (McGee et al. 2000). Examples of frequently used metrics include the image entropy (Atkinson et al. 1997) or gradient

entropy (Cheng et al. 2012). Overall their minimisation favours images with reduced blurring and distinct edges. Autofocus has been applied in areas outside the brain, such as the shoulder (Manduca et al. 2000), wrist (Lin et al. 2007) and abdomen (Cheng et al. 2012). These methods are attractive as they can be readily applied without additional hardware or modifications to the pulse sequence, but the optimisation stage can be very time consuming. Nevertheless, recent improvements in this area tackle this problem by reducing the search space of possible motion estimates through the use navigators (Cheng et al. 2012) or performing a more efficient gradient-based optimisation (Loktyushin et al. 2013). The problem of non-rigid motion correction has also been addressed recently by Cheng *et. al* by approximating non-rigid motion as localized linear translations (Cheng et al. 2012; Cheng et al. 2015). This framework was also employed for motion correction in angiography (Ingle et al. 2013). Loktyushin also expanded on their original work (Loktyushin et al. 2013) to allow for non-rigid correction (Loktyushin et al. 2015).
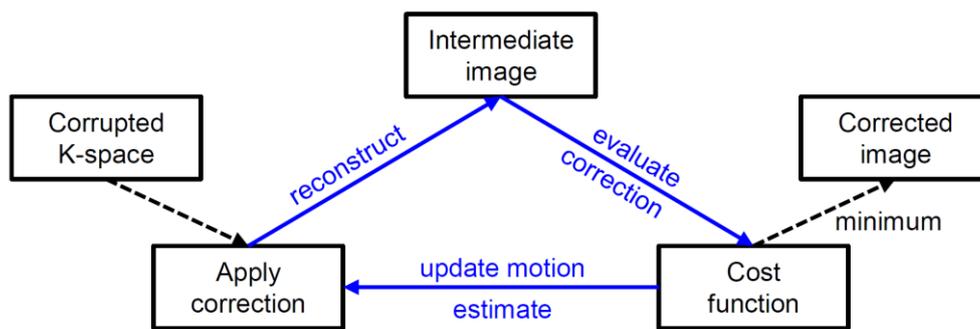
Figure 3-2 Autofocus loop

3.3 Motion correction using data redundancy

Modern MRI systems acquire data in parallel using independent receiver coil elements. This permitted the advent of parallel imaging in the second half of the 1990s (Sodickson & Manning 1997; Pruessmann et al. 1999), constituting an important technical advancement in MRI technology. Initially devised to speed up MR image acquisition, parallel imaging can also be useful to improve image quality (e.g. decreasing blurring in single-shot scans) and reduce artefacts (e.g. due to motion). An example shown by Larkman *et. al* (Larkman et al. 2004) consists of splitting a corrupted dataset into two in a way that makes each of them amenable to a parallel imaging reconstruction. After reconstruction, the two can immediately be averaged, which may mitigate image artefacts, or an image registration algorithm may be used in the first place to achieve a greater level of data consistency before the averaging procedure. Parallel imaging can be used to detect and reject corrupted parts of k-space, and if the corruption is due to a short-lived event, the generalised simultaneous acquisition of spatial harmonics (SMASH)

approach (Mark Bydder et al. 2002) can be used to regenerate the incomplete k-space (M Bydder et al. 2002). Consider a fully sampled dataset where the acquired phase-encoded line n is used to generate line $n + 1$. Because the data is fully sampled, two "versions" of line $n + 1$ are available, the sampled data and the data generated using parallel imaging. If a comparison of both lines reveals differences consistent with a certain model of motion, the data can be corrected. This forms the basis of the SMASH navigators approach by Bydder *et. al* (Bydder et al. 2003). The fact that data is corrected instead of rejected is beneficial in terms of the signal-to-noise ratio (SNR) of the final image. A drawback is that errors may end up propagating through k-space as the different lines are processed. In the Data convolution and combination operation (COCOA) method (Huang et al. 2010), a synthetic k-space is produced and then compared to the acquired data for the detection of corrupted data. An optimised combination of the two then yields an image with reduced artefacts. A different approach proposed by Atkinson *et. al* (Atkinson et al. 2004) uses multi-coil data in an optimisation scheme whereby the artefact source is parameterised and the parameters iteratively adjusted to maximise consistency of the data from different coils. Also in the image space, Winkelmann *et. al* (Winkelmann et al. 2005) perform a voxelwise consistency check to locate artefacts which are then removed with an extended Sensitivity Encoding (SENSE) reconstruction algorithm. More recently, Hu *et. al* (Hu et al. 2011) used multi-coil data to develop a self-gating approach for free-breathing cine MRI.

3.4   Navigators

Navigators are fast pulse sequence elements incorporated into a "parent" sequence. They can provide motion-related information (bulk motion, respiratory/cardiac phases) and/or hardware-related information (e.g. B_0 changes (Hess et al. 2012)). This information is used either prospectively or retrospectively, respectively to guide the image acquisition or reconstruction to obtain a higher quality image. A main assumption is that negligible motion occurs between the navigator and the acquired data to be corrected. These can be divided into three categories: i) k-space, ii) image and iii) free induction decay (FID)-based navigators. Ehman *et. al* introduced the concept of navigator echo with the 1D linear navigator (Ehman & Felmlee 1989). This is simply a non phase-encoded line that crosses the centre of k-space and allows monitoring of translations in the direction of the corresponding readout gradient. To resolve in-plane rotational motion, Fu *et. al* proposed orbital navigators (Fu et al. 1995). Welch *et. al* extended this concept to three dimensions, introducing the so-called spherical navigators which are able to capture translations and rotations in 3D (Welch et al. 2002). These improvements require more data to be acquired which increases the time required for not only for the navigator

acquisition but also its processing to extract the motion information. Nevertheless, recent studies addressed this issue by introducing a template-based approach to speed up the processing of the spherical navigator data (Liu & Drangova 2011; Johnson et al. 2016). Alternatively, some methods focus on reducing the time necessary for acquisition of the navigator data, such as the cloverleaf navigators (Van Der Kouwe et al. 2006) which limit the amount of data acquired to 3 orthogonal lines crossing the centre of k-space (to track translations) and 3 quarter-circles (to track rotations), taking less than 5ms to be acquired. A different approach is to perform an image-based estimation of the motion parameters during the scan. Thesen *et. al* (Thesen et al. 2000) addressed the problem of motion correction in fMRI (functional-MRI) by registering successive multi-slice echo planar imaging (EPI) datasets in a time-series to a reference volume in real-time and using the obtained motion parameters to update the imaging coordinates for the following acquisitions. Other image-based navigators include PROMO (White et al. 2010), which has been used in brain ASL (Zun et al. 2014), and volumetric navigators (Tisdall et al. 2012), which have been shown capable of reducing the detrimental effects of motion in brain morphometry analyses (Tisdall et al. 2016). Again, as the amount of data in the navigator increases, more time is required for its acquisition which may limit their application especially in clinical sequences, which are optimized to have little dead time. Fat navigators (Nguyen et al. 2003) have been combined with highly accelerated parallel imaging for motion correction in neuroimaging (Skare et al. 2015; Gallichan et al. 2016; Gallichan & Marques 2017). Finally, FID-navigators are simplest type of navigator, sampled for a very brief period while no gradients are being applied. Initially proposed by Brau et. al. (Brau & Brittain 2006) to gate abdominal acquisitions, these have more recently been used to detect motion and trigger data re-acquisition in the brain (Kober et al. 2011; Kober et al. 2012) and for MRI of the carotid arteries (Dyverfeldt et al. 2014). Another possibility, given their speed is to sample them frequently throughout the scan (e.g. at least every TR) to detect motion and if necessary trigger the acquisition of a more complex navigator as previously discussed (Waszak et al. 2016). Extracting motion information from FID navigators is a current topic of research (Babayeva et al. 2015).

3.5   Image Registration

Image registration refers to a class of image processing algorithms whose aim is to align structures within images. This is achieved by determining a spatial one-to-one mapping between the pixels in the images. In the medical imaging context, the successful application of image registration is able to produce datasets which correspond to what would have ideally been acquired had the patient not moved. These have been increasingly used to compensate for

patient motion during image acquisition, which degrades the quality of the data. Image registration algorithms have been found useful for a variety of applications (Sotiras et al. 2012; Oliveira & Tavares 2014), including registration of similar types of images (intra- or mono-modality registration) and images obtained with different techniques (inter- or multi-modality registration). Registration of data from different subjects is also common, as is registration to certain types of atlases, allowing group analyses and comparisons across studies to be performed. These and many other applications have solidified image registration as an ubiquitous pre-processing step in the MRI literature, available as part of widely used software packages (Ashburner 2012; Jenkinson et al. 2012; Klein et al. 2010).

This section aims to summarise the main components within an image registration algorithm. These are depicted in Figure 3-3 and include: a type of *transformation* applied to the input image to match a reference image, an image *similarity measure*, which quantifies the degree of alignment between two images and an *optimiser*, involved in the estimation of the transformation which maximises the similarity between the images. A particular focus is given to the discussion of image similarity measures, in particular to those based on information theory, as some of the involved concepts (e.g. image entropy) are not only useful for image registration but also in other motion correction strategies relying on image quality measures. An example of these is the autofocus method (Atkinson et al. 1999; Atkinson et al. 1997) which was previously discussed and developed within the framework of this thesis (see Appendix A).

Comprehensive reviews on this topic can be found in (Maintz & Viergever 1998; Hill et al. 2001; Pluim et al. 2003; Zitová & Flusser 2003).
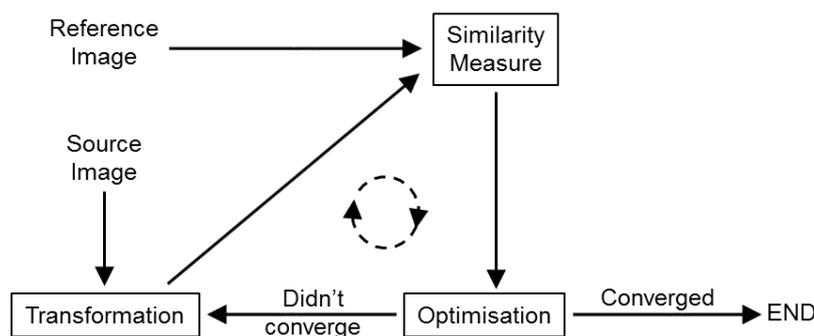
Figure 3-3 Typical image registration workflow (simplified)

3.5.1 Transformation

Image registration involves the computation of a transformation. This can be thought as the spatial mapping T which maps the position and corresponding intensity from the pixels in one image to another, making them spatially aligned. The latter remains stationary, being commonly

referred to as reference, or target image. The image which will be transformed to match the reference image is typically referred to as moving or source image.

The transformation is characterized by its degrees of freedom. Factors influencing the choice of transformation include the dimensionality of the data (2D images or 3D volumes) and whether it is expected that the object of interest will undergo deformations or distortions during imaging. Another consideration is the resolution of the data which influences the required accuracy of the registration process.

In general, we can group transformations in two main families: those based in *parametric* models and *non-parametric* approaches. The latter type of transformations employ a larger number of parameters characterising a deformation vector for each pixel location*.* A classic example is the Demons algorithm (Thirion 1998; Pennec et al. 1999). All registrations performed within the context of this thesis belong to the group of parametric approaches. Therefore, non-parametric transformations will not be discussed any further.

The majority of medical image registration tasks involve registration of image volumes (i.e. 3D images). A standard approach, in particular when dealing with data from a single subject, is to apply transformations in which there are six degrees of freedom, allowing for three translations and three rotations. This is known as a "rigid body" transformation, which is characterized by preserving all distances between any pair of points. Additional degrees of freedom may allow scaling and skews in addition to translations and rotations. When these are incorporated, we are dealing with affine transforms, which preserve parallel lines. In 3D, they are represented by 4x4 matrices (using homogenous coordinates). It is worth noting that rigid body transforms are a special case of affine transforms (unit scaling and zero skews). These have a relatively narrow range of applications where rigid body transforms do not suffice. This is the case since when tissue deformations exist, these are usually complex in nature and not typically limited to stretching and shearing. Nevertheless, affine transformations find use generally in initializing non-rigid transforms, as it is helpful from an optimisation perspective to ensure a good initial alignment for complex registrations. Additionally, these have been used to correct MRI hardware-related issues, such as image distortions (Hill et al. 2001).

Non-rigid deformations are particularly useful for inter-subject registrations. An illustrative example depicting the application of different types of transformations in image registration is shown in Figure 3-4. Non-rigid transforms allow to model much more complex deformations. Unlike the previous types of transforms, these have a higher number of degrees of freedom, sometimes on the order of the number of pixels in the images. A common approach for the

parameterisation of the deformation fields is using cubic b-splines (Rueckert et al. 1999; Klein et al. 2010). A comprehensive review of deformable registration methods can be found in (Sotiras et al. 2012).
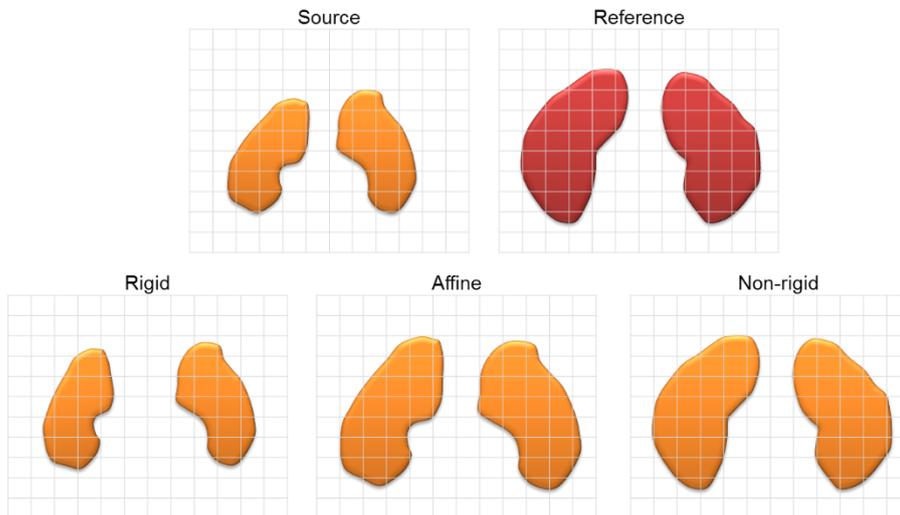
Figure 3-4 Different types of image registration transformations. In this illustrative example, the transformations were applied independently for each of the objects, an approach which was used throughout this thesis, as the two kidneys can move independently.

3.5.2   Similarity Metrics

Image registration algorithms rely on specific image characteristics for determining the transformation that best aligns them. Examples of these include points, lines or curves defining edges, corners or other types of geometric features. Alternatively, the image intensity levels can be used directly. We focus on registration algorithms which are based on image intensities. As such, the transformation we recover aims to optimize some similarity measure which takes into account the intensities of the two images. Specifically, the role of image similarity metrics is to quantify how well the target image matches the reference image after the target image has been transformed. This typically involves an iterative process where, in each iteration, the similarity of the reference image and the moving image (being brought into alignment with the reference image) are compared. Several types of image similarity metrics have been developed over time, and their suitability to a given application is intimately dependent on the types of images involved.

3.5.2.1   Intra-modality registration

When registering images of the same modality, where any changes between them can be attributed to motion, and to a smaller extent, noise, we are referring to *intra-modality* registration. Therefore, if two images are aligned, their subtraction yields an image with reduced

signal throughout the field of view. This is the rationale behind one of the most common intra-modality similarity metrics, the mean of squared differences (MSD), which is defined as:

$$MSD = \frac{1}{N}\sum_{i=1}^{n}(I_R(i) - I_M(i))^2 \qquad \text{Equation 3-3}$$

Where $i = 1, \dots, n$ are the indexes of the pixels within the field of view of both images and $I_R(i)$ and $I_M(i)$ are the image intensities of the reference and moving images, respectively, at the pixel location i. Also widely used are the less restrictive correlation-based techniques, such as the correlation coefficient (CC):

$$CC = \frac{\sum_{i=1}^{n}(I_R(i) - \overline{I_R})(I_M(i) - \overline{I_M})}{\{\sum_{i=1}^{n}(I_R(i) - \overline{I_R})^2 \sum_{i=1}^{n}(I_M(i) - \overline{I_M})^2\}^{1/2}} \qquad \text{Equation 3-4}$$

Where $\overline{I_R}$ and $\overline{I_M}$ are the mean intensities of the reference and moving images, respectively. The CC allows registration of objects whose intensities are related by linear transformations.

However, these metrics face challenges when where there is not necessarily a direct relationship between the intensities of the source and reference images. This is the case when registering images of different image modalities, or in the case of MRI, different pulse sequences and/or MR parameters. An important class of metrics, which are better suited for this purpose have their underpinnings in information theory.

3.5.2.2 Inter-modality metrics

The concept of entropy is frequently borrowed from information theory for use in image processing as a statistical measure of the amount of information or uncertainty in a given process. The most widely used definition of entropy was introduced by Shannon in his 1948 seminal paper (Shannon 1948) and defined as:

$$H = -\sum_{i=1}^{n} p_i \log_2 p_i \qquad \text{Equation 3-5}$$

for a process whose events $e_1, \dots, e_n$ occur with probabilities $p_1, \dots, p_n$.

Entropy can be interpreted in various ways. It is maximized when the probabilities of all events occurring are equal. This corresponds to the state of highest uncertainty regarding the outcome of the given process. On the other hand, when the probability of a single event occurring dominates over all others, we are increasingly certain of the outcome of the process, and the value of entropy decreases. Therefore, entropy can be thought of as a measure of uncertainty.

A practical example is shown in Figure 3-5, where a number of processes, each characterized by 5 events which occur with different probabilities, were simulated. The entropies of each of the processes were computed. The outcome of process 1, whose events have the most similar relative probabilities of occurring, is the most difficult to predict. The entropy of this process is therefore the highest of the set.

	Process 1	Process 2	Process 3
p_1	0.1012	0.0514	0.0809
p_2	0.2541	0.5652	0.0473
p_3	0.2545	0.0996	0.0143
p_4	0.2924	0.1859	0.8013
p_5	0.0977	0.0978	0.0562
H	2.1858	1.7963	1.0790
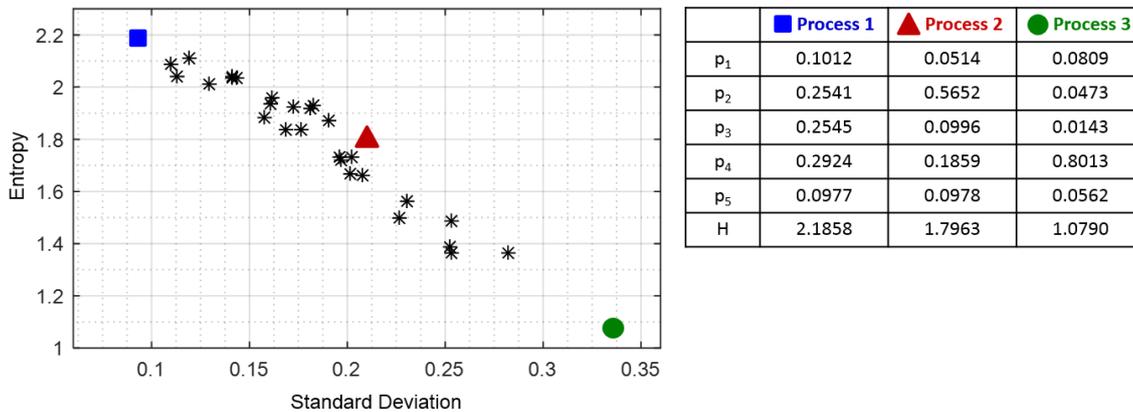
Figure 3-5 Entropy of general processes. Each data point in the plot corresponds to a simulated process, each with 5 possible events which occur at different probabilities. Three processes are highlighted (Processes 1, 2 and 3), and the probabilities of their events are shown in the table on the right-hand side. The higher the uncertainty of the outcome of each process, the higher the entropy (Process 1). In this context, entropy can be thought as a measure of dispersion of the probabilities of the events. Although not the same as standard deviation, entropy *tends* to decrease as the spread in the probabilities increases.

Another common interpretation of entropy is as a measure of information. The closer we get to a process with a single event dominating in terms of its probability of occurrence (Process 3 in Figure 3-5), the lower the amount of information. This can be understood as in the limit where a single event is 100% likely, entropy will be zero ($-1 \times \log_2 1 = 0$). Indeed, the amount of information in this process is zero as we already knew its outcome *a priori* with certainty before it took place.

In the context of image processing, we focus on the probability distribution of the grey levels within an image when computing its entropy. The entropy is maximized for an image whose grey levels occur at equal probabilities (Image A in Figure 3-6). An image with a smaller "amount of information" (i.e., a larger proportion of voxels converging at a similar intensity) will have a smaller entropy (Image B in Figure 3-6). Image entropy therefore decreases as the histogram of an image gets "sharper". This is used as the basis of an image quality metric in the autofocus work. Similarly to as discussed above, in the limit where all the pixels of an image have the same intensity, entropy is zero.

Image A

3	5	9
8	4	1
2	6	7

$H_A = -(1/9) \times \log_2(1/9) \dots$
$\quad -(1/9) \times \log_2(1/9) \dots$
$\quad \dots$
$\quad -(1/9) \times \log_2(1/9)$

$H_A = 3.2$

Image B

3	3	3
3	9	3
3	3	9

$H_B = -(7/9) \times \log_2(7/9) \dots$
$\quad -(2/9) \times \log_2(2/9)$

$H_B = 0.8$
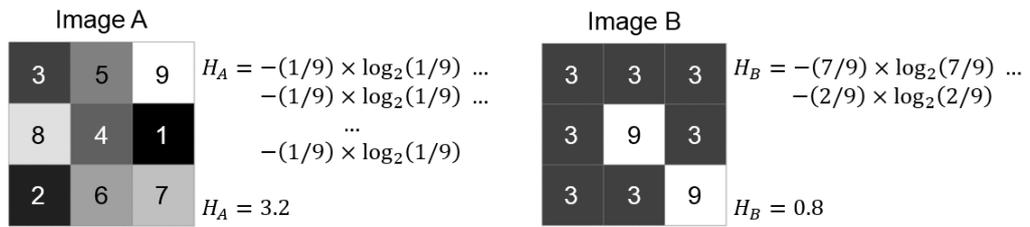
Figure 3-6 Simplified example of image entropy. In image A, all intensities occur with the same probability, which maximises the information content of the image, and thereby the value of entropy. When a larger proportion of pixels converge around a similar intensity level (Image B), the entropy value is minimised.

Shannon also introduced the concept of joint entropy (Shannon 1948). In the context of image registration, joint entropy provides a measure of the amount of information in the combined reference and moving images. This can be conveniently represented by a joint histogram. Each axis in the joint histogram contains the range of possible intensities for each image. Each entry is obtained by counting the number of times a given combination of grey values occurs in corresponding locations of both images (Figure 3-7).
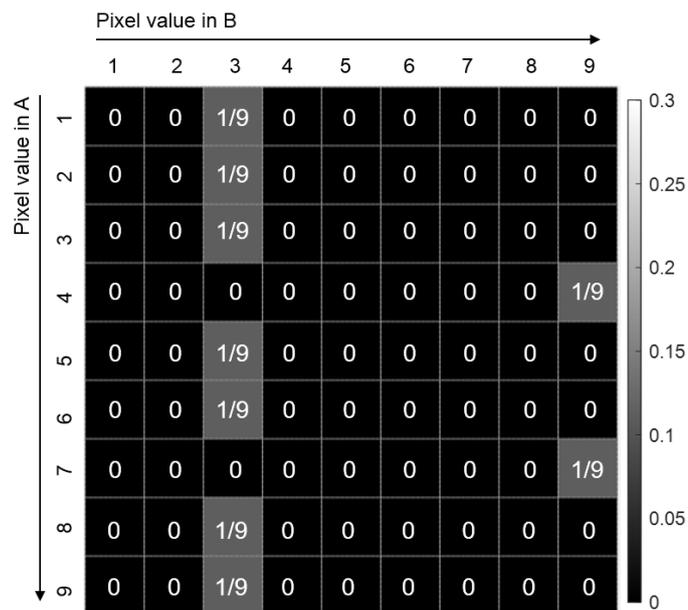
Figure 3-7 Joint histogram corresponding to images A and B in Figure 3-6. In this instance, the value of the joint entropy equals the value of the entropy of image A as all combinations of intensities in the 3x3 images occur only once.

This is particularly useful when attempting to align two images, and the basic rationale for this is depicted in Figure 3-8. Taking a renal MRI image as an example, I generated an additional set of images by applying translations of different amounts in the head-foot direction to the original image. The degree of misalignment between the images increases with the absolute value of the translation applied. This is easily visible in the checkerboard images combining one pair of

original and translated images. As the degree of misalignment increases, the locations of corresponding anatomical structures in both images no longer match. This translates to joint histograms with increased blurring, as the likelihood of finding the same combinations of image intensities in corresponding locations of the two images decreases. In a way analogous to the entropy of a single image, the joint entropy increases as the intensities in the joint histogram spread from the well-defined clusters which exist when the images are aligned. Even though the particular example in Figure 3-8 deals with intra-modality registration, it is important to realise that these clusters will exist regardless of the relationship between the intensities of corresponding structures in both images. This is the property that makes joint entropy-based information theoretic measures suitable for inter-modality registration.

The joint histogram can be normalized by dividing it by the total number of voxels. The result is a joint probability density function $p(i,j)$ of images A and B, which is used for the calculation of the joint entropy (H_{AB}):

$$H_{AB} = -\sum_{i,j} p(i,j) \log_2 p(i,j)$$

The joint entropy was first proposed as an image similarity metric for image registration in 1995 by both Collignon (André Collignon et al. 1995) and Studholme (Studholme et al. 1995).

A known issue with using joint entropy for this purpose is that optimizing it (i.e. minimizing the joint entropy of two images) may in fact result in a failed registration. This is the case because minimization of joint entropy occurs when the all entries in the joint histogram cluster around a single value. Depending on the field of view and structure being imaged, a large proportion of the voxels in both images may correspond to air. A joint-entropy based optimization may therefore result in a maximisation of the correspondence of air in the two images at the expense of ensuring the alignment of other structures, as this would yield the sharpest joint histogram.

This is known as the overlap problem and other metrics have been proposed which perform better and are somewhat less sensitive to this problem. One of these metrics is mutual information, another information-theory concept proposed as a similarity metric for image registration in (A Collignon et al. 1995) and (Viola & Wells 1997). It can be defined as:

$$I_{AB} = H_A + H_B - H_{AB}$$

Where H_A and H_B are the marginal entropies of images A and B, computed within their overlapping volume. Their values will be low when background intensities make up a large proportion of this volume.
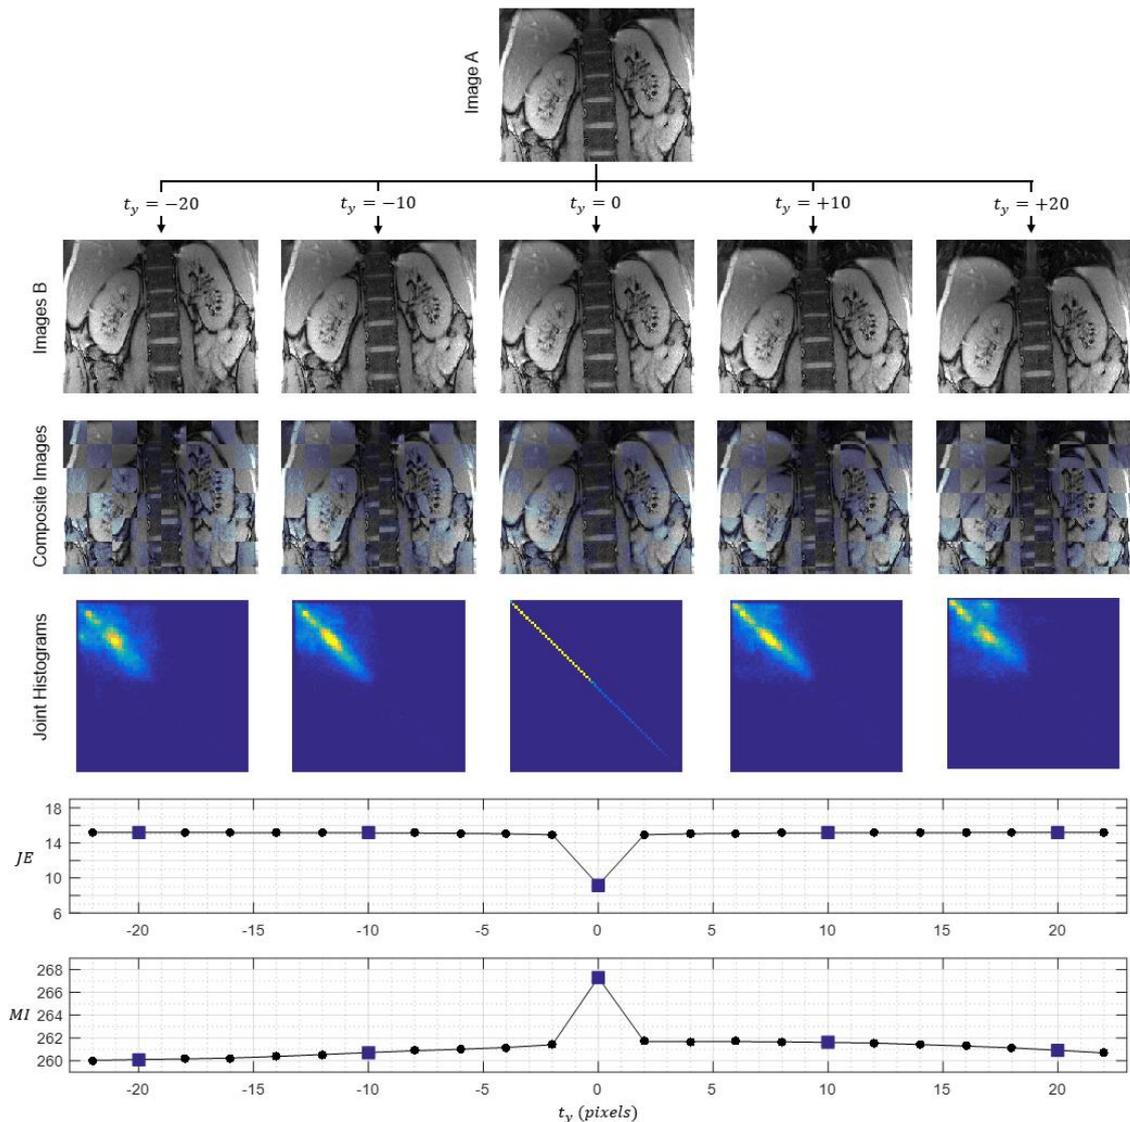
Figure 3-8 Influence of increased image misalignment on image theoretic measures (joint histogram, joint entropy and mutual information). Image A, a coronal oblique anatomical image of the kidneys, was translated in the in plane *y* direction from -22 to 22 pixels with increments of 1 pixel. This generated a new set of images (Image B) which were overlaid to image A. Joint Entropy and Mutual information between the two was computed (plots). For the translations corresponding to multiple values of ten, a checkboard image combining image A and B (neighbouring squared tiles displaying voxels from different images), and the joint histogram of the two images is displayed. Note that for the case of no translation (i.e. identical images) all grey level counts lie on the diagonal of the joint histogram. Although these images are of the same subject / modality, which is a case which could be handled by one of the simpler similarity metrics, this example suffices to illustrate the inner workings of information theoretic based similarity metrics.

A maximisation of mutual information therefore will correspond to a scenario characterized simultaneously by:

- *High marginal entropies:* both images will contain a relatively high amount of
  information, avoiding images dominated by background.

- *Low joint entropy:* large amount of shared information between the two images, implying they are aligned.

This explains why mutual information *per se* is less sensitive to the overlap problem when compared to the joint entropy alone. Further robustness to changes in image overlap was demonstrated by Studholme et al. (Hawkes et al. 1999) who proposed a normalization scheme for the mutual information, which outperformed the standard mutual information measure in inter-modality MR-CT and PET-MR registrations:

$$NMI_{AB} = \frac{H_A + H_B}{H_{AB}}$$

A comprehensive review of mutual information based registration applications can be found in (Pluim et al. 2003).

3.5.3 Optimisation

Optimisation is effectively a stepwise process that seeks to find the "optimal" transformation which brings the structures of interest into alignment. In practice, a given transformation is applied to the moving image and the corresponding output is compared to the reference image using an image similarity metric (see section above). This result informs the registration algorithm with regards to the transformation to apply in the next iteration, if convergence has yet to be reached. A key component of this step is the optimisation strategy. An evaluation of the complete parameter space by performing an exhaustive search is rarely performed as a multitude of efficient numerical optimisation techniques can be employed to tackle this issue. This is of crucial importance as the computational burden of a registration task increases rapidly with the complexity of the transformations (i.e. their number of degrees of freedom). Examples of such techniques include: Gradient Descent, Levenberg-Marquardt, Newton-type methods (Gauss-Newton and Quasi-Newton), Robbins-Monro, Downhill Simplex method, Powell's method, stochastic and evolutionary algorithms, among others. A detailed description of many of these algorithms can be found in (Press et al. 2002), in (Maes et al. 1999), and in (Klein et al. 2007). In the latter, Klein et al. (Klein et al. 2007) evaluated a subset of the existing optimisation strategies within the framework of a cubic *B*-spline based deformation model and a mutual information similarity measure. Importantly, they've shown that it is possible to reduce the computational time considerably without compromising the accuracy of the registration results by using fewer voxels to calculate the cost function and its derivatives, especially when using stochastic gradient-based algorithms, which do not require the exact derivatives to be computed. An example of algorithms amenable to this type of approach is the Robbins-Monro

(Robbins & Monro 1951). In fact, Klein et al. (Klein et al. 2007) found this approach to perform the best among the methods assessed in their work.

In 2009, Klein et al. further proposed an automatic approach for an adaptive determination of step size along the search direction, to be used in conjunction with the Robbins-Monro optimisation approach (Klein et al. 2009). This was the optimization technique used for all registration tasks in this thesis, as implemented in the open-source registration toolbox *elastix* (Klein et al. 2010).

4 Investigation and minimisation of the sources of image artefact in 3D-GRASE ASL of the kidney

Part of this work was presented at an European MRI conference (Nery et al. 2015a) and a National MRI conference (Nery et al. 2015b).

4.1 Introduction

As discussed in Section 2.3.6, renal ASL using a segmented, background-suppressed 3D-GRASE acquisition scheme with respiratory triggering has been shown to be reproducible in healthy volunteers (Cutajar et al. 2012), allowing for multi-slice studies with a range of TIs. However, when scanning less compliant subjects, motion represents a significant hurdle that needs to be overcome before ASL can be routinely applied in the clinic.

This chapter describes an investigation into the sources of image artefacts due to motion in multi-shot 3D-GRASE ASL of the kidneys. The primary aim was a thorough overview of the image acquisition and reconstruction scheme to identify the major sources of artefacts due to movement during the image acquisition. Even though motion-related artefacts that corrupt the ASL blood flow maps may arise at different steps of the acquisition and processing pipeline, here the focus was on the quality of the individual ASL perfusion-weighted images (PWIs). Further issues related to alignment of different sets of image data required for Renal Blood Flow (RBF) calculation (e.g. M_0 and T_1) were addressed elsewhere in the thesis (see Chapter 6).

We scanned healthy volunteers who were instructed to move and change their respiration patterns during the scan to examine the major sources of motion and their effects on image quality. Different schemes based on the 3D-GRASE pulse sequence were compared. A retrospective analysis of the effectiveness of respiratory triggering during irregular breathing was performed by synchronizing the respiratory waveform recordings to the data readout instants, allowing identification of instances where respiratory triggering was unable to ensure data acquisition at a consistent point in the respiratory cycle.

The 3D-GRASE pulse sequence was chosen as the basis for renal ASL experiments within this thesis as it constitutes an attractive image readout module for ASL. Reasons for this include:

- Intrinsically higher signal-to-noise ratio (SNR) of 3D imaging, which is particularly important for a low SNR technique such as ASL. This is also advantageous in paediatric imaging due to the lower amount of signal available (comparatively smaller structures to image);

- Fast whole-kidney coverage;

- Consistent background suppression (BS) throughout all slices: all slices are excited at the same time so BS is equally effective in all of them, unlike in 2D multi-slice methods;

- Exactly the same inflow time for all slices, for the same reasons as in the previous point;

- Reduced SAR when compared to FSE, due to the lower frequency of spin-echo refocussing, particularly important for high field strength imaging;

- Lower sensitivity to field inhomogeneity/susceptibility effects than gradient echo echo planar imaging (EPI) due to spin rephasing at the k-space centre in each partition.

4.2 Motion-sensitivity in segmented 3D-GRASE ASL

As discussed in Section 2.2.4.5, acquisition schemes based on the 3D-GRASE pulse sequence either consist of a single-excitation (single-shot imaging) or require multiple excitations for completely filling k-space (multi-shot imaging). In both cases, *intra-shot motion* (i.e. motion occurring during the imaging readout) is assumed to be negligible given the sampling time for each shot, which is necessarily short (typically less than 300ms) to avoid severe degradation of the image quality due to signal decay during the echo train. On the other hand, in the case of multi-shot imaging, the time between the acquisitions of the different segments corresponding to a given 3D k-space is much longer. Within these larger time intervals, *inter-shot motion* is bound to occur when scanning less compliant subjects, especially in the paediatric population. When an object (e.g. a kidney) changes position between the acquisitions of different segments, phase inconsistencies will be present among the phase-encoding lines of the fully sampled k-space, as discussed in Section 3.1. This will generate artefacts along the phase-encoding direction in the reconstructed images. Unless the amount and type of motion was precisely the same during acquisition of the ASL control and labelled images, the level of artefacts will differ among them. Consequently, the quality of the PWI (resulting from a subtraction of the two) is severely affected (see Figure 4-1).

One obvious strategy to reduce inconsistencies in the position of the kidneys related to breathing effects is simply to acquire data during subject breath-hold. Nevertheless, this is of limited applicability in paediatric MRI given the difficulty that children have in consistently holding their breath for a reasonable amount of time (Chavhan et al. 2013). To allow for free-breathing scans with some degree of robustness to breathing motion, respiratory bellows are typically used in abdominal MRI. These record a waveform corresponding to diaphragmatic motion and as such offer a proxy of the kidney movement due to respiration.

a) Slice 3/12; TI = 1700ms

b) Slice 6/12; TI = 2700ms
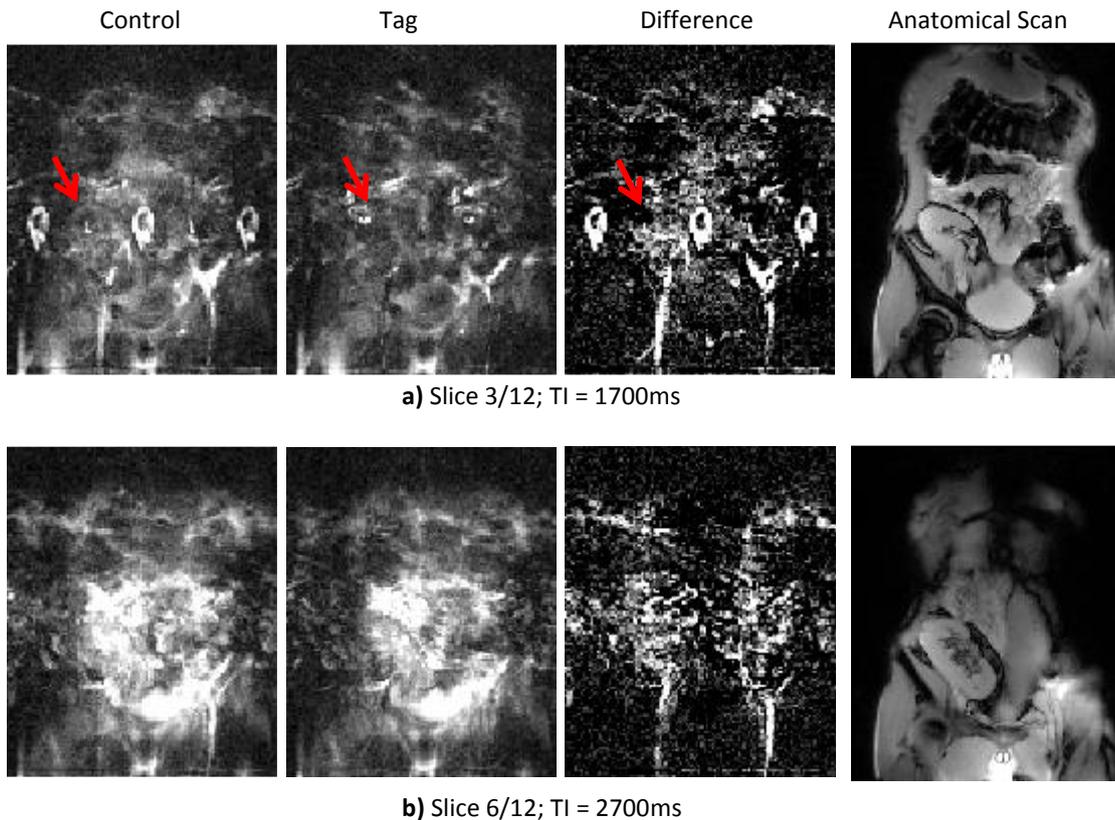
Figure 4-1 Severely corrupted images from a multi-shot 3D-GRASE ASL dataset (transplant patient). Motion between the acquisitions of the different segments of k-space results in a combination of blurring and ghosting after image reconstruction. The red arrows point to the kidney, almost invisible in these slices due to a combination of BS and motion artefacts.

Throughout all the experiments in this thesis, respiratory bellows were used to prospectively trigger the non-selective and slab-selective inversion FAIR ASL pulses at the end-expiration stage of the respiratory cycle. This ensures all inversions to be consistently applied at the same point in the respiratory cycle. Furthermore, this also increases the probability of different shots at the same TI to be acquired in the same part of respiratory cycle, significantly decreasing breathing-related motion artefacts. The drawback with using triggering is the increase of the acquisition time, since the true acquisition TR varies according to the duration of the respiration cycle.

An additional source of error exists whenever multiple image volumes are acquired and combined during the processing stage, either to enhance SNR (such as the case of averaging several measurements at the same TI) or for multi-TI perfusion quantification using compartmental models. Despite the individual images being artefact-free, the position of the target tissue may vary throughout the time series (see Figure 4-2). In fact, this is always a potential issue in ASL since two image volumes are required to obtain a perfusion weighted-

image. Different techniques can be used to address this problem, such as image registration (see Section 3.5) and implemented in Chapter 6 or the use of retrospective sorting algorithms, as reviewed in Section (see Section 2.3.6) and implemented in Chapter 5.
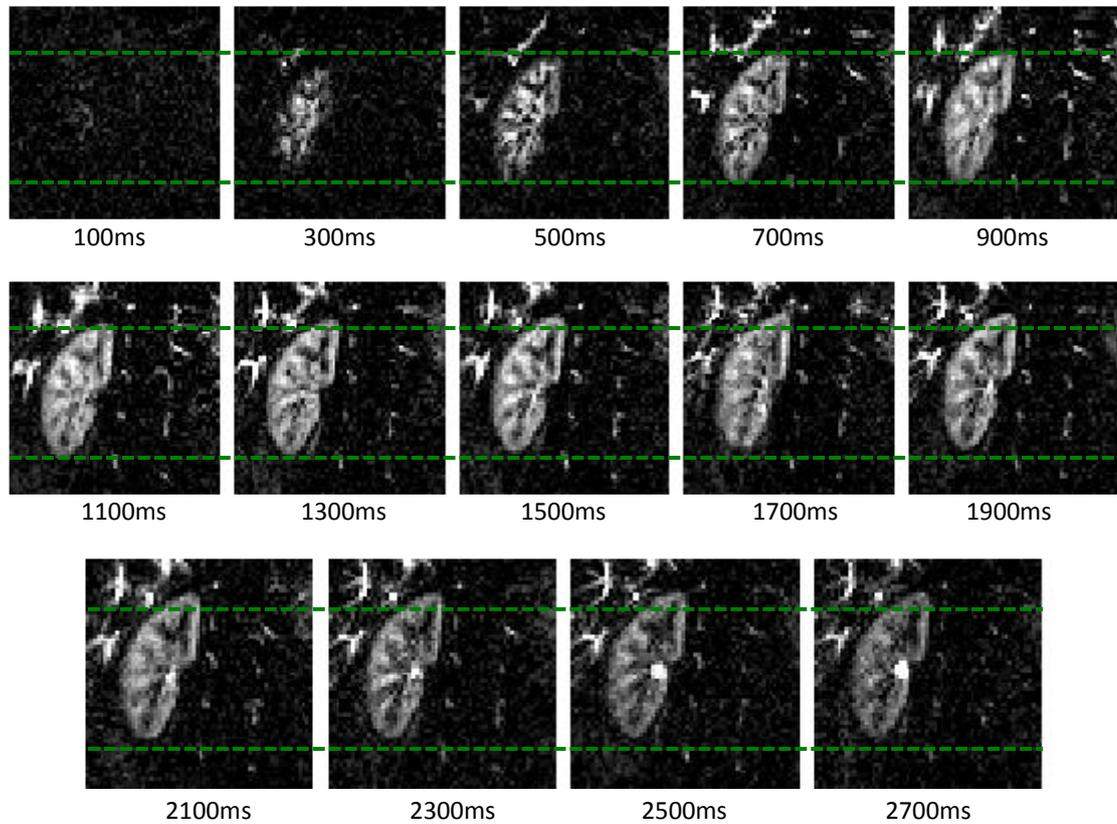
Figure 4-2 Renal inter-scan translational movement. The kidneys are largely at consistent positions during the acquisition of each shot in a given TI, yielding mostly artefact-free images. However, the different TI lengths result in acquisitions at varying parts of the respiratory cycle. The values below each image correspond to the TI used for the corresponding acquisition.

4.3 Methods

Renal ASL data, in the form of reconstructed DICOM images, was available at the start of this project which had been previously acquired with a multi-shot 3D-GRASE sequence. A careful visual inspection of these data revealed a non-trivial amount of motion artefacts. These included blurring and ghosting (appearance of replicas of the object or its structures along the PE direction). The nature of these artefacts makes it challenging to confidently discern their cause, whether it is involuntary movement due to breathing, bulk body motion, flow or peristalsis. Image-based registration algorithms are not suitable for correcting these types of artefacts, since the data corruption occurs at the k-space level. Consequently, appropriate data correction algorithms must also act at this level. As such, access to the raw k-space data is paramount. Furthermore, even though respiratory bellows had been used, physiological information (i.e. the

respiratory trace) for these data was not recorded. These three main reasons motivated the acquisition of a new renal ASL datasets, in healthy volunteers, where both respiratory waveforms and raw k-space data were recorded. Acquiring data in compliant healthy volunteers is advantageous as it is expected that motion will be kept to a minimum. Furthermore, tasks may be assigned to attempt to replicate some of the common artefacts and thus understand their origin.

Recording raw k-space data was made a mandatory step in the standard operating procedures for an (at the time) ongoing study where renal ASL data was acquired in transplant patients also using a multi-shot 3D-GRASE sequence. As such, some of these data is also shown in this chapter, even though I was involved only in scanning, but not recruiting these patients.

4.3.1   Optimisation of the acquisition loop structure

Image reconstruction software was developed to serve as a basis for future k-space based motion correction algorithms. For each scan, the raw data is essentially a file which strings together all outputs from the analog-to-digital converter (ADC) during data acquisition (in the order they were acquired) together with specific headers for each k-space line containing information useful for the reconstruction algorithm. The order in which data is acquired (as implemented in the pulse sequence software), on its own, can have a dramatic influence in the motion sensitivity of a multi-shot multi-TI ASL sequence. From an implementation point-of-view, acquiring all k-space volumes is an iterative process including four main loops:

1. **Segments (or shots) loop**: iterates to completely fill a 3D k-space in segmented acquisitions;
2. **ASL loop**: alternates from the control to tag condition;
3. **Inflow time (TI) loop**: iterates to acquire all TIs specified in the scanning protocol;
4. **Averaging loop**: iterates to acquire all measurements specified in the scanning protocol (for averaging purposes);

The averaging loop will no longer be discussed for two reasons: first, no averaging was performed in all experiments in this chapter. Secondly, this was the outermost loop in the sequence implementation and thus did not directly contribute to an increased sensitivity to motion artefacts. Through an analysis of the raw data files, it was found that the loop structure used in the acquisition scheme was sub-optimal from a motion sensitivity perspective. Firstly, the innermost loop in the image acquisition is the ASL loop. Previous work supports this choice to reduce ASL subtraction errors (Berr et al. 1999; Alsop et al. 2015). After one iteration of this loop, the first k-space segment at the first TI is acquired for both control and tag conditions. In

the following steps, instead of acquiring the remaining shots at the first inflow time, all first segments corresponding to all TIs were acquired. Only then the remaining k-space segments were acquired. For a hypothetical 2-segment 2-TI acquisition, the ordering of the scans is thus:

C.S1.TI1---T.S1.TI1---C.S1.TI2---T.S1.TI2---C.S2.TI1---T.S2.TI1---C.S2.TI2---T.S2.TI2

Where C/T correspond to Control and Tag scans, S1/S2 refer to segments 1 and 2, and TI1/TI2 refer to the two specified TIs. This results in a long interval of time during which the patient can move and thus introduce k-space inconsistencies which manifest as motion artefacts. This time depends largely on the length of the perfusion time series (i.e. number of TIs in a given protocol). This loop structure (from now on referred to as loop structure #1, LS#1) is depicted in Figure 4-3, as well as an example from an actual acquisition which shows that in this case (6 TIs), separate segments corresponding to the same k-space were acquired approximately 1 minute apart.

We modified the sequence loop structure (now referred to as LS#2) by modifying the sequence source code so that the outermost loop is the "TIs" loop (see Figure 4-4). Following the previous example, the new ordering is:

C.S1.TI1---T.S1.TI1---C.S2.TI1---T.S2.TI1---C.S1.TI2---T.S1.TI2---C.S2.TI2---T.S2.TI2

The time between the acquisitions of the different segments corresponding to the same k-space is significantly reduced (from approximately 1 minute to less than 10 seconds in the example in Figure 4-3 and Figure 4-4). In addition, a major advantage of LS#2 is that any given bulk body motion event no longer potentially affects all TIs in the acquisition but instead motion artefacts are now "localised" to the TI where the motion took place.
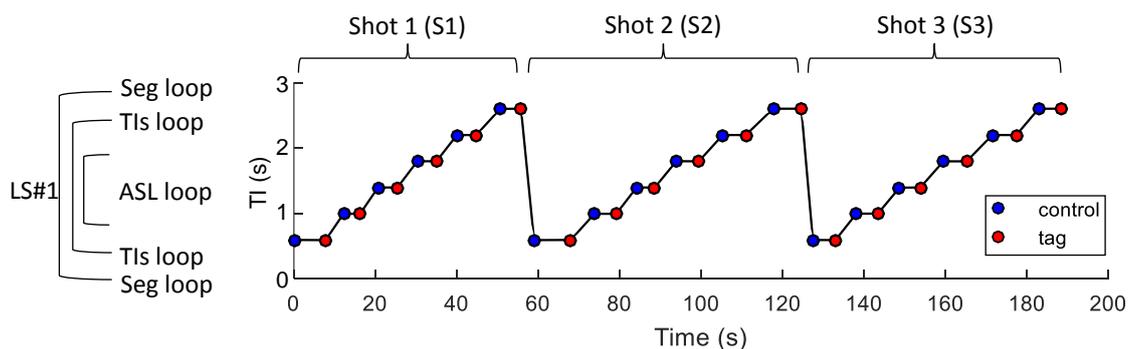
Figure 4-3 Original loop structure (LS#1) used in multi-shot multi-TI 3D-GRASE ASL acquisitions.
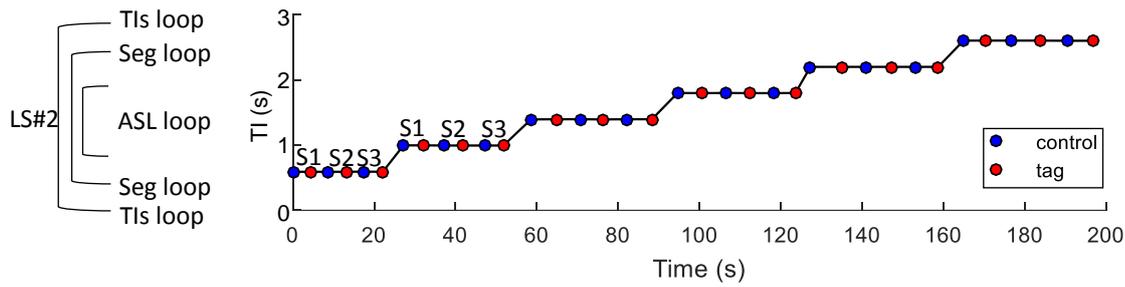
Figure 4-4 Modified loop structure (LS#2) for multi-shot multi-TI 3D-GRASE ASL acquisitions.

4.3.2 Retrospective investigation of motion during acquisition

As mentioned in the start of this chapter, potentially useful motion information is available in the respiratory bellows waveform. This is used during the data acquisition but not recorded by default for a retrospective analysis outside the scanner. This is the case for the software version which interfaces with the MR scanner used in this work (Siemens NUMARIS/4 VB17A). I enabled the automatic recording of the respiratory trace amplitude, together with the triggering instants, by modifying the sequence source code. Combining this with the timestamps of the acquisition of the different k-space segments (which can be parsed from the raw data headers), one can create plots as shown in Figure 4-5. This allows one to precisely identify where in the respiratory cycle the image readouts occur, and as we will show later in this chapter, provides a handle into understanding common motion artefacts in this type of sequence. This particular example shows that the end-expiration "static" portion of the respiratory cycle varies throughout the scan. In addition, due to the different delays between the trigger and data acquisition, longer TIs can be acquired at different parts of the respiratory cycle when compared to the shorter TIs.
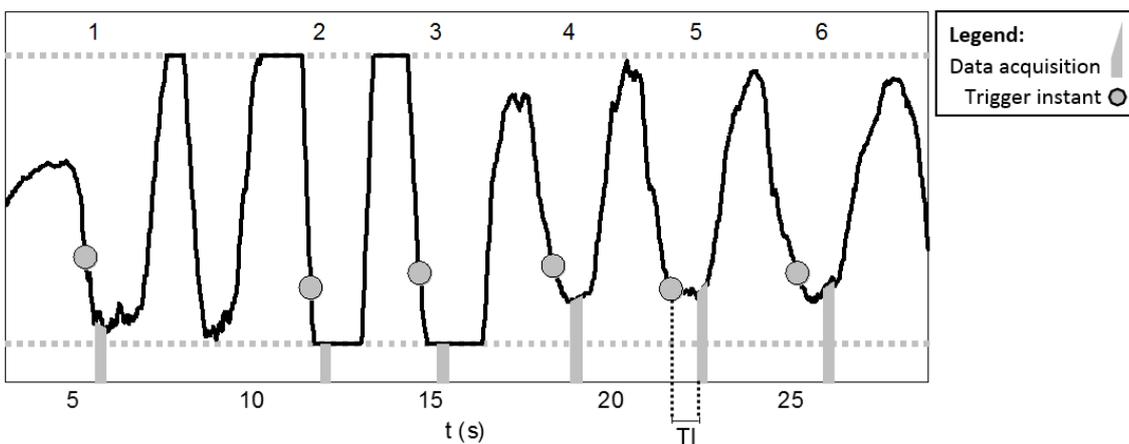
Figure 4-5 Example of a respiratory waveform recorded with bellows, with highlighted instants of triggering (gray dots) and data acquisition (gray columns). Note that the columns highlighting data acquisition are in their actual time scale (readout durations of 237ms). The numbers in the top of the plot indicate the shot index in the acquisition.

94

Another potential source of motion information is embedded in the 3D-GRASE pulse sequence. Non-phase encoded lines are acquired immediately before the first refocusing pulse (see Figure 4-6). These lines are used to correct phase errors related to the data sampling under alternating polarity readout gradients (EPI modules within 3D-GRASE pulse sequences). A Fourier transform of any of these reference lines yields a projection of the object along the orthogonal direction of the readout gradient. In our acquisition, the readout gradient is along the Head→ Foot direction. Therefore, our hypothesis is that these data intrinsically contain motion-related information along this direction, which is the direction along which breathing-related kidney movement is predominant.
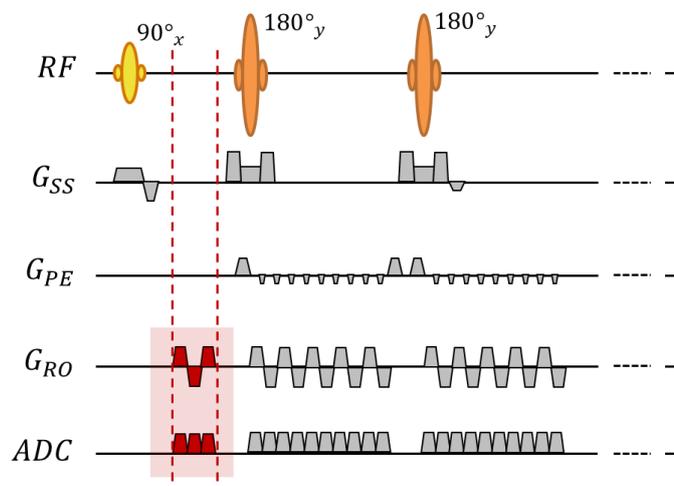
Figure 4-6 Non-phase encoded scans embedded in the 3D-GRASE pulse sequence for EPI Nyquist ghost correction. A Fourier Transform of these yields a projection of the object along the orthogonal direction of the readout gradient (Head→Foot in our MR protocol)

Finally, as a last tool to investigate the source of motion artefacts, we performed zero-filled reconstructions of the individual k-space segments. The rationale behind this is that even though aliasing artefacts will be present in the reconstructed images, because the low number of segments (one third of the full k-space data acquired per excitation), the approximate position of the kidneys during acquisition of each segment will be possible to discern. An example of this approach is given in Figure 4-7, which depicts a simplified motion simulation, as well as undersampled reconstruction of k-space segments. The vertical translations applied to each of the segments are clearly visible in these reconstructions.

| | fully sampled recon | shot1 zero-filled recon | shot2 zero-filled recon | shot3 zero-filled recon |

no motion

intershot motion

shot2 shift = 30 pixels   shot3 shift = -15 pixels
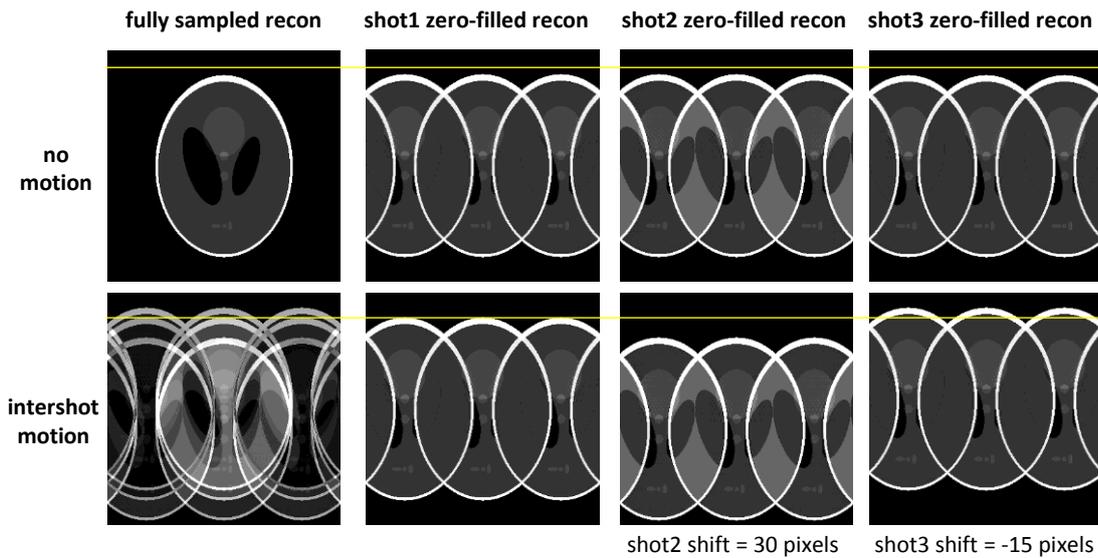
Figure 4-7 Simulation of the effects of multi-shot motion and zero-filled reconstruction of each segment. The second and third segments where shifted in the up-down direction. This results in ghosting along the phase-encoding direction in the fully-sampled reconstruction.

4.3.3 MR acquisition and Subjects

MR data was acquired in a Siemens Avanto 1.5T scanner (Siemens Healthcare, Erlangen, Germany) running the Siemens NUMARIS/4 VB17A software version. Both the spine coil embedded into the scanner table and a dedicated abdominal array coil were used. The latter is composed of two clusters, each with 3 coil elements in the left-right direction, whereas the spine coil has eight clusters also with 3 elements each in the left-right direction (of the eight clusters, three or four are active, depending on the FOV). A multi-TI FAIR 3D-GRASE pulse sequence (see Figure 2-11) was used, with segmentation (number of shots = 3) in the phase-encoding direction (left-right). Partial Fourier (factor = 0.75) was used in the second phase-encoding (or partition-encoding) direction (anterior-posterior). Slice-selective pre-labelling WET saturation pulses (Ogg et al. 1994) as well as a single post-labelling sinc saturation pulse were used. Each coronal-oblique volume is comprised of 12 partitions with 5mm thickness. The field-of-view was 400mm (Head → Foot) × 325mm (Left → Right). TR and TE were, respectively, 3000 and 26.8ms. In-plane resolution was 3.1mm × 3.1mm. The bandwidth and echo-spacing were, respectively, 2170Hz/Px and 0.6ms. Fourteen control-tag ASL pairs were acquired at TI values ranging from 100ms to 2700ms with 200ms increments. No data averaging was performed. Respiratory bellows were used to capture a waveform corresponding to the breathing cycle with a sampling frequency of 50Hz, allowing triggering of the ASL inversion pulses at end-expiration. Background suppression is implemented as described in (Günther et al. 2005) (see Figure 2-14), resulting in a consistent suppression across a range of TIs (see Figure 4-8). Spectral fat saturation was used. All 3D-GRASE

images were reconstructed off-line using programs written by myself in MATLAB (Matlab R2013a (The MathWorks, Inc.)).
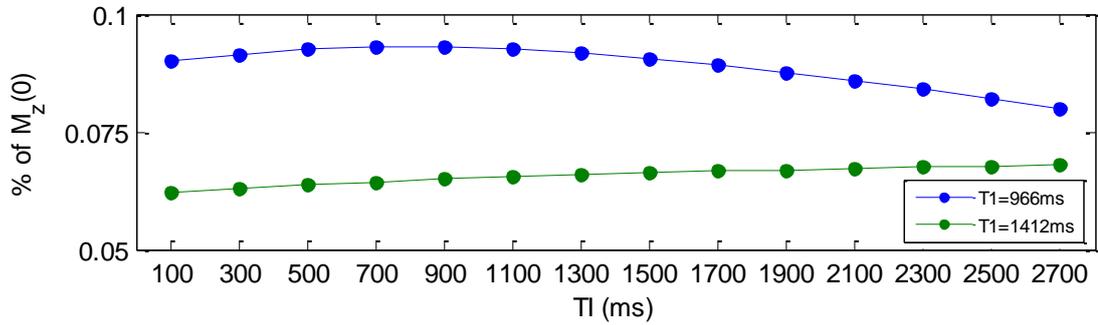
Figure 4-8 Longitudinal magnetisation available at the time of the RF excitation pulse over a range of TIs used in our measurements. Simulation performed for T_1 relaxation times of 966ms and 1412ms, corresponding respectively to the renal cortex and medulla (de Bazelaire et al. 2004), at 1.5T.

To plan the different experiments to be performed within a single MR scanning session (around 60 minutes), a preliminary analysis of the time penalty due to the use of respiratory triggering was performed. The dataset consisted of 58 3D-GRASE ASL scans from 15 subjects (age (years) = 31±18 (mean±std), range=14-71). These were data acquired outside of the scope of this project with a similar protocol as described above. The nominal (i.e. without respiratory triggering) scan time ($TA_{nominal}$) of a 3D-GRASE ASL acquisition is given by:

$$TA_{nominal} = 2 \times TR \times N_{shots} \times N_{TIs} \times N_{measurements}$$

Where the 2 factor accounts for the acquisition of two distinct sets of images (control and tag). For a multi-shot ($N_{shots} = 3$) acquisition with 14 TIs, TR of 3 seconds and no averaging, $TA_{nominal}$ is 252 seconds. The true acquisition time (including triggering) ($TA_{trigger}$) of each of the 58 scans was determined by calculating the time elapsed from the acquisition of the first to the last k-space lines using the timestamps of the raw MR data. The ratio $TA_{trigger}/TA_{nominal}$ was calculated for each scan. A boxplot summarising the $TA_{trigger}/TA_{nominal}$ results is shown in Figure 4-9.
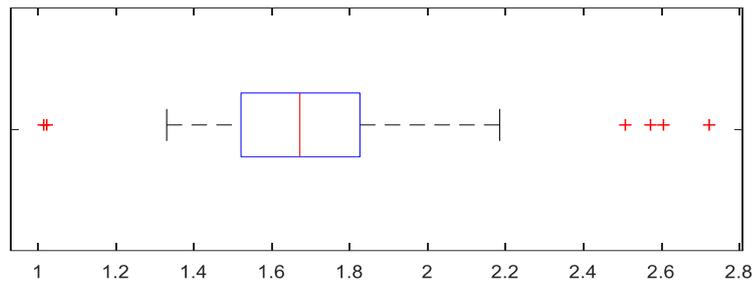
Figure 4-9 Triggering time penalty boxplot in 58 3D-GRASE ASL scans, calculated as the ratio $TA_{trigger}/TA_{nominal}$. The central line corresponds to the median, the edges of the box to the 25% and 75% percentiles (q1 and q3, respectively). Points larger than q3+1.5(q3 − q1) or smaller than q1−1.5(q3 − q1) are considered outliers.

The testing protocol is summarised in Table 4-1 (localiser scans are not included).

Table 4-1 Volunteer 3D-GRASE ASL scans

Scan#	Task	Loop Structure	N_{slices}	N_{shots}	N_{TIs}
1	Still	LS#1	12	3	12
2	Still	LS#2	12	3	12
3	Still	*	10	1	12
4	Still	*	10	1	12
5	Bulk Motion	LS#1	12	3	6
6	Bulk Motion	LS#2	12	3	6
7	Resp. Motion	LS#1	12	3	6
8	Resp. Motion	LS#2	12	3	6

* Loop structure LS#1 and LS#2 are identical for the case of single-shot imaging. The difference between scans #3 and #4 was the voxel size (see Table 4-2 for details)

The main goals for these acquisitions were:

- Compare images from both loop structures in still (Scans #1 and #2) and motion (#5/#6 and #7/#8) conditions;

- Assess the image quality of single-shot scans (Scans #3 and #4);

- Assess to what extent motion artefacts due to breathing correlate with the information from the respiratory trace together with the data readout intervals.

- Investigate if the EPI phase correction projections can provide useful "navigator" information.

In scans #5 and #6, the volunteers were instructed to move into a new position two times during each scan, at the instants indicated in Figure 4-10. They were asked to attempt to limit bulk body movement to these instants, and not to return to their original position after the intentional

movement events. These corresponded to the moment in LS#1 after the 1st and 2nd segments had just been acquired for all TIs in the experiment. Visual instructions were displayed in a screen inside the scanner room, visible through a mirror inside the scanner bore (see Figure 4-11). All the volunteers moved at the intended time and were asked to perform movements of similar amplitude in both scans. For scans #7 and #8, they were asked to speed up their respiratory rate for 5+ respiratory cycles in varying parts of the scan. The goal was to force data acquisition during different parts of the respiratory cycle to assess breathing motion effects, in an attempt to replicate what is likely to occur when subjects with very irregular respiratory rates are scanned.

Four healthy adult volunteers, with previous experience in MRI (age (years) = 33±5 (mean±std), range=25-37) were scanned using this acquisition protocol. In one of the volunteers, Scans #6 and #7 could not be performed due to time constraints.
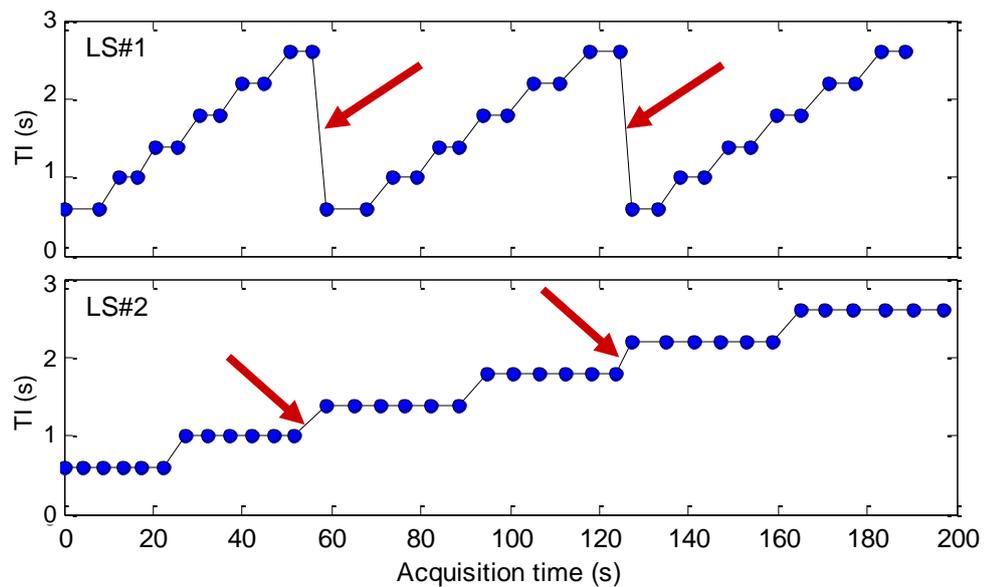
Figure 4-10 Instants when the volunteers were instructed to move during the scan (arrows).
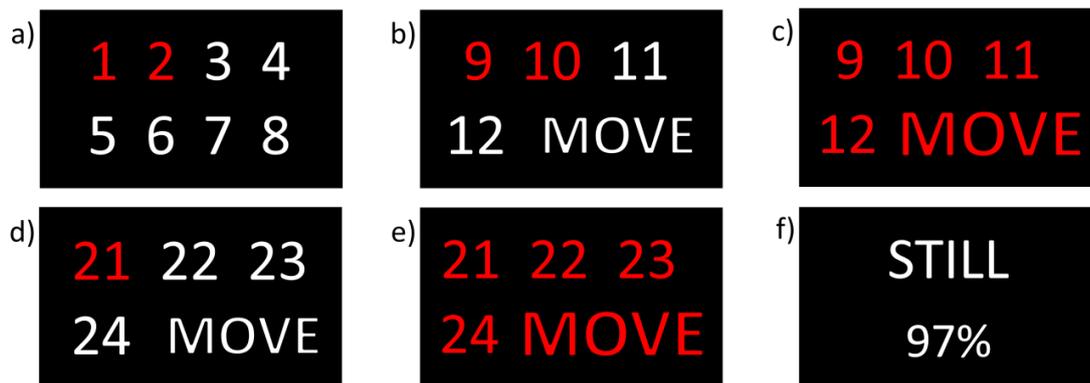
Figure 4-11 Visual commands displayed inside the scanner to inform the patients about the voluntary movement instants. Each of the figures a)-f) corresponds to what the patient could side inside the bore. The patients were instructed to move after shots 12 and 24), in the instants indicated by the arrows in Figure 4-10.

4.4   Results and Discussion

In the presence of motion, the loop structure order greatly impacts image quality. A comparison of the images acquired with both loop structures is shown in Figure 4-12. As would be expected, given the instants within the acquisition where the subjects moved, all data at six TIs acquired with LS#1 shows severe motion artefacts despite the fact that only two motion events took place during the entire acquisition. This is no longer the case with the data acquired with LS#2. In theory, assuming that i) all movement within the scan was limited to the instants when the subjects were instructed to move and ii) that the subjects would be capable to remain in a static position after moving, all images acquired with LS#2 should be free of motion artefacts, and only show the kidneys in different positions. This is not the case, and a few reasons may explain this. On one hand, a consistent observation by the volunteers is the difficulty in remaining in a static position after voluntarily moving into a new position following the scanning protocol. Therefore, even though the largest movements within the scan were those when the subjects were instructed to move, it is likely that these were not the only motion events during the whole acquisition. Furthermore, it was challenging to instruct the volunteers to move by a predetermined amount. Therefore, the amplitude of movement between the volunteers varied widely. A visual inspection of the data suggests the largest amount of movement was seen in volunteer two (see Figure 4-12, images at 1300ms and 1700ms acquired with LS#2 in this volunteer). Such amplitudes of movement will have an effect on the ASL labelling efficiency, which can also explain the reduced perfusion signal in the PWIs. In any case, the benefits of using LS#2 are evident, in particular if we consider the images at the lower TIs, which for the case of LS#2 were completely acquired before any motion event, when the volunteers were able to remain still given that they were in their default, more comfortable position. On the other hand

even at these lower TIs, movement artefacts are already evident in the data acquired with LS#1, as the later segments are acquired after the patient has moved into a different position.

a) Control images (Volunteer 1)

b) Difference images (Volunteer 1)

c) Control images (Volunteer 2)

d) Difference images (Volunteer 2)
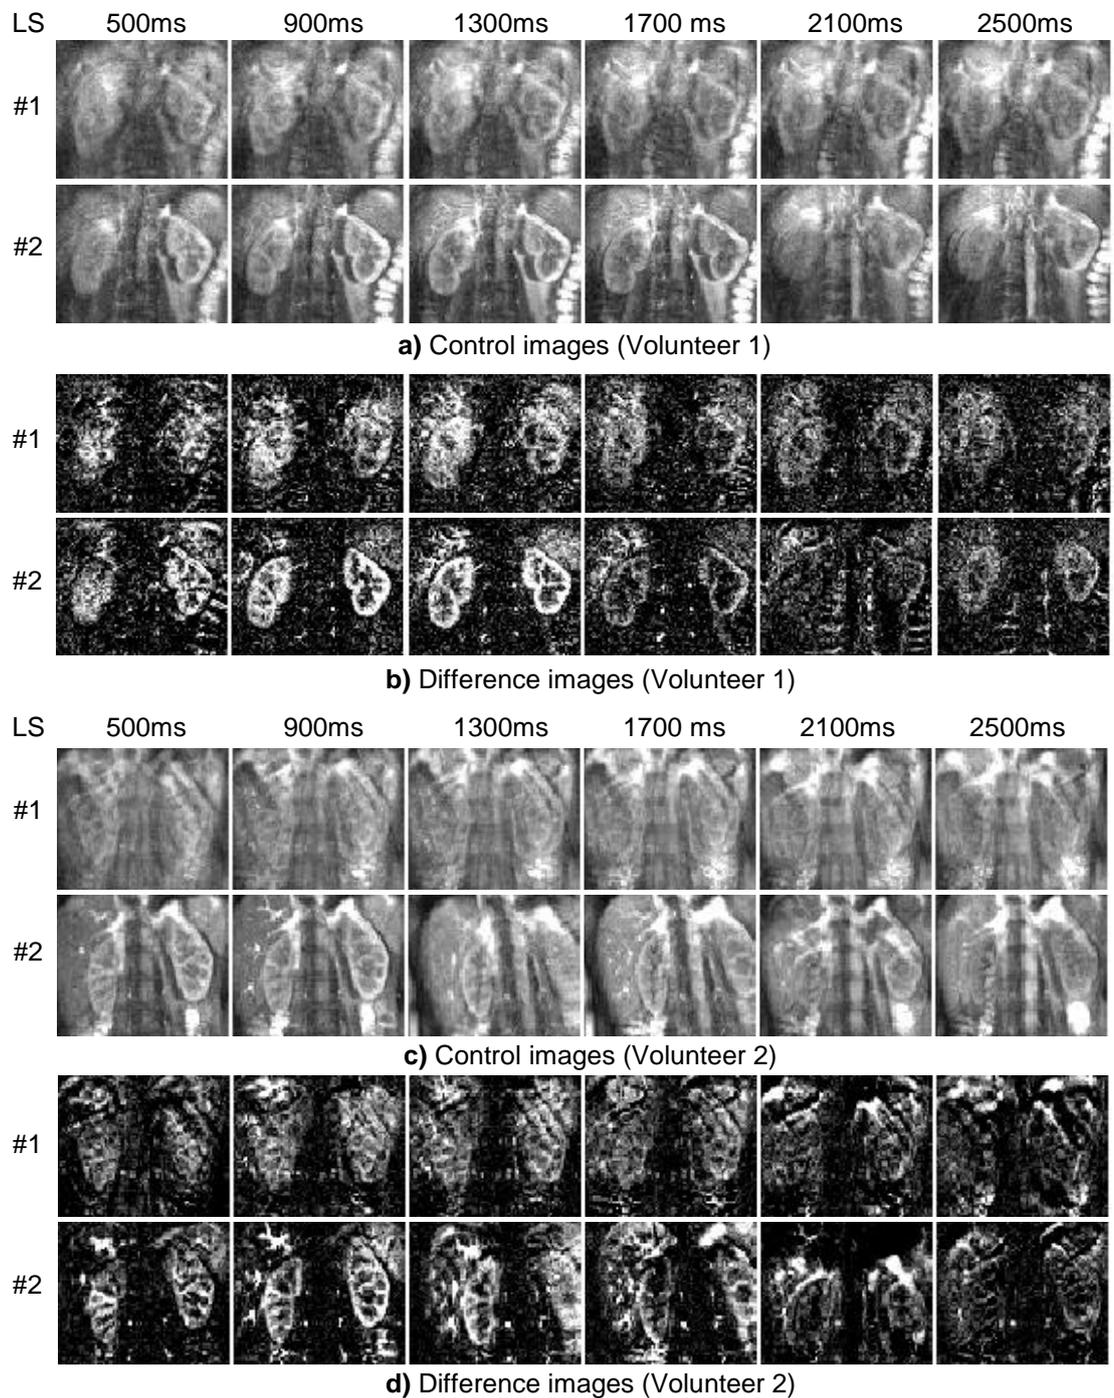
Figure 4-12 LS#1 and LS#2 comparison in two volunteers

A visual inspection of the scans where the subjects exhibited a voluntary irregular respiration revealed that the appearance of motion artefacts is consistent with the information from the respiratory waveform. Specifically, motion corrupted images result when the segments of their corresponding k-space are acquired in different parts of the respiratory cycle. An example is

shown in Figure 4-13 c). As can be seen in the bellows waveform (Figure 4-13 a)), the first shot corresponding to the motion-corrupted image (shot #19 in the acquisition) was acquired at a significantly different part of the respiratory cycle when compared to the remaining two shots corresponding to the same k-space (shots #21 and #23). According to the respiratory waveform, shot #19 was acquired following the inspiration of a deep breath, which causes the kidneys to move to an inferior position. This explains the artefact in c), which will be further confirmed in Figure 4-17 e) and h). A similar example, from a different volunteer can be seen in Figure 4-14, which in addition shows how these artefacts, which are initially seen in the control and tag images, propagate to the PWIs which are used for RBF quantification.

b) Motion-free image   c) Motion-corrupted image
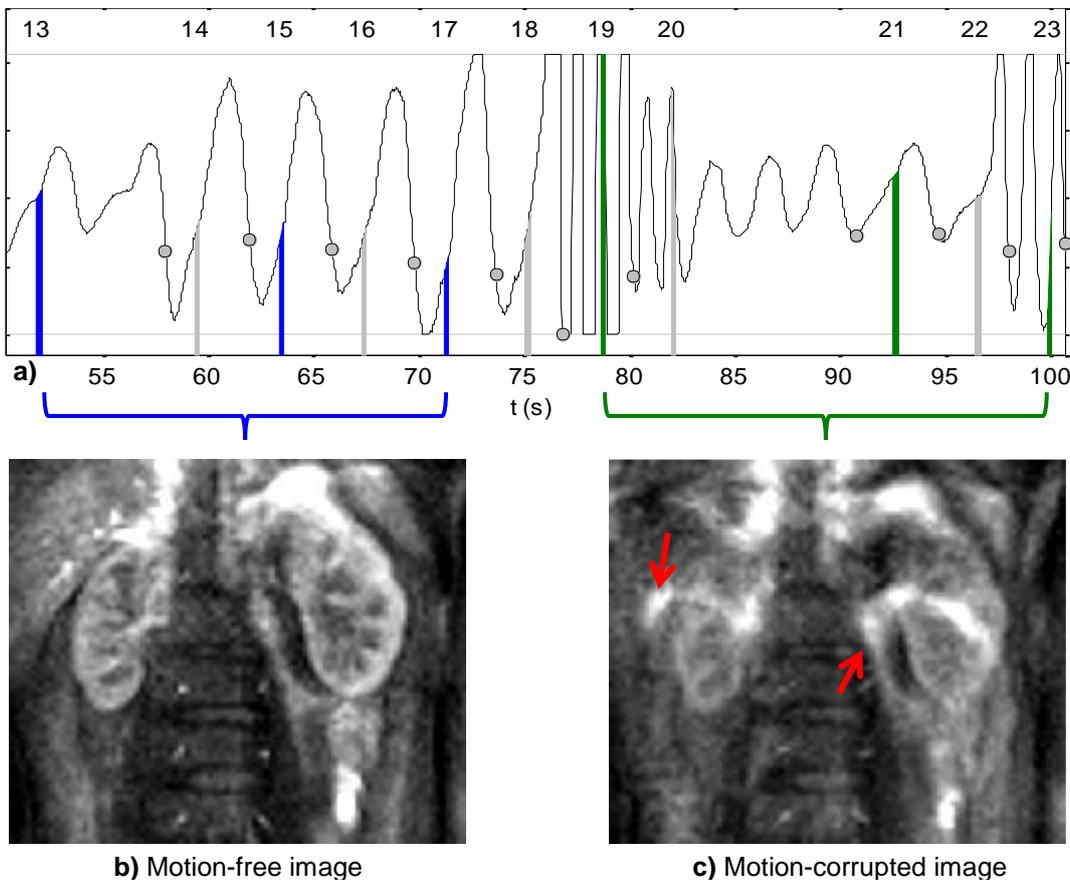
Figure 4-13 Effects of inter-shot motion due to respiration (I) in multi-shot background-suppressed 3D GRASE acquisitions. a) Respiratory trace recorded with bellows. The blue readouts correspond to the acquisition of the motion-free image (b)) and the green readouts correspond to the acquisition of the motion-corrupted image (c))

control (TI = 1300ms)   PWI (TI = 1300ms)   control (TI = 1700ms)   PWI (TI = 1700ms)
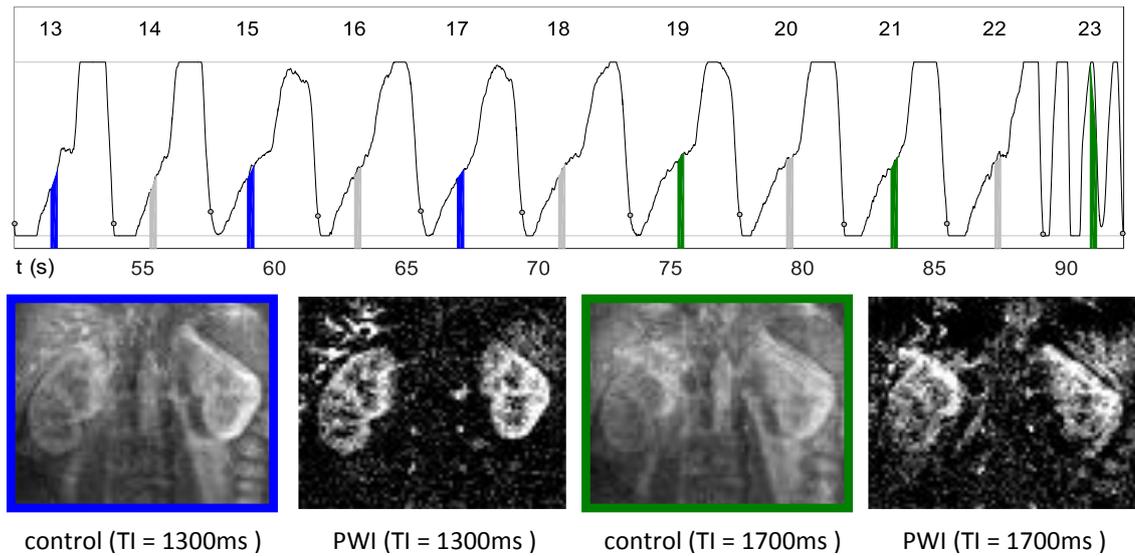
Figure 4-14 Effects of inter-shot motion due to respiration (II) in multi-shot background-suppressed 3D GRASE acquisitions and corresponding difference images. Readout #23, i.e. the third shot of the control image at the TI of 1700ms, was acquired in a significantly different part of the respiratory cycle.

Consistent with the existing literature, the increased motion sensitivity of non-background-supressed data was confirmed. Image artefacts are sometimes encountered even when respiratory triggering appears effective, ensuring the segments of a single k-space volume are acquired in comparable parts of the respiratory cycle (Figure 4-15 b). ASL subtraction errors also appear more prevalent in non-background suppressed data, especially at longer TIs (Figure 4-16).

b) Motion-corrupted image **c)** Motion-free image
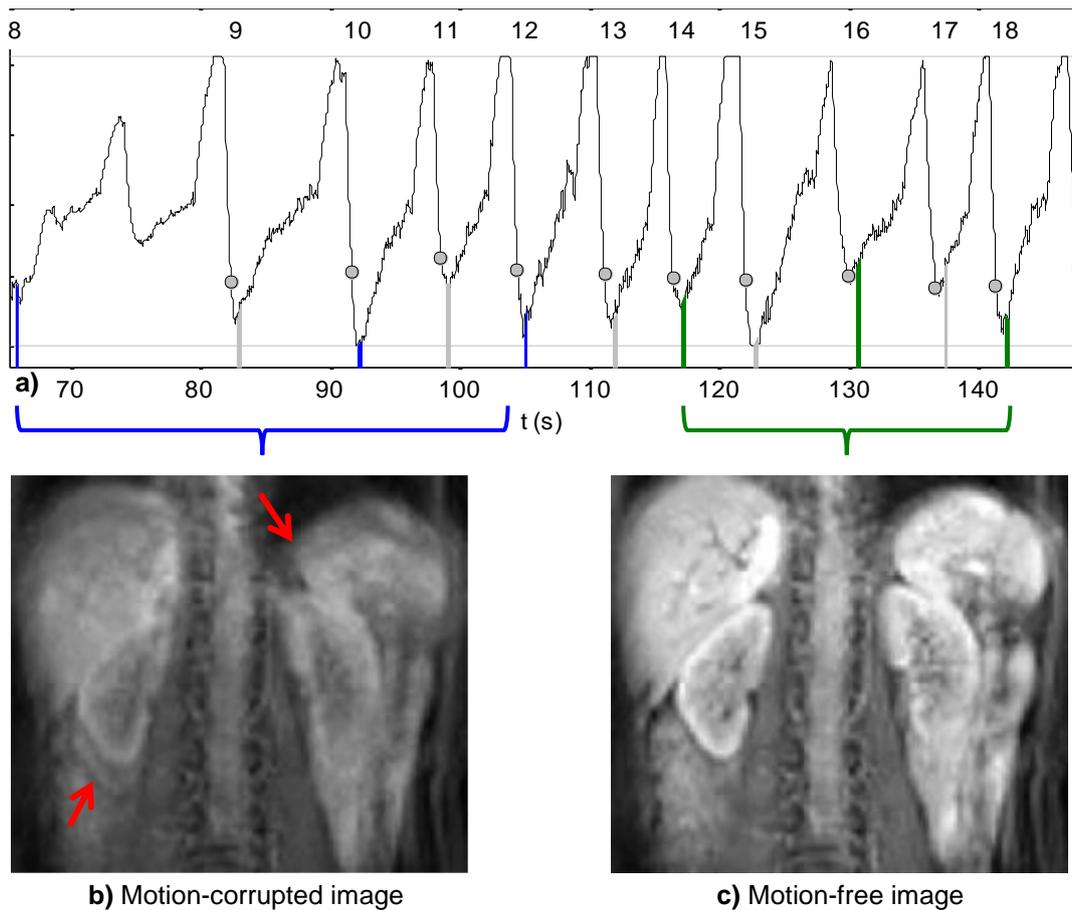
Figure 4-15 Effects of inter-shot motion due to respiration (III) in multi-shot non background-suppressed 3D GRASE acquisitions. a) Respiratory trace recorded with bellows. The blue readouts correspond to the acquisition of the motion-corrupted image (b)). Unlike in Figure 4-13, each readout occurs at consistent points in the respiratory cycle, but breathing-related artefacts are still evident.

The projections obtained from the FT of the first phase correction reference line of each shot corresponding to the images shown in Figure 4-13 and Figure 4-15 are shown in Figure 4-17 d),h) and Figure 4-18 d),h), respectively. In both cases, it can be seen that the projections corresponding to the k-space segments (e)-g) in both figures) giving origin to the corrupted images are the less consistent. This suggests that to a certain extent, there is motion-related information embedded in these projections, which could be used as a proxy for kidney movement. The question of whether this information can be used for detecting or correcting motion artefacts is outside the scope of this chapter, but will be revisited in Chapter 5.

a) Control images without background suppression

b) Difference images without background suppression

c) Background-suppressed Control images

d) Background-suppressed Difference images
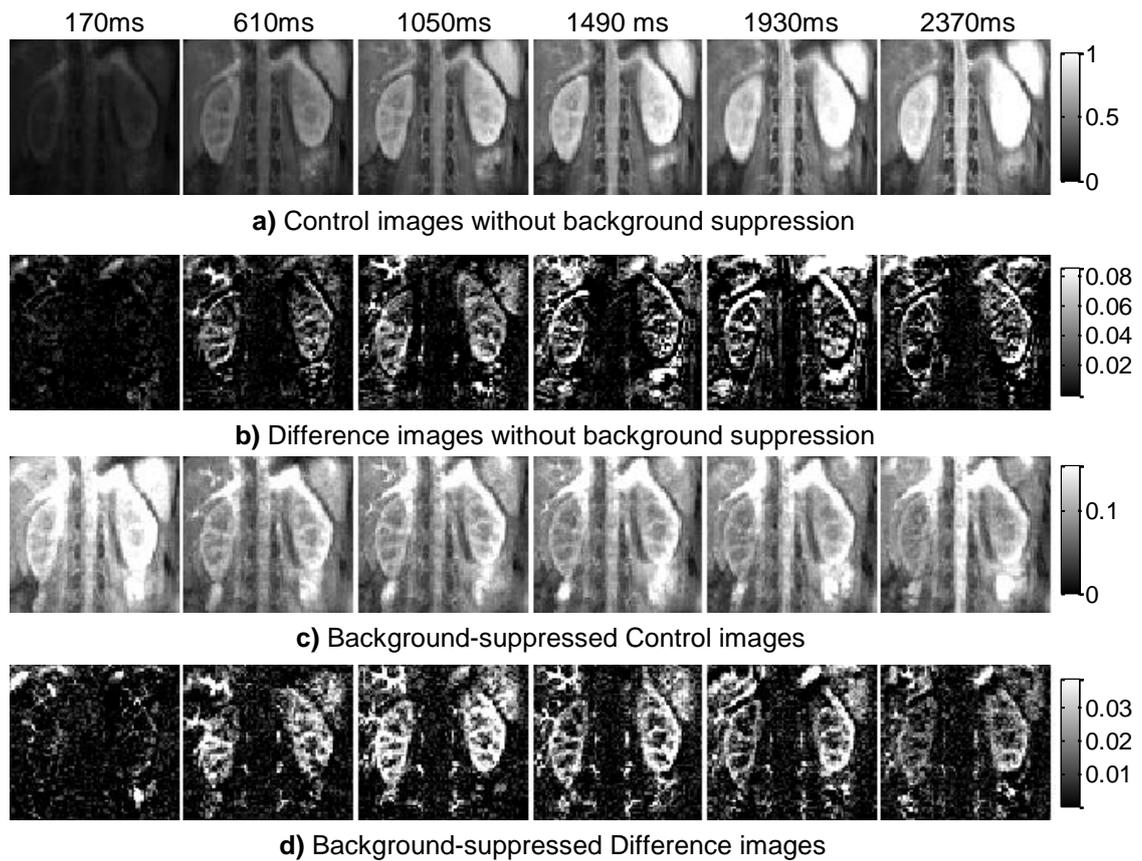
Figure 4-16 Comparison of background-suppressed (BS) and non-BS data acquired in the "still" condition in volunteer 2. BS is effective in reducing subtraction errors, especially at longer TIs when the effectiveness of respiratory triggering is reduced.

The undersampled reconstructions of each shot corresponding to each of the previously referred projections are also shown in Figure 4-17 and Figure 4-18. The respiratory waveform in Figure 4-13 a) indicates that shot #19 was acquired at end-inspiration. The zero-filled reconstruction of this shot (Figure 4-17 e)) is consistent with this observation, showing the kidneys (which overlap due to aliasing) in a significantly more inferior position when compared to the other two shots (Figure 4-17 f) and g)). These shifts in position cause phase inconsistencies in the fully sampled k-space giving rise to the artefacts in Figure 4-13 c) (highlighted by the arrows). Furthermore, the projections shown in Figure 4-17 h) also show a greater variation among shots when compared with the projections in Figure 4-17 d) acquired immediately before, in the motion-free scan (Figure 4-13 b)).

The origin of the artefacts in Figure 4-15 b) can also be explained by respiratory related shifts as suggested by the undersampled reconstructions of Figure 4-18 e), f) and g) and the projections in Figure 4-18 h). In this case however, the renal shifts in position among the shots were smaller, and hard to detect based only on the respiratory trace plot in Figure 4-15 a).

The zero-filled reconstructions of each shot making up the control images in Figure 4-1 are shown in Figure 4-19s. In a), the undersampled reconstructions reveal different structures among the segments (arrows). This, together with potential bowel movements between the acquisition of the different shots results in ghosting in the fully sampled reconstruction. These data were acquired before enabling the recording of the respiratory data. However, in these cases, breathing-related motion is probably less problematic given the location of the transplant kidney (Lanzman et al. 2010). To reduce bowel movement, a more important source of artefacts when scanning more inferior regions, fasting for 4-6h before the scan or using anti-peristaltic drugs may be helpful (Romero et al. 2014). Nevertheless, one might argue that in these severe corruption cases, data rejection and/or reacquisition are possible alternatives if correction is a challenge.
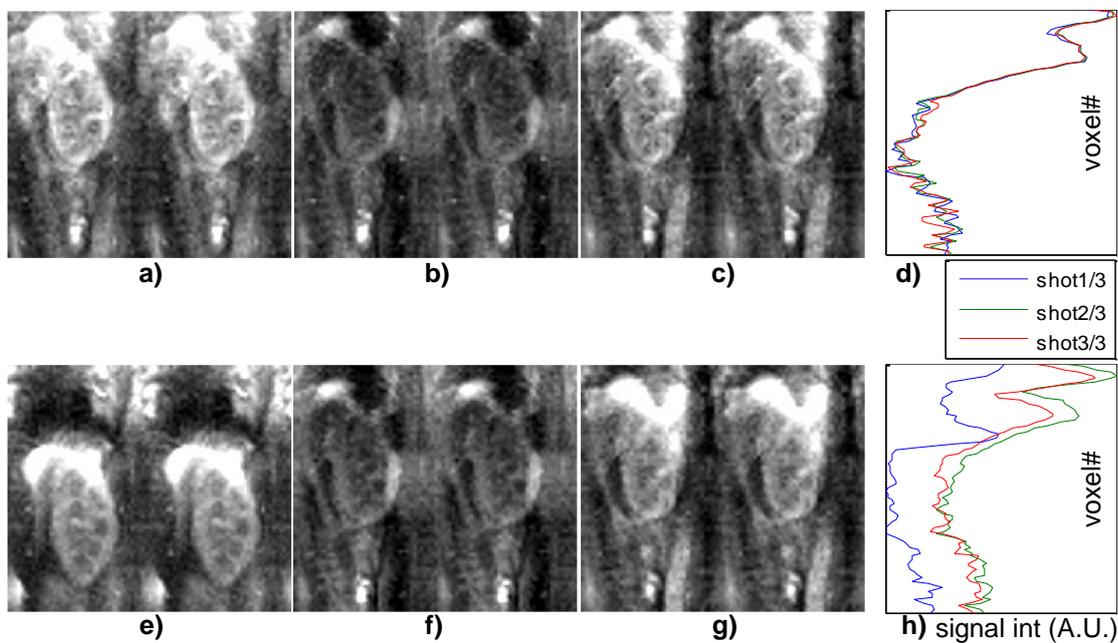
Figure 4-17 Zero-filled reconstruction of each shot from data in Figure 4-13 and corresponding projections obtained from the EPI phase correction data. a), b) and c) (which correspond respectively to shots #13, #15 and #17 in Figure 4-13) suggest that the acquisition of each shot took place at a similar point of the respiratory cycle. Conversely, both the zero-filled reconstructions e), f), g) (which correspond respectively to shots #19, #21 and #23 in Figure 4-13 and the projections suggest that shot #19 (i.e. projection #1 in h)) was acquired when the kidneys were in a significantly lower position compared to the other two shots. This matches what can be seen in the bellows waveform in Figure 4-13 a), providing evidence for the origin of the motion artefacts in Figure 4-13 c).
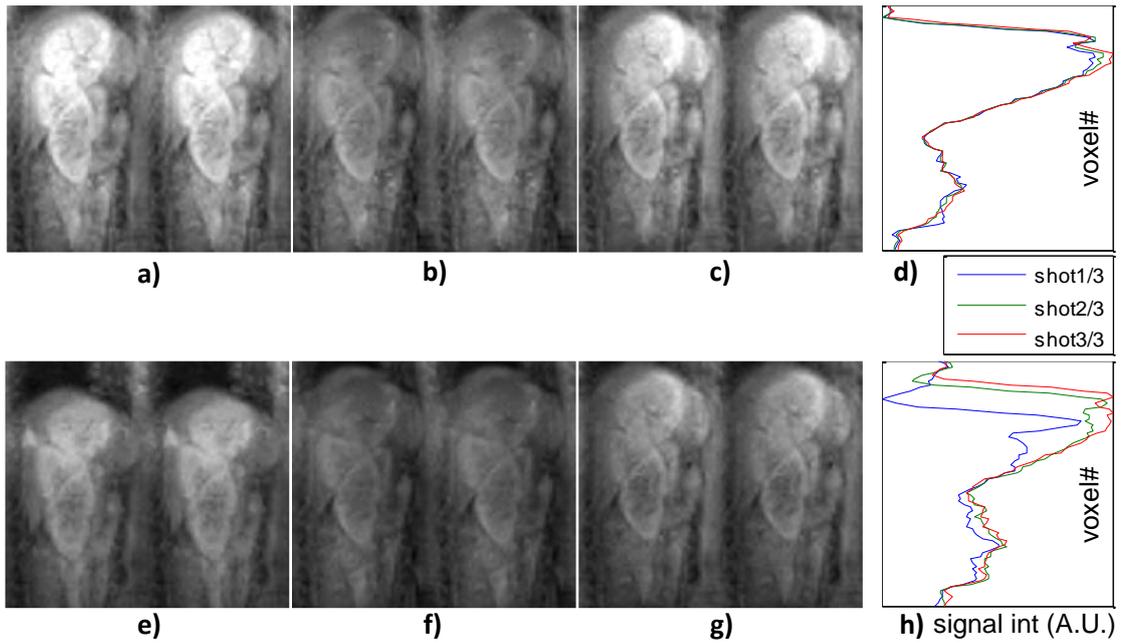
Figure 4-18 Example analogous to Figure 4-17. Both e) and the projections in h) show the kidneys acquired in a different part of the respiratory cycle compared to the other shots, giving origin to the artefacts in Figure 4-15 b)
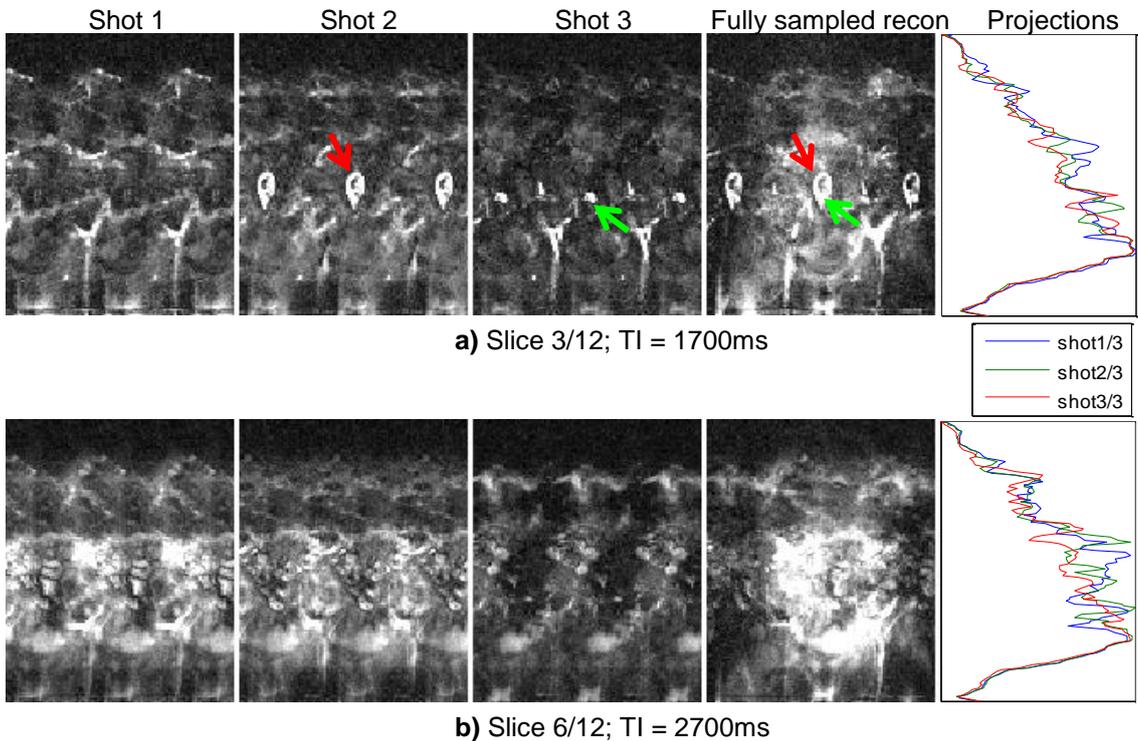
Figure 4-19 Zero-filled reconstruction/projections from corrupted data. These data correspond to each shot making up the control images shown in Figure 4-1.

Another alternative is to use single-shot acquisitions. These are more robust to motion, given that the entire k-space is sampled after a single excitation. In this scenario, care must be taken

so that the readout duration is kept within acceptable limits. On one hand, this is crucial to ensure that the assumption of negligible intra-shot motion still holds. In addition, this is a requirement to avoid excessive blurring due to T_2 decay which broadens the point-spread function of the system, decreasing the effective resolution of the scans. Examples of PWIs acquired with single-shot scans are shown in Figure 4-20. Relevant image readout specific parameters are detailed in Table 4-2.

Table 4-2 Single-shot and multi-shot image readout parameters

Scan#	Resolution (mm)	TE (ms)	BW (Hz/Px)	Echo spacing (ms)	EPI factor	Readout duration (ms)
3	4.2×4.2×6	38.24	2741	0.4	78	302
4	6.3×6.3×6	26.46	2790	0.4	52	208
All others	3.1×3.1×5	26.8	2170	0.6	40	237

One of the disadvantages of the single-shot scans is the lower image resolution. However, from an SNR point-of-view, a lower image resolution is beneficial. This is especially important in an inherently low SNR technique such as ASL. In single-shot acquisitions, one can trade the extra time previously used to acquire multiple segments for additional single-shot measurements (multiple averages), which not only is beneficial in terms of SNR but further increases the robustness of the measurement to motion. Data acquired at 6.3mm in-plane resolution, with the shortest readout duration of the set, is of high quality, but the coarse image resolution results in a severe loss of cortico-medullary differentiation. At this resolution level, partial volume effects will likely have a pronounced effect on the quantification of cortical perfusion. Nevertheless, the high SNR and speed of acquisition of these scans can render them useful for transit delay mapping (Dai et al. 2012). An improved resolution of 4.2mm in-plane can still be obtained in a single-shot condition (Figure 4-20 b), at the expense of a 27% increase of the readout duration (237ms to 302ms). The resulting readout duration is still within the limit of what is acceptable (with respect to the total amount of signal decay occurring during the echo train) according to a consensus publication with recommendations of pulse sequence parameters for 3D-GRASE ASL (Alsop et al. 2015). Even though this was targeted to the brain, the relatively similar T_2 values of the brain (e.g. 92.1±2.6ms for gray matter (Deoni et al. 2003)) and kidneys (e.g. 87±4ms for cortex (de Bazelaire et al. 2004)) implies that this particular recommendation is arguably translatable to renal imaging.

|170ms | 610ms | 1050ms | 1490 ms | 1930ms | 2370ms |

a) Original 3-shot acquisition (In-plane resolution: 3.1mm isotropic, Slice thickness: 5mm)

b) Single-shot acquisition (In-plane resolution: 4.2mm isotropic, Slice thickness: 6mm)

c) Single-shot acquisition (In-plane resolution: 6.3mm isotropic, Slice thickness: 6mm)
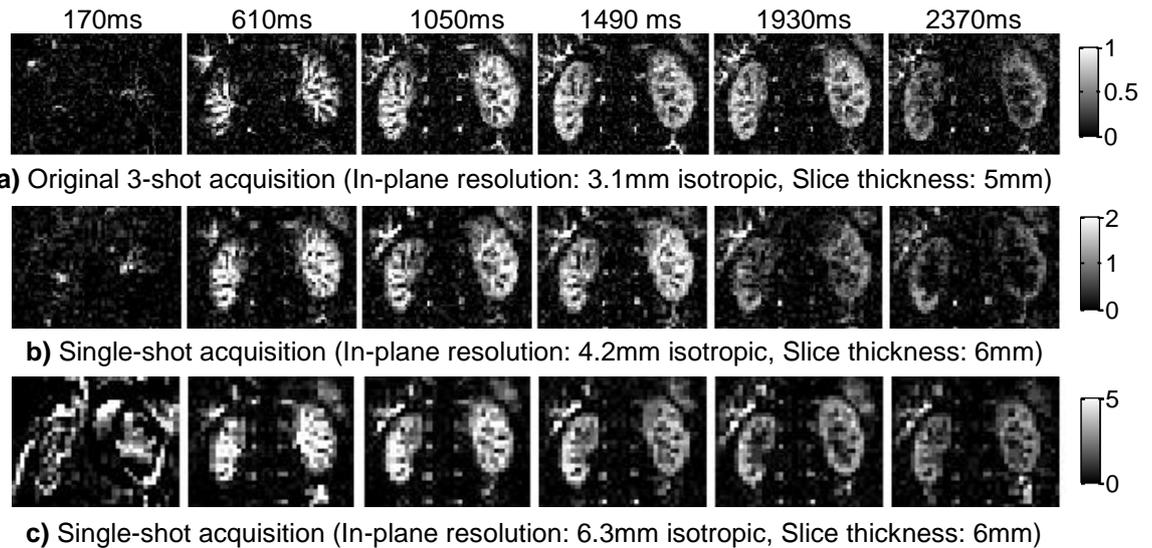
Figure 4-20 Multi-shot and single shot ASL difference images. Comparison of PWIs acquired with the optimised multi-shot sequence and a single-shot sequence, the latter at two different in plane resolution settings. Scans acquired during free breathing, without any extra intentional motion by the subject.

4.5 Conclusion

ASL is intrinsically a motion sensitive technique since is dependent on an image subtraction step (control-tag) to generate a perfusion weighted image. Imaging in the abdomen presents additional challenges, motion in particular, when compared to the brain, where most of the ASL research is conducted. Respiratory triggering has shown to be useful, even when using advanced image registration algorithms (Seif et al. 2015). However, it is unable to solve all respiratory-related problems because of the time delay between the trigger (ASL inversion pulse) and the imaging readout. Additional motion sensitivity is introduced when k-space is filled in a segmented fashion, with significant time gaps between the acquisition of the different RF excitations. I have shown the importance of reducing these time delays, in particular when acquiring data at multiple TIs. I have shown that our acquisition scheme can intrinsically provide motion-related information, by using EPI phase-correction lines in a navigator-like fashion. Whether this information can be used for rejection of corrupted PWIs, due to displacements between control and tag images was investigated elsewhere in this thesis (see Chapter 5). Alternatively, inconsistencies between the projections corresponding to different shots in segmented acquisitions will likely allow for detection of corrupted data, even before control-tag subtraction. Our preliminary volunteer data suggests that there is value in further investigating single-shot acquisition schemes. One may hypothesize that these are potentially better suited in the routine clinical environment, where motion is more likely to be an issue, especially in the paediatric population. With optimised sequence parameters, image artefacts due to motion

should be avoided in the control/tag images simply because of the fact that all data is acquired in a very short period of time, effectively freezing any motion. Nevertheless, both subtraction errors and shifts in position by the tissue of interest throughout the multi-TI data will still be a source of error in the perfusion measurements. Image registration algorithms will be used to tackle these issues in Chapter 6. In summary, 3D-GRASE is an efficient acquisition scheme for ASL but must be optimised when applied in the abdomen to overcome motion-related issues before it can be routinely used in the clinic.

5 Robust kidney perfusion mapping using a single-shot 3D-GRASE readout and retrospective noise-reduction methods

Part of this work was presented at an European MRI conference (Nery et al. 2016b) and (Nery et al. 2016a)

5.1 Introduction

The ultimate aim with this thesis is to develop a renal ASL method for clinical application in paediatric patient cohorts. This specific application demands for a robust approach capable of delivering usable data in the face of potentially uncompliant subjects. Therefore, the renal ASL approach developed in this chapter focuses on speed and robustness to motion.

This chapter is divided into two parts, building on the conclusions laid out in the previous chapter. The first part consists of a feasibility study of perfusion quantification in the kidneys of healthy volunteers where for the first time single-shot 3D GRASE was employed as an imaging readout module. Unlike in (Cutajar et al. 2012), where ASL data was acquired at multiple TIs, here we take the conservative approach of using a single TI. The disadvantage of this approach is that it relies on the assumption that at the chosen TI the labelled bolus has arrived at the imaging region, which can be problematic in situations where the renal haemodynamics deviate from normal. Nevertheless, this approach is in line with our goal of achieving a fast acquisition scheme resilient to motion. Reasons for this include the fact that using a single TI results in:

- *Shorter scan times:* the scan time is proportional to the number of TIs in the ASL protocol, which if combined with signal averaging may result in a prohibitively long scan. On the other hand, in single-TI single-shot 3D-GRASE imaging the nominal scan time remains most likely under 3 minutes, even when acquiring a large number of ASL control-label pairs;

- *Possibility of performing signal averaging:* due to the shorter scan time;

- *Higher signal-to-noise ratio (SNR):* resulting from signal averaging, particularly useful in cases of reduced perfusion, as is expected in patients with impaired renal function;

- *Possibility of rejecting corrupted data:* given the protocol's built-in data redundancy resulting from the acquisition of multiple ASL pairs at the same TI. This may be particularly beneficial in background-suppressed acquisitions, given that the reduced static tissue signal may affect the effectiveness of image registration algorithms;

- *Increased effectiveness of respiratory triggering:* since even though the triggering event doesn't directly initiate data acquisition, the delay between the trigger and actual

readout is constant for all acquisitions. This increases the likelihood of acquiring the data at consistent points in the respiratory cycle;

Another reason why single-shot scans are beneficial is their reduced motion sensitivity when compared to non-background suppressed multi-shot scans. This is an issue even if these scans are not used to compute perfusion-weighted images (PWIs), but rather for calculation of baseline M_0 and T_1 maps, used in the quantification of Renal Blood Flow (RBF).

Inconsistencies in k-space that translate to artefacts such as those seen in section 4.2 are no longer expected when data is acquired in a single fast readout (assuming no intra-segment motion). However, movement will still play a detrimental role as after control-label subtraction it gives rise to spurious signals which instead of reflecting tissue perfusion are caused by subtraction errors. This is particularly important when imaging the kidneys since they are subject to breathing-related movement. Furthermore, standard respiratory triggering implementations may face ASL-specific challenges (see Figure 5-1).

Several methods have been developed which reduce the contribution of corrupted ASL difference images to a mean PWI (Tan et al. 2009; Shirzadi et al. 2015; Tanenbaum et al. 2015). In the work described in the second part of this chapter, I tackled this problem using two main strategies to either outright reject or downweigh the contribution of corrupted measurements. First, we developed and implemented an algorithm that makes use of the "navigator"-like Nyquist-ghost phase correction projections (as discussed in Chapter 4) for identifying and rejecting ASL pairs corrupted by motion.

Furthermore, we apply for the first time in the kidneys two threshold-free weighted averaging approaches, based on (Tanenbaum et al. 2015). All motion compensation strategies are evaluated in terms of the quality of the final averaged PWI, using metrics such as the SNR and image entropy as potential image quality proxies.
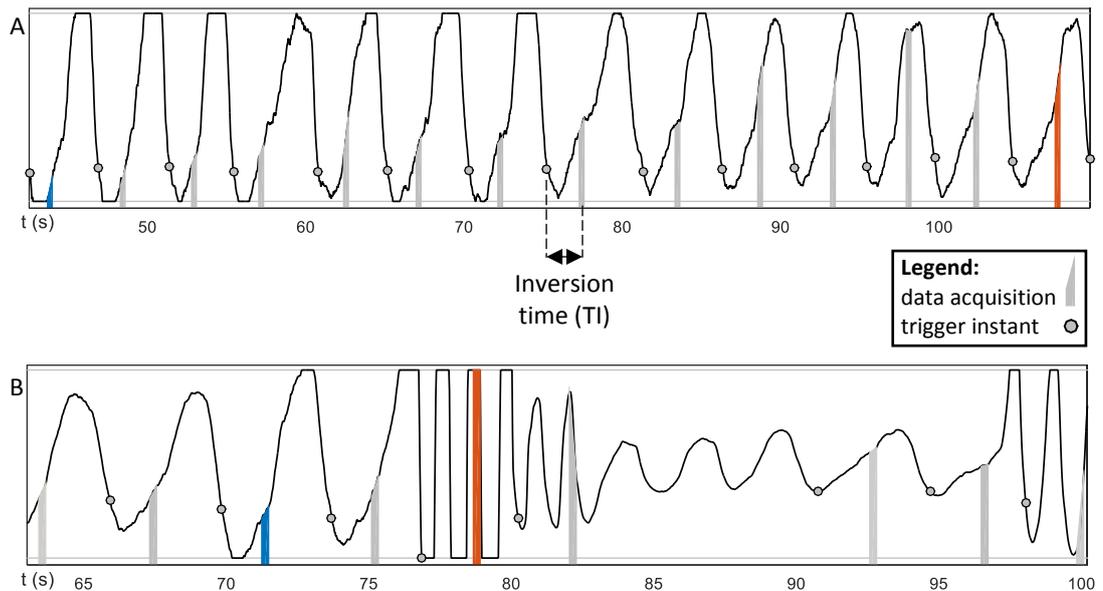
Figure 5-1 Challenges with respiratory triggering. A – Varying inflow times may result in different acquisitions being performed in different parts of the respiratory cycle, even in subjects with reasonably regular respiratory patterns. B – Irregular respiratory rates. Note how the highlighted acquisitions occur at significantly different parts of the respiratory cycle, and consequently the kidneys will be in different positions.

5.2 Methods

5.2.1 MR acquisition and Subjects

MR data was acquired in a Siemens Avanto 1.5T scanner (Siemens Healthcare, Erlangen, Germany) running the Siemens NUMARIS/4 VB17A software version. Both the spine coil embedded into the scanner table and a dedicated abdominal array coil were used. The latter is composed of two clusters, each with 3 coil elements in the left-right direction, whereas the spine coil has eight clusters also with 3 elements each in the left-right direction (of the eight clusters, three or four were active, depending on the FOV). A single-shot FAIR Q2TIPS 3D-GRASE pulse sequence (see Figure 2-11 and Figure 2-14) was used. Partial Fourier (factor 3/4) was applied along the partition direction, resulting in an ETD of 248 ms (for the brain, it is recommended an echo train shorter than 300 ms (Alsop et al. 2015) to avoid excessive through-plane blurring). In order to maintain the highest possible SNR (which is crucial for maximising the accuracy of measurement for the low perfusion values expected in disease), it was decided not to use a parallel imaging acquisition method. The g-factor SNR penalty resulting from the small number of coil elements in the PE (L-R) direction (3 for both the abdominal and spine array coils, as described above) would severely compromise image quality, even if a relatively modest parallel imaging acceleration factor (e.g. R=2) was employed, and so only partial Fourier methods were used in this work. Each coronal-oblique volume is comprised of 10 partitions with 6mm thickness

(8 sampled, 2 zero padded). The field-of-view was 288mm (Head → Foot) × 288mm (Left → Right). TR and TE were, respectively, 3000 and 31.5ms. In-plane resolution was 4.5mm × 4.5mm. The bandwidth and echo-spacing were, respectively, 2790Hz/Px and 0.4ms. Fifteen control-tag ASL pairs were acquired at a TI of 1200ms, allowing for signal averaging. The bolus duration (τ) was fixed to 900 ms. The nominal scan time was 90 seconds. Respiratory bellows were used to trigger the application of the ASL inversion pulses at end-expiration. Background suppression is implemented as described in (Günther et al. 2005) (see Figure 2-14). Slice-selective pre-labelling WET saturation pulses (Ogg et al. 1994) as well as a single post-labelling sinc saturation pulse were used.

Five healthy adult volunteers, with previous experience in MRI (age (years) = 32±5 (mean±std), range=26-37) were scanned using this acquisition protocol, under two different conditions. During the first scan, the volunteers were instructed to remain still. No further instructions were necessary as the acquisition was performed during free-breathing. In the second scan, the volunteers were asked to change their respiratory rate/amplitude during approximately five out of the fifteen ASL pairs to intentionally reduce the efficacy of respiratory triggering. To assess the reproducibility of the technique, one of the subjects was scanned in the "no-motion" condition on three different occasions, with the second and third scan occurring 7 and 21 days after the first.

5.2.2 Image analysis

Perfusion quantification was performed using a single-compartment model (Song et al. 2010), assuming a T_1 value for the kidney cortex of 0.966s (de Bazelaire et al. 2004). Because at the time of writing this chapter, registration algorithms able to register ASL images (background-suppressed in this protocol) to the reference non-background suppressed data (reference M_0 images) had not been validated yet, the M_0 value used for quantification is the median M_0 in whole-kidney regions of interest (Fig. 1), obtained from a non-background suppressed 3D GRASE reference scan. This ensures that in the case of the ASL and the M_0 data being misaligned, perfusion quantification values close to the boundaries of the kidney (i.e. in the cortex) will not be biased by using M_0 values from outside the kidneys. Renal cortical ROIs were manually drawn in the reference non-background suppressed data. Average cortical perfusion values were calculated within these ROIs in each of the subjects.

5.2.3 Reducing the effects of motion corruption – theory

Several methods were implemented with the aim of minimising the effect of large movements, which corrupt data beyond what is generally feasible for image registration methods to

compensate for. This is highly relevant in background-suppressed renal ASL data, as on one hand the background-suppression (BS) reduces the effectiveness of image registration algorithms and on the other hand extreme abdominal motion may be characterised by non-rigid transformations which further increase the complexity of the registration problem.

5.2.3.1 Strategy 1: Data rejection

As discussed in Chapter 4 immediately before the image readout portion of the 3D-GRASE pulse sequence, non-phase encoded lines are acquired for the purposes of phase correction. The 1D Fourier transform of each non-phase encoded echo from the phase correction reference scan corresponds to a projection of the object along the orthogonal direction of the readout gradient (Head→ Foot in all imaging protocols in this thesis). We have shown in Chapter 4 that these projections contain intrinsic motion information. A simple algorithm was developed and implemented to make use of this information to identify misaligned ASL control-tag pairs, which after image subtraction yield a motion-corrupted ASL PWI. These corrupted intermediate PWIs are then rejected before data averaging. An overview of this algorithm is depicted in Figure 5-2. This figure shows a practical example of acceptance and rejection of PWIs, respectively in a case where the position of the underlying control and tag images is consistent and a case where the two are misaligned. This algorithm can be divided into three steps:

- Step 1: definition of the ROI (within the projection data) within which shifts caused by motion will be tracked;
- Step 2: determination of the shift (in pixel units) which maximises the similarity between control and tag projections (*bestShift*);
- Step 3: acceptance/rejection of an ASL pair according to *bestShift*.

The aim with Step 1 is to determine the region within the projections which ideally provides the maximum amount of information related to kidney motion. Following a visual inspection of the projection data and their variability especially in inconsistent ASL pairs, a set of heuristics were defined to determine this region in an automatic way. A consistent observation was that the region with highest signal in the projection data is located at the level of the liver region. For the purposes eliminating highly corrupted data (assuming a more fine motion correction to be feasible with image registration), it was assumed that tracking this region signal would suffice to detect significant breathing-related mismatches between the position of the kidneys in the control and tag images of a given ASL pair. As such, the ROI for motion tracking was defined to be the 9cm-wide region with the highest signal intensity in the sum of the control projections (see Figure 5-2.a)). The rationale for the size of this ROI is to provide a large enough margin to

likely include the lung-liver interface (proxy for breathing-related motion) within the tracking region in the second step of the algorithm.
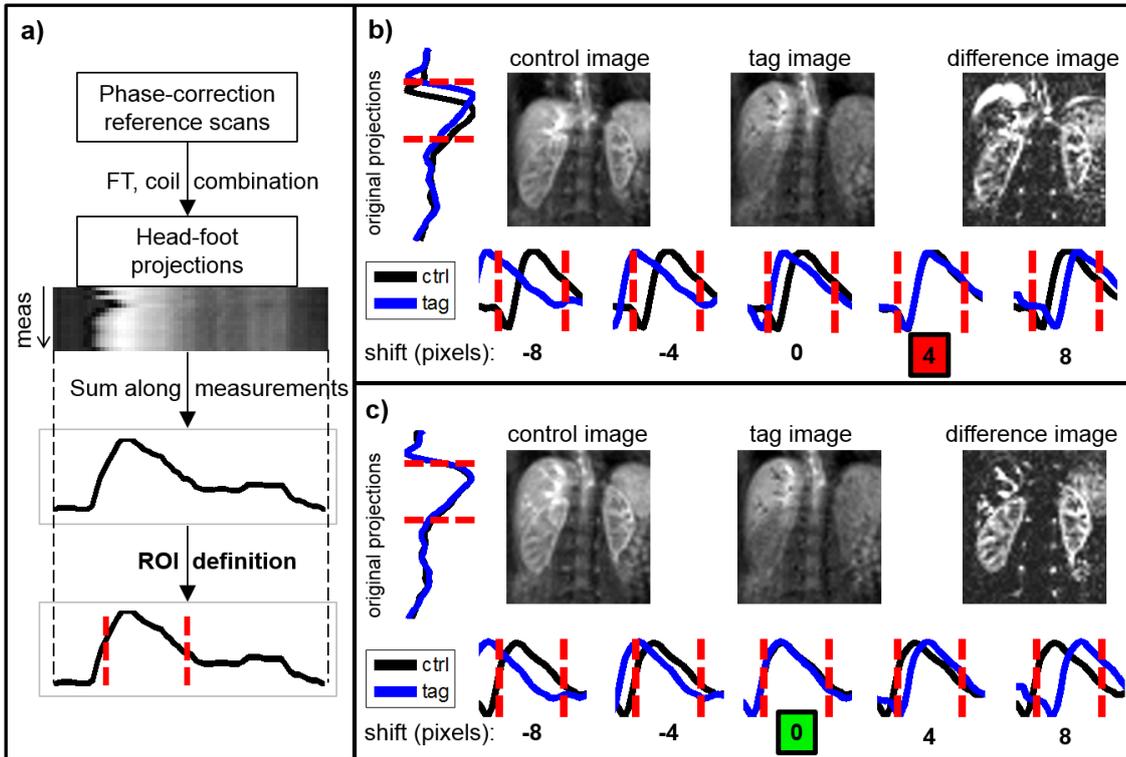
Figure 5-2 Automatic retrospective rejection in ASL using Nyquist-ghost elimination data. a) Definition of the ROI from which changes in position are tracked; b) example of rejection of a motion-corrupted difference image; c) example of acceptance of a motion-free difference image

In the second step of the algorithm (see Figure 5-2.b) and c) for a practical example), for each ASL pair, the projection from the control acquisition is taken as a reference, and the projection from the corresponding tag acquisition is iteratively shifted from -4.5cm to 4.5cm (-10 to 10 pixels at the resolution at which these data were acquired) in the head-foot direction, relative to its original position. For each shift, the sum of squared differences between the control projection and the shifted tag projection is computed within the previously defined ROI. The *"bestShift"* is defined as the shift which minimizes this sum of squared differences.

Finally, in step 3 ASL pairs are accepted/rejected according to *bestShift*:

- $|bestShift| > 1 \rightarrow$ reject ASL pair
- $|bestShift| \leq 1 \rightarrow$ accept ASL pair

The threshold of 1 was chosen to avoid false positives (i.e. rejection of motion free data), as a small shift in the ROI may not necessarily translate to a significant mismatch between the positions of the kidneys at the time of acquisition of the control and tag image volumes. It is

116

assumed that in the event of the algorithm failing to reject motion corrupted data whose absolute value of *bestShift* is smaller or equal than 1, the underlying mismatch between control and tag image is small enough to be amenable for correction using image registration.

The biggest advantage of this method is that the non-phase encoded data is acquired immediately before the actual image readout. Therefore, the motion information within these data matches exactly the position of the different organs exactly at the point of image acquisition, which is not the case with triggering, due to the delay between the triggered labelling pulse and the image acquisition.

5.2.3.2 Strategy 2: Weighted averaging

An alternative strategy is to weight the contribution of the different measurements in different proportions according to noise estimates derived from the ASL data itself. The contribution of measurements deemed corrupted by motion (and potentially other artefacts) to the final PWI can then be significantly reduced. The advantage of these approaches is that they largely require no user input (if the user wishes to restrict the FOV region over which the noise estimates will be calculated then the only step that requires user interaction at this stage is the definition of masks for specifying these), do not depend on arbitrarily defined thresholds, and can be readily implemented for any type of ASL sequence with little computational cost. The disadvantages include the fact that they are a purely retrospective method and unlike registration, they do not correct corrupted data.

In general, if ϵ_p is a vector of noise estimates for each measurement (e.g. an ASL pair in this case), the vector of weights can be obtained by:

$$w_p = \frac{1}{c\epsilon_p}$$

Equation 5-1

Where:

$$c = \sum_{p=1}^{N} \frac{1}{\epsilon_p}$$

Equation 5-2

And N is the number of measurements. As such:

$$\sum_{p=1}^{N} w_p = 1$$

Equation 5-3

The calculation of the mean PWI ($\bar{I}$) is simply the dot product:

$$\bar{I} = \sum_{p=1}^{N} w_p \cdot I_p \qquad\qquad \text{Equation 5-4}$$

Where I_p is the PWI corresponding to the control-tag pair p.

We explored two approaches for computing the noise power ϵ_p for each measurement.

In the first, noise is estimated by computing the so called DVARS measure (Smyser et al. 2010; Power et al. 2012). In this approach, applied in a previous neuroimaging ASL study (Tanenbaum et al. 2015):

$$\epsilon_{p_A} = \frac{1}{M} \sum_{r=1}^{M} \left\{ \left[I_p(r) - I_{p-1}(r) \right]^2 + \left[I_{p+1}(r) - I_p(r) \right]^2 \right\} \qquad \text{Equation 5-5}$$

Where $I_p(r)$ corresponds to the intensity of the PWI corresponding to the ASL pair p at the locations (voxels) r. The total number of voxels over which the noise metric is calculated is M.

In this case, the DVARS measure is only defined if "temporally" adjacent measurements exist. As such, both the ASL pairs 1 and N are discarded, regardless of them being corrupted by motion or not. This method will henceforth be referred to as *wMeanA*.

I propose a second approach (*wMeanB*), where the difference between each individual measurement and the mean of the PWI time series is computed:

$$\epsilon_{p_B} = \frac{1}{M} \sum_{r=1}^{M} \left\{ \left[I_p(r) - I_{mean}(r) \right]^2 \right\} \qquad \text{Equation 5-6}$$

Unlike when using ϵ_{p_A}, no data needs to be discarded for the computation of the weights for each ASL pair. One might expect this to be particularly important for multi-TI ASL. This is because such protocols necessarily include a smaller number of measurements due to time constraints, and as such, avoiding the rejection of uncorrupted data is crucial. Furthermore, this modification allows the possibility of running this algorithm in an iterative way. In the first iteration, I_{mean} corresponds to a typical ASL averaged image where each measurement is assigned the same weight $(1/N)$. In subsequent iterations, I_{mean} corresponds to the output weighted mean image in the iteration that preceded it. As such, in the second and later iterations the individual PWIs are "compared" to an already improved mean image, further decreasing the weight assigned to corrupted measurements (see Table 5-1). Our preliminary tests show that after a small number

of iterations (less than 10) the algorithm converges (the change in individual weights across iterations becomes negligible).

Both approaches allow the regions over which the noise estimates are calculated to be specified. The user therefore has control over the desired area for optimisation. In the particular case of renal ASL, subtraction errors due to respiration can also occur in organs/regions other than the kidneys. As such, each weighted averaging algorithm was tested in two conditions: 1) including all voxels within the FOV, to assess whether more robust noise control is achieved when accounting for possible corrupted data outside the kidneys; and 2) restricting the processing to voxels within renal cortical regions of interest only, where the perfusion would be subsequently calculated.

Table 5-1 Pseudo-code 1: Iterative weighted averaging

```
For each iteration of "N_ITERATIONS"
    If (first iteration)
        % output of first iteration → unweighted mean image
        meanImage = mean(pwImages, 4);
        out{iteration} = meanImage;
    Else
        % mean image in iteration "n" → output of iteration "n-1"
        meanImage = out{iteration};
    EndIf
    For pwiImage of "N_MEASUREMENTS"
        Compute $\epsilon_{p_B}$;
        Compute $w_p$;
    EndFor
    % compute weighted mean image using weights $w_p$
    out{iteration+1} = mean(pwImages, 4, $w_p$);
EndFor
```

In summary, five methods for data rejection were tested (see Table 5-2):

Table 5-2 Summary of implemented data rejection / noise-reduction methods

Method #	Name	Noise estimate	Mask	No. of iterations
1	*rej*	---	---	---
2	*wMeanA*	ϵ_{p_A}	none	1
3	*wMeanA-mask*	ϵ_{p_A}	cortex	1
4	*wMeanB*	ϵ_{p_B}	none	10
5	*wMeanB-mask*	ϵ_{p_B}	cortex	10

5.2.4 Reducing the effects of motion corruption – quantitative evaluation

To quantitatively assess the effects of the motion/noise correction algorithms, four metrics were employed:

I. *Cortical signal-to-noise ratio (SNR_{ctx})*
- Ratio between the cortical signal intensity and the standard deviation of the signal intensity in a background ROI. The background ROI consisted of a region of 7×7×3 voxels without discernible perfusion signal.
- $SNR_{ctx} = \bar{I}_{ctx}/\sigma_{noise}$

II. *Cortical temporal standard deviation ($tSTD_{ctx}$)*
- Standard deviation of the signal intensity time series in cortical voxels
- $tSTD_{ctx} = \sigma_{ctx}^{t}$

III. *Cortical temporal SNR ($tSNR_{ctx}$)*
- Ratio between the cortical intensity and the standard deviation of the corresponding signal intensity time series
- $tSNR_{ctx} = \bar{I}_{ctx}/\sigma_{ctx}^{t}$

IV. *Image (volume) entropy (H_{vol})*
- Image entropy (as discussed in Section 3.5.2.2) computed throughout the entire image volume.
- $H_{vol} = -\sum_{b=1}^{B} p_b \, log_2 \, p_b$

Where:

- $\bar{I}_{ctx}$: perfusion-weighted (ΔM) signal in a cortical voxel
- σ_{noise}: standard deviation of the signal intensity in a background ROI
- σ_{ctx}^{t}: standard deviation of the cortical time series in a cortical voxel
- B: number of bins for entropy calculation (B = 256)
- p_b: normalised intensity histogram counts (probabilities) at bin b.

Note that for the evaluation of the weighted averaging methods, σ_{ctx}^{t} is calculated as the weighted standard deviation using the vector of weights w_p (whose sum equals one) obtained during the weighted averaging methods. MATLAB's implementation of the weighted standard deviation was used, which for a vector x with elements $x_1, x_2, \ldots, x_N$ and using weights $w_1, w_2, \ldots, w_N$ is defined as:

$$\sigma_x^w = \sqrt{\frac{\sum_{i=1}^{N} w_i (x_i - \bar{x})^2}{\sum_{i=1}^{N} w_i}}$$

Where $\bar{x}$ is the weighted mean of x.

All metrics except H_{vol} are computed in a voxel-wise basis. Results, where presented in numeric form, are expressed as mean±standard deviation considering all voxels within the cortical ROI.

Two-tailed paired t-tests were used to assess for statistically significant differences in each of the metrics after applying each motion correction method when compared to the "No correction" condition. The threshold for statistical significance was chosen to be 0.05.

5.3 Results

5.3.1 Perfusion quantification in the still condition

In the still condition, all ASL scans were free of motion artefacts (see Figure 5-3). Note that given the nature of the readout used during acquisition (3D), all slices are acquired simultaneously and therefore the level of motion corruption (and the effect of the motion suppression algorithms) is identical for all slices. For this reason, and for simplicity, only a central slice from each subject is shown when visually evaluating the effects of the motion suppression algorithms. Good image quality was obtained in a short nominal scan time (90 seconds) with clear contrast between renal cortex and medulla. Perfusion maps for subject 3 are shown in Figure 5-4, demonstrating good reproducibility of the technique. The mean cortical perfusion across all subjects/scans was 295 ± 97 ml/100 g/min (range: 245–343 ml/100 g/min), which agrees with the range of renal cortical perfusion values obtained in healthy volunteers in previous studies (see Section 2.3.6). Some degree of heterogeneity of the perfusion values can be see within different cortical regions in the kidneys, a finding which is repeatable in the case of subject #3 (see Figure 5-4). Intra-kidney variations in perfusion have been found in the literature with other MRI techniques such as DCE-MRI (Sourbron et al. 2008; Cutajar et al. 2014) and may reflect different metabolic demands according to the kidney region. Further reasons that may explain some of the variability in RBF include the fact that a single constant TI was used for all subjects. Even though care was taken in the choice of TI (1.2s is standard across single-TI renal ASL literature), an intravascular contribution in some subjects can be an inevitable consequence of choosing a single TI which may be suboptimal for some subjects. Furthermore it must also be acknowledged that partial volume effects are likely non-negligible given the chosen voxel size, which ultimately is tied to the trade-off of sensitivity (high SNR) and accuracy of the measurements (high resolution and thus minimisation of partial volume effects).
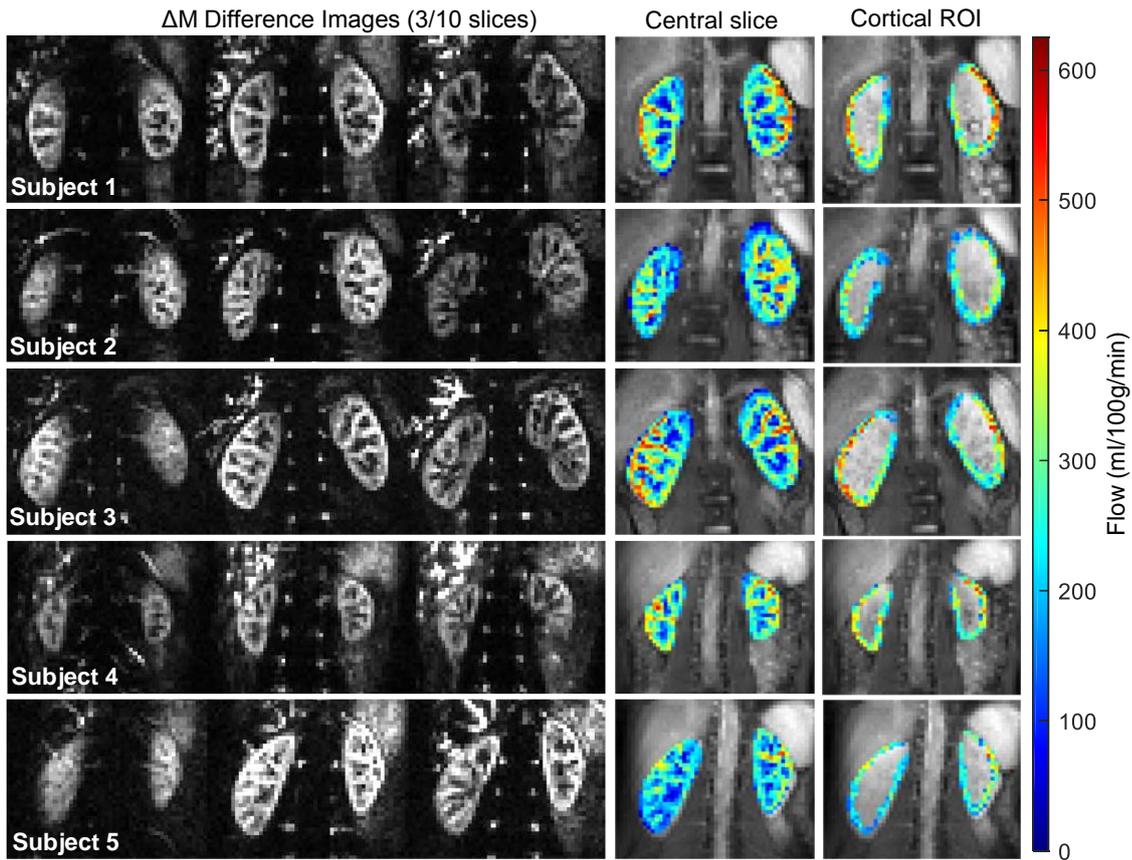
Figure 5-3 PWIs and perfusion maps for the 5 volunteers in this study. Whole-kidney and cortex perfusion maps are shown. For each subject, the perfusion maps shown on the right correspond to the central slice in the difference image set. The cortical ROI was obtained by manual segmentation.
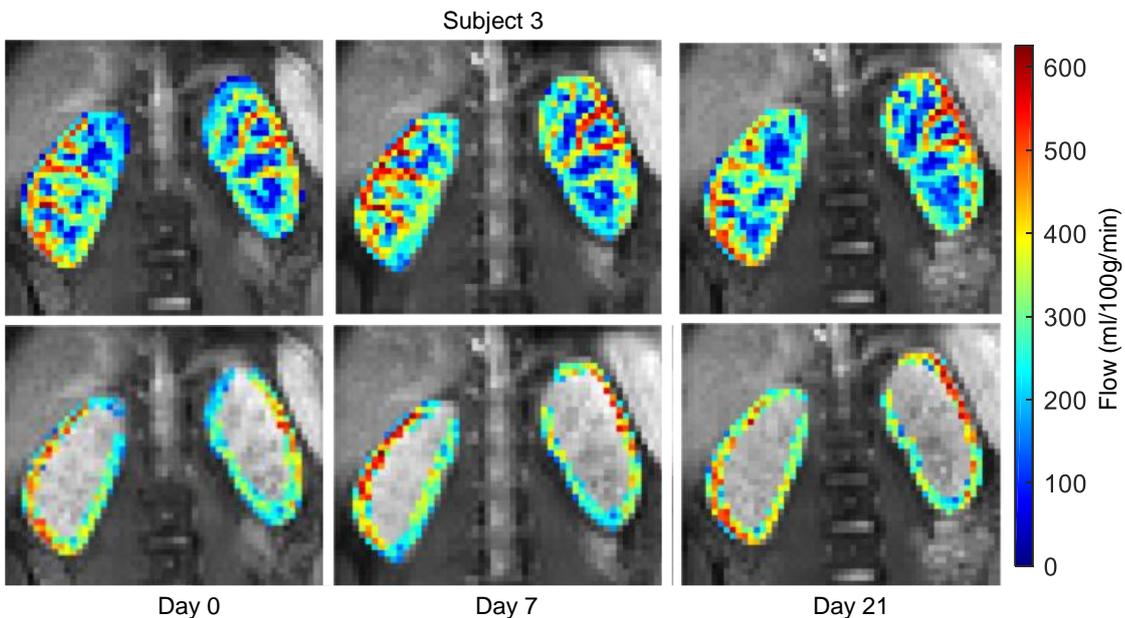
Figure 5-4 Perfusion maps from subject 3 from three scanning sessions. Top row: whole kidney; Bottom row: cortical region of interest. The mean cortical perfusion for days 0, 7 and 21 was respectively 305±100, 320±110 and 343±101 ml/100g/min.

5.3.2 Effect of the noise-reduction methods

5.3.2.1 Qualitative evaluation

All methods were implemented successfully and reduce the effect of the corrupted measurements in the final PWI, albeit to a different extent. The boundaries of the automatically defined motion-tracking regions (ROI in Figure 5-2 a) for the *rej* method are shown in Figure 5-5 for all subjects.
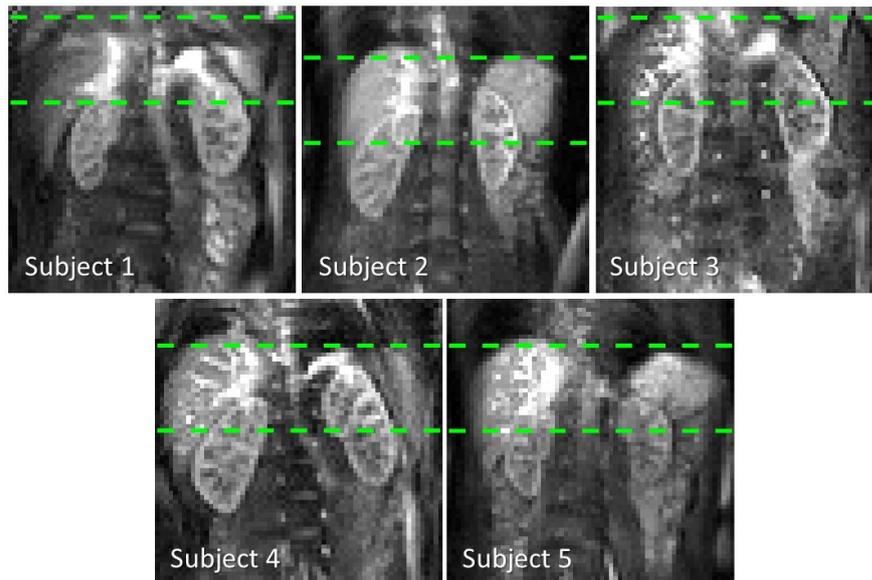
Figure 5-5 Boundaries of the motion-tracking region detected automatically in the first step of the data rejection algorithm (corresponding to step a) in Figure 5-2) overlaid on a central slice from the first control BS volume for each subject.

Figure 5-6 allows a qualitative evaluation of the result of applying the *rej* method. Rejection of corrupted data reduces blurring artefacts caused by the inclusion of motion-corrupted difference images in the averaging procedure. This increases the contrast between the renal cortex and medulla and will translate to a more accurate perfusion quantification. The algorithm is effective at preserving good data, thereby avoiding a SNR loss in motion-free datasets. This can be seen in subject #3, who remained motionless and whose regular respiratory cycle allowed for effective respiratory-triggering resulting in no motion artefacts. No data was rejected in this case. Subject #4 was asked to breathe irregularly for 2/3 of the scan to induce a greater amount of motion artefacts. Consequently, almost half of the measurements were rejected (7 out of 15). The decrease in the temporal standard deviation of the perfusion-weighted data is also evident, particularly in the more severely affected data from subjects #1 and #4, and as expected, not restricted to the kidneys. A visual analysis of the individual PWIs together with the measurement indices indicated for rejection by the *rej* method suggests that this method is very effective at rejecting severely corrupted data and in not rejecting motion-free data. These two factors are

particularly important if one wishes to combine this method with other motion correction approaches, such as image registration. In fact, the most important challenge faced by the *rej* method appears to be instances where data is not severely motion corrupted and fails to be rejected. However, these are the cases where one would expect an effective image registration algorithm to be able to correct the data.

Subject #	1	2	3	4	5
No. rejected pairs	4	6	0	7	3
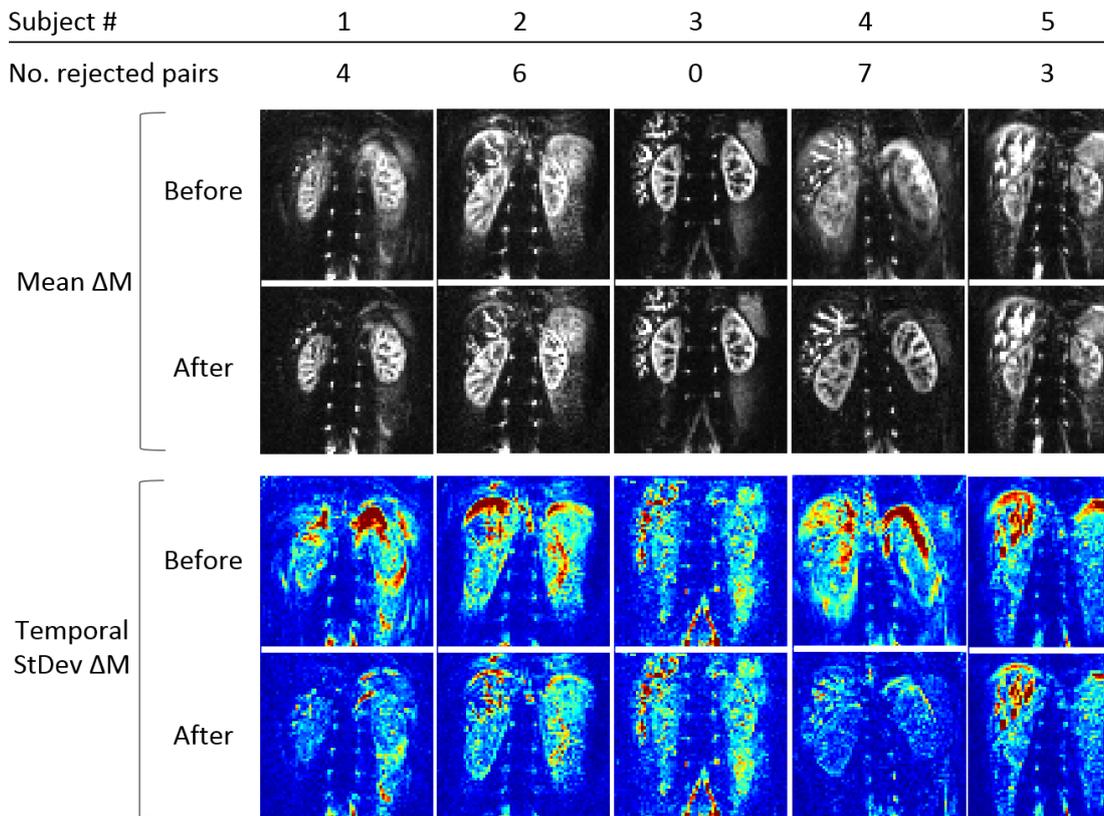
Figure 5-6 ASL PWIs and tSTD maps before and after automatic data rejection. Central slices are shown. The intensity display range is the same for each set of images of each subject. For the tSTD maps, the intensities range from 0 to the 99% percentile.

An instance where *rej* failed to reject one motion corrupted measurement can be seen in Figure 5-7. This figure shows the individual PWIs from subject #2, during the "motion-task" scan. As expected, the last 5 measurements different significantly from the first 10, showing prominent motion artefacts resulting from the kidneys (and adjacent organs) changing position from the control to the label condition. The table accompanying the figure shows the weights assigned to each measurement by methods *rej* and *wMeanB-mask*. As can be seen, measurement #13 (which clearly shows subtraction errors) was not rejected by *rej*. All the remaining motion corrupted measurements were rejected, and all of the motion free good quality data was preserved. Also in this figure, at the end of iteration #1 of the *wMeanB-mask* method, the last 5 corrupted measurements contribute to 8.3% of the final perfusion image. Running additional iterations of the algorithm appears beneficial, as after 3 iterations, this contribution is reduced

to 3.5%. Therefore, the contribution of the set of corrupted images to the final difference image is approximately one order of magnitude lower than in the standard unweighted averaging method (3.5% vs. $\approx$33% (5/15)).
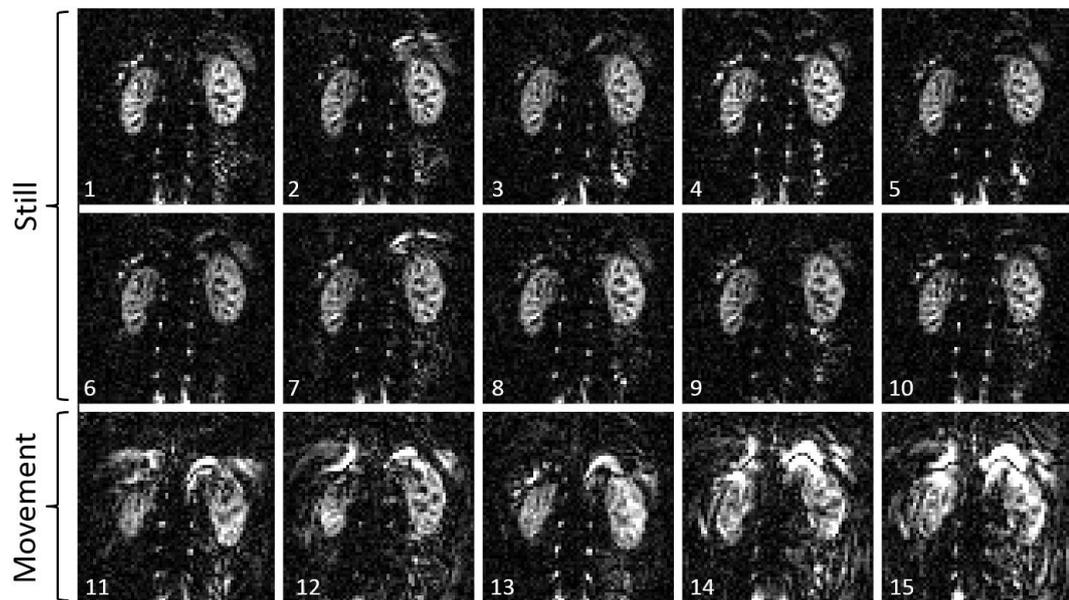
Method		#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12	#13	#14	#15
rej		0.091	0.091	0.091	0.091	0.091	0.091	0.091	0.091	0.091	0.091	0	0	0.091	0	0
	It #0	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067	0.067
wMeanB	It #1	0.075	0.097	0.125	0.089	0.105	0.091	0.076	0.084	0.107	0.068	0.021	0.022	0.02	0.013	0.007
mask	It #2	0.065	0.099	0.137	0.094	0.131	0.102	0.073	0.081	0.111	0.067	0.01	0.011	0.009	0.006	0.003
	It #3	0.062	0.096	0.138	0.095	0.14	0.106	0.073	0.077	0.109	0.067	0.009	0.01	0.008	0.005	0.003

Figure 5-7 Effect of noise reduction algorithms on the contribution of each PWI (before averaging) to the ASL averaging step (subject #2). Top: Individual PWIs; Table: Weights attributed to each individual PWI, by both the data rejection method and one of the weighted averaging methods (4 iterations shown, where the first corresponds to a standard unweighted averaging averaging). Both methods are able to reject, or greatly down-weigh the contribution of the later measurements in the time series, which are corrupted due to severe motion.

Central slices for each of the subjects before and after application of all methods are shown in Figure 5-8. The largest visual difference among the set of images is between the "No correction" condition and any of the motion suppression methods. Notable differences between the motion suppression methods can be seen in subject #4, where the method *wMeanA* (both with and without applying the cortical mask) shows the increased blurring which reduces the differentiation between cortex and medulla, when compared to *rej* and *wMeanB* (with and without mask). Although difficult to perceive visually, *wMeanB-mask* was the most effective method in reducing blurring in subject #5.

Figure 5-8 Final averaged PWIs (central slice) for all subjects (following application of each of the proposed motion/noise suppression methods).

5.3.2.2   Quantitative evaluation

The absolute values of all image quality metrics averaged across all subjects are shown in Table 5-3. In order for these metrics to better reflect the effect of the motion suppression methods on the actual perfusion measurements, only voxels in the cortex where considered during their calculation. The cortical masks used for each subject can be seen in Figure 5-9, overlaid onto a central slice from the averaged PWI dataset. The exception is H_{vol} which was calculated as a proxy for the image quality over the entire field of view (see Section 3.5.2.2 for more details on the use of entropy in image processing). The first metric being considered is SNR_{ctx}, calculated using the "background ROI" approach. The results from this metric are unexpected, as it decreases in all methods. Two possible explanations are 1) the fact that motion artefacts often cause the signal within the region of interest to increase (most obvious in subject 4) and 2) the reduced noise suppression in the background ROI resulting from the inclusion of a smaller number of measurements (or an increase in the variance in the weights assigned to each

measurement) in the averaging procedure. A more detailed discussion of this can be found in Section 5.3.3.

Table 5-3 Image quality metrics before/after each proposed noise suppression methods. For the first 3 metrics (SNR_{ctx}, $tSTD_{ctx}$ and $tSNR_{ctx}$) which are calculated on a voxelwise basis, first the mean and standard deviation of the metric values within the ROI are computed, for each subject. The values in the table correspond to the mean across the subjects of the mean and standard deviation within their corresponding ROIs. The values of H_{vol} correspond to the mean $\pm$ standard deviation across subjects of the entropy calculated across the entire image volume.

	SNR_{ctx}	$tSTD_{ctx} \times 10^8$	$tSNR_{ctx}$	H_{vol}
No correction	48.1$\pm$14.5	4.38$\pm$1.99	3.46$\pm$1.41	5.33$\pm$0.37
rej	43.4$\pm$13.8	3.43$\pm$1.42	4.29$\pm$1.94^b	4.86$\pm$0.17
wMeanA	42.7$\pm$13.3^{a,b}	3.71$\pm$1.50	3.95$\pm$1.66	5.02$\pm$0.23
wMeanA-mask	41.3$\pm$13.0^{a,b}	3.60$\pm$1.46	4.09$\pm$1.73	5.00$\pm$0.27
wMeanB	46.0$\pm$13.6	3.46$\pm$1.30	4.18$\pm$1.69^{a,b}	4.85$\pm$0.17
wMeanB-mask	42.4$\pm$14.0	3.14$\pm$1.09^{a,b}	4.50$\pm$1.81^{a,b}	4.83$\pm$0.21^b

^{a}statistically significant difference when compared to the "No correction" case
^{b}statistically significant difference when compared to the "No correction" case (excluding subject 3)
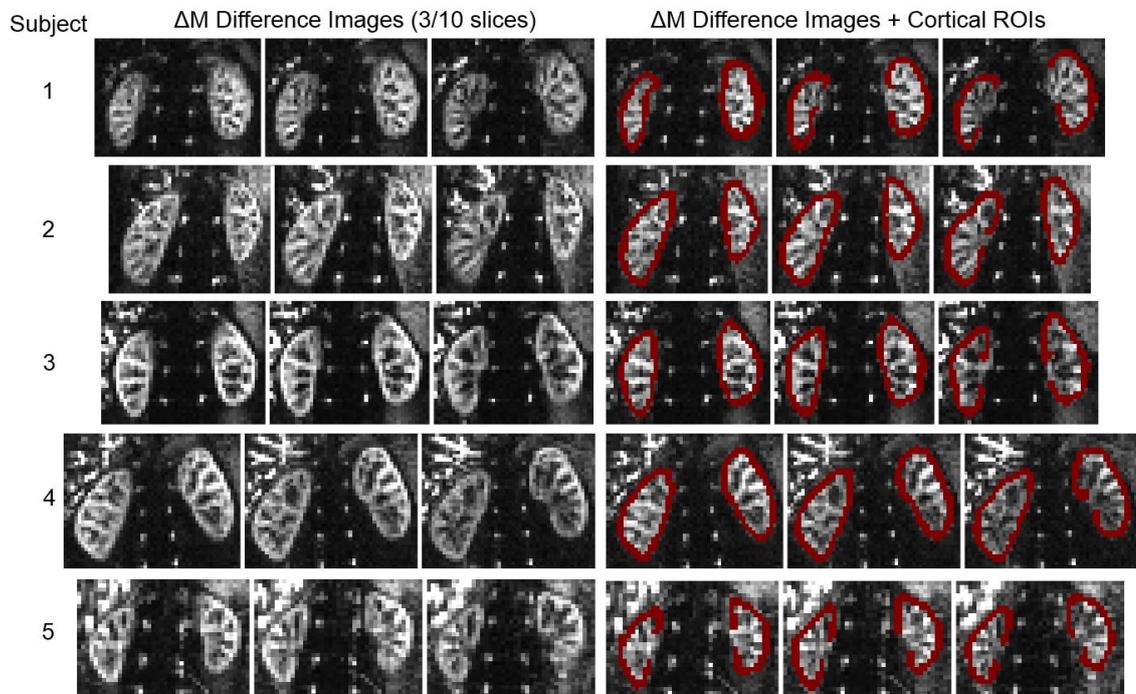
Figure 5-9 Visual depiction of the manually drawn cortical ROIs used for quantitative evaluation of the motion suppression methods. Averaged difference images shown after application of the *wMeanB-mask* method.

Assuming that the specific haemodynamics of each subject remains fairly constant over the few minutes the data takes to be acquired, in ideal data without motion and other spurious artefacts, the variation in signal among the different measurements should be minimal. Therefore, any motion correction algorithm should reduce this variation. A metric that directly reflects this is the temporal standard deviation calculated along the 4th dimension of the data (time, i.e. measurements), in this case calculated accounting for cortical voxels ($tSTD_{ctx}$). Examples of $tStd$ maps before and after applying the *rej* method can be seen in Figure 5-6. All methods were successful in decreasing this metric. The decrease was statistically significant in the case of *wMeanB-mask* even when including subject #3 in the analysis (whose data is practically motion-free). Not only the mean values of this metric decreased, but also the mean of the standard deviation of the perfusion weighted signal within the ROIs approximately halved from the "No correction" case to the *wMeanB-mask* case.

An alternative definition of the SNR is the temporal SNR (tSNR) over the cortical ROI ($tSNR_{ctx}$). Because in the definition of this metric the numerator is the mean of values within the cortical ROI, this metric is also sensitive to positive signal outliers (which similarly to the case of SNR_{ctx} may result in an increase value of the metric for motion corrupted data). However, an increase of $tSNR_{ctx}$ was achieved with all methods, a testament of the effectiveness in reducing the variability of the signal across the measurements. Statistically significant differences were found between the "No correction" condition and the *wMeanB* method (with and without masking) and between the "No correction" condition and the *rej* method when subject 3 (motion-free data) was excluded from the analysis. Voxelwise maps of $tSNR_{ctx}$ for each subject can be seen in Figure 5-10 for the "No correction", *rej*, and *wMeanB-mask* cases. The increase in $tSNR_{ctx}$ is greater as the level of motion corruption increases. Finally, the image (volume) entropy H_{vol} decreased in all methods, with a statistically significant decrease in the case of *wMeanB-mask* if subject 3 is not accounted in the analysis (as it is expected motion correction methods to have negligible effect in this data as it is free of motion).
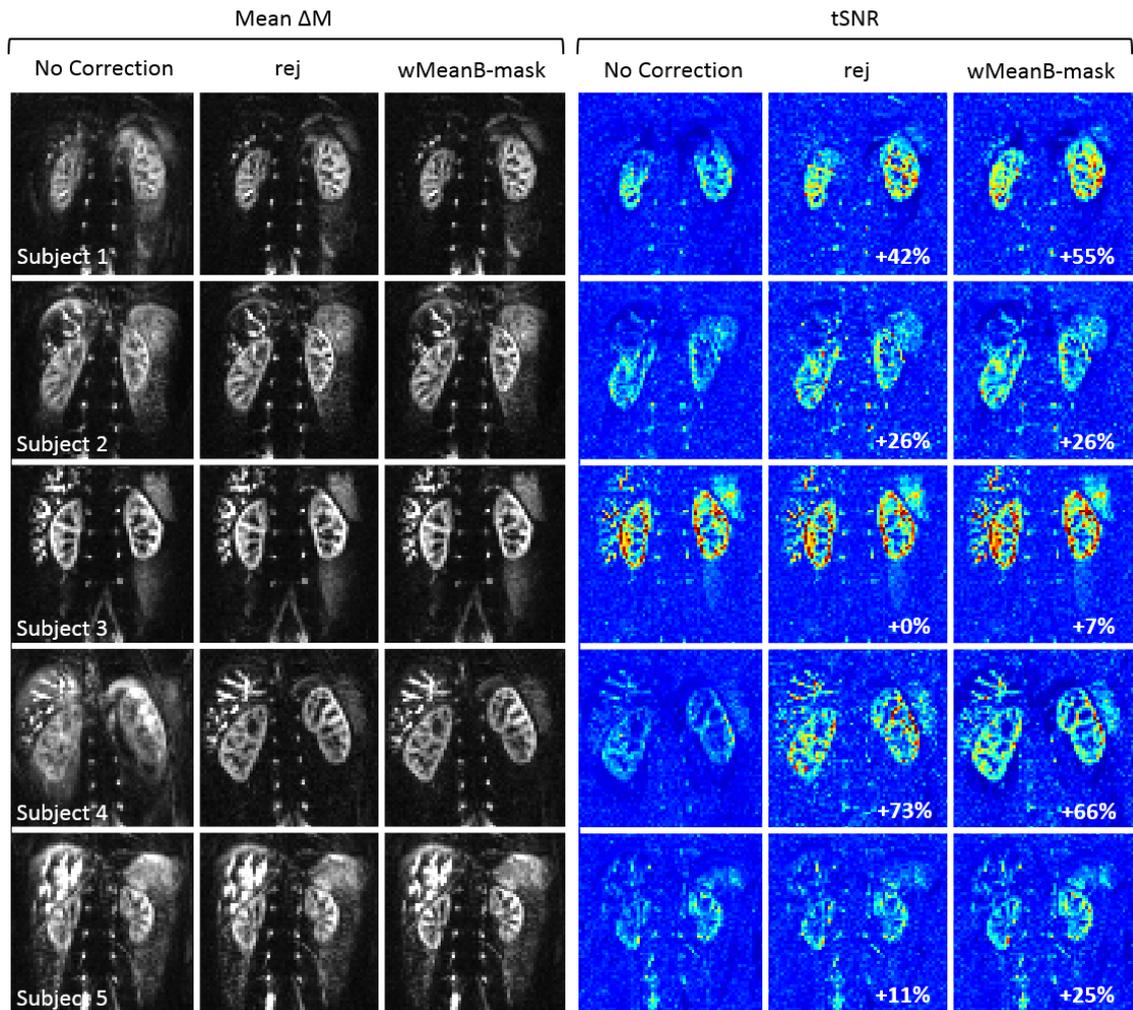
Figure 5-10 Averaged PWIs and tSNR maps before/after noise-correction. The data improves significantly after using both methods, with the degree of improvement being greater for highly motion corrupted data such as that from Subject #4. The numbers in tSNR maps correspond to the mean increase in this metric in the cortical ROI when compared to the "No correction" case.

The change in all metrics following the application of the data rejection and weighted averaging methods, with respect to the standard averaging case is summarised in the boxplots displayed in Figure 5-11, Figure 5-12, Figure 5-13 and Figure 5-14. Focusing on $tSNR_{ctx}$ and H_{vol} one might argue the best performing of the weighted averaging methods is *wMeanB-mask*, which is closely followed by the *rej* method. The "motion task" performed by the subjects was limited to changing the respiratory cycle. As such, bulk body motion is not expected, limiting the majority of the kidney motion to occur along the head-foot direction. This is arguably the best case scenario for the *rej* method, as this is the direction along which the reconstructed projections carry motion information.
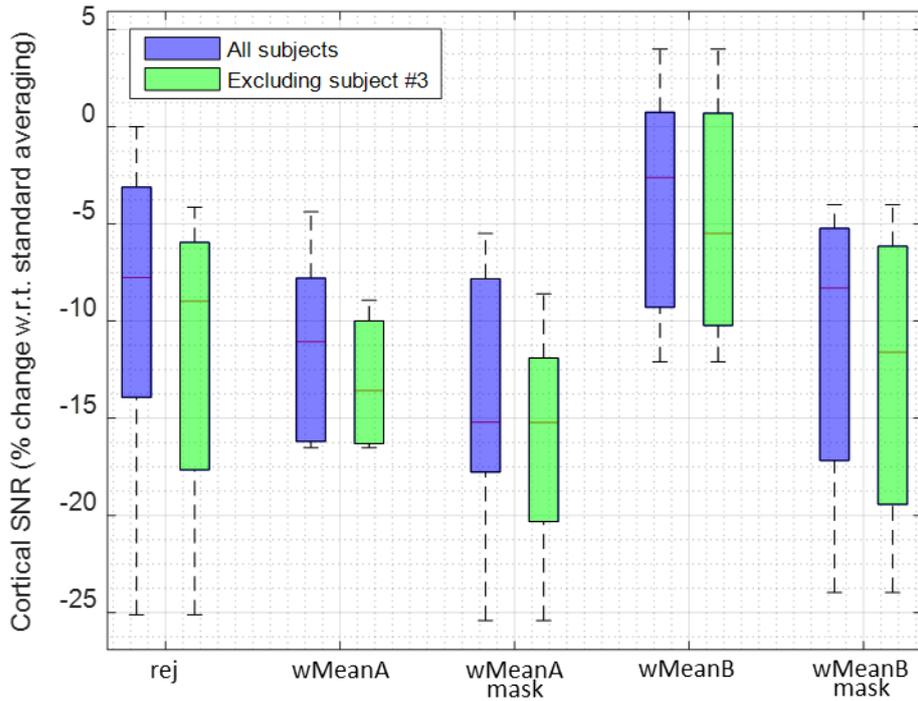
Figure 5-11 Change (%) in the cortical SNR of the motion corrected averaged PWI signal relative to the corresponding non motion corrected PWI signal. The central line corresponds to the median, the edges of the box to the 25% and 75% percentiles (q1 and q3, respectively).
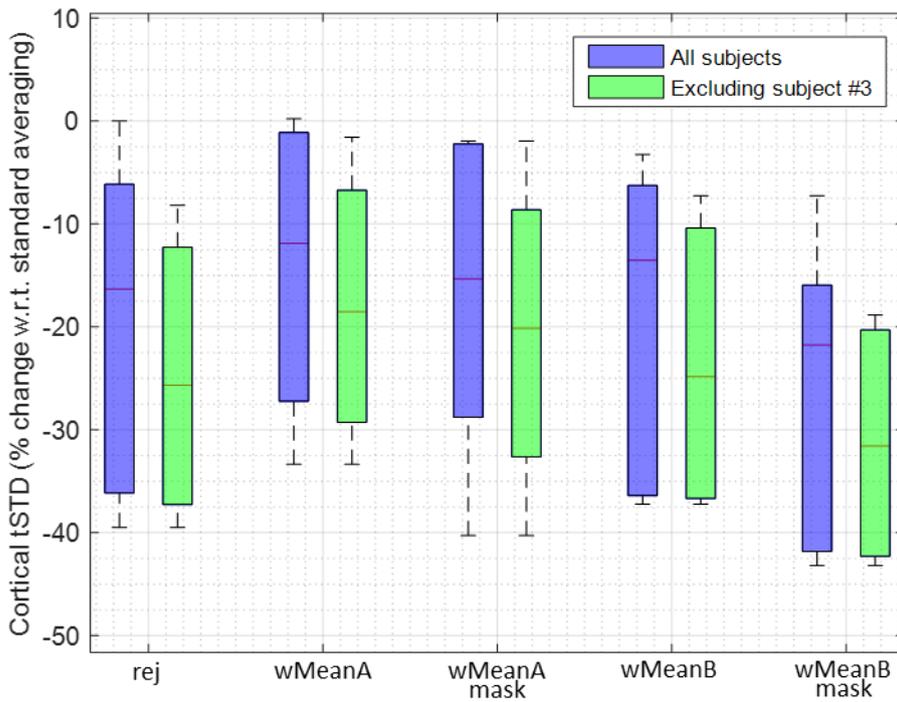
Figure 5-12 Change (%) in the cortical tSTD of the motion corrected averaged PWI signal relative to the corresponding non motion corrected PWI signal.
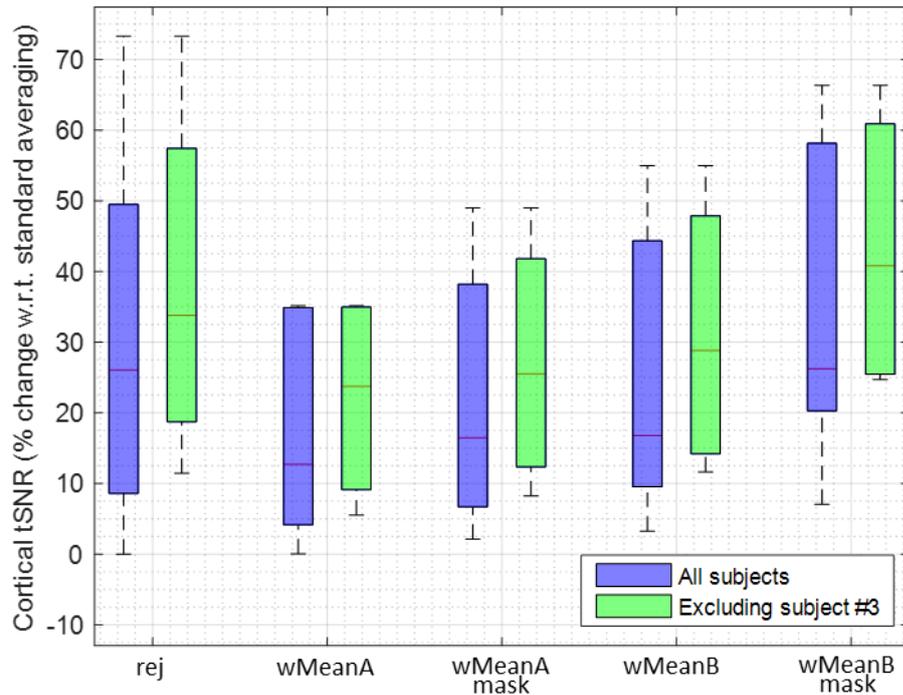
Figure 5-13 Change (%) in the cortical tSNR of the motion corrected averaged PWI signal relative to the corresponding non motion corrected PWI signal.
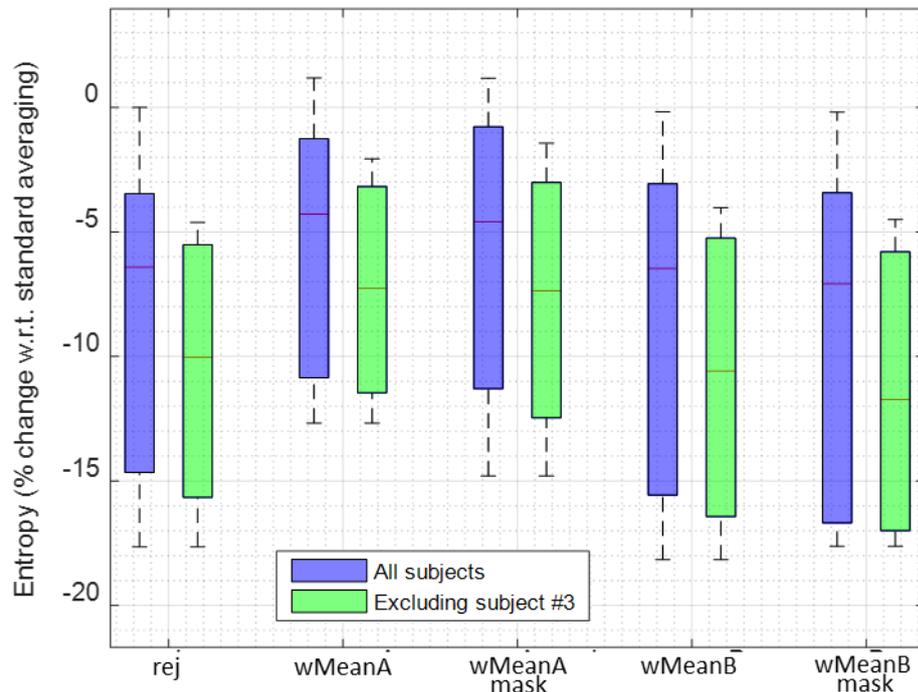
Figure 5-14 Change (%) in the entropy of the motion corrected averaged PWI signal relative to the corresponding non motion corrected PWI signal.

The performance of the *wMeanB-mask* method should be independent to the type of motion which occurred during the scan (and thus not limited to head-foot motion as *rej* is). The fact that

this method performed similarly to the *rej* method in this head-foot motion condition, together with its simplicity and ease of implementation across different ASL sequences makes it very attractive and arguably superior to *rej*, at least in the case of each method being used in isolation. However, given the differences in the motion detection framework of both methods and that they use different information, a combination of both methods may further improve the results. A possibility for this could be to use the *rej* method as a pre-processing step applied before the data averaging, to allow the first iteration of the weighted averaging methods to start from a less corrupted PWI. However, I believe the biggest advantage of the *rej* method is that it provides information on whether the perfusion-weighted data is corrupted independently for each measurement. This is unlike the weighted averaging methods, which for each measurement derive a noise estimate by using adjacent measurements or the complete time series. This advantage is what would allow *rej* to be implemented for real-time data rejection/reacquisition, and this should be the scenario in which *rej* is most effective, while the weighted-averaging methods would be used retrospectively for further improving the data quality.

5.3.3   Effect of asymmetric weighting on background-based SNR measures

The results obtained with the background-noise ROI-based SNR measure (SNR_{ctx}) were unexpected. One would expect that an increase in the data quality by rejection of corrupted measurements would increase the SNR of the data. Qualitatively, a visual inspection of the data following the motion/noise corruption methods clearly shows improvement in the final averaged PWI (reduction of blurring/edge artefacts and improvement in the cortico-medullary differentiation). However, the values from this metric decreased following the application of the motion/noise reduction methods. There can be several explanations for this. First, motion artefacts in the uncorrected data may result in artificially enhanced perfusion-weighted signal (positive signal outliers), which if occurring within the region of interest will positively bias the mean cortical signal, resulting in an increase of the SNR measure. The second reason is that all methods under examination will skew the weight of certain measurements during the averaging process which one may hypothesize reduces the noise suppression effect in regions of the background, specifically increasing the standard deviation of the signal in the voxels in this background region. This quantity corresponds to the denominator in the definition of the SNR_{ctx} metric. As such, an increase in this quantity results in a decrease of SNR_{ctx}. A basic simulation was implemented to test this hypothesis.

The simulation begins with the initialisation of a synthetic noise-free image (Shepp-Logan phantom; Figure 5-15 a)). To generate a set of measurements in an analogous way to the

multiple measurements (ASL pairs) obtained in an ASL experiment, Rician noise is simulated to make several copies of this image Rician distributed, each with a slightly different noise profile Figure 5-15 b)). This step was achieved by using a freely available MATLAB toolbox (Ridgway 2008). Then, each of this set of measurements is averaged multiple times, each with each measurement contributing to the final image in different proportions (different averaging weights). The SNR of each resulting averaged image is computed as the ratio of the mean of the voxels in ROI 1 and the standard deviation of the voxels in ROI 2 (see Figure 5-15 c)). This entire process is repeated a number of times specified by the user (see Table 5-4). For each run, only one weight set corresponds to the unweighted average condition (i.e. equal weights for each measurement). A plot of the resulting SNR values obtained for each set of weights against their variance is shown in Figure 5-16.

a)   b)   c)
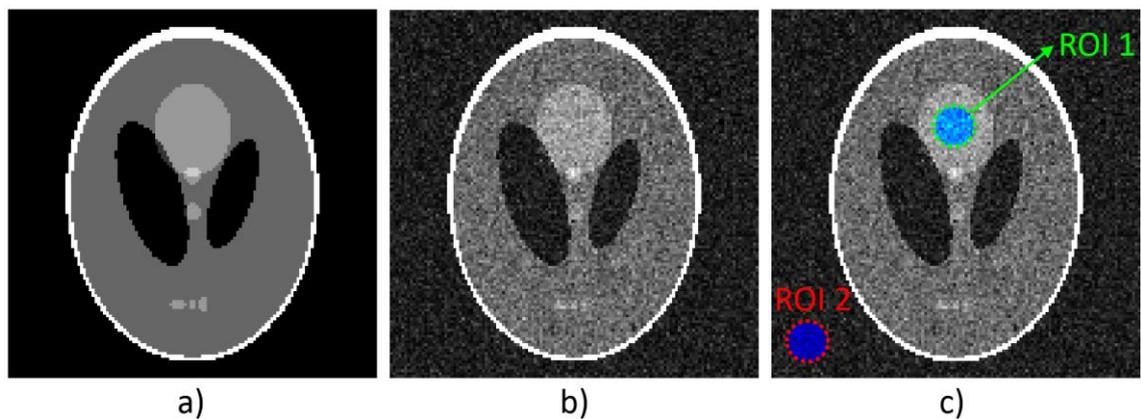
Figure 5-15 Synthetic noise-corrupted data and ROIs. a) Noise-free synthetic image; b) Rician noise-corrupted version of a); c) Highlight of the ROIs used for a simulation of the effect weighted averaging on background noise-based SNR measures – ROI 1: signal; ROI 2: background.

Table 5-4 Pseudo-code 2: Simulation of the effect of weighted averaging on SNR(calculated as the ratio between the mean signal intensity at the tissue ROI and the standard deviation of the signal intensity in a background ROI)

```
Initialise noise-free test image
Initialise SignalROI, NoiseROI
For each run of "N_RUNS"
    Generate "N_MEASUREMENTS" Rician noise-corrupted images
    For each weightSet of "N_WEIGHT_SETS"
        Perform weighted averaging of Rician noise-corrupted images*
        Compute SNR of averaged image (SNR=mean(SignalROI)/std(NoiseROI))
        Compute Variance of weightSet
    EndFor
    Plot SNR vs. Variance of weightSet
EndFor
```
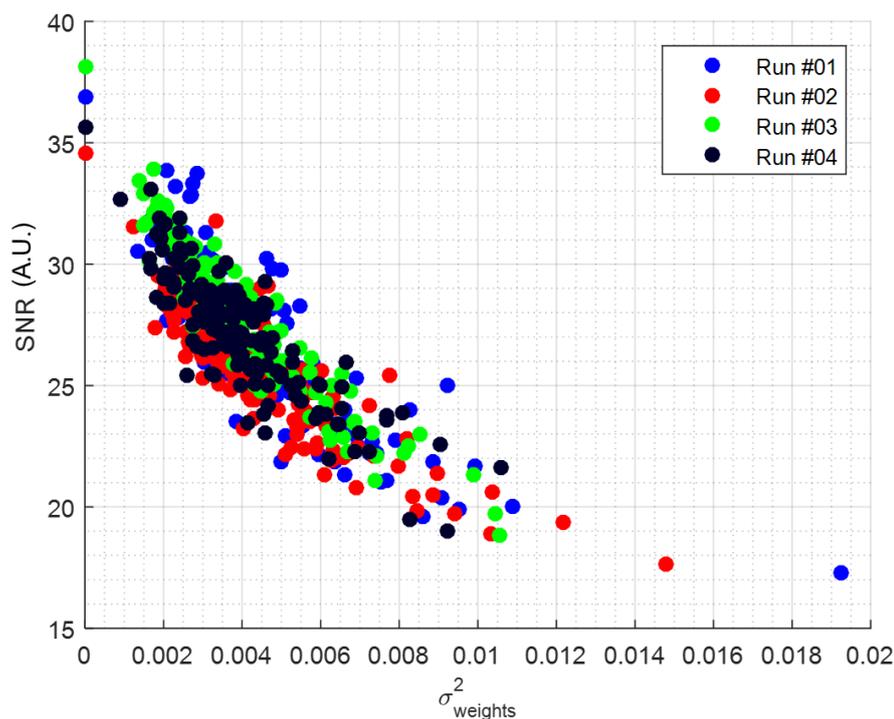
Figure 5-16 Effect of increased weight variance on SNR of the weighted averaged data (simulation results). SNR is maximum when all data is equally weighted due to a better suppression of noise.

As can be seen, for each of the runs of the experiment, the maximum SNR is obtained for a set of measurements whose variance is zero (i.e. unweighted averaging condition). Furthermore, the SNR tends to decrease as the variance of the weights increases. The decrease in SNR_{ctx} following the application of the proposed correction methods can likely be explained by this fact, together with the previously mentioned potential contribution of positive signal outliers resulting from motion artefacts. As such, the results of this simulation suggest that this particular SNR measure is not necessarily associated with a more reliable ASL measurement. Arguably more appropriate for this purpose is the $tSNR_{ctx}$, as it better reflects the actual improvement the implemented noise suppression methods have in the data of interest and how motion would affect the quantitative perfusion measurements derived from the ASL data. Another benefit of that approach (because only the cortical region of interest is taken into account) is that it is not biased by varying noise profiles across the FOV, owing to the different coil sensitivities throughout the FOV as an array coil was used for data acquisition.

Figure 5-16 was obtained by initialising the settings constants to:

- N_RUNS = 4;
- N_MEASUREMENTS = 15;

- `N_WEIGHT_SETS = 150;`

5.4 Conclusion

ASL is a unique MRI technique which uses blood water as a non-invasive freely diffusible tracer to quantify tissue perfusion. ASL benefits from 3D acquisition methods, which in the kidney have employed segmented readouts. These have provided high-quality data with good spatial resolution in healthy volunteers (Cutajar et al. 2012). The major drawback of these approaches is their high sensitivity to motion, potentially limiting their applicability in the clinic.

To improve the robustness to motion of our ASL technique, tests with single-shot data readouts were conducted in healthy volunteers. In this work, we have demonstrated the use of single-shot 3D GRASE in the kidney for the first time, and obtained perfusion values which are in accordance with previous renal ASL studies. This encourages further investigation and optimisation of single-shot ASL acquisition schemes to provide quantitative assessment of renal perfusion in clinically reasonable scan times.

Despite the use of single-shot acquisition schemes, ASL is inherently a motion-sensitive technique due to the fact that PWIs are obtained by subtraction of two image datasets. Retrospective sorting is commonly performed in ASL to reject measurements where spatial mismatches exist between the ASL labelled and control images (e.g. (Gardener & Francis 2010)). We implemented and tested several approaches to reduce the negative impact of corrupted measurements to the final averaged ASL PWIs. First, an automatic algorithm was developed to identify and reject corrupted measurements in renal ASL. This algorithm uses "navigator-like" information extracted from reference scans already embedded in the acquisition, so no pulse sequence modifications were required and no extra data needs to be acquired. Even though this algorithm was used for retrospective data rejection, its simplicity and low computation burden makes it amenable for implementation at the scanner. This would allow real-time rejection of highly corrupted data, potentially followed by reacquisition until a set number of motion free measurements are obtained to ensure sufficient SNR. Furthermore, weighted averaging schemes were implemented which down-weigh the contribution of corrupted measurements during the averaging operation. This is achieved by calculating pairwise noise estimates without resorting to thresholds or arbitrarily fixed constants. Rejection of corrupted data reduces blurring artefacts caused by the inclusion of motion-corrupted PWIs in the data averaging procedure. This increases the contrast between the renal cortex and medulla and should translate to a more accurate and repeatable perfusion quantification. A quantitative evaluation

of the quality of the perfusion-weighted data using several image quality metrics was performed which demonstrates the benefit of using these methods.

6 Optimised post-processing pipeline for renal 3D-GRASE ASL

Part of this work was presented at an International MR conference: (Nery et al. 2017a) and (Nery et al. 2017b).

6.1 Introduction

The previous chapter focused on the development of non image registration-based retrospective approaches to account for patient motion during ASL acquisitions. The rationale behind the strategies implemented in Chapter 5 is the assumption that in routine clinical scanning, a non-trivial amount of data will be severely corrupted beyond what registration-based retrospective motion correction methods will be able to correct. Hence, these data need to be discarded before the subsequent quantification stage. This is feasible because traditional ASL sequences typically acquire several measurements with identical scanning parameters for the purposes of signal averaging. On another hand, while less severe motion by the patient can still induce strong image artefacts (particularly in ASL as it relies on image subtraction), the data might still be amenable for correction, particularly if the underlying patient motion did not introduce k-space inconsistencies such as those described in Chapter 4. The literature available on the correction of these types of artefacts in ASL is very sparse, even for brain ASL which is likely to pose fewer challenges than abdominal ASL. In fact, an initial foray into the development of such methods where I implemented an autofocus-based approach for multi-shot motion correction in ASL can be seen in Appendix A. It is important to note that respiratory triggering is only able to prevent artefacts due to motion associated with the patients' respiration. To account for bulk body motion, other strategies are required. As such, a fundamental component of any state-of-the-art MR image processing pipeline is an image registration framework to correct for patient movement by retrospectively aligning the images (or organs of interest therein) acquired throughout the ASL experiment (Ashburner 2012; Jenkinson et al. 2012).

The number of steps and complexity of an ASL image registration pipeline depend on several assumptions related to the experimental protocol and the nature of the data itself. Even though in this chapter we will focus on single-TI ASL, many of the developments presented here are readily extendable to multi-TI acquisitions. First, the ASL time series needs to be realigned. This consists of images with potentially varying contrasts, as the acquisition typically alternates between the label and control condition. Before applying any model-based approach for quantification of blood flow, it is then necessary is to ensure all these data are in the same space as the ASL "calibration" scans. These might consist of a simple PD image (as was the case in Chapter 5) or an additional image series, which itself needs to be realigned. The latter is the case

if one wishes to obtain patient-specific T_1 maps, which is the approach undertaken in this chapter. An artefact-free voxelwise quantification of blood flow is only possible when all these image datasets are realigned.

Here, I implement a framework aiming to fulfil this requirement and thereby enable accurate mapping of Renal Blood Flow (RBF) using ASL. This chapter is divided in two parts. First, two approaches for registration of saturation recovery time series are implemented. Their effect is then assessed by comparing the repeatability of the T_1 estimates derived from these data as well as by quantifying the impact they have on the underlying saturation recovery fit quality. Using a saturation recovery sequence ensures that the longitudinal magnetisation state is consistent between TRs (i.e. nulled) regardless of the duration of the TR which is variable in respiratory triggered applications. The realignment of the ASL time series is then tackled by implementation of a customised image processing pipeline. Similarly, the effect of its application on the repeatability of RBF is evaluated. Furthermore, we also assess the effect of incorporating one of the noise-reduction approaches discussed in Chapter 5 in the processing pipeline. The impact of this approach, as well as image registration on the temporal standard deviation and temporal signal-to-noise ratio (tSNR) of the background-suppressed ASL PWI time series, is evaluated. The amount of renal translational motion which occurred during the acquisition of this time series is calculated based in the transformation matrices output from the image registration steps. Furthermore, this is correlated to the degree of improvement in the image quality metrics after applying the motion correction pipelines. All these experiments were performed using data from a cohort of paediatric patients with severe CKD (stages 3B-5).

6.2   Methods

6.2.1   Patient recruitment

Twelve children with CKD were recruited by myself from the chronic renal failure clinic at Great Ormond Street Hospital NHS Foundation Trust after being referred by a nurse or a consultant nephrologist. Ethical approval for this study had been obtained. The scanning procedure was explained to the parents and children, initially by myself and also through 2 movies (one for children and one for the parents). Of these, one of the patients was scanned once but found the experience of MR scanning intolerable due to the closed environment of the MR scanner, resulting in the patient being excluded from the study. The parents were given the choice to remain in the scanner room in the vicinity of the child. Eleven patients (age (years): 12±3 (mean ± std), range: 7-17, median 12 - 8 males) were successfully scanned on two different occasions for assessing the intra- and inter-session reproducibility of the RBF measurements using the

single-shot protocol described in the previous chapter. Scans were performed on a 1.5T Avanto scanner (Siemens Healthcare, Erlangen). The time between scans was: 23±10 (mean ± std), range: 7-35, median 21 days).

6.2.2   Image acquisition

MR data was acquired in a Siemens Avanto 1.5T scanner (Siemens Healthcare, Erlangen, Germany) running the Siemens NUMARIS/4 VB17A software version. The ASL acquisition consisted of the same protocol described in Chapter 5 (section 5.2.1). The exception was that in this case, twenty-five instead of fifteen ASL pairs were acquired. This modification was performed with the aim of allowing a further improvement in SNR, as we expected renal perfusion to be reduced in CKD patients. A larger number of measurements may also be beneficial in terms of motion robustness, insofar as the number of motion-free scans is likely to increase with the total number of measurements. The ASL scan was repeated up to 2 times in each session with a ≈40min gap between acquisitions. Importantly, the patient remained inside the scanner during this period, while other MR sequences were run. Scanning each patient two times in each MR session allowed and assessment of intra (i.e. within)-session reproducibility. Furthermore, as will be detailed in section 6.2.3.1, a saturation recovery (SR) dataset was separately acquired to allow for voxelwise T_1 mapping on a subject-by-subject basis. A SR sequence was used with 9 post-saturation delays (TD) (range: 100-2500ms; increment 300ms; saturation pulse spatially non-selective) (see Appendix B for an example of a T_1 map in one healthy volunteer obtained with this protocol). The imaging module consisted of a single-shot 3D GRASE readout with main parameters matched to the accompanying ASL imaging module: matrix size 64x64x10, voxel size 4.5x4.5x6.0mm, Partial Fourier (factor 3/4), TR/TE=3000ms/31.54ms. Data acquisition was repeated up to 3 times in each session (dependent on patient tolerance) which across all patients resulted in a total of 52 SR datasets (2 excluded due to extreme motion). A proton-density (PD) volume was also obtained with similar scanning parameters (without saturation pulse) to serve as reference for image registration and functional renal parenchyma ROI drawing during data analysis. Respiratory triggering was used to trigger the saturation pulse at end-expiration. Therefore, the scans were performed under free-breathing and no respiratory coaching was necessary. Even though the PD acquisition required no saturation pulse, a delay of 1200ms from the triggering instant to the data acquisition was implemented to match the delay corresponding to the inflow time in the ASL acquisitions. This was implemented as it was verified that if no delay was present, the acquisition of the PD image and the ASL data rarely took place in a similar part of the respiratory cycle which further increased the misregistration of the two datasets.

6.2.3 Processing pipeline

6.2.3.1 General remarks

The main theoretical and practical factors, as well as assumptions that guided the development of the ASL image processing pipeline described in this chapter, are as follows:

I. *T_1 maps required on a subject-by-subject basis:* In Chapter 5, one of the assumptions during RBF quantification was that all voxels considered part of the renal cortex were assigned a fixed T_1 value, obtained from a previous study (de Bazelaire et al. 2004). The rationale supporting this choice was that only data from healthy volunteers was used, and the range of variation of cortical T_1 in this group of subjects is expected to be relatively small, compared to patient cohorts (Huang et al. 2011). When scanning tissue with well characterised T_1 values (e.g. renal cortex in healthy volunteers), using a literature T_1 simplifies and shortens the image acquisition and post-processing stages, and may reduce image artefacts. This is because it eliminates the use of T_1 maps which themselves may have artefacts (resulting from misalignment of the underlying time series from which they are derived), which propagate to the RBF maps. Furthermore, issues of alignment of T_1 maps to the ASL/M_0 data are no longer a concern. For these reasons, this approach is common in many studies (Shimizu et al. 2017; Niles et al. 2016; Robson et al. 2016), among many others. However, T_1 is an established indicator of pathophysiological tissue status (Huang et al. 2011), and a voxelwise quantification of this quantity on an individual subject basis may be especially important in clinical studies as the underlying pathophysiology may cause T_1 values to vary across subjects. Indeed, recent studies have shown increases in T_1 with higher degrees of renal function impairment (as assessed by the estimated glomerular filtration rate (eGFR)) (Breidthardt et al. 2015; Gillis et al. 2016). With this in mind, a dedicated T_1 mapping scan (saturation recovery (SR)) was included in the imaging protocol (as discussed in section 6.2.2). This scan is itself sensitive to patient movement due to the need to collect images at multiple recovery times. Therefore a registration module dedicated to the realignment of these data was implemented in the image processing pipeline.

II. *Varying contrast throughout the T_1 mapping time series:* Given that these data are acquired in a way so that multiple points throughout the longitudinal relaxation time curve are sampled, the contrast of the underlying images changes as a function of the post-saturation delay (TD). As such, three strategies (including two registration

approaches) were evaluated for pre-processing the data before model fitting to obtain the T_1 maps:

A.   No motion correction;

B.   "Chain" registration;

C.   "Direct" registration;

In the "chain" registration approach, images were successively registered to the adjacent image in the SR time course, starting from the longest TD. This implementation aims to minimize contrast differences between the target and source image volumes throughout the recovery time series. Starting from the longest TD ensures the images with relatively higher SNR (due to longer T_1 recovery time) are used as reference for the multiple registration steps. In the "direct" registration approach, all images in the time series are directly registered to a fixed reference scan (chosen to be the PD acquisition). A schematic of both approaches is depicted in Figure 6-1.

$$[TD_n] \longrightarrow \boxed{r(PD)} \longrightarrow [rTD_n]$$

$$[TD_{n-1}] \longrightarrow \boxed{r(rTD_n)} \longrightarrow [rTD_{n-1}]$$

$$\vdots \qquad\qquad\qquad\qquad \vdots$$

$$[TD_2] \longrightarrow \boxed{r(rTD_3)} \longrightarrow [rTD_2]$$

$$[TD_1] \longrightarrow \boxed{r(rTD_2)} \longrightarrow [rTD_1]$$

a)

$$[TD_1,\ldots,TD_n]$$
$$\downarrow$$
$$\boxed{r(PD)}$$
$$\downarrow$$
$$[rTD_1,\ldots,rTD_n]$$

b)
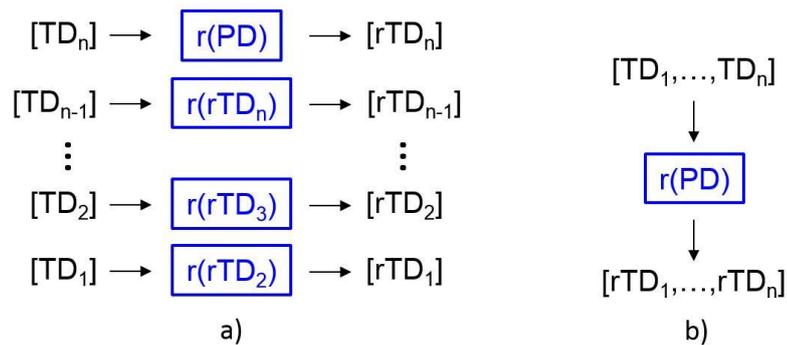
Figure 6-1 SR registration approaches. a) "chain" method: start by registering the longest TD in the time series to the PD image, then every other TD is registered to its closest volume in the time series in decreasing order of post-saturation delay time; b) "direct" method all TDs are directly registered to a reference PD image

III.   *Assumption of rigidity of the renal movement:* At this stage we focused on accounting for changes in position and orientation of the kidneys by performing a six degrees-of-freedom 3D rigid body registration. The main reason behind this assumption is that the background suppression (BS) module which is part of the ASL acquisition scheme used in this work significantly reduces the static tissue signal. The resulting lack of anatomical specificity in the unsubtracted ASL data renders highly complex transformations with thousands of degrees of freedom such as those required for non-rigid deformations unfeasible. Nevertheless, at the relatively coarse resolution we are imaging (4.5x4.5mm in plane), non-rigid behaviour such that caused by kidney pulsation should present less of a problem due to the fact that the renal capsule has a relatively rigid, inelastic nature

(Meyers et al. 2011) (particularly in CKD where the capsule is expected to be more fibrous and consequently even more rigid). Furthermore, no intra-subject registrations were performed (where it can be argued that non-linear transformations play a more important role).

IV.   *Separate registration of the left and right kidneys:* This is required due to the fact that the movement of both kidneys is not the same (Song et al. 2011; Siva et al. 2013). Therefore, since no single rigid transformation can simultaneously account for the displacements of both kidneys, rectangular masks were manually drawn in all slices to allow separate registration of left and right kidneys (see Figure 6-2 for an example of such masks). Several studies have also used similar approaches (Nathan S. Artz et al. 2011; Fenchel et al. 2006).
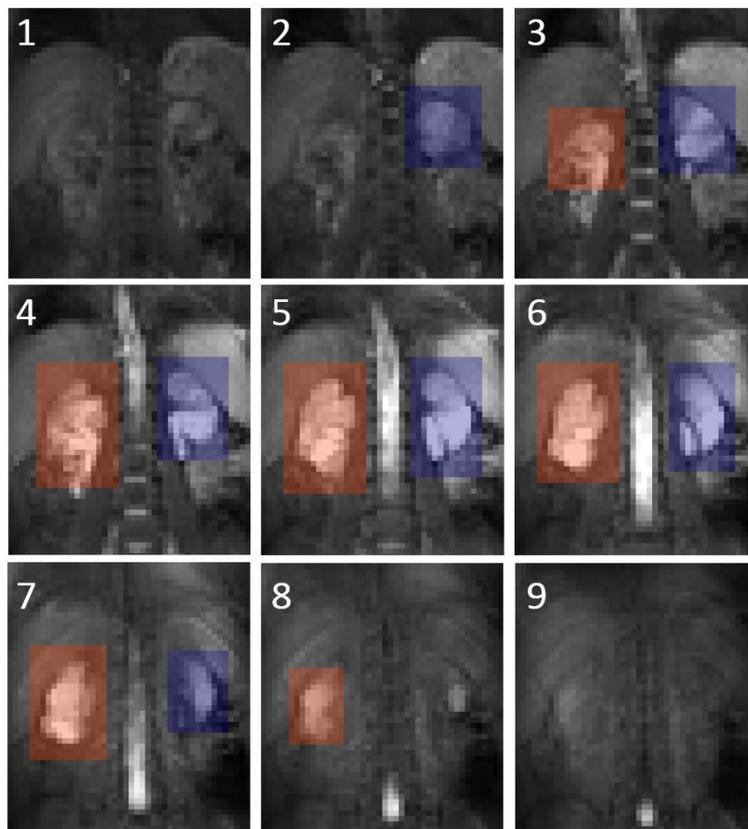
Figure 6-2 Left/right masks for separate registration of each kidney. The masks are overlaid in a reference proton-density volume from a child with CKD. Nine slices out of the ten acquired are shown. The numbers correspond to the slice indices, increasing towards the posterior part of the abdomen.

V.   *Proton-density image used as "absolute" reference for registration:* Even though a M_0 dataset can also be obtained from the same fitting operation from which the T_1 maps are obtained, these are not used in the quantification. This is because M_0 maps can be obtained from a single separate PD acquisition. As such, these M_0 maps should contain

less artefacts as they do not depend on the assumption of a completely aligned time series. In addition to being used for calibration of the ASL data, these M_0 maps are used as the "absolute" reference to which ultimately all data is aligned. That is, even though the intermediate registration steps in the pipeline use several types of images as reference, ultimately all T_1 maps and ASL data will be aligned to the PD-derived M_0 map. This accounts for the fact that even if the ASL time series is aligned, severe image artefacts will still propagate to the blood flow maps if the T_1/M_0 data is not in the same space as the ASL data.

VI. *Separate registration of control and label image series:* As mentioned in section 6.1, the contrast of adjacent images in the ASL time series changes if the label and control conditions alternate throughout the acquisition (which is the case for standard ASL pulse sequence implementations). Importantly, the absolute difference in signal between these conditions is amplified if BS is used (which is the case for all ASL acquisitions in this work). This factor is known to possibly affect registration results as these differences in contrast may be erroneously assumed to result from patient movement and consequently be minimised by the registration algorithm (Wang et al. 2008). Nevertheless, this is an often overlooked factor in the registration of ASL data, with few brain studies (and none in the kidneys) explicitly implementing this step (e.g. (Wang et al. 2008; Di Cataldo et al. 2011)). The proposed registration pipeline was designed so that only one control-label registration is performed per ASL dataset. The reason for why one registration step of this type was performed will be detailed below.

6.2.3.2 Implementation

All motion correction pipelines were developed and implemented by myself using MATLAB, using the open-source library *elastix* (Klein et al. 2010) for the registration tasks. Three-dimensional rigid body registrations were performed using masks for independent registration of both kidneys (as discussed above), employing a mutual information-based image similarity measure (see all registration settings in Appendix B). An *elastix* MATLAB wrapper was written to create a simple interface providing all customised options necessary for this application (e.g. output of single 3D volumes following registration with multiple masks, "chain" registration). In addition to the T_1 pre-processing strategies discussed in section 6.2.3.1 II, three ASL processing pipelines were evaluated:

A. No motion correction
B. Motion correction using image registration

C. Motion correction using image registration combined with a weighted averaging scheme (corresponding to *wMeanB-mask* introduced in Section 5.2.3.2).

Two factors are worth highlighting regarding the complete ASL processing pipelines:

- For B and C, the SR registration module was included into the pipeline using the approach deemed more robust following the "chain" vs. "direct" registration comparison.

- For C, the masks used for the weighted averaging method corresponded to the rectangular masks drawn for separate registration of the left and right kidneys (an example is shown in Figure 6-2)

Diagrams depicting pipelines B and C are shown in Figure 6-4 Note that in these diagrams, images or image series are referred to within brackets (i.e. the symbols "[" and "]"), whereas operations acting on images are enclosed in rectangular boxes. Both pipelines start by registering each individual control [SS_1,…,SS_n] and label [NS_1,…,NS_n] image series separately to the mean of all images with the same labelling condition ([μSS] and [μNS], respectively). The resulting registered individual control/label images ([rSS_1,…,rSS_n] and [rNS_1,…,rNS_n]) are then averaged for improved SNR. Both resulting images ([μrSS] and [μrNS]) may still not be in the same space and therefore a direct subtraction of the two might result in subtraction errors. As such, the label image ([μrNS]) is registered to the control image [μrSS] (resulting in [$r\mu rNS$]). Note this is a single registration step using the averaged images (higher SNR) and the control image is used as reference given it has more signal as the blood had not been labelled. Finally, the control image ([μrSS]) is registered to the reference PD image. This yields the "final" control image ([fSS]). Because the label image is in the same space as the control image, simply using the transformation from this step (T_{SS}) will bring the label image into alignment with the PD image, resulting in the "final" label image ([fNS]). This choice has to do with the significantly reduced tissue contrast in the label image (see Figure 6-3), which compromises a direct registration of this image to the PD image.
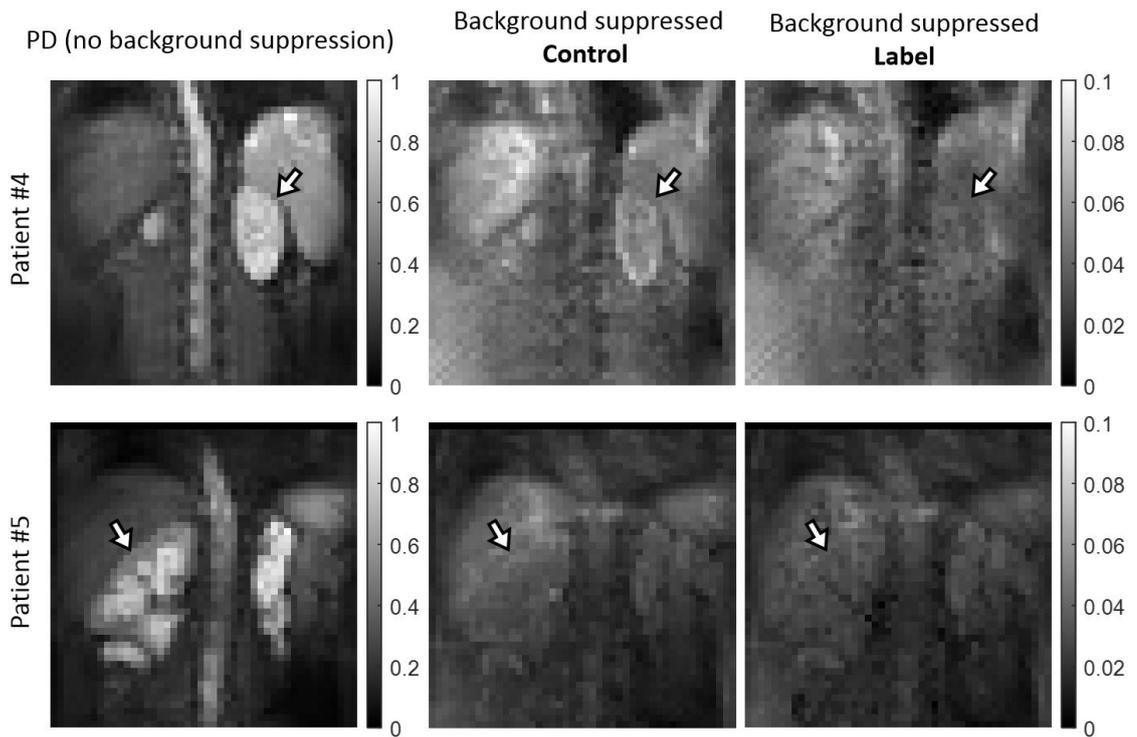
Figure 6-3 Reduced anatomical specificity in background-suppressed data. The loss of contrast is particularly pronounced in the label images, where the renal tissue (arrows) is almost indistinguishable from the adjacent structures. The control and label images are shown in the same intensity range.

The difficulty of performing accurate image registration on background-suppressed data has been highlighted in several ASL studies (e.g. (Cutajar et al. 2014; Alsop et al. 2015), among others). Note that each image volume in the SR recovery time series was also registered to the same reference PD image for ensuring alignment between the T_1 maps and the perfusion-weighted datasets (not shown in Figure 6-4).
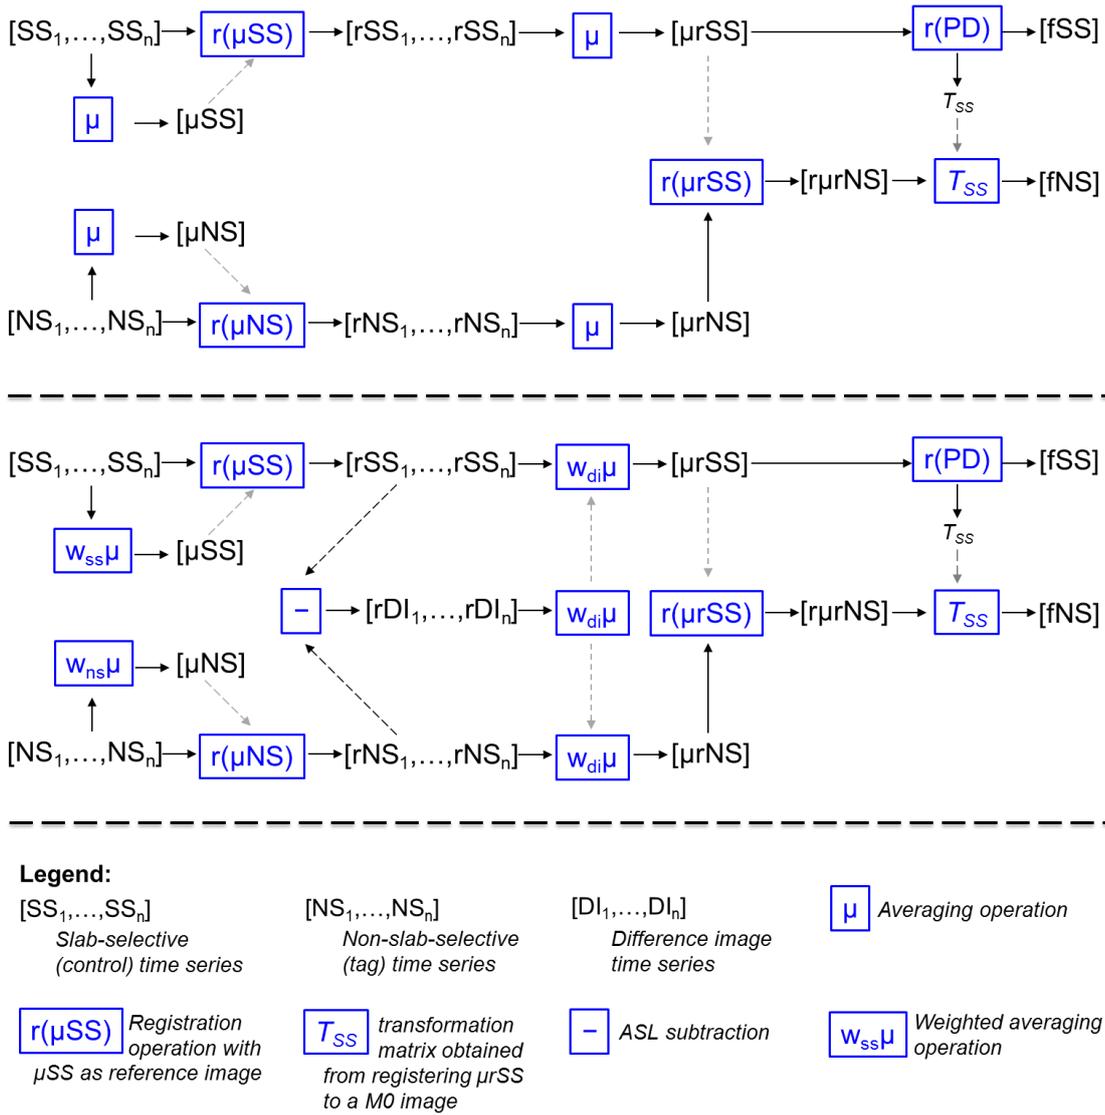
Figure 6-4 Motion correction pipelines for ASL data. Top: Processing pipeline B; Centre: Processing pipeline C (note that processing pipeline A corresponds to the "no registration" case and therefore is not shown in this figure. The main difference between the two is that in processing pipeline C, all image averaging steps were performed assigning weights to each of the individual images, as implemented in section 5.2.3.2 (smaller weights for individual images which deviate more from the unweighted average of the scans).

6.2.4 Image Analysis

All analysis software was implemented by myself using MATLAB. Voxelwise fits of the saturation recovery data to the Bloch equation describing the recovery of the longitudinal magnetisation (Equation 6-1) were performed to obtain T_1 and M_0 maps (note that the M_0 maps from this fit were not used for RBF quantification, for reasons described above). A three-parameter fit was performed using the following expression:

$$M_z(t) = M_0[1 - (1 - \cos(\theta))e^{-t/T_1}] \qquad \text{Equation 6-1}$$

ASL RBF maps were derived from the single-TI data using the single-compartment modelling approach described in the previous chapter (Section 5.2.2). The major difference is that in this chapter (for reasons detailed above), rather than using fixed T_1 and M_0 values throughout the entire ROI, voxelwise T_1 and M_0 values were used instead.

Regions-of-interest (ROIs) were manually drawn on the PD images for calculating ROI-based statistics related to the quantitative MR measures. The aim was to select all voxels deemed part of the functional renal parenchyma, with care not to include high intensity regions (corresponding to a dilated collecting system or cysts) in the ROIs. Several studies (e.g. (Robson et al. 2016; Shimizu et al. 2017), among others) use perfusion information in the difference images to identify and segment cortical regions of interest. This is arguably a valid assumption in healthy volunteers as the high perfusion results in strong signal in the cortical region. However, we believe this will likely not hold in certain patient groups as the cortical perfusion (and thus signal) is markedly reduced. Therefore, we opted by using the PD images (used as reference M_0 maps) for drawing our ROIs given their higher amount of anatomical information. This should avoid biases related to drawing ROIs in PWIs in impaired kidneys, such as failing to include anatomically present cortex which is undetectable in the PWIs due to reduced perfusion. This is another reason why image registration is paramount for these datasets. Having a multi-slice dataset proved crucial to allow tracking of the cysts across the volume to assist in the boundary definition.

6.2.5   Evaluating the impact of motion correction

Validation of the motion correction approaches was performed in two steps: separately for the registration of the SR data and for the complete pipeline integrating SR and ASL motion correction methods.

6.2.5.1   Saturation recovery

The impact of the registration algorithms on the intra- and inter-session repeatability of the T_1 estimates was evaluated using the intra-class correlation coefficient (ICC) and the within-subject coefficient of variation (WSCV). Furthermore, the quality of the underlying SR fits obtained following the different motion correction approaches was assessed employing the root mean squared error (RMSE) as a goodness of fit metric. Both M_0 and T_1 maps obtained for all patients in this work are shown both in the "no motion correction" condition and also after employing image registration to align each image in the time series. The quality of the M_0 maps is also compared to that of the reference PD acquisition which is relatively insensitive to motion. Therefore, success by the image registration algorithms in aligning the SR time series can be shown by an increased similarity between the M_0 maps and the reference PD images.

6.2.5.2   ASL

Repeatability of RBF estimates was evaluated as assessed by the intra-class correlation coefficient (ICC) and the within-subject coefficient of variation (WSCV). The effect of registering the individual control/tag images (before averaging) was evaluated by computing the temporal standard deviation (tSTD) and tSNR of the ASL perfusion-weighted signal (metrics defined in Chapter 5 (Section 5.2.4)) in the functional renal parenchyma ROIs. As previously in this work, the main assumption that allows these metrics to relate to the amount of movement in the ASL datasets is that no significant change in the renal haemodynamics takes place during the time required for acquisition of the ASL data sets ($\approx$ 3 minutes, accounting for delays associated with respiratory triggering). As such, noise, and especially motion are considered the major source of variability in the perfusion-weighted signal across ASL pairs.

To summarise the "amount of motion" which took place during the acquisition of the ASL data for each patient (i.e. background suppressed control and label time series), a total translational motion index was calculated as the square root of the sum of the squared displacements relative to the reference image, obtained from the transformation matrices output from the registration algorithm. This index will be referred to as Euclidean distance (ED), and has been previously used for this purpose (e.g. in (Studler et al. 2010)). This metric is used as the most significant type of movement by the kidneys is translational in nature, caused by respiration which predominantly induces superior to inferior displacements. Correlation analyses were performed the between the improvement in the image quality metrics and the amount of motion in the corresponding scans, as well as the age of the subjects.

6.2.5.3   Proton-density/perfusion-weighted (PD/PW) hybrid images (PW-PDI's)

The PD and ASL acquisitions share the same scanning parameters and image readout (with the exception of the labelling and BS modules which are only used for the latter). Therefore, the ASL information can be overlaid directly onto the PD image to generate a "hybrid" proton-density-perfusion-weighted image (PD-PWI). The resulting red, green and blue (RGB) image displays a mixture of anatomical (PD) and functional (PWI) information, respectively represented by linear grayscale and a colour scheme (specified by its associated colormap). As such, in way analogous to established multi-modal hybrid techniques such as PET/CT and PET/MRI and other MR approaches such as fMRI/structural MRI, this allows for a superior spatial localization of the information contained in the functional measurements. In the case of ASL, this may be particularly useful for the evaluation of areas of heterogeneous or abnormal perfusion. Furthermore, this also allows for an efficient way of visually assessing the alignment between

the underlying PD and PWI datasets. Hybrid images of this type are shown for each patient, using PWI datasets both before and after motion correction. To obtain the PD-PWI image, a simple voxelwise approach was used. A sum of the intensities of each image (each independently normalised to the range $[0 - 1]$) is performed, weighting this sum with two factors α and β, which refer, respectively, to the contribution of the PWI and the PD information in the final image. The following expressions describe this operation for a given pixel of coordinates (x,y):

Table 6-1 Pseudo-code 3: Fusion algorithm for a given voxel (x,y)

hybridSlice(x,y,1) = α .× pwiSlice(x,y,1) .× map(c,1) + β .× pdSlice(x,y,1) ← R channel

hybridSlice(x,y,2) = α .× pwiSlice(x,y,2) .× map(c,2) + β .× pdSlice(x,y,2) ← G channel

hybridSlice(x,y,3) = α .× pwiSlice(x,y,3) .× map(c,3) + β .× pdSlice(x,y,3) ← B channel

where map is a $C \times 3$ colormap matrix (C is the number of colors in the colormap, each represented by a RGB triplet) and c is the index in the colormap which specifies the mapping of the weighted contribution of the intensity in the PWI (pwiSlice) to its corresponding RGB triplet in the colormap. Furthermore, hybridSlice is the resulting $X \times Y \times 3$ matrix corresponding to the hybrid PD-PWI image, where X and Y are respectively the number of pixels along each direction of the slice (i.e. 2D image). The optimal choice of values for α and β will be explored later in this chapter.

6.3   Results

6.3.1   General remarks on image registration

Visual assessment of the registration results is an integral part of the development and validation of registration algorithms (Hill et al. 2001). In addition to the criteria established in section 6.2.3.1, the development of the image registration pipeline (as well as registration parameters) was guided by a visual assessment of the results at each registration step in the pipeline. Preliminary tests suggested that the greatly increased number of degrees of freedom offered by nonlinear registration severely decreased our confidence in the registration results, as non-linear registration applied directly to the background-suppressed control/label images resulted in unacceptable deformations. This can likely be attributed to the effect of BS which reduces the static tissue signal and thus the lack of anatomical detail rendered it unfeasible to obtain an accurate deformation field.

6.3.2   Post-processing pipeline – T$_1$ mapping

Figure 6-6 and Figure 6-7 provide an overview of the results of T_1/M_0 mapping in a single slice across the centre of the kidneys for all patients enrolled in this study (each row in the figures

corresponds to data from a single patient). These were derived from the SR dataset, both with and without motion correction before the model fitting stage. All scans in these figures correspond to the first run of the first day for all patients. The three first columns (*A*, *B* and *C*) show M_0 maps. *A* was obtained by a single PD acquisition with readout matching the main ASL scan but without any labelling or BS modules. *B* and *C* were the output of SR fitting, respectively without and with motion correction ("*direct*" registration). The improvement in image quality of the M_0 maps in column *C* (i.e. after registration) is evident when compared to the equivalent maps obtained without motion correction (i.e. column *B*). This improvement can be seen as a marked reduction of artefacts close to the boundaries of the kidneys. Because the set of images in column *A* is obtained in a single fast acquisition (readout duration of less than 300ms) without need for any fitting, it is largely immune to motion artefacts. It can therefore serve as the gold standard to which the M_0 maps obtained by SR fitting can be compared. The similarity between this set of PD images (column *A*) and M_0 maps in column *C* is a testament to the effectiveness of the "*direct*" registration approach in compensating for the kidney displacements throughout the SR time series, despite the intrinsic contrast changes due to the varying post-saturation delays.

Columns *D* and *E* show T_1 maps, obtained respectively before and after application of the "*direct*" registration pipeline, overlaid onto the PD image shown in column *A*. Note that the pixels depicting T_1 values are restricted to rectangular boxes, which were those used for independent left/right kidney registration. The same T_1 information is shown, respectively in columns *F* and *G*, but restricted to the voxels assigned to the functional renal parenchyma ROIs, for a better appreciation of the variation in T_1 values in the regions from which the mean T_1 estimates were obtained. As expected, when uncompensated for, the renal movement gives origin to artefacts, especially prominent close to the boundaries of the kidneys. Image registration is capable of suppressing these in most cases as can be seen by comparing the images from column *D* and *E*. After registration, the T_1 maps are also generally smoother and less noisy, as the number of voxels with large spurious variations in T_1 is reduced. At the same time, the maps show better anatomical specificity, preserving the variations in T_1 values throughout different types of tissue, not only when comparing the kidneys to adjacent organs (e.g. liver) but also within the kidney, especially when lesions are visible. This is especially visible in patients #1, #5 and #10, where the kidney dilations visible as hyperintensities in the PD image are characterised by significantly higher T_1 values. Columns *F* and *G* demonstrate how the use of image registration ensures artefacts are suppressed in the cortical regions. This is crucial for ROI-based quantification as many of the most severe artefacts in the T_1 maps occur close to the

boundaries of the kidneys, which are precisely the regions from which the mean T_1 estimates are derived (see artefacts highlighted by the arrows in column *F* in both Figure 6-6 and Figure 6-7). Without motion correction, these artefacts result in severely inaccurate T_1 values that are then used for perfusion quantification, which under these circumstances yield erroneous RBF maps. Despite the success of the proposed motion correction approach, variations in the T_1 estimates across the different subjects still remain. This can likely be attributed to the fact that the patient cohort under examination is very heterogeneous with respect to the underlying pathophysiology (see Chapter 7).

A summary of the quantitative T_1 values obtained across all scans of all patients (total number of scans = 50) with approaches A (no registration), B ("*chain*" registration) and C ("*direct*" registration) can be seen in Table 6-2.

Table 6-2 Cortical T_1 descriptive statistics before/after SR motion correction prior to T_1 mapping using the SR time series. Data from entire cohort (all patients/days/runs) (all values in seconds)

$n_{\text{SR time series}} = 50$	A no registration	B "chain" registration	C "direct" registration
Mean	1.89	1.70	1.65
Standard deviation	0.29	0.34	0.21
Maximum	3.02	3.55	2.14
Minimum	1.40	1.24	1.26

Note that the maximum mean T_1 value obtained following approach B was 3.55s, and corresponded to the scan where this approach failed (see below). Excluding this outlier, the range of T_1 values using approach B was 1.24−2.13s, similar to that obtained using approach C. Statistically significant differences were found between the mean T_1 values (averaged across the "functional renal parenchyma" ROI), in all patients (all runs/days) when comparing approaches A vs. B (two-tailed paired t-test, p<0.001) and A vs. C (two-tailed paired t-test p<0.00001) but not when comparing the approaches B vs. C (two-tailed paired t-test p=0.072). This suggests that in the presence of uncorrected renal motion the T_1 values suffer from a significant positive bias. This can be problematic especially as increases in T_1 have been associated with presence of parenchymal damage (Breidthardt et al. 2015).

Both registration approaches improved the repeatability of the scans, with the "*direct*" registration method outperforming the "*chain*" method according to all repeatability measures (see Table 6-3). Propagation of registration errors in the "*chain*" method introduced severe

artefacts in the T_1 map of one SR time series (approximately doubling the mean T_1 in this scan). Therefore, repeatability indexes were computed for the full dataset and excluding this scan (see Table 6-3). Both registration methods are similarly effective except for the case where the "*chain*" approach failed resulting in a highly corrupted T_1 map.

Table 6-3 Intra/inter-session repeatability measures for the mean cortical T_1 estimates. WSCV: within-subject coefficient of variation; ICC: intra-class correlation coefficient. Propagation of registration errors in the "chain" method introduced severe artefacts in the T_1 map of one SR time series. "Chain" registration repeatability measures computed with and without inclusion of this outlier.

	Intra-session		
	A - no registration	**B - registration (chain)**	**C - registration (direct)**
WSCV	0.123	0.182	0.065
WSCV (excluding outlier)		0.080	
ICC	0.085	0.331	0.750
ICC (excluding outlier)		0.632	
	Inter-session		
	A - no registration	**B - registration (chain)**	**C - registration (direct)**
WSCV	0.130	0.162	0.074
WSCV (excluding outlier)		0.064	
ICC	0.254	0.291	0.614
ICC (excluding outlier)		0.692	

Using both proposed registration approaches, the mean RMSE of the SR fits was found to be significantly lower throughout all patient/runs compared to the "no motion correction" approach (Figure 6-5 a)) (paired t-test p<1e-10). Interestingly, the RMSE decreased in the later runs. A likely explanation for this is that the fact that the only patients able to tolerate 3 runs were probably the more compliant ones, and as such their respective SR time series should be less affected by motion. There was no significant difference between the mean RMSE obtained the two registration approaches (paired t-test p=0.47). Further evidence of this fact can be seen in the histogram in Figure 6-5 b), obtained by pooling the value of the RMSE of all ROI voxels of all patients/days/runs.

These results indicate a comparable performance by both registration methods but that the "*chain*" approach may be more susceptible to artefacts caused by the propagation of errors throughout the registration of the time series due to the fact that the reference image for each registration step is the result of the registration that preceded it (except for the longest TD image). On the other hand, in the "direct" approach, all registration steps are independent as

each image is individually aligned to the same reference image. This suggests that having a fixed high SNR reference image should be prioritized over minimizing differences in contrast throughout the time series using this particular approach. Furthermore, this encourages us to establish the "*direct*" registration strategy as standard for motion correction of the SR time series in our overall ASL processing pipeline.
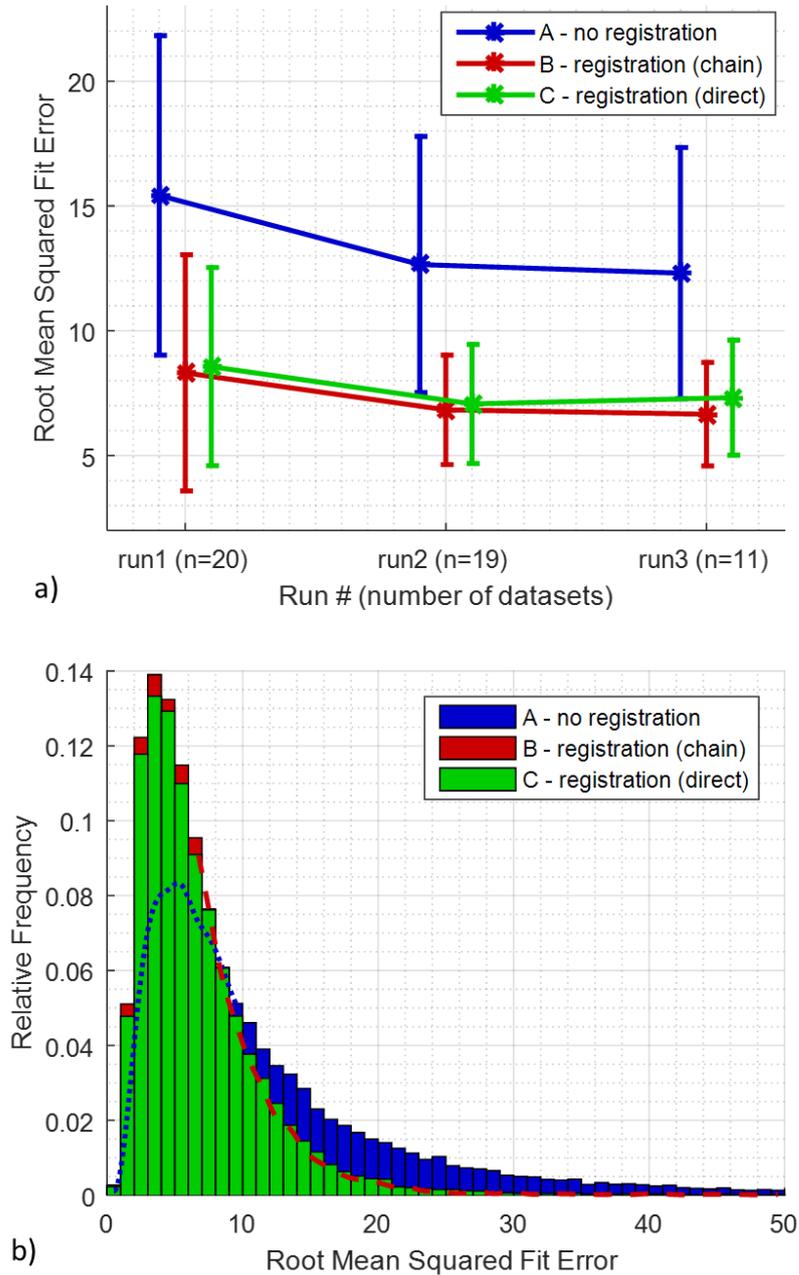
Figure 6-5 Effect of proposed registration approaches on the RMSE of the SR fits on the functional renal parenchyma ROIs. a) Average ± standard deviation (error bars) RMSE across all subjects for each run of a given day. b) Histogram of the RMSE of individual voxels for all patients/days/runs. The dashed lines correspond to the obscured sections of histograms A and B.
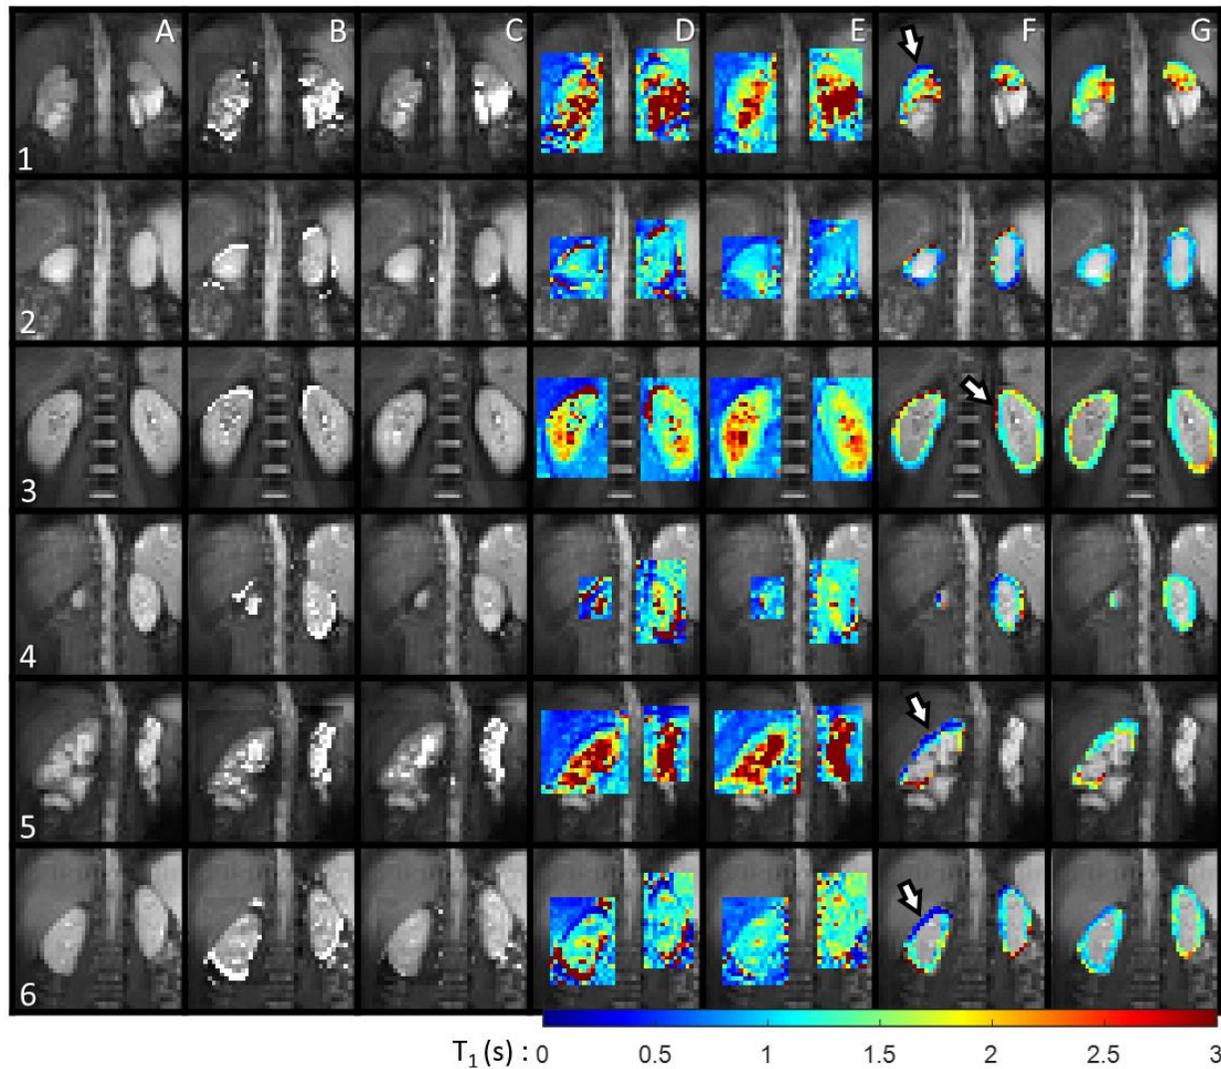
Figure 6-6 Overview of results of SR fitting (patients #1-#6), with and without applying image registration. Each row corresponds to data from a central slice of a single patient (the number in each row corresponds to the patient index). Each column shows images from a given type. A – proton density images (single acquisition); B – M_0 map resulting from SR fitting (no motion correction); C – Same as B but using motion correction ("direct" image registration); D – T_1 map obtained from SR fitting (without motion correction) overlaid on the PD images; E – Same as D but using motion correction ("direct" image registration); F and G: Respectively the same as D and E but T_1 map values restricted to the functional renal parenchyma ROIs, from which mean T_1 estimates were derived. The white arrows in column F highlight artefacts in the T_1 maps caused by misalignment during the time series before applying image registration All data from the 1st run of the 1st scan of each patient.
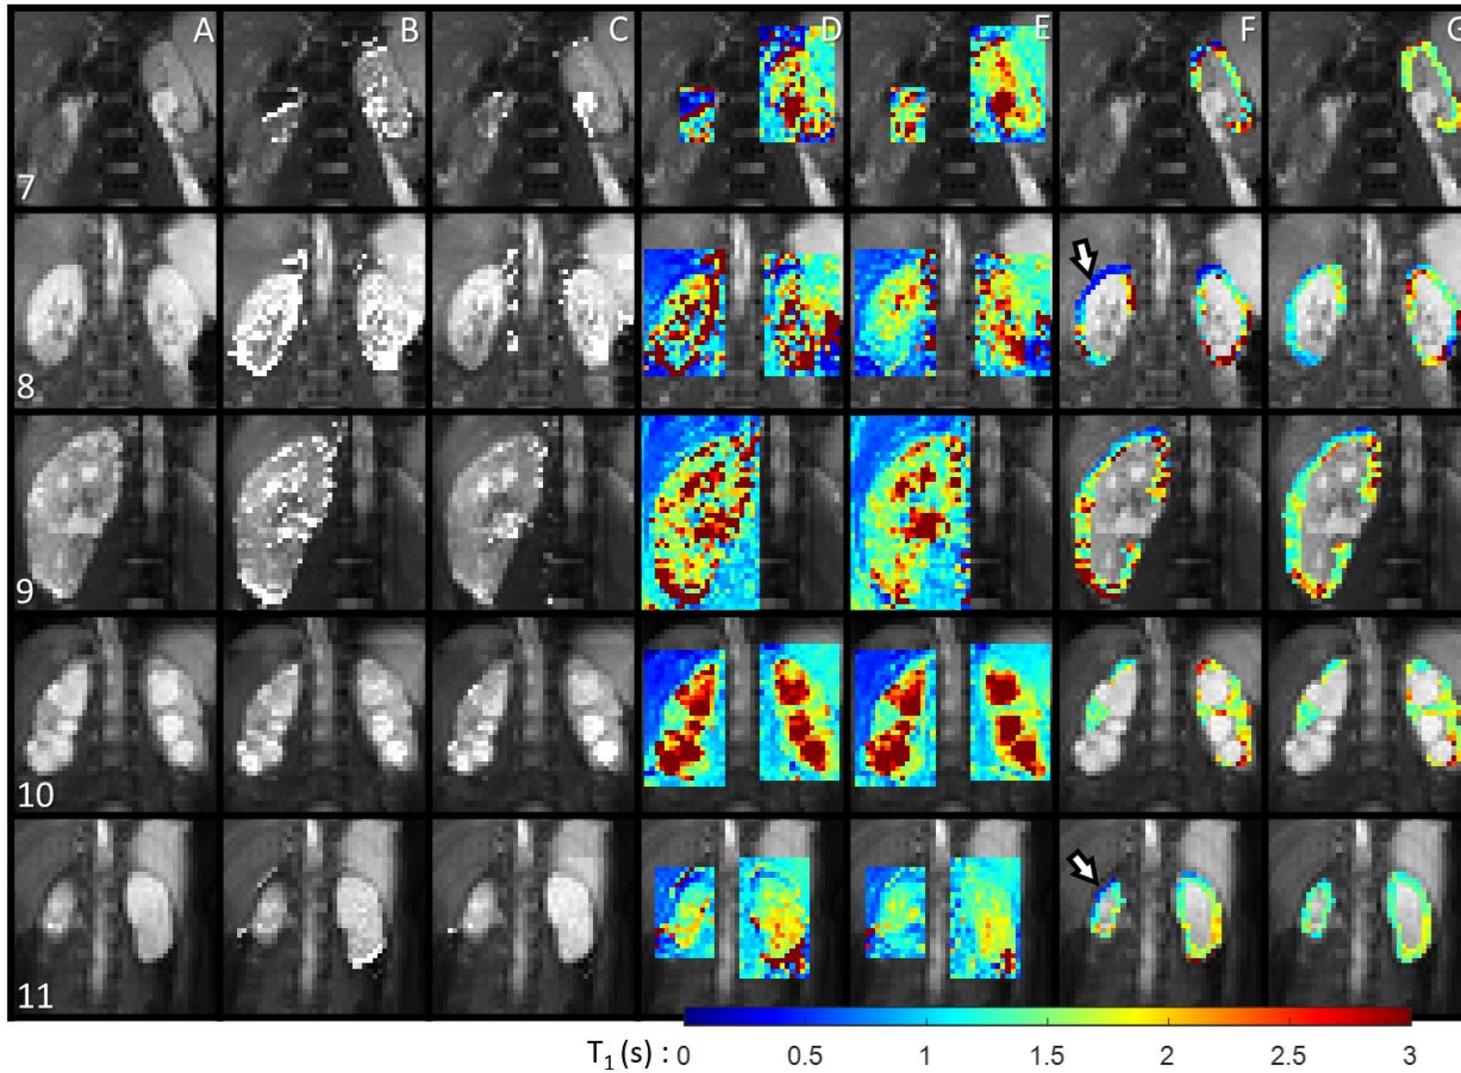
Figure 6-7 Overview of results of SR fitting (patients #7-#11) (figure organisation analogous to Figure 6-6)

$T_1 (s)$: 0   0.5   1   1.5   2   2.5   3

6.3.3 Post-processing pipeline: ASL

6.3.3.1 Registration of background-suppressed time series

The mean change in temporal standard deviation (tSTD) of the perfusion-weighted time series, averaged across after applying the ASL processing pipelines B ("*registration only*") and C ("*registration + weighted averaging*") relative to the "*no motion correction*" case (pipeline A) was -15.6±10.5% (range: -37.8% – 1.8%) and -24.3±14.6% (range: -57.2% – -1.0%) respectively. Analogously, the mean variation in temporal SNR (tSNR) after applying pipelines B and C, relative to A was, respectively 19.2±15.7% (range: 1.3% – 65.6%) and 31.4±25.0% (range: 1.1% – 87.9%). Note that these values relate to all scans in the dataset, including scans during which the amount of motion was smaller (where less room for improvement in the data exists). If only scans during which the average amount of total renal translational motion was larger than 50mm (or ≈1mm per volume) are considered, tSTD further decreases by 21.6±10% and 35.1±11.9 when comparing pipelines B and C with A, respectively. Analogously, the tSNR further increases by 29.1±15.8% and 49.6±24.3%. The percent change in these metrics after applying pipelines B and C, relative to A for each individual run of all patients, as well as averaged across runs for each patient are shown in Table 6-4 (tSTD) and Table 6-5 (tSNR). For enhanced visualization, each entry in the tables is colour-coded according to the degree of improvement in the metrics for each scan, on a colour scale ranging from red (smallest improvement) to green (greatest improvement). For tSTD, improvement corresponds to a decrease in this metric, assuming that during the length of the scan (approximately 3 minutes, with the delays associated with respiratory triggering), the haemodynamics should remain constant and therefore any change in perfusion weighted signal should be attributed to noise and especially motion. On the other hand, improvement in tSNR corresponds to an increase in this metric (see Chapter 5 (Section 5.2.4) for the definition of these metrics). Note the more frequent "green" colour coding on the right hand side of the tables, suggesting the added value of combining weighted averaging with image registration. In fact, applying this pipeline (C), resulted in a decrease in tSTD and an increase in tSNR in all scans in the dataset.

Statistically significant differences were found between the tSTD and tSNR of the perfusion-weighted time series in the "no motion correction" and the "registration only" as well as the "registration + weighted averaging" cases. In addition, these metrics also differed significantly between the "registration only" and the "registration + weighted averaging" cases (two-tailed paired t-tests, all p-values $\leq 1\text{x}10^{-5}$). All runs were considered individually in these tests (i.e. the metrics were not averaged across the four runs for each patient). This suggests a benefit in using

the motion correction pipelines and the usefulness of using a noise-reduction weighted averaging method in addition to standard image registration.

Table 6-4 tSTD change after image registration relative to the no motion correction case (in %). In this table, the Euclidean distance (ED) was averaged across all runs for each patient.

Patient	B – registration only					C – registration + weighted averaging					ED (mm)
	Day 1		Day 2		Mean	Day 1		Day 2		Mean	
	Run 1	Run 2	Run 1	Run 2		Run 1	Run 2	Run 1	Run 2		
1	-25.2	-28.8	-27.3	-20.2	-25.4	-27.2	-34.7	-31.0	-48.3	-35.3	69.9
2	-12.4	-21.7	-28.3	-21.4	-20.9	-16.2	-26.7	-32.6	-30.3	-26.4	36.3
3	-12.1	-18.3	-22.3	-17.5	-17.6	-25.9	-31.5	-28.8	-22.8	-27.3	47.9
4	-5.4	-4.2	-7.4	-4.0	-5.3	-9.5	-8.4	-21.6	-5.1	-11.1	22.1
5	-32.8	-4.8	-26.2	-27.3	-22.8	-39.1	-9.6	-41.5	-38.0	-32.0	83.6
6	-7.7	N/A	-15.5	-10.2	-11.1	-9.4	N/A	-22.2	-21.4	-17.7	33.1
7	-8.9	-14.0	-6.9	-10.5	-10.1	-10.3	-22.0	-11.1	-16.5	-15.0	55.0
8	-5.4	-25.9	-3.0	-17.1	-12.9	-7.7	-35.3	-5.4	-22.8	-17.8	37.5
9	-1.6	-9.1	-0.4	-3.3	-3.6	-2.5	-9.9	-1.0	-3.3	-4.2	33.5
10	-37.8	-37.2	-30.4	N/A	-35.1	-48.2	-57.2	-34.7	N/A	-46.7	80.3
11	-13.8	-13.5	-15.4	1.8	-10.2	-25.5	-30.5	-44.5	-51.4	-37.9	74.7

Table 6-5 tSNR change after image registration relative to the no motion correction case (in %). In this table, the Euclidean distance (ED) was averaged across all runs for each patient.

Patient	B – registration only					C – registration + weighted averaging					ED (mm)
	Day 1		Day 2		Mean	Day 1		Day 2		Mean	
	Run 1	Run 2	Run 1	Run 2		Run 1	Run 2	Run 1	Run 2		
1	29.7	39.9	30.5	28.8	32.2	32.9	48.9	35.0	82.8	49.9	69.9
2	11.8	21.9	32.2	27.7	23.4	16.1	29.2	40.4	43.5	32.3	36.3
3	8.8	22.3	21.6	16.5	17.3	25.4	35.0	26.2	18.8	26.3	47.9
4	5.2	3.7	8.6	5.4	5.7	10.2	8.3	25.7	6.1	12.6	22.1
5	46.3	5.0	37.7	48.9	34.5	52.6	8.3	62.4	76.1	49.9	83.6
6	5.9	N/A	16.7	9.1	10.5	8.5	N/A	25.2	18.6	17.4	33.1
7	11.0	11.8	6.0	10.4	9.8	12.1	23.0	9.9	15.1	15.0	55.0
8	5.5	35.3	2.7	21.2	16.2	6.9	48.0	4.4	25.5	21.2	37.5
9	1.3	7.9	1.4	1.8	3.1	1.3	7.5	1.1	3.0	3.2	33.5
10	65.6	46.0	40.3	N/A	50.7	82.1	87.9	49.0	N/A	73.0	80.3
11	14.4	10.4	21.2	7.6	13.4	27.2	31.5	76.9	70.7	51.6	74.7

As long as the data has not been corrupted to such a degree that motion correction is not feasible, one would expect a greater improvement in tSNR for datasets where motion is more prevalent. As such, the total renal translation motion (summarised by the Euclidean distance metric), averaged across runs for each patient is also shown in Table 6-4 and Table 6-5. A larger value in this metric suggests an overall larger amount of motion by the corresponding patient. In fact, statistically significant correlations between the amount of renal translational movement in each scan and the improvement in tSNR after applying image registration were found (R = 0.59, p = 4.6×10^{-5}) and a combination of registration and weighted averaging (R = 0.80, p =

2.3×10^{-10}) in the corresponding perfusion weighted time series. Figure 6-8 shows the corresponding correlation scatter plots, where R is the Pearson correlation coefficient.
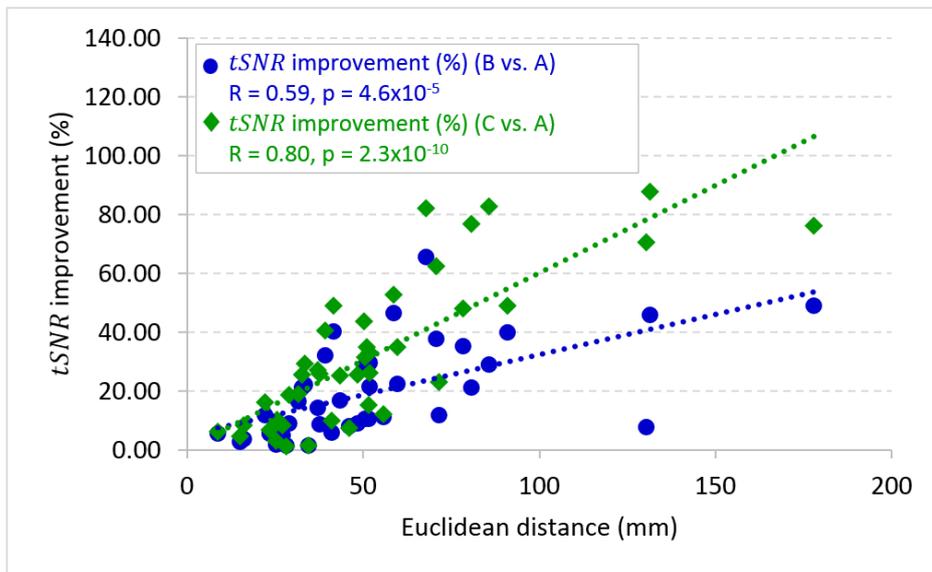
Figure 6-8 Improvement in tSNR after motion correction methods. scatter plot displays the percent increase (i.e. improvement) in tSNR in the PWI time series after applying the motion correction processing pipelines B and C relative to the "no motion correction" case (pipeline A) according to the amount of renal translation motion, for each individual ASL run (n=42). Linear regression lines as well as the Pearson correlation coefficient and p-value are shown. A significant, positive correlation was found between the improvement in tSNR and the total amount of renal translation motion, as determined by the registration algorithm.

The amount of renal translational movement (as determined by the rigid registration algorithm), averaged across runs, differed significantly between runs 1 and 2 (44.2±17.3mm vs. 61.4±43.8mm, two-tailed paired t-test p=0.035, data pooled from the two days) but not between days 1 and 2 (46.3±20.7mm vs. 54.8±39.8mm, two-tailed paired t-test p=0.392, data pooled from the two runs). This finding appears consistent with the practical expectation that the likelihood of patient movement increases with the amount of time spent inside the scanner. This further suggests that for a given protocol, the most motion-sensitive sequences, such as ASL, should be acquired near the beginning of the scanning session. Table 6-6 shows the Euclidean distance metric obtained in each individual run in the dataset, as well as mean and standard deviation of this metric, both across runs and patients.

Table 6-6 Euclidean distance descriptive statistics across the whole cohort (all patients/days/runs)

Patient	Age (years)	Euclidean distance (mm)					
		Day 1		Day 2		mean	std
		Run 1	Run 2	Run 1	Run 2		
1	10	51.7	91.1	51.1	85.6	69.9	21.4
2	14	22.2	33.5	39.1	50.2	36.3	11.7
3	12	48.5	59.6	51.9	31.6	47.9	11.8
4	10	25.5	16.2	37.7	8.8	22.1	12.5
5	7	58.6	27.2	70.8	177.8	83.6	65.4
6	17	26.9	N/A	43.5	28.9	33.1	9.1
7	15	55.7	71.5	41.1	51.6	55.0	12.6
8	13	23.6	78.4	15.1	32.7	37.5	28.2
9	16	34.6	46.0	28.2	25.3	33.5	9.2
10	12	67.8	131.4	41.7	N/A	80.3	46.1
11	9	37.0	50.6	80.8	130.4	74.7	41.4
mean	12.3	41.1	60.5	45.6	62.3	52.4	
std	3.1	16.0	34.2	18.2	53.7	30.5	

A weak, albeit statistically significant correlation was found between the age of the patient and the amount of renal translational movement in each of the patients' ASL runs (see Figure 6-9).
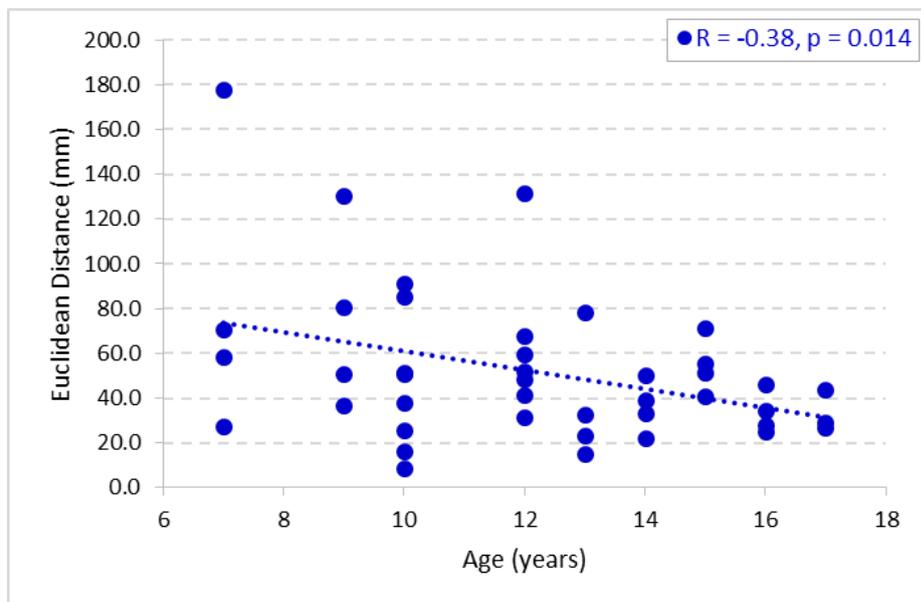
Figure 6-9 Correlation of renal motion and patient age. This scatter plot displays the total amount of renal translational motion (as determined by the rigid registration algorithm) according to the age of the patient. The linear regression line as well as the Pearson correlation coefficient and p-value are shown.

Furthermore, weak but statistically significant negative correlations were also found between the improvement in tSNR (%) and the age of the subjects when comparing the "no motion correction" and the "registration only" pipelines (R = -0.37, p=0.015) and the "no motion correction" and the "registration + weighted averaging" pipelines (R = -0.49, p=0.0009).

6.3.3.2 Complete pipeline

Figure 6-11 and Figure 6-12 provide an overview of the ASL results. Similarly to Figure 6-6 and Figure 6-7, a single slice across the centre of the kidneys is shown for all patients enrolled in this study and each row in the figures corresponds to data from the same patient. Results are shown for "no motion correction" condition and after applying the full ASL motion correction pipeline, including "*direct*" registration for the reference SR time series and the "*registration + weighted averaging*" strategy for the ASL data. All scans in these figures correspond to the first run of the first day for all patients. Column *A* shows the PD scans, used as the baseline M_0 scan for calibration of the ASL data and also shown here to serve as the anatomical reference. Columns *B* and *C* show the hybrid PD-PWI images, respectively without motion correction and after applying the ASL motion correction pipeline.

One immediate observation is the fact that the PW dataset is generally well aligned with the PD dataset, even before applying motion correction. At least three reasons likely explain this finding. First, the SR dataset was acquired immediately before the main ASL acquisition. Second, the fact that each ASL acquisition consisted of 25 identical measurements for averaging which confer motion robustness to the final averaged PWI even if the individual images before averaging may occasionally shift with respect to each other due to changes in the respiratory cycle. Third, the fact that even though no labelling is performed in the PD acquisitions, a delay following the application of the respiratory trigger was added in this acquisition to match the delay that exists between the respiratory trigger and the ASL acquisition (to allow the inflowing blood labelled in the aorta to reach the kidneys). This therefore increases the likelihood of the PD images and the ASL images to be acquired at the same point in the respiratory cycle. This demonstrates that a careful setup in the acquisition stage goes a long way to provide high quality data, even without motion correction.

An example of one scan where a large misalignment (approximately 1cm) between the PD and the ASL data volumes can be seen, as well as the effect of applying ASL pipeline C ("image registration + weighted averaging") is shown in Figure 6-10.
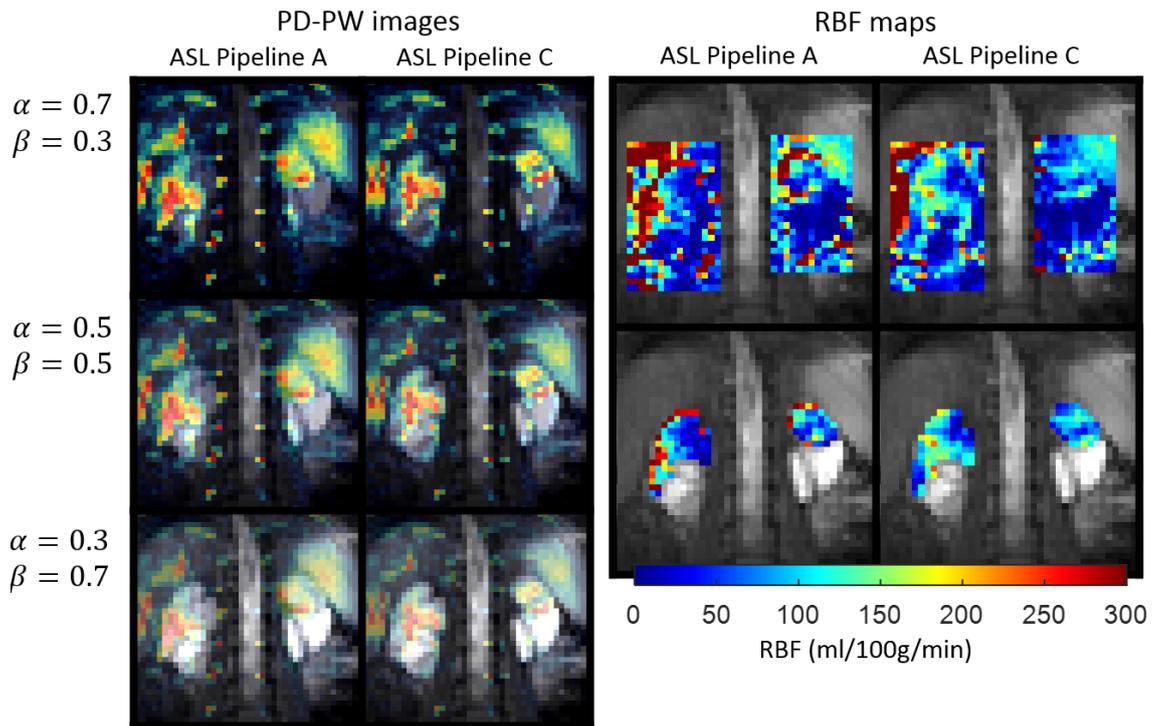
Figure 6-10 Motion correction of a large misalignment between PD/ASL volumes (approximately 1cm). This misalignment is evident when looking at the proton-density perfusion-weighted (PD-PW) images without motion correction (ASL pipeline A), particularly when generating these images with with varying levels of α and β weights. After image registration + weighted averaging (ASL pipeline C), a better match is seen between the PD and ASL datasets and the corresponding RBF maps are markedly improved. This scan corresponds to the second run of the first day of patient #1.

Overall, the hybrid images are of high quality, and show patterns of renal perfusion which vary considerably when comparing to healthy volunteer data (such as shown in Chapter 5) but also between patients, according to the underlying pathophysiology. For example, the data from patient #4 suggests relatively high perfusion with a spatial distribution matching that of healthy volunteers (higher in the renal cortex), patient #3 appears to have a marked reduction in perfusion throughout the whole kidney and patient #1 shows local perfusion defects matching regions of kidney dilations as seen in the PD images. Columns D and E show the resulting quantitative RBF maps computed within the rectangular regions used for masking the image registration step, respectively following ASL pipeline A ("no motion correction") and ASL pipeline C ("image registration + weighted averaging"). Here, the beneficial effects of motion correction are more perceptible, largely because several separate image datasets are used in conjunction to generate the RBF maps, and any misalignments between these give origin to severely biased RBF measurements. Similar to when no motion correction was applied before SR fitting for T_1/M_0, artefacts are seen close to the boundaries of the kidneys in the RBF maps if image registration is not used (highlighted by the arrows in the subfigures in column F). The existence

of such artefacts depends on the relative positions of the kidneys in the ASL and the "calibration" data (PD images / SR time series), which may change between different ASL runs. Therefore, they directly influence the repeatability of the RBF measurements. The proposed registration approach is able to improve the alignment between background-suppressed ASL data and the reference M_0 maps, as well as T_1 maps obtained by SR fitting. The boundary artefacts are accordingly reduced. This can be seen in Figure 6-11 and Figure 6-12, by comparing the corrupted data with artefacts highlighted in column *F*, with the images from the same patients after applying the correction (column *G*). A particularly striking case where this is also the case can be seen in Figure 6-10. The RBF values averaged across the functional renal parenchyma ROIs obtained in each of the four scans in all patients enrolled in this study are shown in Table 6-7.

Table 6-7 Mean functional renal parenchyma RBF values obtained for each patient in all of the four ASL scans (two runs (R1 and R2) in each of two days), together with the mean and standard deviation of the RBF values of each patient averaged calculated across the four scans. The colour coding for each entry in the table relates to the mean RBF values (lowest: red; highest: green) and the standard deviation of the mean RBF values obtained in the 4 scans for each patient (low STD → high repeatability: green; high STD → low repeatability: red). N/A[a]: discarded due to fat shift artefact; N/A[b]: discarded due to failure in respiratory triggering.

	A "no motion correction"						B "registration only"						C - "registration + weighted averaging"					
	Day 1		Day 2		Mean	STD	Day 1		Day 2		Mean	STD	Day 1		Day 2		Mean	STD
Patient	R1	R2	R1	R2			R1	R2	R1	R2			R1	R2	R1	R2		
#1	128	93	86	80	97	22	115	91	82	88	94	15	116	90	82	90	95	15
#2	187	169	154	130	160	24	168	160	188	166	171	12	168	160	188	170	171	12
#3	53	54	67	70	61	9	54	53	52	51	52	2	52	51	51	50	51	1
#4	242	261	192	152	212	49	179	182	163	87	153	45	179	181	162	87	152	44
#5	344	276	125	99	211	118	144	87	117	87	109	28	141	91	114	87	108	25
#6	418	N/A[a]	178	146	247	148	169	N/A[a]	169	142	160	15	170	N/A[a]	169	142	160	16
#7	132	146	106	91	119	25	76	69	73	77	74	4	75	67	73	75	73	4
#8	114	101	72	79	91	19	72	77	52	57	65	12	73	76	51	56	64	12
#9	38	41	43	50	43	5	37	41	42	49	42	5	37	41	41	48	42	5
#10	74	61	155	N/A[b]	97	51	78	67	94	N/A[b]	80	13	76	67	94	N/A[b]	79	14
#11	77	82	78	104	85	13	75	78	46	47	61	17	73	76	46	41	59	18
Mean	127.4						95.3						94.7					
STD	81.8						46.9						47.3					

Overall, across all scans of the entire cohort, the mean±std of the RBF obtained with each of the ASL pipelines was 127.4±81.8, 95.3±46.9 and 94.7±47.3 ml/100g/min, respectively using ASL pipelines A ("no motion correction"), B ("registration only") and C ("registration + weighted averaging"). Overall, the standard deviation of the measurements decreased after applying the motion correction pipelines. A higher standard deviation of the 4 runs of each patient indicates larger variability in the mean RBF obtained in each measurement, which could be related to an incomplete or unsuccessful movement correction. The highest standard deviation is seen in

patient #4, whose last scan (day 2 run 2) yielded an RBF measure of approximately half of the first 3. A retrospective look through all the intermediate steps in the pipeline revealed that the final registration step (registration of the averaged background-suppressed control image to the reference PD scan) failed. What caused registration to fail was the existence of a high-intensity region in the vicinity of the kidneys in the background suppressed control images. Compared to the background-suppressed kidneys, this high intensity region had a greater similarity to the non-background supressed kidneys in the PD dataset. Consequently, the registration algorithm brought this high intensity region into alignment with the kidneys in the PD reference image instead of the actual background suppressed kidneys. Given the rigid transformation, the alignment of the kidneys in the registered background suppressed data to the PD reference dataset ended up being worse after this final registration step that with no registration, which was responsible for a corrupted perfusion map and hence the greatly different RBF estimate in this particular scan compared to the first three. This highlights the need for data quality control checks to avoid propagation of errors/artefacts throughout an image processing pipeline, particularly as its degree of automation increases. Repeatability measures regarding the RBF estimates obtained using the different processing pipelines are summarised in Table 6-8. Note that no data was rejected in this chapter (other than scanning failure). All the scans which were not completely corrected were included in the repeatability analyses for a fair assessment of the success of the registration methods.

Table 6-8 Intra/inter-session repeatability measures for the mean cortical RBF estimates. Pipeline A: "no motion correction"; Pipeline B: "registration only"; Pipeline C: "registration + weighted averaging"; WSCV: within-subject coefficient of variation; ICC: intra-class correlation coefficient.

Pipeline	Intra-session repeatability			Inter-session repeatability		
	A	B	C	A	B	C
	Pooled days			Pooled runs		
WSCV	0.145	0.186	0.179	0.506	0.201	0.204
ICC	0.936	0.857	0.873	0.372	0.833	0.834
	Day 1			Run 1		
WSCV	0.141	0.151	0.139	0.531	0.132	0.135
ICC	0.951	0.898	0.916	0.370	0.927	0.925
	Day 2			Run 2		
WSCV	0.149	0.218	0.214	0.453	0.282	0.285
ICC	0.853	0.818	0.834	0.329	0.645	0.657

Image registration was particularly successful in improving the inter-session repeatability of the measurements. The intra-session repeatability without applying image registration was unexpectedly high. Because the SR dataset was acquired only once at the beginning of the scanning session, systematic errors due to consistent misalignment between the SR and the ASL dataset in both ASL runs resulted in a moderately repeatable mean RBF despite severe

corruption in the RBF maps in the no-registration case. The lowest repeatability was found to be the inter-session repeatability of the second scan of each day, a somewhat anticipated fact given the expected decreasing levels of compliance by the subjects the longer they remain in the scanner. In addition to the challenge faced in the registration of the second scan of the second day of patient #4 (discussed above), arguably the worst repeatability was found in patient #5 (see Figure 6-13). In this case, the final RBF maps of the second run of both days of this patient remain affected by motion artefacts after the complete pipeline. In the "day 1 - run 2" scan, this is due to misalignment of between the PD and the final PWI (highlighted by the white arrows in the second row of Figure 6-13). This resulted in an underestimation of perfusion of the inferior pole of the kidney in (see column G), lowering the mean RBF with respect to that obtained in "day 1 - run 1". In the case of "day 2 - run 2" extreme motion (note this was the scan with largest amount of motion according to the Euclidean distance metric) caused the background suppression (BS) to PD registration to fail (see bottom of Figure 6-13). Note that despite the movement, an improvement in quality of the PWI after BS to BS registrations can be seen.

Indeed, ensuring alignment of the ASL and PD datasets was found to be the most challenging step in the pipeline because of their dissimilar signal intensities and contrast (due to the use of BS in the ASL scans). We therefore would recommend acquiring a SR dataset ($\approx$1min using our acquisition protocol) immediately before/after any ASL acquisition, to minimise the likelihood of significant motion between acquisitions, particularly in challenging patient populations, such as the one in this study. Additionally, not only motion but other types of artefacts can be present in the unsubtracted background-suppressed data. Avoiding these, particularly in the neighbouring regions of the kidneys, is key to enable an already challenging registration task (registering background-suppressed to non-background-suppressed data).
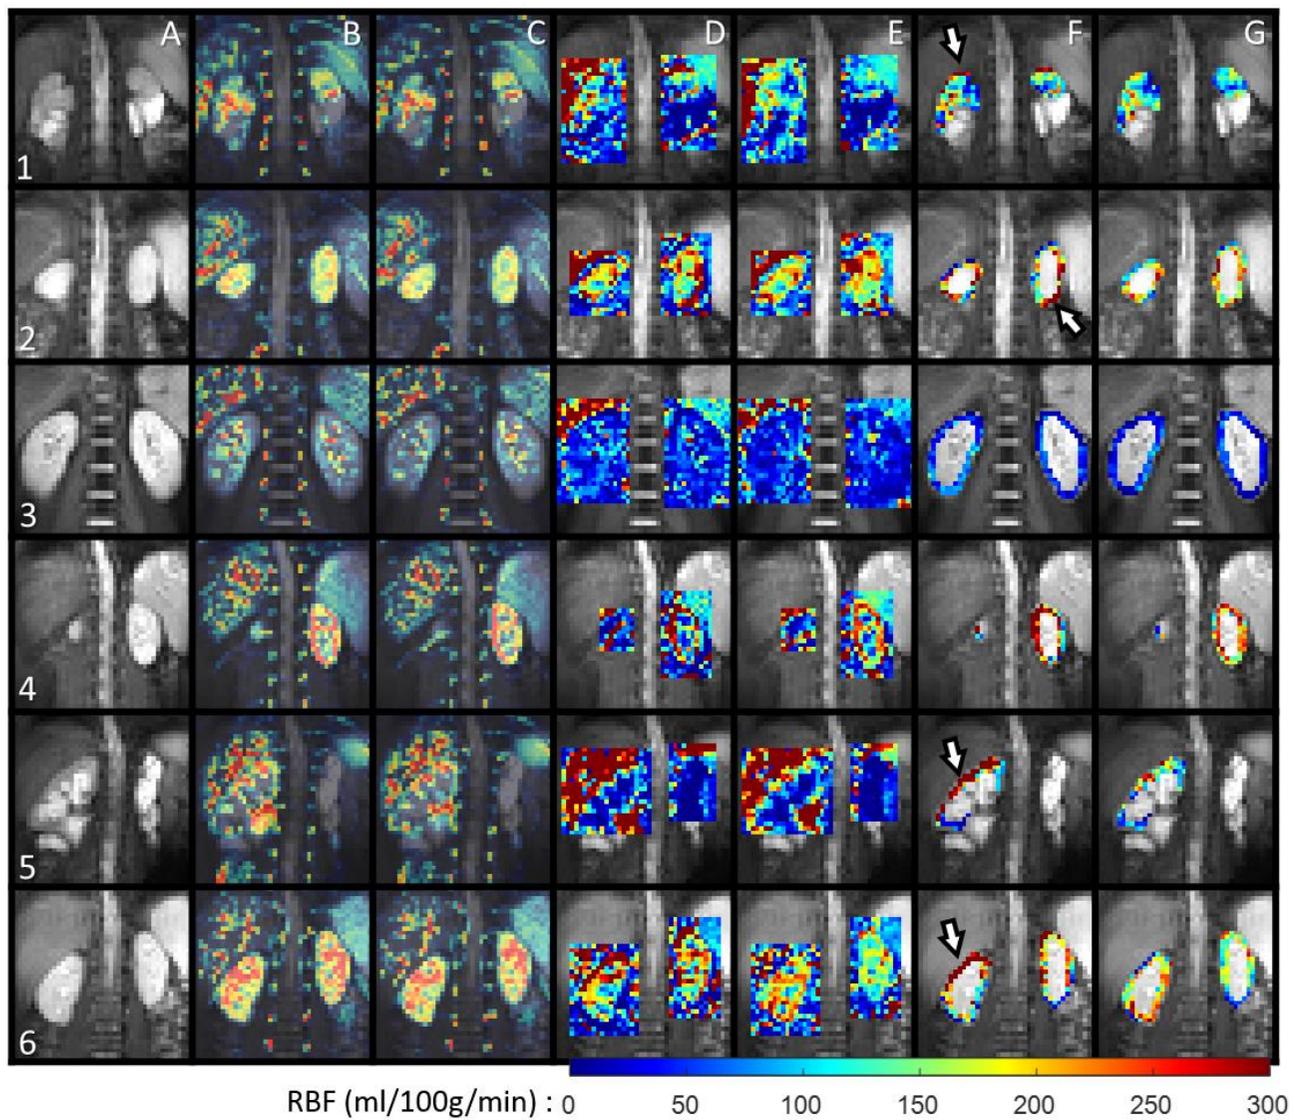
Figure 6-11 Overview of ASL results before/after motion correction (patients #1-#6). Each row corresponds to data from a central slice of a single patient (the number in each row corresponds to the patient index). Each column shows images from a given type. A – Proton density images; B – PD/PWI fusion image, where the grayscale data corresponds to the PD images, and the coloured data corresponds to the averaged perfusion weighted image (no motion correction). C – Same as B but using motion correction ("registration + weighted averaging"); D – RBF map obtained without motion correction overlaid on the PD images; E – Same as D but using motion correction ("registration + weighted averaging"); F and G: Respectively the same as D and E but RBF map values restricted to the functional renal parenchyma ROIs, from which mean RBF estimates were derived. All data from the 1st run of the 1st scan of each patient.

RBF (ml/100g/min) : 0   50   100   150   200   250   300
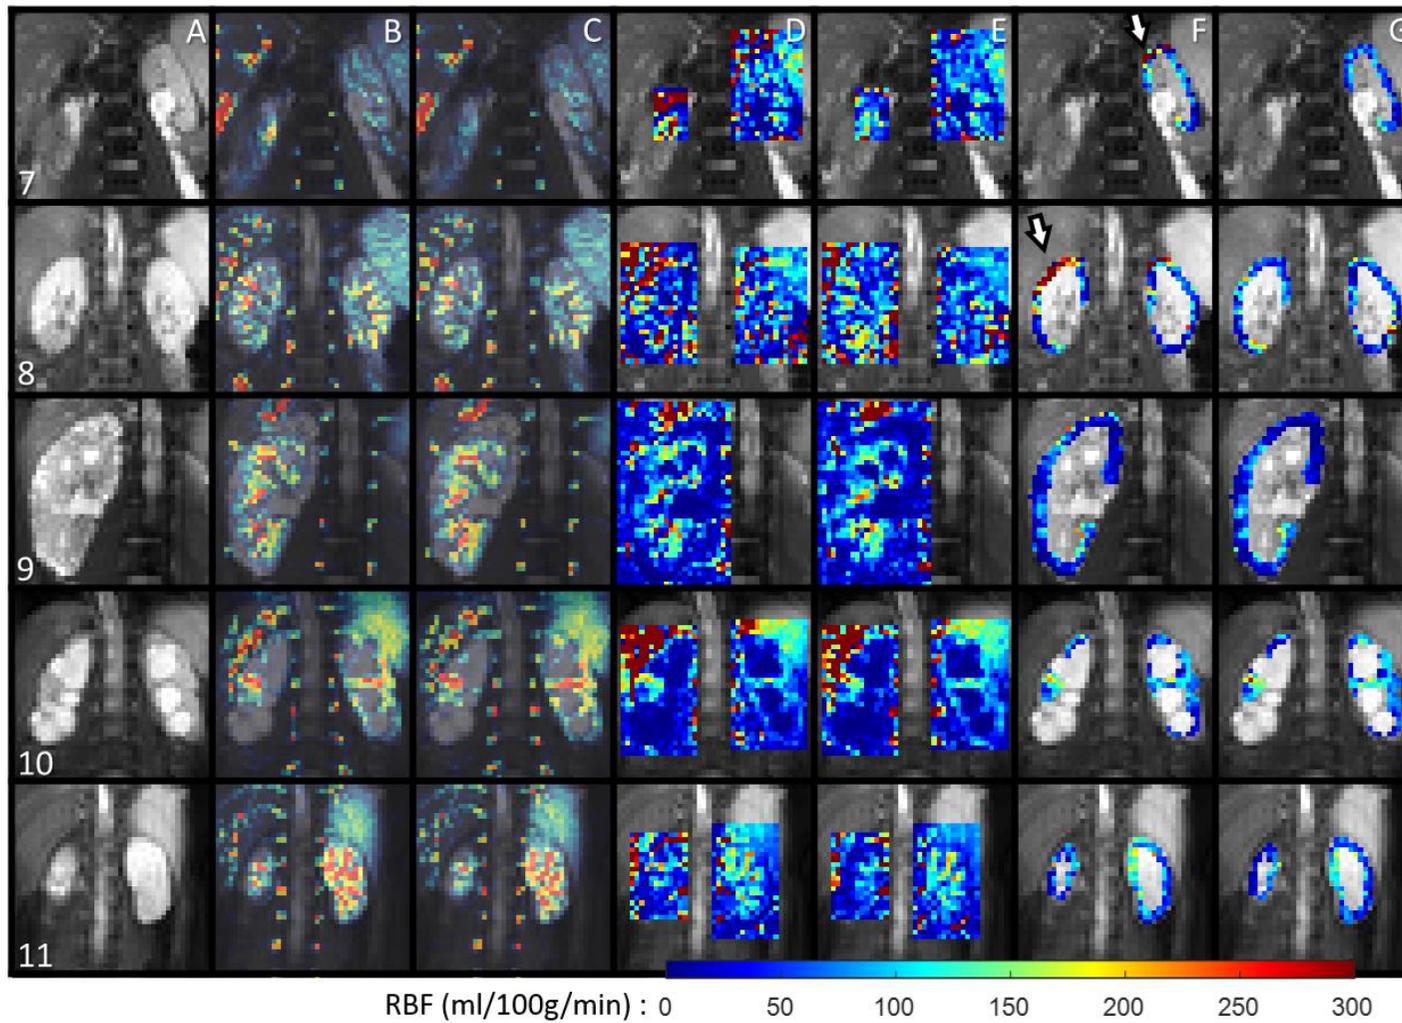
Figure 6-12 Overview of ASL results before/after motion correction (patients #7-#11)(figure organisation analogous to Figure 6-11)

RBF (ml/100g/min) : 0   50   100   150   200   250   300
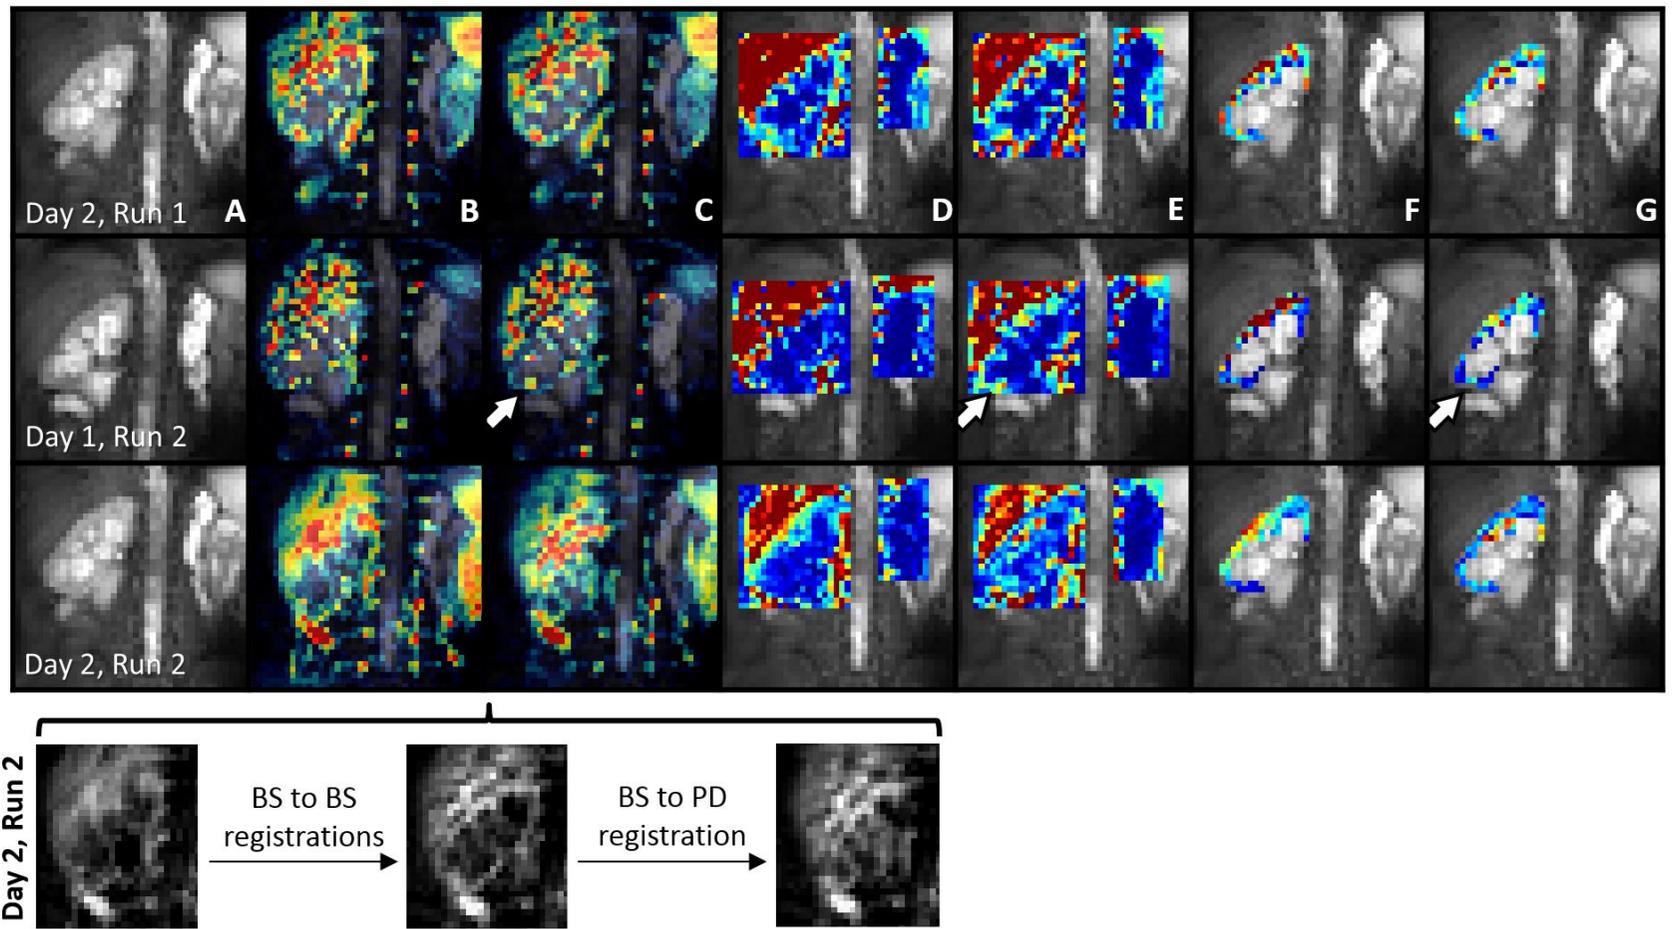
Figure 6-13 Worst repeatability case due to unsuccessful BS to non-BS (i.e. PD) registration (patient #5).

6.4 Correlation of RBF and T_1

As discussed in Section 2.3.6 of this thesis, several studies have shown decreased RBF in CKD. Furthermore, an increase in T_1 has been associated with pathology-induced structural changes to the renal parenchyma.

A statistically significant negative correlation was found between cortical RBF and T_1 (R = −0.57, p = 0.006) in this patient cohort. The corresponding correlation scatter plot is shown in Figure 6-14, where R is the Pearson correlation coefficient.
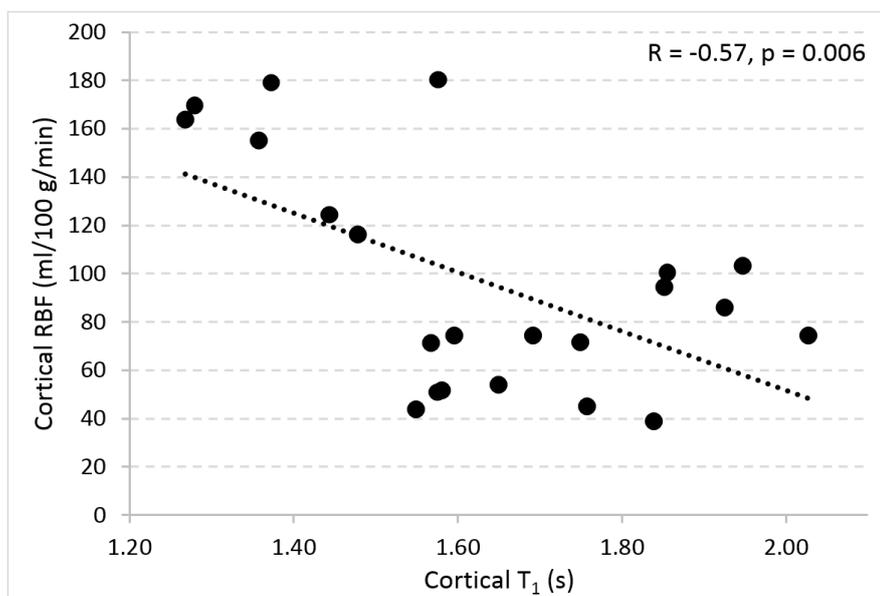
Figure 6-14 Correlation of cortical RBF and T_1. This scatter plot displays the cortical RBF and corresponding T_1 for each patient in each of the scanning sessions (ASL runs from each day are averaged since both were quantified using the same saturation recovery series). The linear regression line as well as the Pearson correlation coefficient and p-value are shown.

This correlation reflects the fact that both T_1 and RBF are MR measures sensitive to pathology. Furthermore, in line with as discussed previously (section 6.2.3.1), supports the need for T1 quantification on an individual subject basis to ensure accurate RBF quantification.

6.5 Conclusion

ASL is a unique MR approach for quantifying tissue perfusion non-invasively. More than two decades of developments enabled great strides in translating this technique from specialized research labs into clinical implementation (Alsop 2012; Alsop et al. 2015). Nevertheless, ever since its inception, ASL has been hindered by its propensity to motion-related artefacts which limit its application in the clinical domain, especially outside the brain (Robson et al. 2009; Gardener & Francis 2010). In this work, we have combined a motion-insensitive ASL acquisition scheme with a specifically tailored image processing pipeline using validated image registration

methods with the aim of enabling repeatable renal perfusion measurements in paediatric patients with kidney disease. Eleven patients were scanned, each twice (40-50 minutes apart) on two different days to assess the reproducibility of the measurements.

Certain assumptions which simplify both the ASL acquisition and data processing in healthy volunteers, such as using literature T_1 times, no longer hold as this parameter may change with the underlying pathophysiology. Therefore, our pipeline includes a SR recovery fitting module allowing patient-specific T_1 maps to be derived. This relies on the acquisition of a series of images at varying post-saturation delays and therefore is also prone to motion artefacts. We evaluated the impact of two SR recovery time series registration approaches on the renal T_1 estimates obtained in this patient cohort. The "chain" method appears susceptible to propagation errors throughout the time series due to the fact that the reference image for the registration step changes in each iteration. This suggests that having a fixed high SNR reference image ("*direct*" approach) should be prioritized over minimizing differences in contrast throughout the time series using this particular "*chain*" implementation. Nevertheless, both registration approaches were successful in improving the quality of the T_1 estimates on a challenging patient population, as assessed by the intra- and inter-scan repeatability and analysis of the RMSE of the SR fits. Furthermore, realignment of the SR time series resulted in an improvement of the quality of the SR fits, as assessed by the RMSE metric (RMSEs of the SR fits were statistically significantly reduced after registration). A visual comparison of the resulting T_1/M_0 maps clearly demonstrates the usefulness of image registration in reducing artefacts caused from misalignment of the individual image volumes from the SR time series. Finally, the mean T_1 estimates from each scan of each patient (averaged across the entire ROI) were found to be statistically different after image registration, compared to the no motion correction case. Across all scans, we obtained a mean T_1 of 1.89±0.29 and 1.65±0.21s (mean±std), respectively in the "no motion correction" condition and after applying the "*direct*" registration method.

As anticipated, from an ASL standpoint scanning paediatric patients with CKD presented additional challenges, especially in terms of motion and reduced signal to noise ratio (SNR) of the perfusion signal. From an acquisition perspective, our data suggests that the optimised protocol was able to cope with this as the images are free of motion-artefacts due to the single-shot nature of the acquisition. The amount of perfusion signal varies considerably among patients as the population is quite heterogeneous in terms of their pathologies. Nevertheless, the perfusion-weighted signal following control-label subtraction is clearly detectable in all patients when compared to the surrounding tissue and background regions with no expected

perfusion. This can be attributed to the moderate resolution at which the ASL scans are performed as well as the multiple averaging step implemented in this protocol, which boosted the SNR. We observe a very close match between visible pathologies in our anatomical scans (e.g. collecting system dilations) and a marked reduction in perfusion signal in the ASL scans. This is especially evident in the PD-PWI hybrid images shown in this chapter.

Despite the single-shot acquisition, the additional bulk body movement and irregular respiratory rates (which reduces the efficacy of respiratory triggering) prevalent in this patient cohort still poses challenges as all the measurements rely on processing multiple image volumes acquired separately. As seen above, registration of the SR time series was successful. In addition, the registration of the background-suppressed data was assessed by computing the temporal standard deviation (tSTD) and temporal SNR (tSNR) of the perfusion-time series. A statistically significant improvement in both metrics was achieved by using image registration, which was further enhanced by integrating the weighted averaging approach into the processing pipeline. The degree of improvement in tSNR positively correlated with the amount of renal translational motion which took place during the acquisition of the ASL time series, as described by the Euclidean distance metric.

Given the large number of averages in the ASL scan, the factor which most critically affected the measured RBF values was the alignment between the averaged ASL data (i.e. the final averaged PWI) and the calibration datasets (T_1 and M_0 maps). Indeed, if these image volumes are not aligned, the resulting RBF maps are corrupted, showing severe artefacts particularly close the boundaries of the kidneys and producing wildly inaccurate RBF values. This is especially problematic as cortical perfusion measurements are typically derived from these regions. The effect of registration in reducing these types of artefacts was demonstrated in Figure 6-11 and Figure 6-12. This resulted in an improved inter-session repeatability of the scans, assessed by the intra-session correlation coefficient and within-subject coefficient of variation. Interestingly, even though the intra-session repeatability did not improve *per se*, a reduction in the systematic errors in the RBF estimates due to consistent misalignment between the SR and the ASL dataset between different runs in the same session was achieved.

7 Feasibility of 3D-GRASE Renal ASL in a Paediatric Cohort with Impaired Renal Function

7.1 Introduction

Having previously established ASL in healthy volunteers (see previous work by our group in (Cutajar et al. 2012) and Chapter 5 within this thesis), this chapter presents a cross-sectional study delving into the question of whether ASL is feasible in paediatric patients with different degrees of renal failure. This chapter examined two hypotheses:

1. Reproducible Renal Blood Flow (RBF) measurements are possible in children with CKD.
2. RBF and the estimated glomerular filtration rate (eGFR) are directly correlated.

To investigate these hypotheses, eleven children with severe CKD of different pathophysiological origin (eGFR (ml/min per $1.73m^2$) = 26 ± 9 (mean ± standard deviation), 12-47 (range)) underwent the same ASL acquisition on two different days to assess the repeatability of the RBF measurements and whether ASL can provide clinically-relevant information. Furthermore, the challenges of applying ASL in a paediatric cohort are described, namely in terms of acquisition (image artefacts, especially as they relate to movement issues) and image analysis (selection of regions of interest for RBF analysis and delayed bolus arrival times).

7.2 Methods

7.2.1 Recruitment

Between March and July of 2016, eleven patients (8 male) underwent MRI scans on two different occasions. This is the same cohort of patients originally introduced in Chapter 6 (section 6.2.1). After an initial interaction with the patients and their parents, where the overall aims of the study were explained, children/parents who agreed to participate in the study were provided with information sheets and videos to explain the MRI procedure. On the day of each scan, both the parents and the child were given the opportunity to ask questions prior to the start of the MRI session. To reduce anxiety and maximise comfort, the children were encouraged to bring a DVD to watch during the scan using a mirror system mounted on a frame positioned over the child's head and a pillow was placed under the legs to minimise lumbar lordosis. Parents were given the option to remain in the scanner room during image acquisition and encourage their child to remain still for the duration of the scan.

7.2.2 MR acquisition and image processing

All data acquisitions were performed on a 1.5T Siemens Avanto system (Siemens Healthcare, Erlangen, Germany). The MR parameters were described in Section 5.2.1 and Section 6.2.2. All children underwent two single-TI ASL scans with an inflow time of 1.2s; one at the beginning of the session and the other at the end of session. In addition to the single-TI ASL scan with an inflow time of 1.2s, some patients (depending on their tolerance level to MRI scanning) had additional ASL runs at longer inflow times of 1.4, 1.6 and 1.8 seconds, albeit with 15 averages rather than the 25 averages that were acquired for the main ASL scan. Quantitative RBF results from these scans with longer TI are not reported in this work; they were only used to confirm cases where it was suspected that the inflow time of 1.2s was not sufficiently long.

Data pre-processing consisted of the methods described in Chapter 6. Motion correction was performed using rigid registration for both the saturation recovery and ASL data, combined with noise/outlier removal using weighted averaging of the ASL data.

Quantification of single-TI data was performed as previously described Chapter 6 using a single-compartment modelling approach and subject-specific voxelwise T_1/M_0 maps obtained from separate saturation recovery / proton-density acquisitions.

7.2.3 Data analysis

7.2.3.1 Region-of-interest (ROI) selection

Several criteria were established for the manual selection of regions-of-interest (ROIs) for RBF quantification:

1. Two types of ROIs were drawn:
    a. Cortical ROIs or functional renal parenchyma (FRP) ROIs (the differences between these are discussed below).
    b. Whole-kidney ROIs.
2. All ROIs were drawn on the PD images (acquired in the same space as the ASL data). This was done so that no renal parenchyma was missed, as it could be the case if the ROIs had been drawn using the ASL data only (given the low perfusion in patients with advanced CKD).
3. ROIs were drawn in all slices of the 3D dataset with clearly distinguishable renal tissue. The drawing of ROIs in the edge slices was done conservatively to avoid including voxels with obvious partial volume effects.
4. The approach to drawing the ROIs differed depending on the presence or absence of significant dilatations in the kidney.

a. For patients without dilations or when dilations were not part of the renal tissue (such as in patient #7), the cortical ROIs were drawn using the same methods as for the healthy volunteers: including only voxels of the outer region of the kidney, approximately 2 to 3 voxels wide (9 – 13.5mm at the native resolution of the scans). In these patients, the whole-kidney ROIs consist of the cortical ROIs plus all the remaining renal parenchyma.

b. For patients with large dilatations, all non-dilated renal parenchyma was included in the regions (henceforth referred to as functional renal parenchyma (FRP) regions). No cortico-medullary differentiation was possible. In these patients, the WK ROIs are equal to the FRP ROIs, as no more parenchymal renal tissue was left after drawing the FRP ROIs.

5. Any dilatations were tracked across the multislice dataset to assist on the determination of their borders with the goal of avoiding their inclusion into both the cortical/FRP and whole-kidney ROIs.

Examples of ROIs can be seen in Figure 7-1 and Figure 7-2, respectively for one patient with dilatations and one without dilatations.
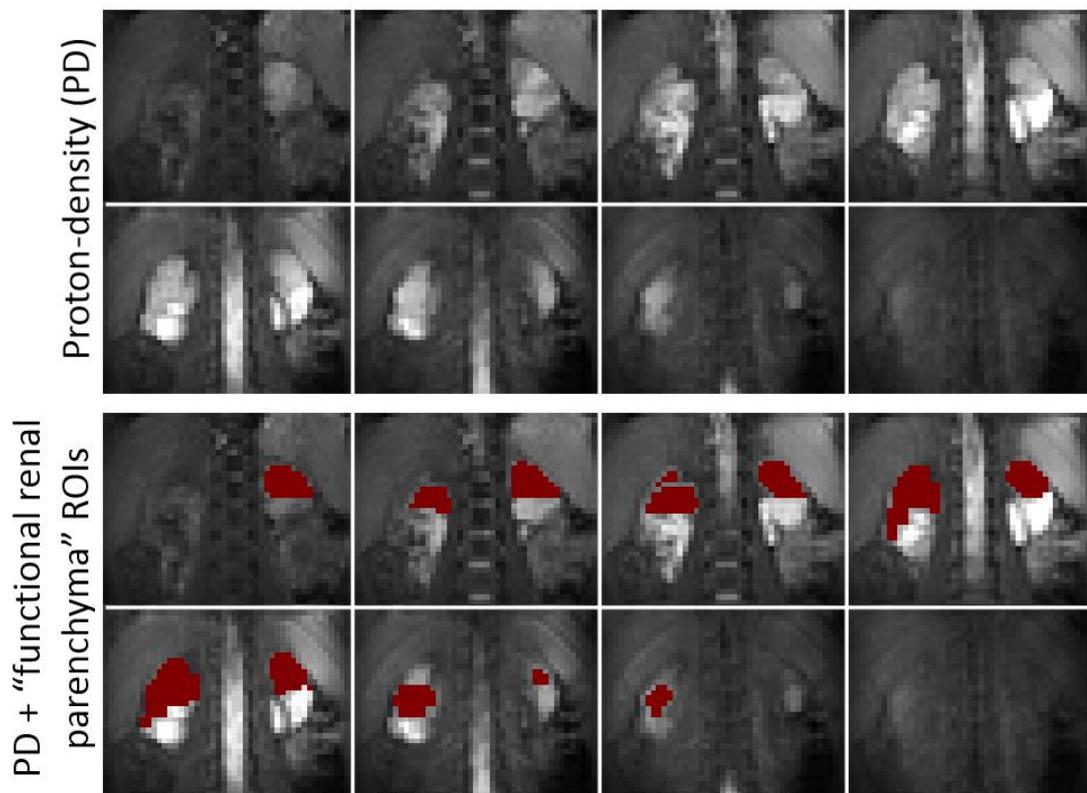
Figure 7-1 Regions of interest (ROIs) in a patient with dilatations. In patients with dilatations, the same ROI was considered for deriving functional renal parenchyma and whole-kidney statistics.
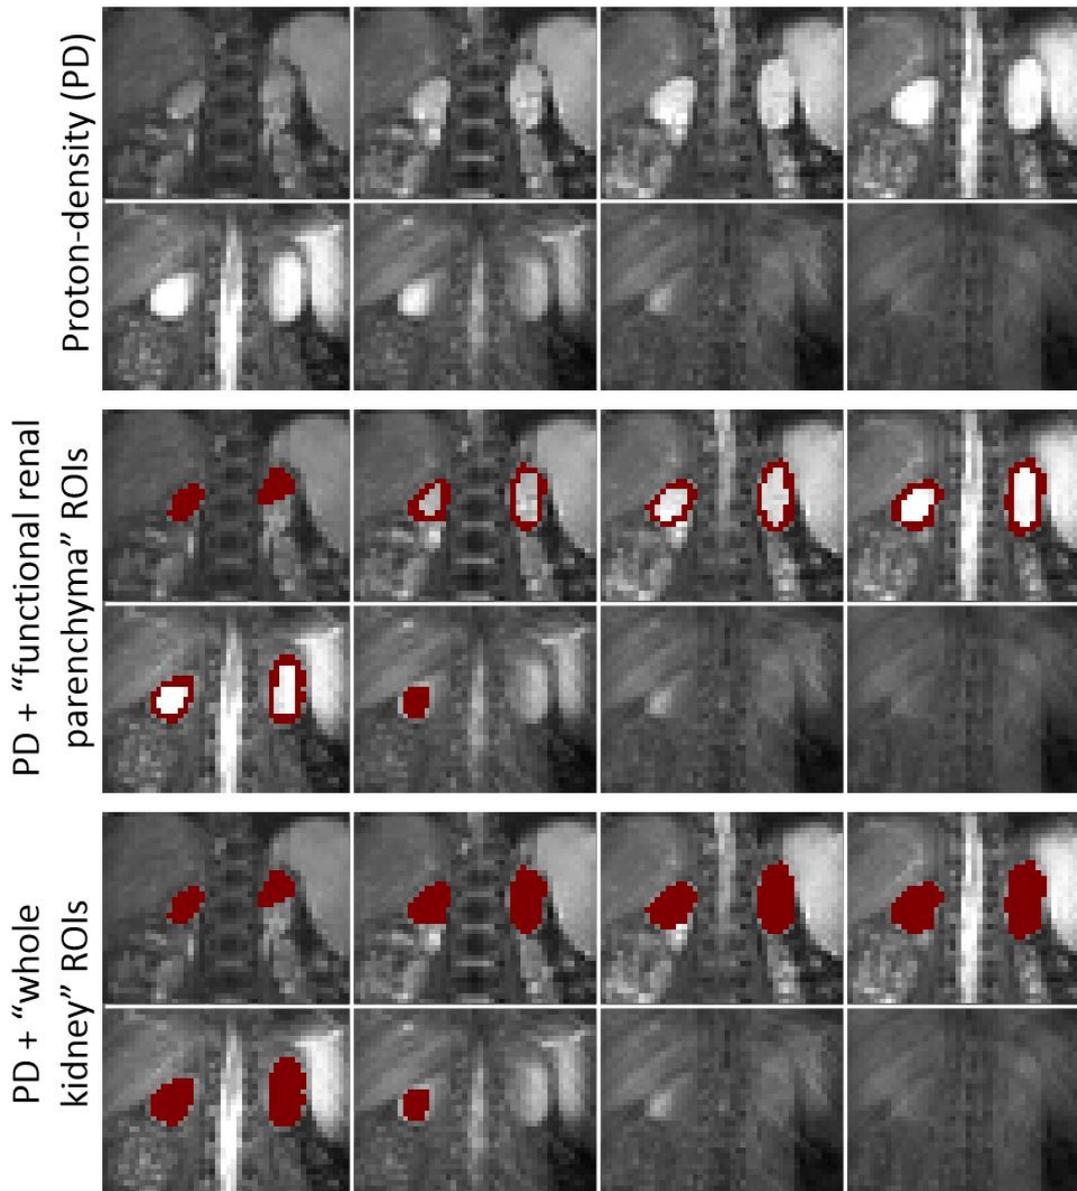
Figure 7-2 Regions of interest (ROIs) in a patient without dilatations.

7.2.3.2 Data quality assessment

Residual uncorrected motion poses one of the main challenges in the assessment of the repeatability of ASL. Therefore, after motion correction, the data quality of all scans was assessed by visual inspection (separate for each kidney) to allow a fair assessment of drawbacks and merits of ASL without being confounded by major movement issues. Only RBF results from datasets of sufficient quality were included in the Bland-Altman and correlation plots (see Section 7.2.5). The inclusion criteria were:

- No obvious artefacts in the final averaged PWIs, T_1 maps and PD images

- Good alignment between T_1, M_0 and perfusion-weighted data confirmed by visual subjective evaluation, based in the types of artefacts found when clear misalignments between these datasets exist (see Chapter 6 for a thorough discussion).

7.2.4   Clinical topics

7.2.4.1   Assessing clinical stability of the patients

Serum creatinine measurements for all patients were obtained from routine blood sampling. Due to logistic reasons, blood samples were only obtained in days when the patients attended clinical appointments, which did not necessarily correspond to the days when MR scanning was performed. The time gap between blood sampling and the MR acquisition is reported below.

Several previous studies have used different heuristics to establish whether subjects are in a clinically stable condition with respect to their renal function between MR visits. For example, Lanzman *et. al.* (Lanzman et al. 2010) considered subjects clinically stable if their variation in serum creatinine was less than 20% in the four months prior to the study, whereas Artz et. al. (Nathan S. Artz et al. 2011) considered patients stable if their serum creatinine varied by 0.3mg/dL, (i.e. 26.5 μmol/L) or less between MR visits (on the order of months).

In this study, the two scans for each patient were performed within a short period (time between scans: 23±10 (mean ± std), range: 7-35, median 21 days). For this reason, and given that for all patients no acute changes in clinical status occurred between visits, we assume all the children were clinically stable between the two MR sessions.

Stability over a longer term was also assessed. Given the variable frequency of blood sampling in our cohort, patient clinical stability was established based on the serum creatinine measurements closest to the scan days (shown in Table 7-1) and measurements available in the months prior to the scans (up to 1 year before), considering patients stable if the variation in serum creatinine was less than 20% following Lanzman's approach (Lanzman et al. 2010). The preliminary results of patient follow-up (up to 1 year after the MRI scans) are also reported where available.

Table 7-1 Demographic and clinical data for each patient in the cohort. [a]Where only one blood pressure value is shown it refers to the diastolic blood pressure. [b]Negative values indicate that the MR was performed before the creatinine measurement.

#	Age (years)	Sex	Height (m)	Weight (kg)	BP (mmHg)[a]	Serum creatinine (SCr) (μmol/L) @ scan 1		@ scan 2		eGFR (ml/min/1.73m^2) @ scan 1	@ scan 2	Primary diagnosis	Dilatation
						SCr	Δt^b (days)	SCr	Δt^b (days)				
1	10	M	1.34	31	99/56	169	5	169	33	29.0	29.0	Posterior urethral valves	Yes
2	14	M	1.52	51	104/80	264	0	278	0	21.0	19.9	Bilateral renal dysplasia	No
3	12	M	1.52	52	126/78	304	0	340	0	18.3	16.3	Nephronophtisis Type 4	No
4	10	F	1.21	32.9	100	226	0	219	-28	19.6	20.2	Bilateral renal dysplasia	No
5	7	M	1.15	24.8	90/50	146	-30	146	-2	28.7	28.7	Posterior urethral valves	Yes
6	17	F	1.51	55	114	117	61	117	82	47.1	47.1	Acute tubular necrosis	No
7	15	M	1.65	61	126/76	436	11	498	-22	13.8	12.1	Posterior urethral valves	Yes
8	13	M	1.68	55.5	142	251	0	273	0	24.4	22.5	Haemolytic uraemic syndrome	No
9	16	M	1.81	79	104	244	42	247	4	27.1	26.8	Renal Cyst and Diabetes Syndrome	No
10	12	M	1.47	34.8	118/78	152	12	152	20	35.3	35.3	Posterior urethral valves	Yes
11	9	F	1.25	23	97/53	158	89	174	-86	28.9	26.3	Dysplasia and neuropathic bladder	No

7.2.4.2 Calculation of eGFR

The estimated Glomerular Filtration Rate ($eGFR$) was calculated from serum creatinine (SCr) (obtained from routine clinical blood sampling) and height (h) measurements using the "Bedside Schwartz" equation, which in SI units is as follows (Schwartz et al. 2009):

$$eGFR = 3654.86 \times \frac{h}{SCr} \qquad \text{Equation 7-1}$$

Where the respective units of $eGFR$, h, and SCr are $mL/min\ per\ 1.73m^2$, m, and $\mu mol/L$.

7.2.5 Statistical Analyses

The agreement of the RBF measurements between runs in a given day and between days was evaluated. Bland Altman plots, as well as to a lesser extent correlation plots, were used to visualize the spread of RBF among the various measurements in each subject. Bland Altman plots were constructed by plotting the difference in RBF measurements against to their means. The 95% limits of agreement are obtained by computing the mean $\pm$ 1.96 $\times$ standard deviation of the difference in the RBF estimates. T-tests were used to test the null hypothesis that the RBF differences come from a distribution with zero mean. The correlation between RBF (in ml/100g/min) measured using a TI of 1.2s and eGFR (ml/1.73m²/min) was evaluated by computing the Pearson's correlation coefficient, where *p* values smaller than 0.05 were deemed to be significant.

7.3 Results and discussion

7.3.1 Recruitment

Despite liaising with the parents and the child before the MRI session, identifying children that were able to cooperate during the scan presented a significant challenge, a finding also described elsewhere (e.g. (Cahoon 2011)). Some centres already used mock MRI scanners to attempt to rule out subjects whose anxiety levels and movement during a MRI practice session suggest the actual imaging experiment is unlikely to provide good quality data (Hallowell et al. 2008). Though the benefits of this approach are clear, we did not have access to a mock scanner at GOSH and so were unable to assess potential volunteers for the study in this way.

Twelve children/parents agreed to take part in the study. One child was unable to undergo MRI scanning as despite careful preparation and setup she was unable to remain in the scanner for more than a couple of minutes. Therefore, eleven children (age (years): 12±3 (mean ± std), range: 7-17, median 12) were scanned on two occasions. Four children had posterior urethral valves, three children had renal dysplasia, one child had nephronophtisis, one had acute tubular

necrosis (ATN), one had haemolytic uraemic syndrome and one had renal cyst and diabetes syndrome (see Table 7-1). In one subject, the MR session had to be prematurely aborted approximately 25 minutes after the session start due to perceived discomfort and anxiety. In this particular case, the MR protocol was slightly adapted to prioritise the ASL sequences, which allowed test-retest data to be obtained despite the shortened MR session.

7.3.2   Patient's clinical status

Demographic and clinical data for each of the 11 patients enrolled in this study are displayed in Table 7-1. The time elapsed between the blood sampling and the MRI scan of the corresponding patients was 24±30 (mean ± std), range: 0-89, median 11.5 days. (see results for each patient/day on Table 7-1.

As discussed in Section 7.2.4.1, the clinical status of all children remained stable between the MR scans. In terms of longer term stability, eight out of the 11 patients (specifically patients #1, #2, #4, #5, #6, #9, #10, #11) were clinically stable based in the criteria defined in Section 7.2.4.1. One year following the MR scans, the 3 patients with progressive CKD (patients #3, #7 and #8) had reached end-stage renal disease and either started preparation for (or were already undergoing) renal replacement therapy.

Across the whole cohort, serum creatinine levels were 231 ± 99 (mean ± std), range: 117-498, median 223 µmol/L corresponding to eGFR levels of 26 ± 9 (mean ± std), range: 12-47, median 27 ml/min per 1.73m^2. The number of patients at each CKD stage across the cohort is: Stage 3A = 1; Stage 3B = 1; Stage 4 = 8; Stage 5 = 1.

Four out of the eleven children presented with renal dilatations. In some patients, significant atrophy and thinning of the renal cortex can be seen, particularly in cases with hydronephrosis, where dilation of the renal pelvis and calyces has occurred. A loss of cortico-medullary differentiation was also apparent in the T_1 maps (see Figure 7-3), similar to clinical ultrasound findings.
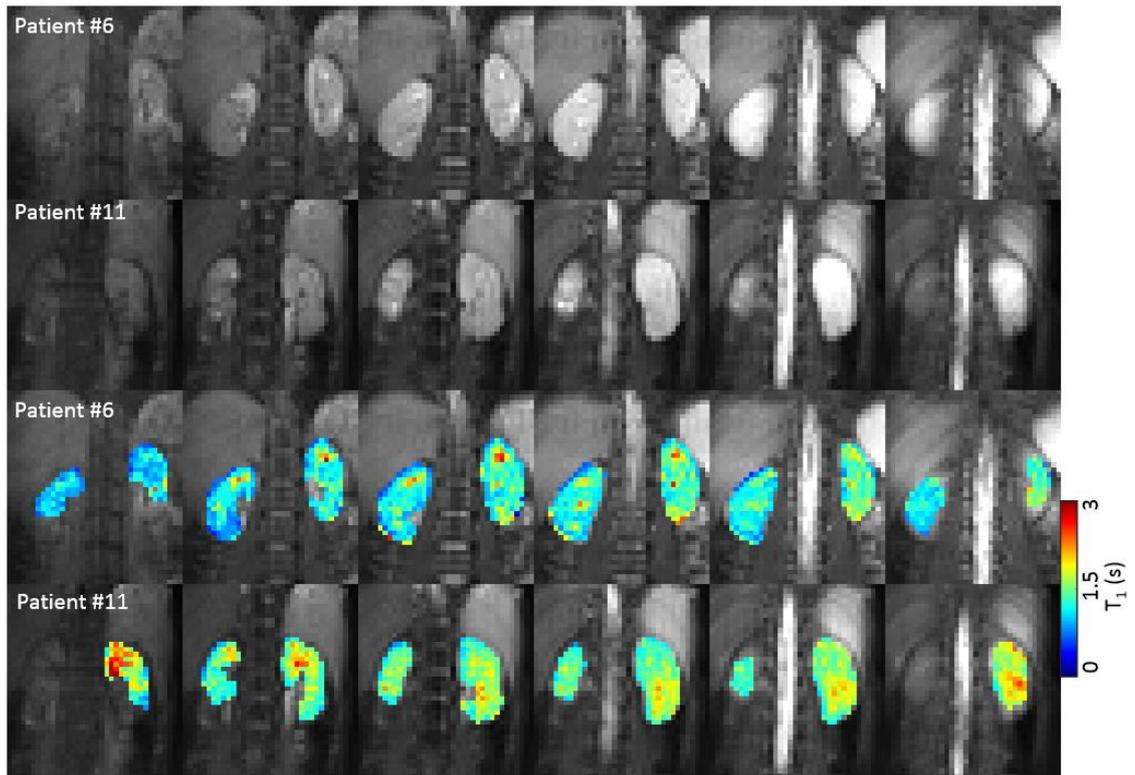
Figure 7-3 Loss of cortico-medullary differentiation in two patients. Rows 1 and 2 (from top): PD images. Rows 3 and 4: T_1 maps in the whole-kidney ROIs overlaid on the PD images.

7.3.3 Data quality assessment and handling of exceptional data

The results of the data quality assessment are summarised in Table 7-2. This table provides information (separately for each kidney) on which scans met or failed the data quality criteria (within the latter are those that failed due to scanning failure or excessive unsalvageable motion). In addition, kidneys which could not be analysed are specified, either because they could not be accurately discerned in the PD images, therefore resulting in a very low degree of confidence in drawing the ROIs (such as the left kidney of patient #9), or because no functioning renal parenchyma could be segmented (e.g. left kidney of patient #5 where only severe dilations are visible due to obstruction and associated reflux). Furthermore, the table is color-coded to highlight which intra and inter-session comparisons were not possible following the exclusion of corrupted data therefore specifying which data could not be included in the plots shown in this chapter.

Table 7-2 Results of visual evaluation of data quality. Only the scans that met the inclusion criteria described in Section 7.3.3 (checkmarked) were included in the plots seen in this chapter.

#	Left Kidney Day 1 Run 1	Left Kidney Day 1 Run 2	Left Kidney Day 2 Run 1	Left Kidney Day 2 Run 2	Right Kidney Day 1 Run 1	Right Kidney Day 1 Run 2	Right Kidney Day 2 Run 1	Right Kidney Day 2 Run 2
1	✓	✓	✓	✓	✓	✓	✓	✓
2	✓	✓	✓	✓	✓	✓	✓	✓
3	✓	✓	✓	✓	✓	✓	✓	✓
4	✓	✓	X	X	NoK	NoK	NoK	NoK
5	NoK	NoK	NoK	NoK	✓	X	✓	X
6	✓	NaN	✓	✓	✓	NaN	✓	✓
7	✓	✓	✓	✓	NoK	NoK	NoK	NoK
8	✓	✓	✓	✓	✓	✓	✓	✓
9	NoK	NoK	NoK	NoK	✓	✓	✓	✓
10	✓	✓	✓	NaN	✓	✓	✓	NaN
11	✓	✓	✓	✓	✓	✓	✓	✓

Legend:
- ✓ Data met inclusion criteria
- X Data failed inclusion criteria
- NoK Kidney not analysed
- NaN Data not available (scan failure)
- (yellow) Intra-session comparison impossible
- (cyan) Inter-session comparison impossible
- (green) Intra/inter-session comparison impossible

In summary, following motion correction, the majority of scans were deemed of good quality and therefore used for further analyses. The averaged PWIs show a satisfactory signal-to-noise ratio (SNR) despite reduced perfusion, which can be attributed to the 3D sequence that ran with multiple averaging at a relatively coarse resolution. Out of the 44 scans (11 patients, 4 runs each), 2 were discarded due to scanning failure: fat shift artefact (patient #6) and triggering failure (patient #10), as detailed previously (Chapter 6). Four scans were rejected due to excessive motion/image registration failure. Details on the failure of image registration in these cases are discussed in Section 6.3.3.2. Therefore this translates to an 86% (38/44) success rate in this cohort of children.

Children with single functioning kidneys included patients #4, #5, #7 and #9. In terms of individual kidneys, in 4 children the function of one kidney was so poor as to be unanalysable on ASL. Nevertheless, all 4 only contributed to a very small fraction of the total renal perfusion in each child, which can be visually confirmed by examining to the perfusion-weighted data, particularly when displayed together with the complementary anatomical information (e.g. using the PD-PWI fusion images). Of the 72 possible kidneys (seven patients with two kidneys and four runs and four patients with a single kidney and two runs), 6 had to be excluded. Specifically, the right kidney of patient #4 is hypoplastic and contributed to less than 10% of the total renal function when assessed by a Nuclear Medicine DMSA scan at 3 years of age. Its small size renders ROI drawing challenging due to partial volume effects and image registration is impaired due to the very small number of voxels included in the rectangular renal mask, which specifies the region of the FOV to be registered. Both the left kidney of patient #5 and the right kidney of patient #7 were considered non-functioning following MAG3 renograms and show

little to no perfusion weighted signal in the ASL data. The left kidney of patient #9, despite being larger than the other rejected kidneys, is still small compared to the contralateral kidney, and was difficult to visualise and segment in the PD images (similar finding in clinical ultrasound) with very low perfusion-weighted signal.

Patient #3's day#1−run#1 scan was the only case where using image registration introduced errors in the data. There was relatively little patient movement during this particular scan, which combined with data averaging resulted in a sufficient quality PWI, therefore allowing the criteria of data quality to be met even without motion correction. Therefore, in all analyses, the results from day#1−run#1 of patient #3 refers to the data without motion correction. This particular ASL run was unique in this regard.

7.3.4 Effect of respiratory triggering on acquisition time

As discussed in Section 4.3.3, the "true" acquisition time (TA) of respiratory-triggered scans ($TA_{trigger}$) is longer compared to the nominal scan time of non-respiratory triggered acquisitions ($TA_{nominal}$) due to the fact that with respiratory-triggering, the actual TR of the sequence depends on the subjects' respiratory cycle. For the ASL scanning parameters used in this study (single-shot single-TI ASL with 25 measurements), $TA_{nominal}$ is 150 seconds ($2 \times 3[s/TR] \times 25[TR]$). Table 7-3 shows the "triggering time penalty factor" ($TA_{trigger}/ TA_{nominal}$) as well as $TA_{trigger}$ (in seconds) for all ASL runs of all patients in this cohort (total number of runs = 42).

Table 7-3 $TA_{trigger}/TA_{nominal}$ and $TA_{trigger}$ for all ASL runs. The underlined values correspond to the mean and standard deviation across all runs of all patients (total number of runs = 42).

| Patient | $TA_{trigger}/TA_{nominal}$ | | | | | | $TA_{trigger}$ (s) | | | | | |
| | Day 1 | | Day 2 | | mean | std | Day 1 | | Day 2 | | mean | std |
	Run 1	Run 2	Run 1	Run 2			Run 1	Run 2	Run 1	Run 2		
1	1.92	1.91	1.43	1.84	1.77	0.23	288	287	214	275	266	35
2	1.40	1.53	1.69	1.68	1.57	0.14	210	230	253	251	236	20
3	1.49	1.51	1.58	1.50	1.52	0.04	223	226	238	225	228	7
4	1.64	1.71	1.35	1.26	1.49	0.22	246	256	202	189	223	33
5	1.75	1.53	1.80	1.90	1.74	0.15	263	230	270	284	262	23
6	1.58	N/A	1.60	1.44	1.54	0.08	236	N/A	240	217	231	13
7	1.40	1.33	1.32	1.48	1.38	0.08	211	200	197	223	208	12
8	1.30	1.57	1.28	1.55	1.42	0.16	194	235	191	233	213	24
9	1.36	1.58	1.31	1.54	1.45	0.13	204	236	196	231	217	20
10	1.53	2.00	2.04	N/A	1.86	0.28	229	299	307	N/A	278	43
11	1.54	1.62	1.74	1.88	1.70	0.15	231	243	261	282	254	22
mean	1.54	1.63	1.56	1.61	1.58		230	244	234	241	237	
std	0.17	0.19	0.24	0.20		0.21	26	28	35	30		31

In summary, the actual scan time of the ASL runs in this study was 237±31s (range = 189-307s). That is, using respiratory triggering doubled the scan time in the worst cases, and at best increased the scan by approximately 25%.

7.3.5 ROI definition

Only the PD images were used for ROI definition. This avoids an underestimation of the number of renal parenchyma voxels due to reduced perfusion signal particularly in the patients where perfusion impairment is most severe, which would bias the RBF quantification results. All PD images were of high quality, free of movement artefacts owing to the single-shot nature of the acquisition. Good tissue contrast is observed particularly between functioning cortex and lesions characterised by a very long T_1 (such as dilatations and cysts), which in several cases allowed them to be excluded from the functional renal parenchyma ROIs provided their size was sufficient to make them resolvable at the scan's set voxel size. Nevertheless, ROI selection proved more challenging than on healthy volunteers as seen in Chapter 5. Despite the volumetric data set, avoiding the inclusion of non-parenchymal regions was challenging in patients, especially as this was done based on the PD images alone and no functional information guided this process (to avoid introducing bias, as described above). A retrospective visual analysis of the PD-PWI fusion images together with the manually drawn ROIs suggests that in some cases (e.g. patient #1) non-parenchymal kidney regions (dilatations) were included in the ROIs. A similar outcome is likely to have occurred in patient 9, where cysts close to the kidney boundaries were included in the cortical ROI. When this happens, longer mean T_1 values are obtained and the mean RBF is biased towards lower results. Drawing the ROIs for volume measurements in the PD images may also result in some errors due to partial volume effects, particularly in the edge slices, due to a combination of the relatively coarse resolution and the blurring characteristic of the single-shot 3D readout.

This raises the question on which image should the ROI be drawn. If the ROI is based on the PD image, then the ROI may include anatomically present renal parenchyma that is not perfused. If the ROI is based on the PWI then all the parenchyma might not be included. The volumes of these two ROIs will be different giving different results of RBF in ml/kidney/min. The importance of the difference between these two ROIs remains an open question.

Furthermore, volume measurements using the PD images are further complicated by a limitation intrinsic to the FAIR labelling scheme. In renal ASL experiments using the FAIR labelling scheme, it is necessary to avoid including the aorta within the slab-selective inversion region. This ensures that inflowing blood is not inverted (labelled) in the control acquisition. Depending

on the specific anatomy of each subject, specifically on the relative position of the kidneys and the aorta, meeting this criterion while imaging the entirety of both kidneys simultaneously may not be possible. Figure 7-4 depicts one case (patient #3) where this is the case. This illustration shows anatomical reference scans (acquired in the three anatomical planes using a TrueFISP sequence) for positioning of the ASL imaging volume. As can be seen, in order to exclude the aorta from the selective inversion slab (which by definition closely matches the imaging slab), the selective inversion region has to be moved posteriorly. This causes a significant anterior portion of the kidneys not to be included in the imaging volume, thus preventing estimation of the whole-kidney volume from these images. Visual inspection revealed that the entirety of both kidneys was not included in the field of view in the more than half of the subjects. Since in this work all ROIs were drawn on the PD images, all RBF results are shown in ml/100g/min, and not ml/kidney/min.
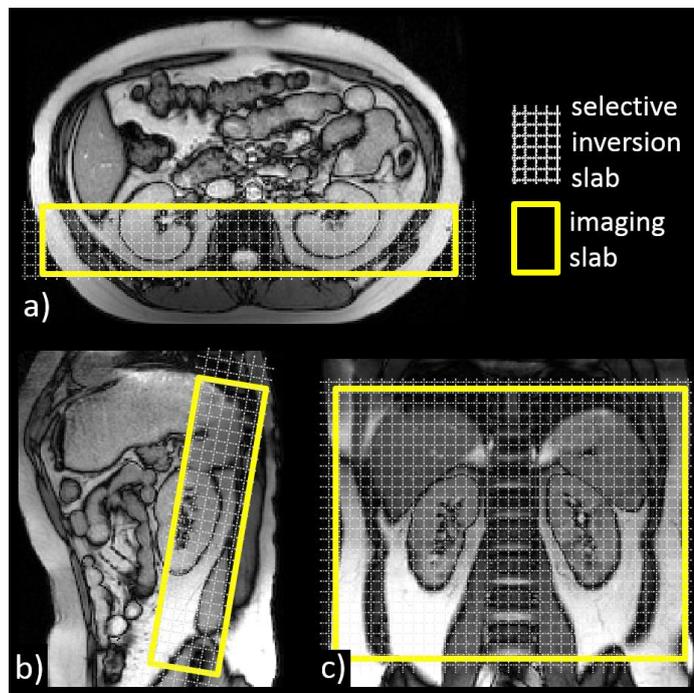
Figure 7-4 Incomplete kidney coverage due to FAIR labelling. Anatomical reference TrueFISP scans used for position of the ASL selective inversion and imaging volumes. a) Axial slice; b) Sagittal slice; c) Coronal slice. The axial dataset is crucial for positioning of the ASL imaging volume and inversion slab excluding the aorta so that inflowing blood is not labelling in the slab-selective (i.e. control condition). Depending on the relative anatomical position of the kidneys and the aorta, and assuming both kidneys are to be included in the imaging volume, then anterior portion of the kidneys may be excluded from the imaging region.

7.3.6   RBF Intra and Inter-session agreement

Bland Altman and correlation plots are shown for the intra and inter-session comparisons, both considering separate kidneys and for the case when perfusion is averaged across all voxels of

the two kidneys (see Figure 7-5 to Figure 7-12). Furthermore, plots were constructed for both the functional renal parenchyma and whole-kidney ROIs. As mentioned above, all scans deemed not of sufficient quality due to uncorrected movement artefacts or scanning failure are excluded from these plots. These are highlighted in the legend of the figures using a cross symbol. A strong correlation was found for all comparisons, particularly for the intra-session condition. Because a good agreement between measurements cannot be deduced from even a high correlation (Bland & Altman 1986), Bland Altman plots are also shown.
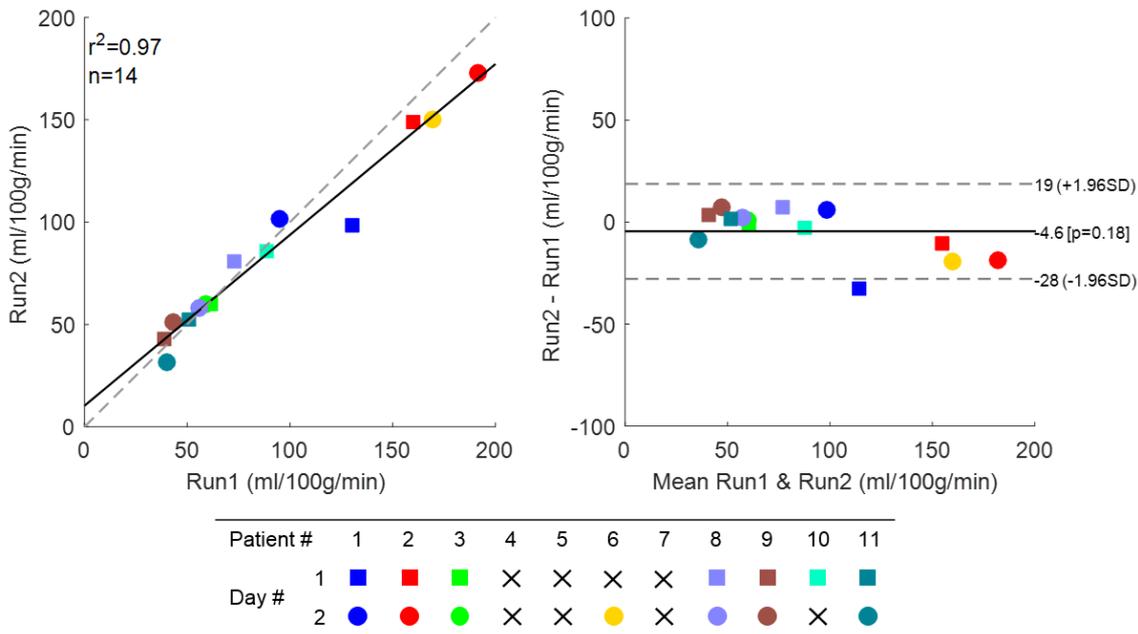
Figure 7-5 Mean RBF intra-session agreement (right kidney – FRP ROIs). Left: correlation between runs in a single day for all datasets of acceptable data quality (together with lines of equality to illustrate the agreement between scans). The scans marked with a cross in the legend table are not shown either because the kidney was not analysed (non-existent or impossible to segment) or the scan was deemed of insufficient data quality. See Table 7-2 for details.
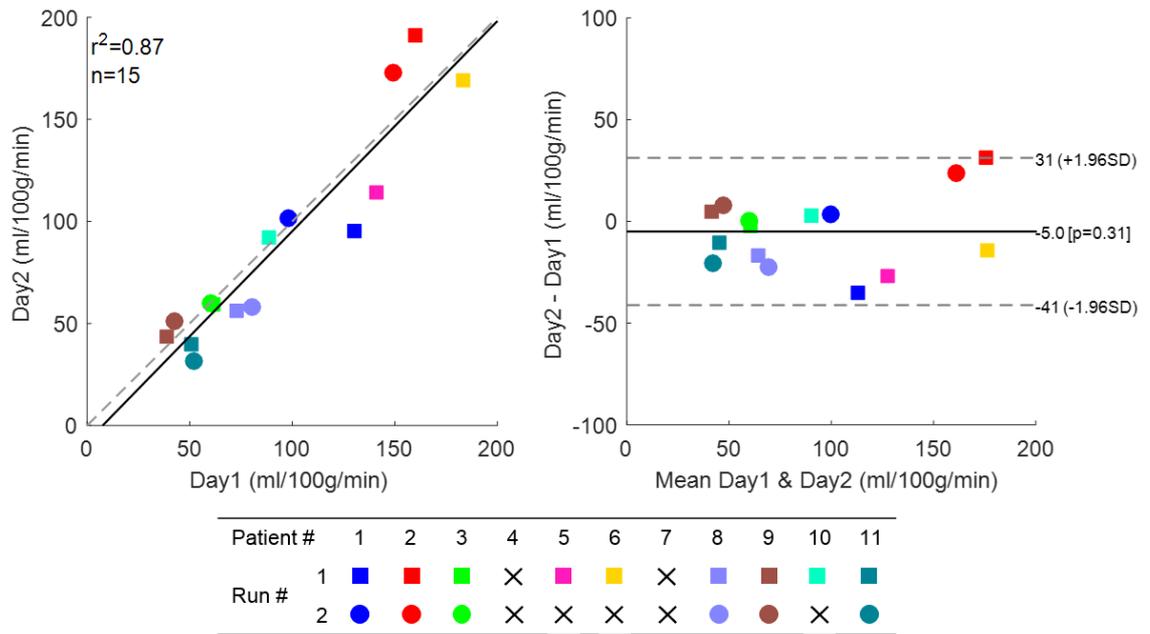
Figure 7-6 Mean RBF inter-session agreement (right kidney – FRP ROIs).
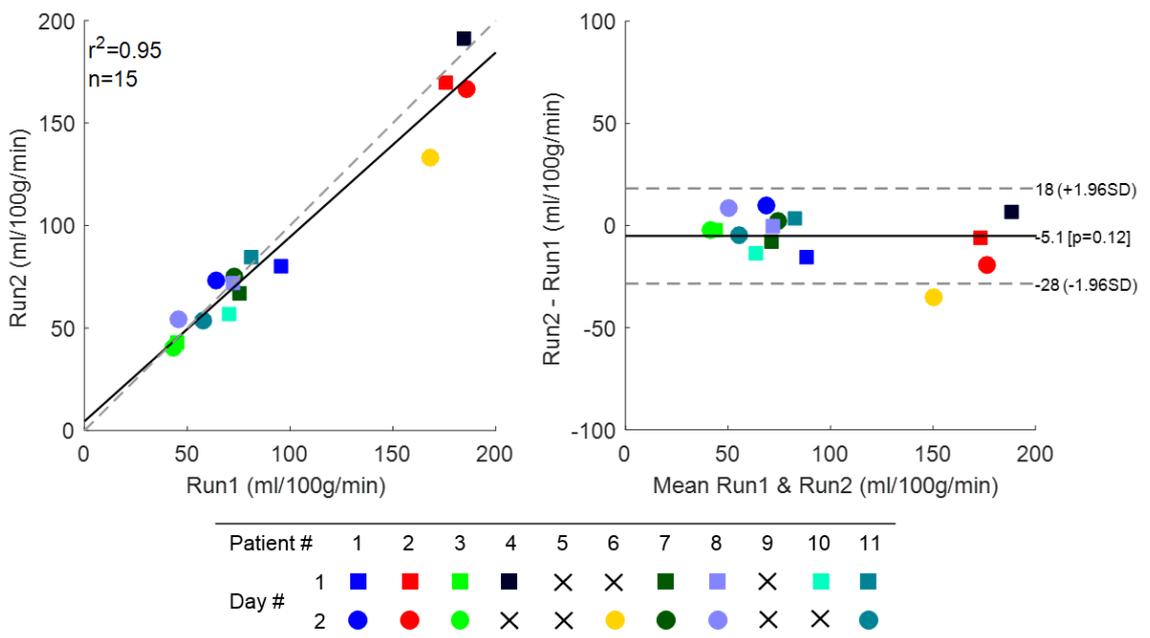
Figure 7-7 Mean RBF intra-session agreement (left kidney – FRP ROIs).
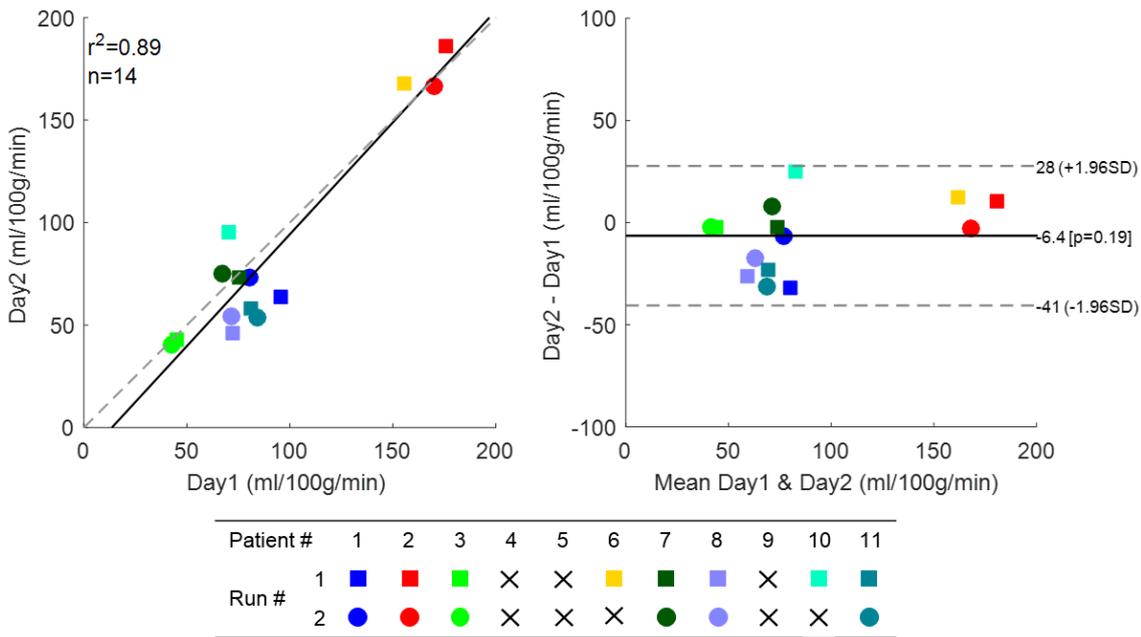
Figure 7-8 Mean RBF inter-session agreement (left kidney – FRP ROIs).
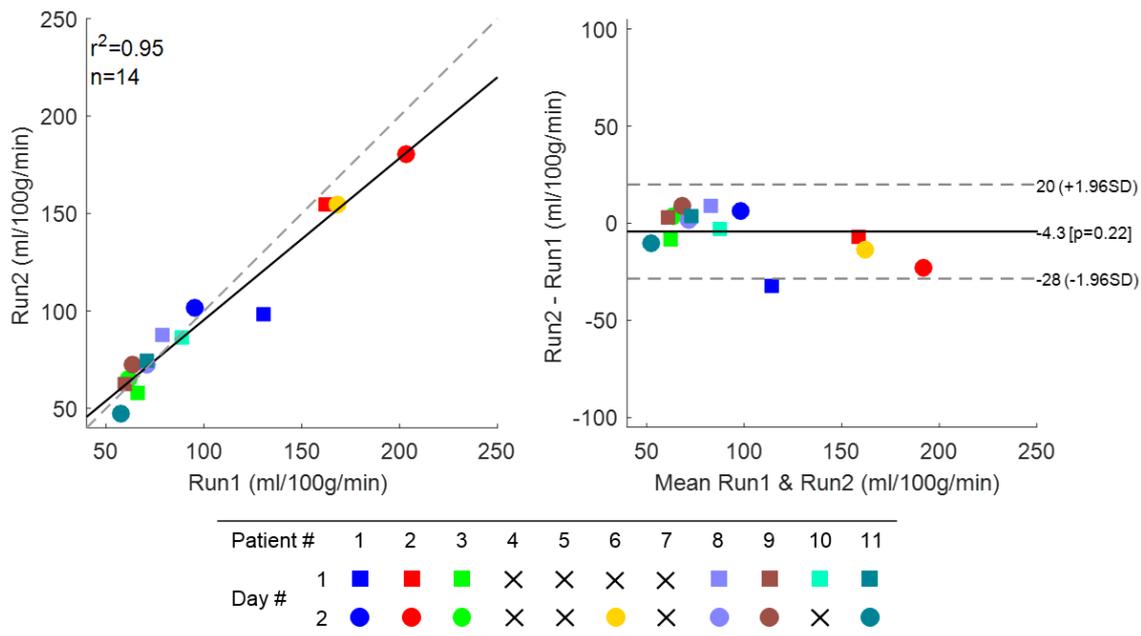
Figure 7-9 Mean RBF intra-session agreement (right kidney – WK ROIs).
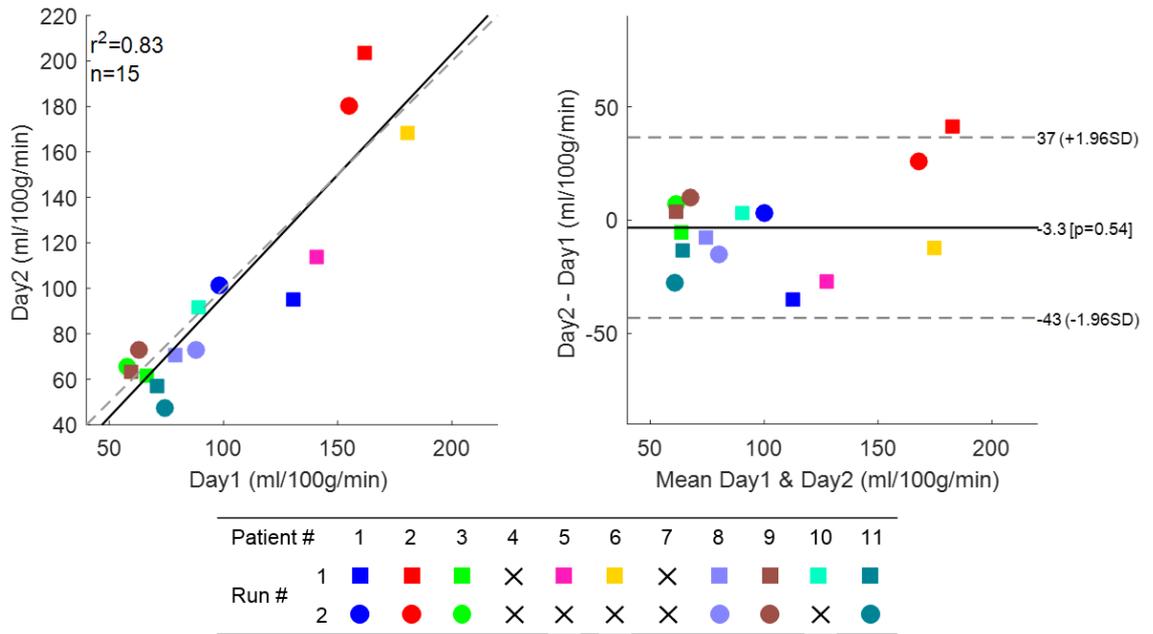
Figure 7-10 Mean RBF inter-session agreement (right kidney – WK ROIs).
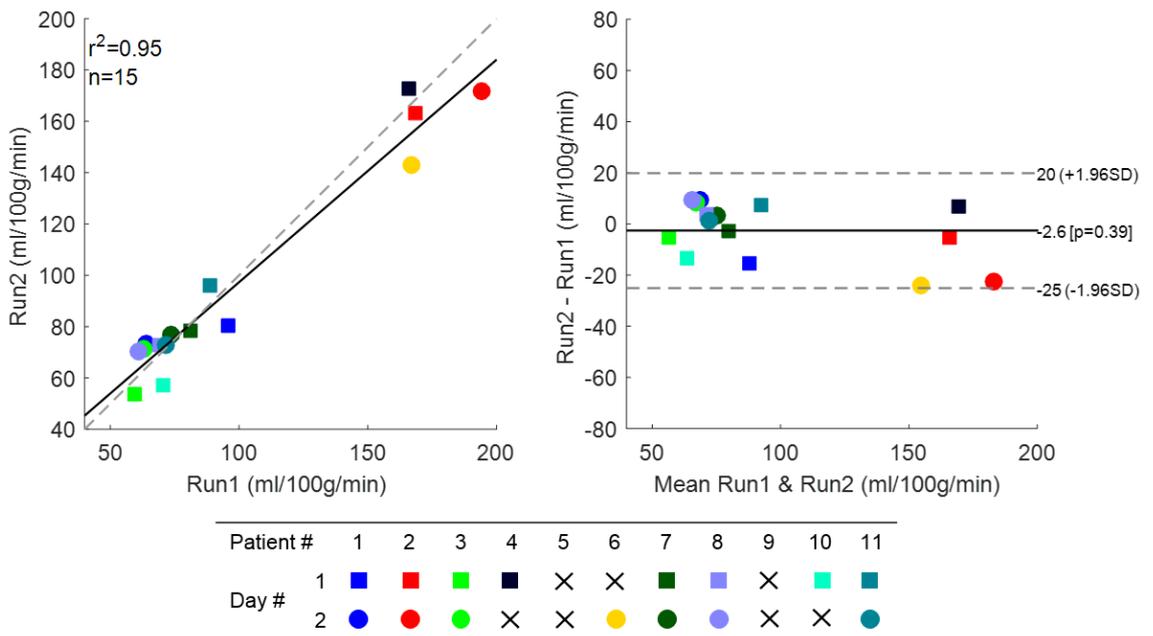
Figure 7-11 Mean RBF intra-session agreement (left kidney – WK ROIs).
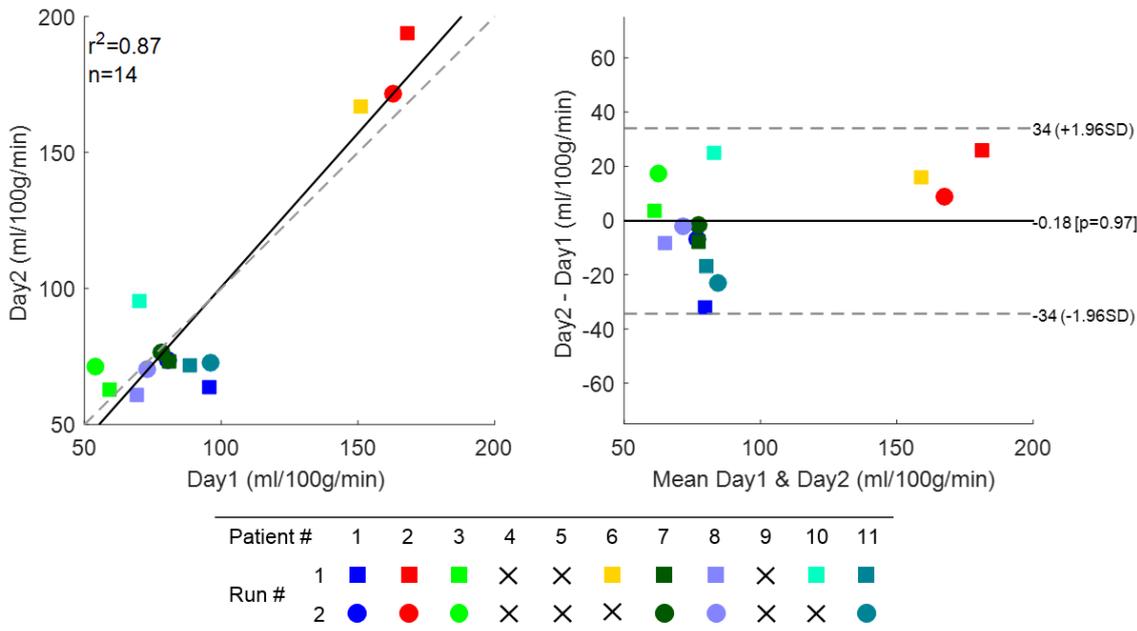
Figure 7-12 Mean RBF inter-session agreement (left kidney – WK ROIs).

A summary of the statistical parameters obtained from the Bland-Altman analysis is shown in Table 7-4, where the limits of agreement (LOA) are shown in ml/100g/min.

Table 7-4 Summary of statistics from the Bland-Altman plots, where LOA is in ml/100g/min

| | Functional renal parenchyma | | | | | | Whole kidney | | | | | |
| | Right | | | Left | | | Right | | | Left | | |
	N	LOA	CV (%)	N	LOA	CV (%)	N	LOA	CV (%)	N	LOA	CV (%)
Intra-session	14	23	14	15	23	13	14	24	13	15	22	12
Inter-session	15	36	20	14	34	20	15	40	20	14	34	18

In all plots, the mean of the differences is relatively small, with the largest bias (6.4 ml/100g/min) being found in the inter-session comparison of the left kidneys considering the FRP ROIs. For each comparison, t-tests have shown that the null hypothesis (RBF differences coming from a distribution with zero mean) is not rejected at the 5% significance level. The limits of agreement are calculated based on the spread of the differences of the RBF measurements (mean $\pm$ 1.96 $\times$ standard deviation of the differences). As such, the majority of points on the plots will lie within these limits, except for outlier cases where the differences between measurements are the greatest. These include the intra-session comparisons of patient#01-day#1 (Figure 7-5 and Figure 7-9) and patient#06-day#2 (Figure 7-7). Since the M_0 and T_1 maps are shared among the intra-session scans, these do not account for the variability of RBF in these cases. In the case of patient#06-day#2, the variability can be at least partly explained due to an exaggerated through plane translation found by the image registration algorithm when registering the left kidney. In the case of patient#01, even though no motion issues are seen in both runs of the first day, the

ASL PWI from the second runs shows reduced PW signal, resulting in the lower mean RBF obtained. Inter-session comparisons that yield differences outside of the limits of agreement include: Patient#02-Run#1 (Figure 7-6 and Figure 7-10) and Patient#10-Run#1 (Figure 7-8), where in both cases RBF increased in the second day. Good alignment between the different datasets in these scans was confirmed, therefore physiological variation may play a role in these variations, especially as the changes in mean T_1 between these days are relatively small (approximately 100ms and 50ms, respectively).

7.3.7   RBF vs eGFR

Plots of the mean RBF in each kidney obtained in all scans vs. eGFR are shown in Figure 7-13 and Figure 7-14, respectively for the FRP and WK ROIs. Analogous plots with the mean RBF averaged across all voxels of both kidneys are shown in Figure 7-15 and Figure 7-16.The mean functional renal parenchyma / cortical RBF at the inflow time of 1.2s in this patient cohort was, respectively 101 ± 52 (mean ± std), range = 51-188 ml/100 g/min, significantly lower than in the cohort of adult healthy volunteers presented in Chapter 5 (295 ± 97 ml/100 g/min). Whole-kidney RBF was 105 ± 44, range = 61-175 ml/100 g/min (see Table 7-5). All patients with progressive CKD (patients #3, #7 and #8), which at 1 year following the MRI scans require renal replacement therapy had low RBF values, in the range of 60-75 ml/100g/min. No statistically significant correlations were found between RBF (in ml/100g/min, averaged across both kidneys) at the inflow time of 1.2s and eGFR in both FRP and WK ROIs considering each run individually, respectively (R = 0.227, p = 0.17) and (R = 0.244, p = 0.14) (see Figure 7-15 and Figure 7-16, respectively). The same is verified after averaging the RBF results of the two runs in a given day, respectively (R = 0.295, p = 0.19) and (R = 0.307, p = 0.18).

Table 7-5 Mean FRP and WK RBF values obtained for each patient in all of the four ASL scans (two runs in each of two days), together with the mean and coefficient of variation of the RBF values of each patient averaged calculated across the four scans. Only scans of good quality were included (see Section 7.3.3). The colour coding for each entry in the table relates to the mean RBF values (lowest: red; highest: green) and the coefficient of variation of the mean RBF values obtained in the 4 scans for each patient (low CV → high repeatability: green; high CV → low repeatability: red).

| Patient | FRP RBF (ml/100g/min) | | | | | | WK RBF (ml/100g/min) | | | | | | %Diff(FRP,WK)* |
| | Day 1 | | Day 2 | | Mean | CV (%) | Day 1 | | Day 2 | | Mean | CV (%) | |
	Run 1	Run 2	Run 1	Run 2			Run 1	Run 2	Run 1	Run 2			
#1	116	90	82	90	95	15	same as FRP						0
#2	168	160	188	170	171	7	165	159	198	175	175	10	-2
#3	52	51	51	50	51	2	57	56	62	68	61	9	-19
#4	185	191	X	X	188	2	166	173	X	X	169	3	10
#5	141	X	114	X	127	15	same as FRP						0
#6	170	NaN	169	142	160	10	167	NaN	168	149	161	6	-1
#7	75	67	73	75	73	5	81	78	73	77	77	4	-7
#8	73	76	51	56	64	19	74	80	66	72	73	8	-14
#9	39	43	44	51	44	11	60	63	64	73	65	9	-47
#10	76	67	94	NaN	79	18	same as FRP						0
#11	73	76	53	47	62	23	84	91	68	66	77	15	-24
Mean					101	12					105	10	
STD					52	7					44	5	

*%Diff(FRP,WK): difference (%) between Functional Renal Parenchyma (FRP) and Whole-Kidney (WK) RBF
*%Diff(FRP,WK) = ((FRP-WK)/FRP)×100
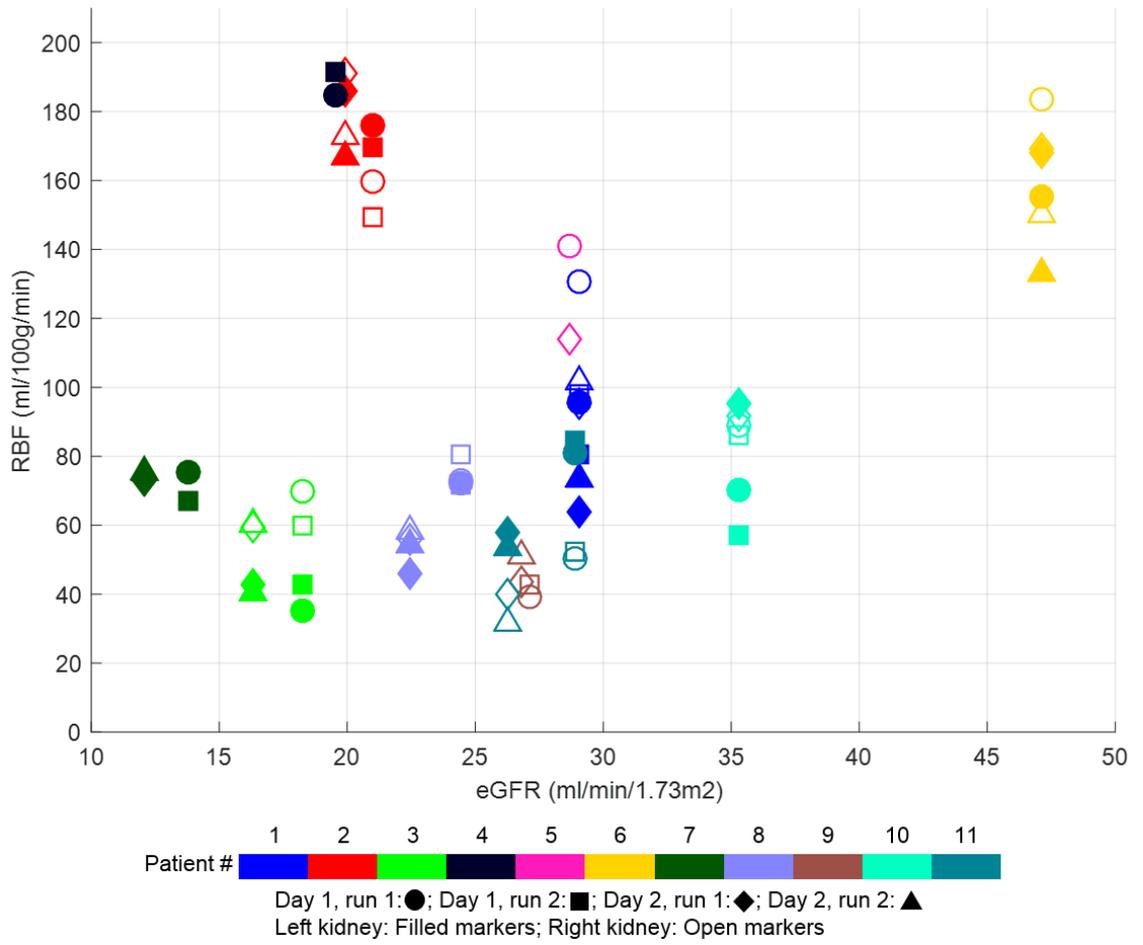
Figure 7-13 Mean RBF (FRP ROIs) vs. eGFR - separate kidneys.
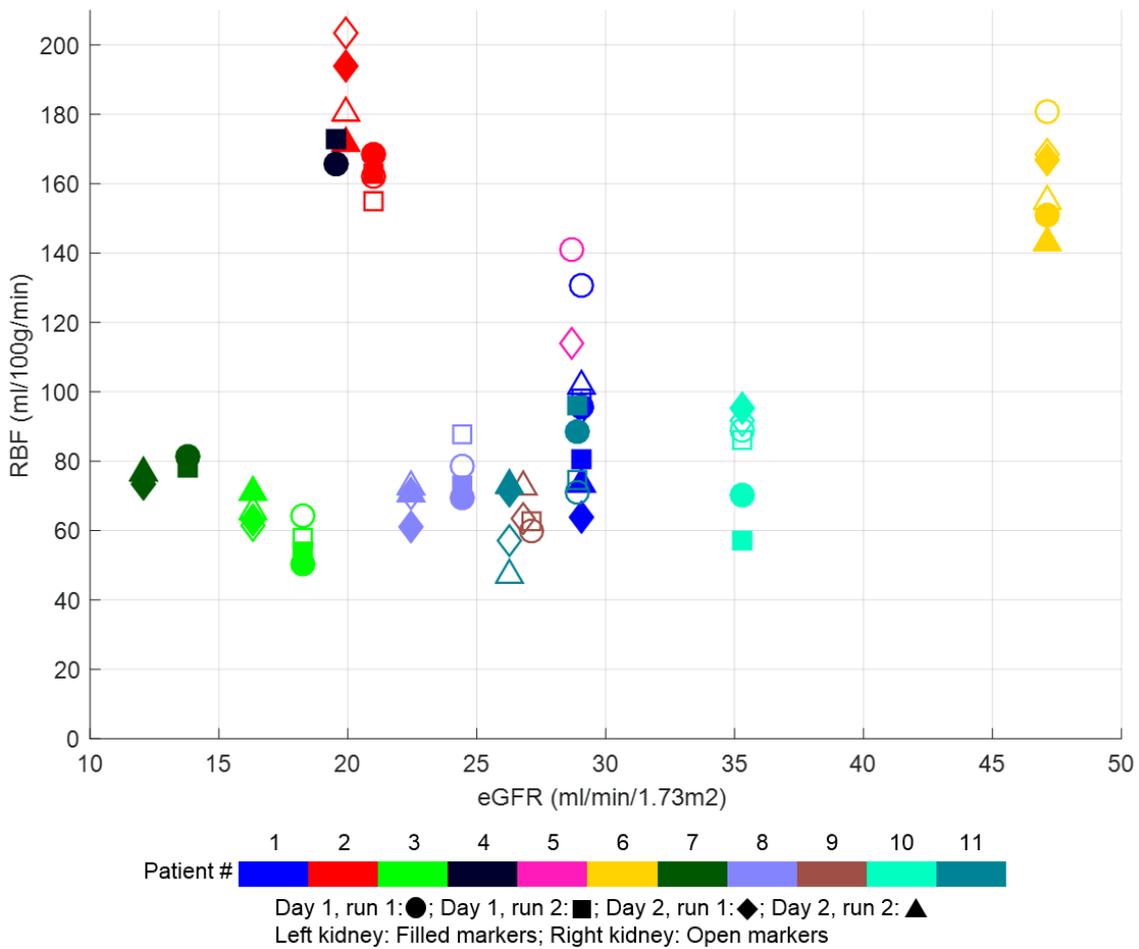
Figure 7-14 Mean RBF (WK ROIs) vs. eGFR - separate kidneys.
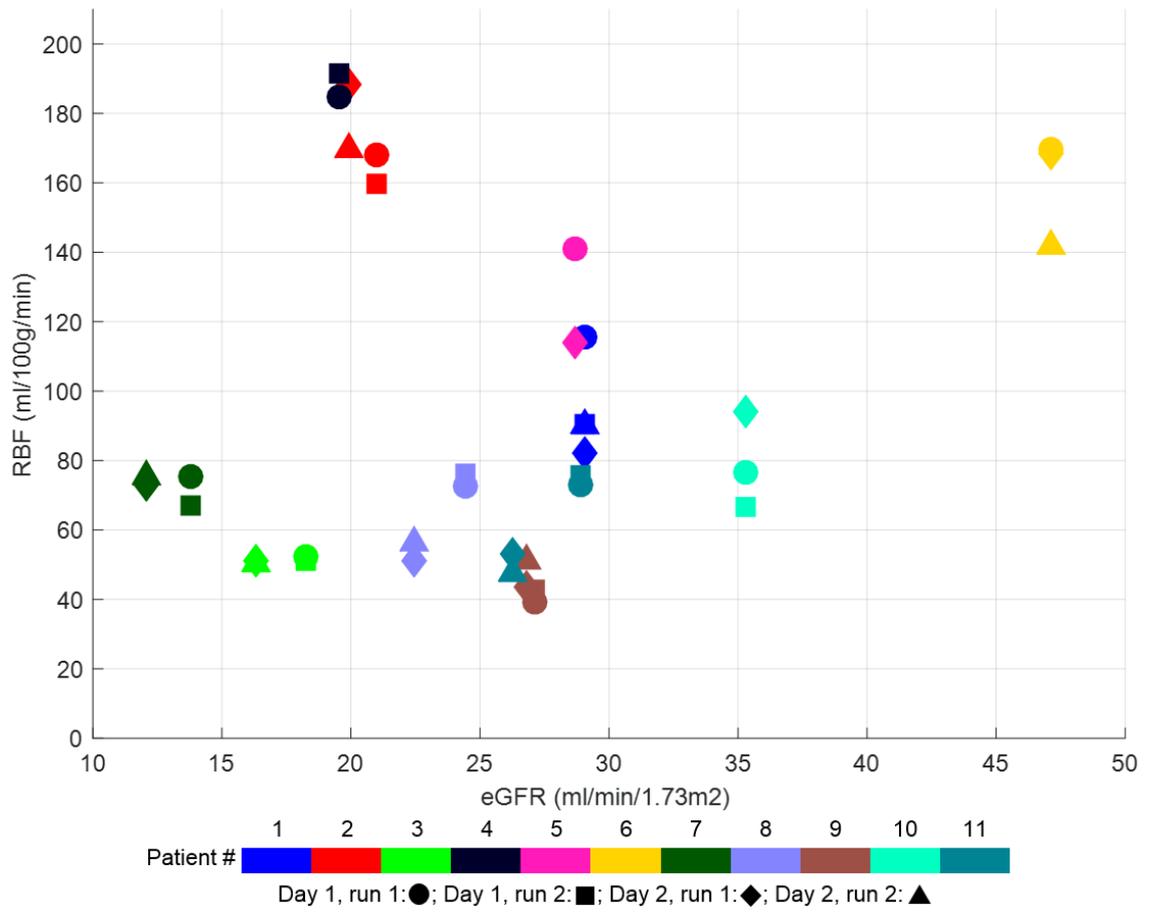
Figure 7-15 Mean RBF (FRP ROIs) vs. eGFR – merged kidneys.
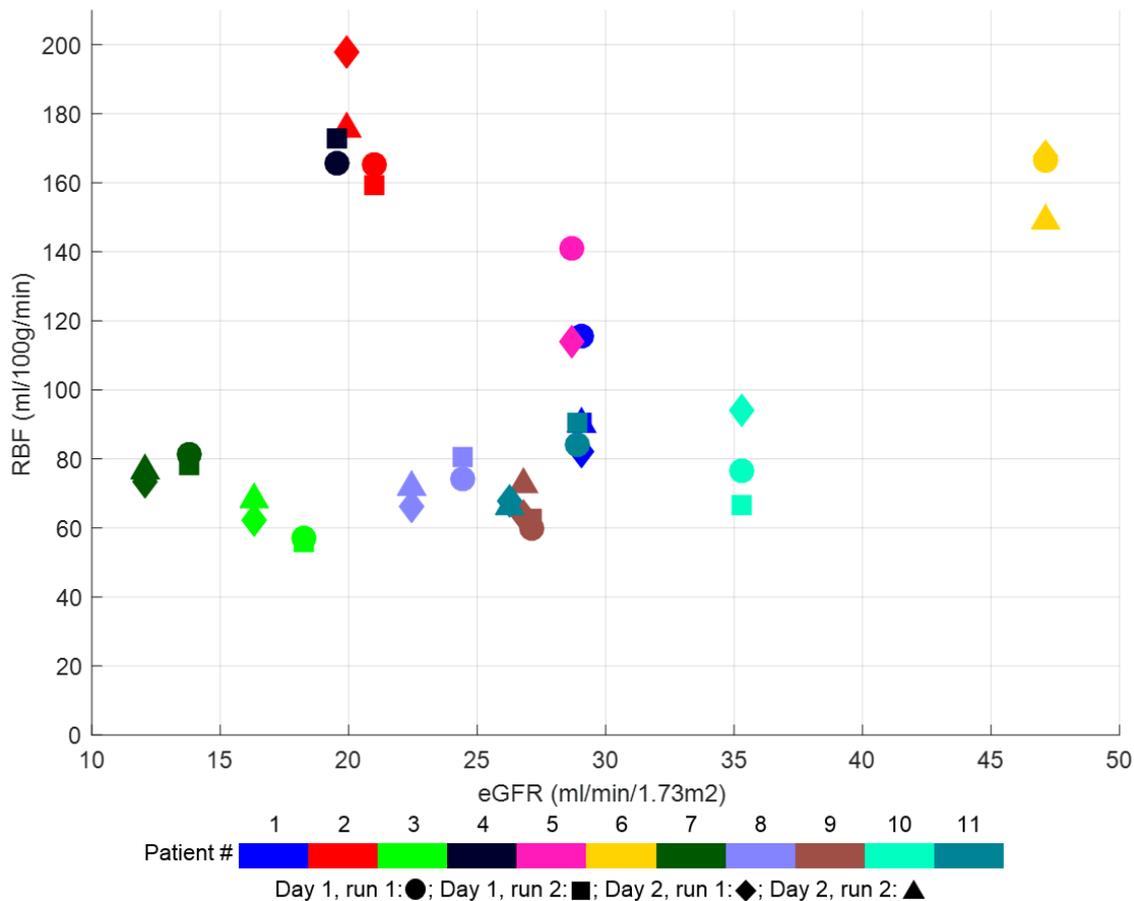
Figure 7-16 Mean RBF (WK ROIs) vs. eGFR – merged kidneys.

The following sub-sections reflect a tentative grouping of the patients based on their underlying pathology.

7.3.7.1   Patients #3, #6, #8

Patients #3 (Nephronophtisis), #6 (ATN) and #8 (haemolytic uraemic syndrome) presented with intrinsic kidney disease without dilatations. Patients #3 and #8 have relatively similar eGFRs, whereas the eGFR of patient #6 is around twice as that of patients #3 and #8. Given the underlying pathophysiology in patients #3 and #8, a progressive decrease in renal function is expected. This was the case, as in the months following the scans, there was a significant progression towards end-stage renal disease leading both patients to require dialysis. Patient #6 presented with ATN. The natural history of patients with ATN is variable, depending on the degree of ischemia suffered by the kidney. The ASL data reveals better perfusion in patient #6 as compared to #3 and #8, consistent with the significantly higher eGFR in patient #6. Unlike patients #3 and #8, patient #6 remained in stable CKD, not needing dialysis. This suggests that ASL may be providing clinically relevant information and might provide insight into different pathological processes.

194

7.3.7.2 Patients #1, #5, #7 and #10

Patients #1, #5, #7 and #10 presented with posterior urethral valves resulting in bilateral dilatations. The PWIs closely match the anatomy: no PW signal is seen in the anatomical locations of the dilatations, and tissue is clearly perfused in the remaining functioning parenchyma (see Figure 7-17). An exception is patient #7 (lowest eGFR in the cohort) whose dilatations occur at the ureter level and the renal parenchyma shows a globally reduced perfusion, which fits what one might expect from a lower eGFR measure.
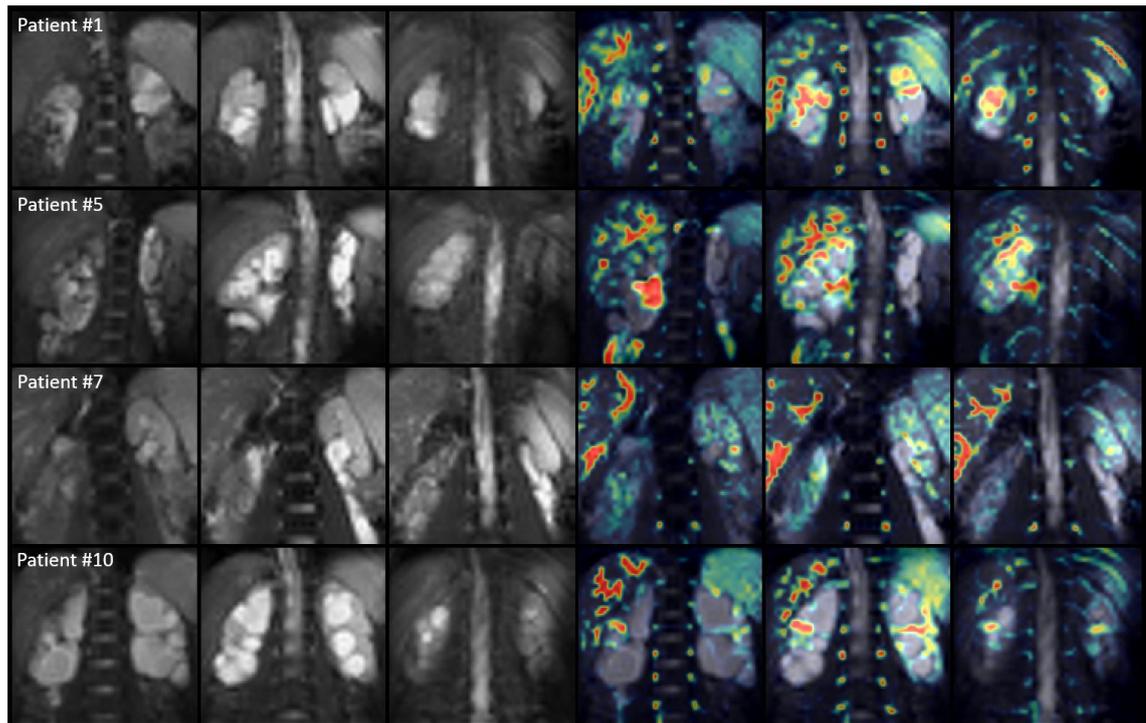
Figure 7-17 PD and PD-PWI fusion images in patients with dilatations (patients #1, #5, #7, #10). Each show shows 3 slices of each patient. There is a very good match of the ASL perfusion signal and the anatomy as seen in the non-background suppressed PD (left) and PD-PWI (right) images. Note the lack of perfusion signal in regions with dilatations, as one would expect. All scans correspond to the first run on the first day. These data were interpolated (bicubic interpolation, factor of 3), as is common in medical imaging visualization software.

7.3.7.3 Patients #2, #4 and #11

Three patients presented with renal dysplasia as primary diagnosis (patients #2, #4 and #11). Patients #2 and #4 show a perfusion image with good signal. Whilst having a low GFR (between 19.6 and 21.0 ml/min/1.73m^2), these two patients have the highest RBF of the entire cohort, respectively 171 ± 12 and 188 ± 5 (mean ± standard deviation across the scans which met the data quality criteria). Despite a similar primary diagnosis, the perfusion obtained in patient #11 is significantly lower than in patients #2 and #4. In addition, the whole-kidney perfusion (per 100g) is higher than the cortical perfusion in the case of patient #11, which would not be

expected if the majority of blood had perfused the cortex at the inflow time used. The RBF results of Patient #11 are surprising (low RBF which does not match with the relatively high eGFR). Therefore, in this particular patient, ASL data at the inflow time of 1.6s was visually inspected and processed analogously to the 1.2s TI ASL scans for a preliminary analysis of the RBF map at this longer inflow time. A visual evaluation of the perfusion-weighted signal at both TIs, overlaid on the reference PD images, suggests that at the chosen TI the complete bolus of labelled blood might not had left the large vessels and reached the renal cortex (see Figure 7-18). Analysis of the quantitative perfusion maps supports this finding as setting the inflow time to 1.6s resulted in a higher cortical RBF than at the TI of 1.2s (see Figure 7-18). This raises an interesting and important observation in relation to inflow times used in ASL. This will be discussed below in the conclusion of this chapter.
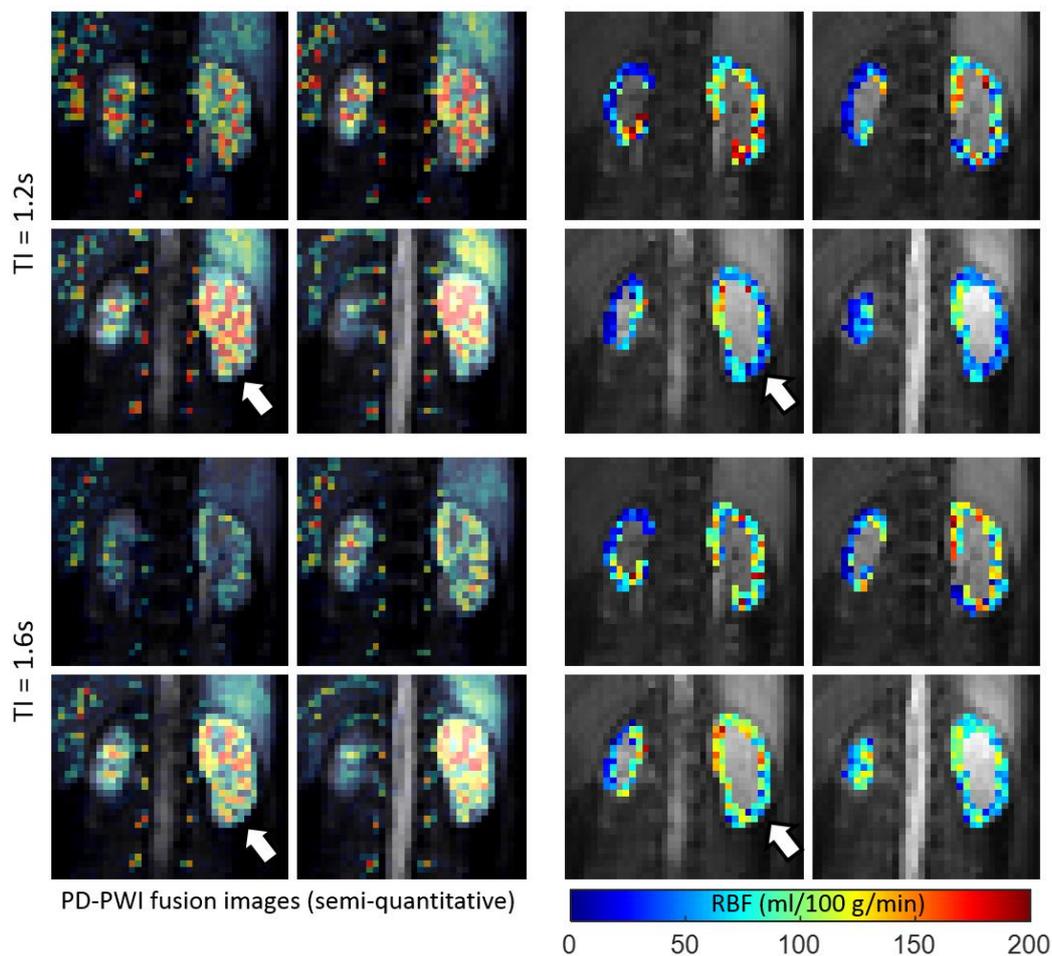
Figure 7-18 Delayed inflow time (patient #11 – 4/10 slices). The PD-PWI fusion images (right) confirm very good alignment between the PD and the PWI dataset. At TI=1.2s, perfusion has yet to reach the cortex, particularly in the left kidney (larger kidney on the right side of the figure). After quantification, this results in a reduced mean RBF at TI=1.2s compared to the longer TI of 1.6s. In the highlighted slice (arrows), the mean perfusion is 95 ± 32 ml/100g/min at TI = 1.6s compared to 71 ± 41 ml/100 g/min at the TI = 1.2s.

7.3.7.4 Patient #9

Patient #9 presented with a markedly enlarged right kidney with multiple cysts visible in the PD images, and extremely reduced perfusion signal in those corresponding regions. The pathophysiology of this patient is unique in this cohort, so forms a group of its own. The data in this patient is of very high quality (high SNR and almost completely free of artefacts – see Figure 7-19). The functional RBF (in ml/100g/min) is low (51±1 and 61±6 ml/100g/min, respectively for the FRP and the WK ROIs, mean ± standard deviation across runs). An inflow time of 1.2s, and to a larger extent the inclusion of cysts (with no perfusion) in the FRP ROIs likely contributes to this. This finding is supported by the fact that despite the large volume of the WK ROI, the difference in mean RBF between this ROI and the FRP ROI is the largest in this patient (with the mean WK RBF being approximately 47% higher than the mean FRP RBF), suggesting the inclusion of non-parenchymal tissue in the FRP ROIs.
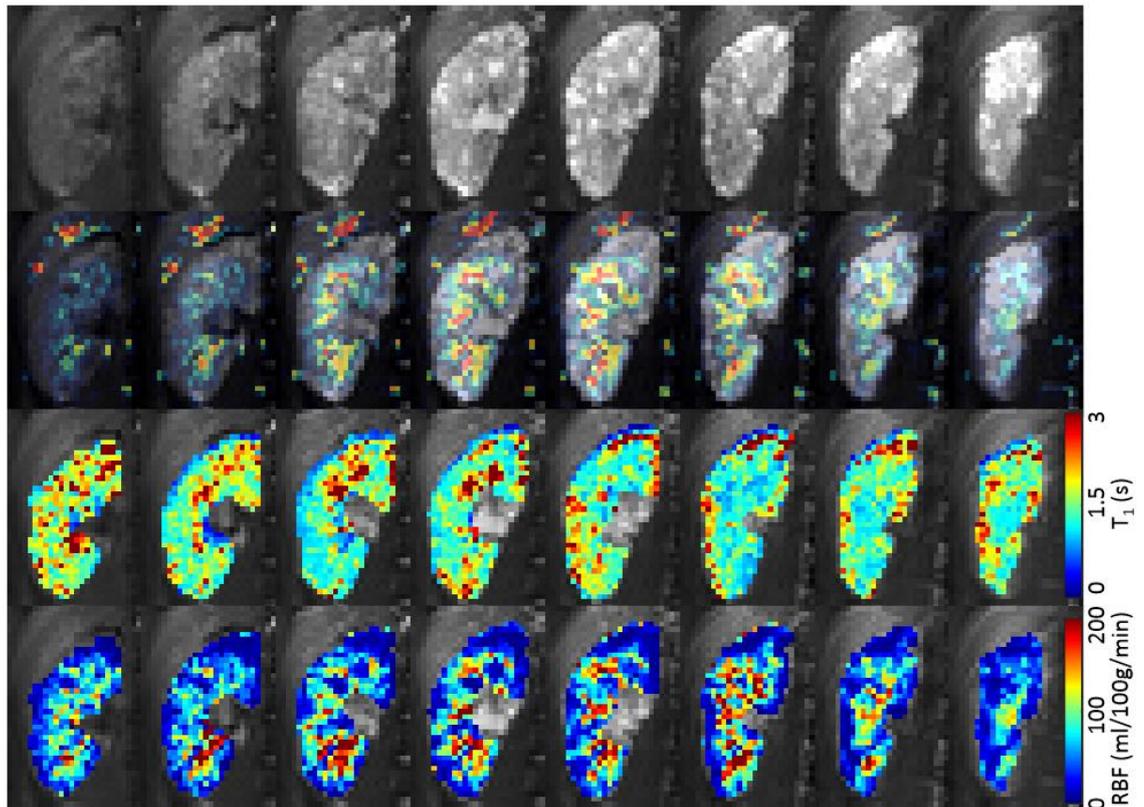
Figure 7-19 PD, T_1 and RBF data from patient #9: PD, PD-PWI fusion images, T$_1$ maps overlaid on PD images and RBF maps overlaid on PD images (from the top to the bottom rows).

7.3.8 Limitations of ASL measurements in children with CKD

This feasibility study has demonstrated several imitations, which are highlighted in the following sections. Despite these limitations, the MR data has shown the potential of renal ASL in

providing clinically relevant information in a challenging patient group in the first feasibility and reproducibility study of this technique on a paediatric cohort with severe CKD.

7.3.8.1 Heterogeneous patient cohort

A small cohort of children enrolled in this study, including patients with different ages and body surface areas, variable degrees of CKD and a wide range of diagnoses, where one might expect that the degree to which perfusion is impaired may be linked to each patients' underlying pathophysiology.

7.3.8.2 Lack of RBF measurements normalised to kidney volume

The kidney volumes vary significantly between subjects with some having only one functional kidney. As such, perfusion measurements normalised per total kidney volume are desirable, but currently not available due to the FAIR labelling imaging field-of-view limitation (see Section 7.3.5), which affects the proton-density ROIs from which whole-kidney ROIs were derived. Normalizing renal perfusion measurements per kidney or per volume/mass of cortical tissue may provide a more complete picture of the haemodynamics of the whole kidney. For example, in (Cutajar et al. 2015), RBF was measured using ASL in kidney donors pre and post-nephrectomy. They found no significant differences in RBF in ml/100 g/min but given an increase of kidney volume post-nephrectomy, RBF in ml/kidney/min did increase significantly, suggesting hyperperfusion at the level of each individual glomeruli. A similar finding was reported by [Eikefjord2016b] who compared single-kidney renal blood flow in healthy controls and living kidney donors, measured with DCE-MRI. No significant difference was found in RBF in ml/100g/min between the two groups. However, when the entire kidney volume was taken into account (significantly higher in the donor group), RBF in ml/kidney/min was found to be higher in the kidney donors.

7.3.8.3 Appropriate inflow times

A visual inspection of the pattern of perfusion-weighted signal enhancement, especially when overlaid on the PD images (PD-PWI fusion images) allows one to assess whether a given scan might have been performed with an inadequate inflow time. A preliminary analysis suggests that in some patients (specifically #3, #7, #8 and #11) an inflow time of 1.2s may not be long enough to ensure arrival of the complete bolus of labelled blood at the renal cortex. This is a well-known limitation of single-TI ASL studies, and is likely to be more prevalent in patients than healthy volunteers (for which an inflow time of 1.2s as used in this study is commonly used in the literature).

To attempt to confirm this hypothesis, exploratory ASL acquisitions were also performed data at an inflow time of 1.6s (0.4s longer than the standard inflow time of 1.2s).

As discussed above, an inflow time of 1.2s seemed insufficiently long in patient #11 for the labelled blood to reach the renal cortex. ASL data from this patient acquired with a 1.6s TI was processed and analysed analogously to the 1.2s TI. Figure 7-17 provides evidence of this fact, where mean cortical RBF calculated at TI=1.6s increased up to 34% in a central slice when compared to the mean RBF obtained at TI=1.2s

 A similar situation occurred in patients #3, #7 and #8, and to a lesser extent in patient #9. However, as seen above, the low mean cortical RBF in patient #9 is also related to the inclusion of cysts close to the boundary of the kidney in the cortical ROIs.

Multi-TI scans can address this issue, but in this preliminary work which primarily addressed the concern of motion in paediatric patients we opted for a conservative approach of using a single-TI. This allows the scan to be completed in a shorter period of time, and also the acquisition of a high number of measurements providing additional motion robustness. Nevertheless, the results presented here provide motivation to move to a more comprehensive multi-TI approach for renal ASL in CKD patients. Under good circumstances (e.g. appropriate SNR and motion-free data), a multi-TI approach can simultaneously reduce errors in the absolute RBF measures and provide additional information on the tissue haemodynamics. In single-TI ASL the main a priori assumption to ensure accurate quantification is that the entire labelled bolus has arrived at the tissue at the chosen TI. Variations (e.g. prolongation) of the arterial transit time, as can be expected in certain pathologies, may render this assumption invalid. Multi-TI techniques, by sampling and fitting the inflow curve across multiple time points, allow for variations in arterial transit time to be accounted for in the RBF calculation. Including measurements at short inflow times also allows for a better characterisation of the bolus arrival time which may in itself be a biomarker of altered perfusion dynamics. Furthermore, by knowing the amount of time spent by the labelled bolus in the intravascular and extravascular spaces, it is possible to account for the different relaxation properties of the label according to where it resides.

7.3.8.4 RBF modelling

A comparatively minor limitation is that similarly to other studies (e.g. (Song et al. 2010; Nathan S. Artz et al. 2011; Niles et al. 2016), the model employed for RBF quantification did not account for possible differences in the T_1 of the blood and tissue). We suspect that in some cases this could have contributed to an underestimation of the perfusion values, for subjects and renal

anatomical regions where the T_1 values obtained from the SR dataset might have exceeded the values typically found for the T_1 of blood.

7.3.8.5 Timing of blood sampling

Ideally blood samples for serum creatinine estimation and thus eGFR calculation in each patient would have been acquired on the day of each MR session, for a more accurate reflection of the patient's biochemical state of the day of the scan. Nevertheless there was a short period between MRI scans and no changes in clinical status occurred between scans.

7.4 Conclusions

This study evaluated the feasibility of renal ASL in a paediatric cohort and was able to show reproducible RBF measurements in children even with severe CKD. RBF was found to be significantly lower in children with CKD than in healthy adult volunteers. The hypothesis that there is a direct correlation between eGFR and RBF in children with CKD has been disproven. The lack of correlation between eGFR and RBF might be a reflection of the different pathophysiological processes in different pathologies. The combination of low eGFR and low RBF may be a biomarker leading to end-stage renal disease: Patients #3, #7 and #8, who had some of the lowest eGFRs in the cohort, had low RBF (<80ml/100g/min), and their clinical status worsened in the year following the scans - these may be index cases. Information on the variation and importance of arterial arrival times needs to the investigated in a larger cohort of children with similar causes of CKD. Arterial transit times could represent a new renal disease biomarker, leading to a better understanding of the underlying pathophysiology and thus natural history of progression or lack of progression of CKD. A prospective study with a larger group of children with a similar diagnosis would be necessary to test this hypothesis.

From a practical perspective, we would recommend acquisition of a SR dataset immediately prior to the main ASL acquisition to minimize the likelihood of large kidney displacements between these, which may be difficult to correct for using image registration. Avoiding non-movement related artefacts, especially those that may introduce regions of high intensity in the vicinity of the kidneys in the background suppressed data (e.g. fat shift or B_1 inhomogeneity) is of paramount importance as these may severely affect the accuracy of image registration (this was seen in patient #4).

Nevertheless, to our knowledge this is first study showing the feasibility of renal ASL in a paediatric cohort with severe kidney disease. Further work will focus on establishing the minimum number of signal averages that allows good SNR and motion robustness properties to be obtained, in an effort to reallocate some of the time currently used to acquire multiple

repetitions of a single TI to acquisitions at different inflow times so that variable haemodynamics can be accounted for in the RBF modelling and quantification.

8 Conclusions and Future Outlook

8.1 Conclusions

Throughout recent years, ASL has been slowly making its way into the clinical realm as technical challenges are addressed allowing the potential of the technique to be realised. This thesis aimed to develop and apply ASL in the context of paediatric CKD. Several challenges were encountered, with patient movement being the most prevalent one. This was addressed at all levels of the ASL framework and applied in the first study showing feasibility of ASL in children with severe chronic kidney disease.

Chapter 4 has shown the importance of optimised multi-shot acquisition schemes to reduce the likelihood of patient movement corrupting the data at the *k*-space level. This is particularly important since such corruption events give rise to artefacts that are impossible to correct using standard post-processing approaches. Nevertheless, even optimised *k*-space filling strategies remain motion-sensitive when they rely on several excitations separated by time gaps long enough so that patient movement is not "frozen". This inherent limitation of multi-shot schemes is demonstrated in a variety of ways. First, a retrospective analysis of the respiratory trace of subjects who underwent multi-shot 3D GRASE ASL scans to which the instants of data acquisition (readout modules) were synchronized. This allowed correlation of the image artefacts to inconsistencies with respect to the phase of the respiratory cycle during which the data was acquired, explaining the source of the image artefacts. These were further confirmed by examining independent undersampled reconstructions of each *k*-space segment, as well as "navigator-like" information obtained from reconstructing projections from the phase correction acquisitions embedded in the 3D-GRASE pulse sequence. This suggests that the use of multi-shot acquisition approaches is likely to hinder the adoption of ASL in routine clinical practice, particularly in patients where compliance may be reduced, such as children. This motivated a preliminary assessment of data acquired with single-shot 3D-GRASE readout schemes. This has shown merit in these types of acquisitions, which were explored in the remaining chapters in this thesis.

Building on the conclusions from Chapter 4, Chapter 5 presents a feasibility study of renal perfusion quantification in healthy volunteers where for the first time single-shot 3D GRASE was employed as an imaging readout module in the kidneys (Nery et al. 2016b). This allowed a robust

and fast (on the order of 2 minutes, depending on the respiratory rate of the subjects) volumetric mapping of cortical perfusion. RBF estimates were obtained in accordance with the existing literature. Furthermore, automatic retrospective approaches to outright reject or reduce the contribution of corrupted ASL PWIs to the final RBF maps were developed. These were tested in the same cohort of healthy volunteers, who had undergone an additional ASL scan which included periods of intentional movement by the subjects. The developed approaches were able to significantly reduce image artefacts and blurring and significantly improved the temporal signal-to-noise ratio of the perfusion weighted time series (up to a 73% improvement in the most corrupted dataset).

After adopting a more robust readout (with respect to movement) and implementing approaches to reduce the effect of corrupted ASL measurements on the final RBF maps, Chapter 6 focused on the development and implementation of a complete ASL quantification pipeline employing image registration for motion correction. The pipeline includes steps for registration of both the saturation-recovery (SR) data (from which renal T_1 estimates are obtained for quantification) and the ASL datasets. Both were independently validated in a cohort of 11 children with severe chronic kidney disease. We show an improvement in the quality of the cortical T_1 estimates, as assessed by intra and inter-scan repeatability and analysis of the residuals of the SR fits. The registration of the individual PWIs (before averaging) resulted in a statistically significant decrease of the temporal standard deviation and increase of the temporal SNR of the PWI time series (i.e. improvement in both metrics). Incorporating one of the noise correction strategies introduced in Chapter 5 (*wMeanB-mask*) into the quantification pipeline resulted in a further statistically significant improvement of both metrics compared to image registration alone. Both the SR and ASL datasets were also registered to the reference PD scan, significantly reducing errors in the RBF quantification caused by misalignment of these different types of data. Considering the complete pipeline with motion correction (i.e. considering the combined effect of SR registration, ASL PWI registration, SR to PD and ASL to PD registration), the repeatability of the RBF estimates obtained in scans performed on different days improved significantly compared to the "no motion correction" case.

The feasibility of ASL in children with severe chronic kidney disease was assessed in Chapter 7, using the optimised methods developed in Chapters 4, 5 and 6. This cohort had arguably the most severe degree of CKD explored so far in the literature with several patients close to end-stage renal failure (eGFR (ml/min per $1.73m^2$) = 26 ± 9 (mean ± standard deviation)). As expected, this patient cohort has shown significantly reduced renal perfusion compared to healthy volunteers. Nevertheless, despite having a limited number of patients, this study

suggests that children with CKD are not a homogenous pathophysiological group with respect to the relationship between RBF and estimated filtration (eGFR). Furthermore, our data suggests that the arterial transit times (ATT) vary in this cohort of patients, which warrants future studies on the optimisation of multi-TI ASL schemes to account for these haemodynamic changes. These will also provide quantitative ATT information, whose variation among patients may provide insight on the underlying pathophysiology. Despite several limitations and highlighting several of the challenges of performing ASL in a clinical scenario in a difficult patient cohort (combination of severe CKD arising from different causes), this study was able to address the fundamental concern of patient movement. This allowed us to obtain reproducible functional renal parenchyma and whole-kidney RBF measurements in the first paediatric cohort to undergo renal ASL, demonstrating feasibility of this technique.

8.2 Future Outlook

With respect to the use of volumetric readouts for ASL, an important avenue for future work is to improve single-shot acquisition schemes. As shown in this work, we believe they are necessary for robust acquisition methods to enable a wide adoption of ASL in the clinic, particularly in applications where movement is of particular concern. Even in the brain, where currently multi-shot acquisition schemes are recommended (Alsop et al. 2015), future improvements in image acquisition should eventually result in a shift to single-shot acquisition methods, as highlighted by Wong *et. al.* (Wong 2014). The greatest disadvantage of volumetric single-shot schemes is their achievable spatial resolution. The spatial resolution of the ASL protocol used to scan children with CKD is not ideal from a theoretical standpoint. Defining an optimal spatial resolution for scanning children with CKD is challenging given that the dimensions of renal cortex vary widely depending on the overall size of the kidneys (age of the subjects) and especially on the underlying pathophysiology (e.g. present or absence of dilatations in the collecting system). In this particular study, a compromise in spatial resolution was necessary to ensure the possibility of whole-kidney single-shot imaging at 1.5T.

Partial volume correction methods may also be beneficial but these rely on accurate segmentations from high resolution anatomical scans and require certain assumptions to be fulfilled (e.g. spatial consistency). Furthermore, the lack of widespread automated image segmentation methods applicable for renal imaging (particularly in disease) makes obtaining such segmentations a laborious process which typically relies on lengthy manual segmentation efforts by experts. As such, these were not explored in this work.

A coarse spatial resolution may cause renal parenchyma voxels to be composed of signal originating simultaneously from both cortical and medullary tissue, particularly in the voxels close to the interface between these two tissues. This is the well-known partial volume effect, which affects all MRI techniques to an extent depending on their achievable spatial resolution. In renal ASL data, this may translate into the difference between the RBF measured in two deemed cortical voxels being caused by an unequal proportion of cortical and medullary tissue in the voxels, rather than a true difference in cortical RBF. This effect is exacerbated by the marked differences between cortical and medullary perfusion, as the former can be up to five times higher than the latter. In certain types of CKD (e.g. vascular diseases), cortical thinning is often observed (O'Neill 2000), which may further amplify errors in cortical perfusion measurements due to partial volume effects. As such, improving the spatial resolution of this technique will be of importance to improving the sensitivity and specificity of renal ASL in detecting perfusion deficits, particularly in focal disease. The approach presented in this thesis has significant room for improvement as no acceleration/undersampling methods (apart from partial Fourier) were used. These can be used to improve the spatial resolution of the scans while leaving the ETD unchanged or reduce the ETD to minimize blurring and thereby improving the effective resolution of the scans (Griswold et al. 1999). Several aspects are relevant when considering the use of parallel imaging in children with CKD at 1.5T. The main factor is related to the SNR of the ASL measurements. First, as previously discussed, ASL is a technique with an intrinsically low SNR, which is further reduced in a subjects where the perfusion is compromised (which is the case in CKD). If parallel imaging is used with the aim of increasing the resolution of the scans, as would be desirable in this particular case (maintaining the same echo-train duration) then the loss of SNR is not only caused by undersampling the data but also from the corresponding decrease in voxel volume. The fact that scans were performed at 1.5T is also not advantageous as compared to at a field strength of 3T where most of the recent renal ASL work is performed. Furthermore, parallel imaging methods were not implemented on the available prototype 3D-GRASE ASL sequence. An in-house implementation of parallel imaging acquisition and reconstruction methods was not pursued as it was superseded by the priority of developing motion correction approaches. The use of parallel imaging will address the main disadvantage of 3D single-shot sequences, which is the long echo train duration, during which relaxation effects are significant and result in a loss of image quality. A reduction in the acquisition window (shot duration) may also have the benefit of reducing physiological noise as shown in (Do et al. 2014). If fact, (Do et al. 2014) et. al. found that the reduction in physiological noise was large enough to offset the increase in thermal noise caused by using parallel imaging resulting in a net

gain in terms of tSNR. However these findings were obtained in a cardiac ASL study using a TrueFISP readout and whether similar gains could be seen in renal 3D-GRASE ASL will require a future study.

Compressed sensing (Lustig et al. 2007) may also be used to improve the spatial resolution or coverage of ASL scans. The latter has been shown by Han et. al., who have used compressed sensing to improve the coverage of a bSSFP-based (i.e. TrueFISP) ASL acquisition scheme (Han et al. 2016). Alternative approaches to cope with blurring include modulating the flip-angles to preserve signal at later stages of the echo train (Liang et al. 2014) and post-processing techniques (Galazzo et al. 2014).

It is important nonetheless to keep in mind that under-sampling techniques come with an associated SNR loss which may be particularly costly for ASL, especially in patients with low RBF as seen in CKD. For this, imaging at a higher field (e.g. 3T) is attractive, given that the SNR of ASL measurements is proportional to the magnitude of the main magnetic field. In ASL, an additional advantage exists at higher fields since the magnetic label takes a longer time to decay, so a smaller amount of the possible signal is lost during the transit of the label. Nevertheless, a methodological approach is required for this investigation as the increased field strength may present additional challenges with respect to magnetic field inhomogeneity and power deposition. Pseudo continuous labelling (pCASL) may present a viable method to further improve SNR but even though this is now the labelling method of choice for brain ASL, further validation is necessary for the kidneys, particularly with respect to the variations in labelling efficiency caused by susceptibility effects and strong blood flows in the aorta (Shimizu et al. 2017). A particularly important and often overlooked fact is the geometric limitations of FAIR when applied in the context of renal imaging (as shown in Chapter 7), which can likely be solved with pCASL.

Both these possible improvements in SNR and spatial resolution are highly desirable in ASL as currently perfusion measurements in the renal medulla are still highly challenging (Nathan S. Artz et al. 2011; Kim et al. 2017), which may also provide important information on the pathophysiology of renal disease (Wentland et al. 2009).

Future research is also necessary on the interplay between background-suppression and image registration, specifically on determining the optimal levels of background suppression to simultaneously reduce the fluctuations of static tissue signal which affect the perfusion measurements while still allowing registration methods to successfully realign the ASL data.

As seen in the renal ASL review presented in Chapter 2, research into multi-TI ASL methods has emerged only recently, and is mostly still limited to adult healthy volunteer studies. Our data shows how using single-TI ASL is advantageous from a motion-robustness perspective, but also that arterial transit times vary in this group of patients. Therefore, it is necessary to work on the compromise between the data redundancy that gives single-TI ASL its motion robustness and the additional haemodynamic information provided by multi-TI scans. Noise and motion compensation methods such as those that were presented in Chapter 5 and Chapter 6 will play an important role in this.

Besides all the potential technical improvements described above, one can imagine that the number of ASL studies in patient populations will continue to increase. Ultimately, multi-centre studies with large patient cohorts will be necessary to ascertain the extent to which ASL (likely in combination with other non-invasive MRI biomarkers) will be adopted in clinical practice. Several scenarios can be envisaged, all of which have the potential to significantly impact the lives of patients with kidney disease: MR biomarkers may help in the stratification of patients with respect to their risk of progression to ESKD, elucidate aspects of renal pathophysiology and in a non-invasive identification of the critical mechanisms of kidney injury in each patient. Furthermore, they may provide valuable information to guide therapy on a patient-specific basis and potentially in the design of clinical trials to assess future therapies for renal disease.

Appendix A - Reduction of motion artefacts in multi-shot 3D GRASE ASL using Autofocus

This work was presented at an International MR conference (Thomas et al. 2016).

Introduction

A brief review of motion correction methods was presented in Chapter 3. Some of these were developed to perform k-space based corrections to correct for artefacts which can not be addressed by traditional image registration strategies such as those presented in Section 3.5. Many of these artefacts are avoided by adopting single-shot protocols, as presented in Chapter 5. Nevertheless, unless advanced methods for accelerating the image acquisition and reconstructing undersampled data are used, the maximum achievable effective resolution of the scans is affected. Despite not being the focus of this thesis, and not constituting a chapter in its own right, I implemented and conducted tests with a k-space-based method (autofocus) for retrospective motion correction of multi-shot data. The results formed the basis of an abstract

(Thomas et al. 2016) and are presented here as a demonstration of merit for future work in this domain.

Methods

It is known that translations in image domain correspond to phase ramps in the frequency domain (Fourier shift theorem), as discussed in Section 3.1. Therefore, the algorithm applies trial phase ramps to the acquired raw data to compensate for the rigid body translations that have occurred between the acquisitions of the different k-space segments. After each shift and following reconstruction, the image quality of the resulting image is quantified using a quality metric, in this case the image entropy (see Section 3.5.2.2). Minimizing the entropy of an image reduces blurring and removes ghosts from otherwise reduced signal intensity regions. This is the basic principle of the autofocus method (Atkinson et al. 1997; Atkinson et al. 1999). To avoid a lengthy brute-force optimization, a basic iterative algorithm was implemented whereby the algorithm starts applying phase shifts in a coarse grid, which is iteratively refined (see Figure A-1). This allows the algorithm to start with a large range of possible phase ramps while simultaneously correcting motion at the subpixel level at later iterations by significantly reducing the step size of the phase increments.

Data from a healthy volunteer were acquired on a Siemens 3T Tim Trio scanner (Siemens Healthcare, Erlangen). Background suppressed FAIR Q2TIPS 3D GRASE ASL volumes were acquired with the following parameters: TE=26.84ms; TI1/TI2=800/2000ms; spatial resolution=3.1x3.1x5mm; acquisition matrix 128x104x12, with 5/8 partial Fourier applied along the second PE direction and segmentation applied in the primary PE direction (3 shots), resulting in an echo train duration (ETD) of 215ms; 140° refocusing pulse flip angle; TR=4s. Data were acquired in batches of 4 repeats: during the first batch, the subject was asked to remain as still as possible; during the second batch, the subject was asked to make translational head movements during the first and third repeat. Raw k-space data were exported from the scanner and processed with MATLAB code developed by myself (The MathWorks, Inc., Natick, MA, United States).

```
---
Settings:
Motion amplitude (mm) = 10.0
Step size  (mm) = 5.0
No. refining grids = 2
---

>> multiGridSearch1D (SHOT 1)
BEST_SHIFT_IDX = 3
#1:
   shiftsPe (mm) = -10.00, -5.00, 0.00  5.00, 10.00
#2:
   shiftsPe (mm) = -5.00, -2.50, 0.00  2.50, 5.00
#3:
   shiftsPe (mm) = -2.50, -1.25, -0.00, 1.25, 2.50

>> multiGridSearch1D (SHOT 2)
BEST_SHIFT_IDX = 4
#1:
   shiftsPe (mm) = -10.00, -5.00, 0.00, 5.00  10.00
#2:
   shiftsPe (mm) = 0.00, 2.50, 5.00, 7.50  10.00
#3:
   shiftsPe (mm) = 5.00, 6.25, 7.50, 8.75  10.00
```
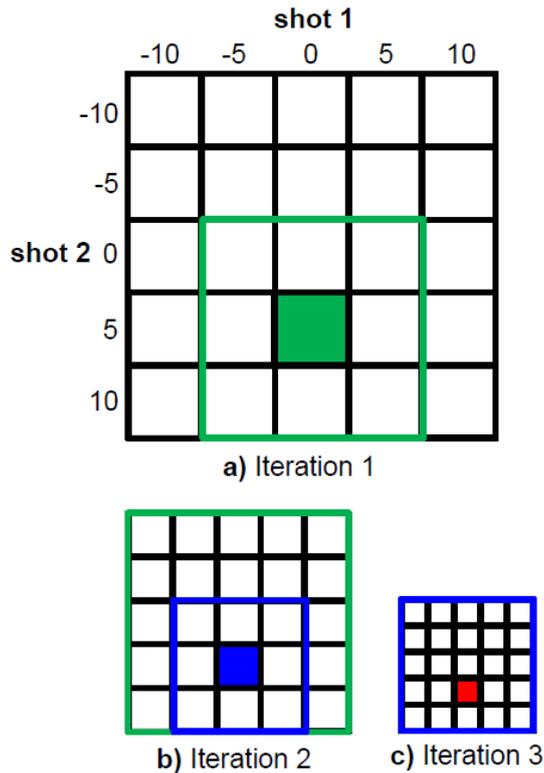
Figure A-1 "Multi-grid" autofocus in a 3-segment acquisition. For simplicity, the figure only considers trial shifts one in direction. I also consider one of the three shots to be in the reference position. As such, trial shifts are only applied to the other two shots.

Results

The autofocus algorithm was able to reduce inconsistencies between the different k-space segments, by applying phase corrections which allowed convergence to an optimal minimum-entropy solution for the 3-shot 3D GRASE data sets. The effect of applying the autofocus algorithm to one of the ASL control acquisitions from a scan with intentional subject movement is shown in Figure A-2. Two main improvements are clear:

i.   Reduction in blurring, visible as an increase in image sharpness when comparing Figure A-2 a) and Figure A-2 b).

ii.   Reduction in image ghosting, demonstrated by Figure A-2 c), which corresponds to a subtraction of the images before and after applying the autofocus method.
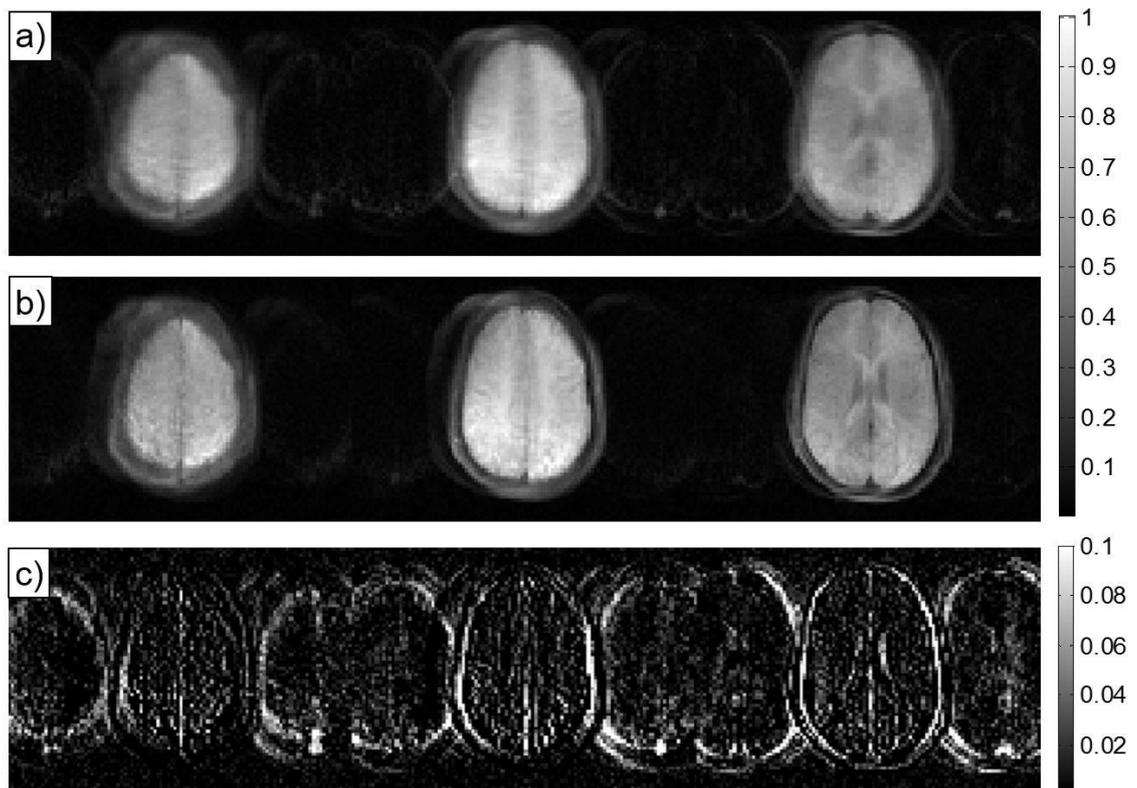
Figure A-2 Result of applying the autofocus algorithm in control multi-shot 3D-GRASE images (3 slices from the same subject). a) Standard reconstruction b) Autofocus reconstruction c) Difference between a) and b)

Conclusion

An autofocus-type method was implemented to correct for k-space inconsistencies resulting from movement occurring during the acquisition of segmented 3D-GRASE datasets. This initial implementation was unable to robustly correct for these types of artefacts in the kidneys. Reasons for this possibly include the non-rigid movement occurring in the abdomen (not necessarily the kidneys). As a proof of concept, we applied this method in brain ASL data (to satisfy the rigid motion condition) acquired from a volunteer which moved during the scan under controlled conditions. Our method was able to significantly reduce image blurring and ghosting artefacts, thus demonstrating the use of autofocus to improve image quality in multi-shot 3D GRASE ASL acquisitions. Autofocus presents advantages such as the fact that it does not require additional modifications to the pulse sequence and can be readily applied to our data. In addition, since in our acquisition multiple k-space lines are acquired within the same excitation, we have an inherent grouping of the k-space data, which significantly reduces the number of unknown motion parameters. These results are promising and encourage further work in optimising our autofocus method for renal acquisitions. Nevertheless, this was not pursued in

the remainder of the PhD due to a shift in focus to single-shot acquisitions, which are less susceptible to movement which directly corrupts individual k-space volumes.

Appendix B – T_1 mapping example in one healthy volunteer

As discussed in Chapters 5 and 6 (sections 5.2.2 and 6.2.3.1), ASL quantification in healthy volunteers was performed assuming a fixed cortical T_1 value obtained from the literature. Two main reasons led to this approach:

- Renal cortical T_1 values in healthy volunteers are relatively well characterised (compared to patient populations) and therefore a literature T_1 can be expected to be representative of the cortex of the healthy volunteer cohort examined in this thesis.

- A practical reason (also discussed in section 5.2.2) was that all motion correction pipelines were developed after the collection and processing of the healthy volunteer data. T_1 mapping data is particularly prone to image artefacts in the absence of motion correction methods due to the varying delay from the respiratory trigger (inversion) to image acquisition. Therefore, the use of corrupted T_1 maps was avoided by using a fixed literature T_1 value.

A T_1 map obtained in a healthy volunteer with the same acquisition (section 6.2.2) and analysis (section 6.2.4) methods (except for the absence of motion correction) employed in the paediatric cohort with CKD can be seen in the following figure.
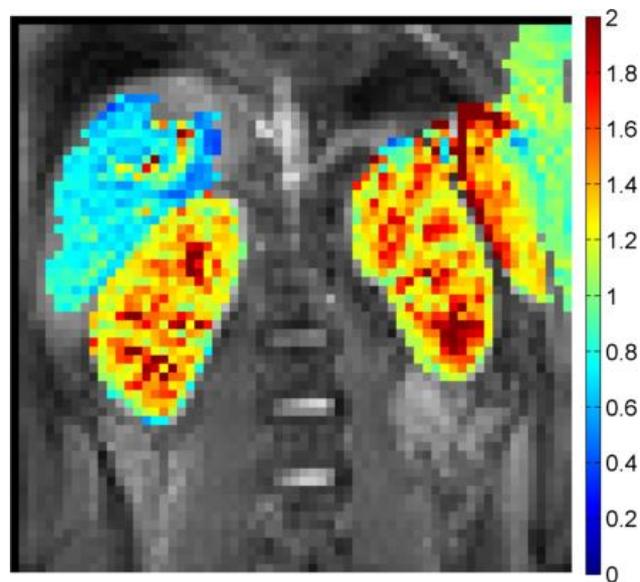
Figure B-1 T_1 map (RBG voxels) overlaid on one of the saturation recovery image volumes (grayscale voxels) from where the T_1 map was calculated.

Despite the lack of motion correction, the minimal amount of movement in the underlying saturation recovery time series allowed a good quality T_1 map to be obtained. Consistency across the cortical regions as well as corticomedullary differentiation can be seen. Furthermore,

the mask used when generating the figure above was obtained by simple thresholding and in addition to the kidneys, T_1 values from the liver can also be seen and are lower than in the kidneys, as expected (de Bazelaire et al. 2004).

Appendix C – *elastix* parameter file

```
// ======================
// ===== Image Types ===== ----------------------------
// ======================
(FixedInternalImagePixelType "float")
(MovingInternalImagePixelType "float")
(FixedImageDimension 3)
(MovingImageDimension 3)
(UseDirectionCosines "true")


// ===========================
// ===== Main components ===== -----------------------
// ===========================
(Registration "MultiResolutionRegistration")
(FixedImagePyramid "FixedSmoothingImagePyramid")
(MovingImagePyramid "MovingSmoothingImagePyramid")
(Interpolator "BSplineInterpolator")
(Metric "AdvancedMattesMutualInformation")
(Optimizer "AdaptiveStochasticGradientDescent")
(ResampleInterpolator "FinalBSplineInterpolator")
(Resampler "DefaultResampler")
(Transform "EulerTransform")


// ========================
// ===== Transformation ===== ------------------------
// ========================
(HowToCombineTransforms "Compose")
(AutomaticTransformInitialization "true")
(AutomaticScalesEstimation "true")


// ============================
// ===== Similarity Measure ===== ---------------------
// ============================
(NumberOfHistogramBins 32)
(ErodeMask "false")


// ==================
// ===== Pyramid ===== -------------------------------
// ==================
(NumberOfResolutions 1)


// ====================
// ===== Optimizer ===== -----------------------------
// ====================
(MaximumNumberOfIterations 500)
(AutomaticParameterEstimation "true")
(UseAdaptiveStepSizes "true")


// ========================
// ===== Image Sampling ===== ------------------------
// ========================
(NumberOfSpatialSamples 4096)
(NewSamplesEveryIteration "true")
(ImageSampler "RandomSparseMask")


// ===================================
// ===== Interpolation and Resampling ===== -----------
// ===================================
(BSplineInterpolationOrder 1)
(FinalBSplineInterpolationOrder 3)

(DefaultPixelValue 0)
(WriteResultImage "true")
(ResultImagePixelType "float")
(ResultImageFormat "nii")
```

9 References

Alsop, D.C., 2012. Arterial spin labeling: its time is now. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 25(2), pp.75–77. Available at: http://dx.doi.org/10.1007/s10334-012-0309-8.

Alsop, D.C. et al., 2015. Recommended implementation of arterial spin-labeled Perfusion MRI for clinical applications: A consensus of the ISMRM Perfusion Study group and the European consortium for ASL in dementia. *Magnetic Resonance in Medicine*, 73(1), pp.102–116.

Alsop, D.C. & Detre, J. a, 1998. Multisection cerebral blood flow MR imaging with continuous arterial spin labeling. *Radiology*, 208(2), pp.410–416.

Andre, J.B. et al., 2015. Toward quantifying the prevalence, severity, and cost associated with patient motion during clinical MR examinations. *Journal of the American College of Radiology*, 12(7), pp.689–695. Available at: http://dx.doi.org/10.1016/j.jacr.2015.03.007.

Archibald, G. et al., 2007. UK Consensus Conference on Early Chronic Kidney Disease--6 and 7 February 2007. *Nephrology, dialysis, transplantation*, 22(9), pp.2455–2457.

Artz, N.S., 2010. *Assessing Renal Perfusion in Native and Transplanted Kidneys Using Arterial Spin Labeling MRI*. PhD Thesis, University of Wisconsin-Madison.

Artz, N.S. et al., 2011. Comparing Kidney Perfusion Using Noncontrast Arterial Spin Labeling MRI and Microsphere Methods in an Interventional Swine Model. *Investigative radiologyInvestigative radiology*, 46(2), pp.124–131.

Artz, N.S. et al., 2011. Reproducibility of renal perfusion MR imaging in native and transplanted kidneys using non-contrast arterial spin labeling. *Journal of Magnetic Resonance Imaging*, 33(6), pp.1414–1421.

Ashburner, J., 2012. SPM: A history. *NeuroImage*, 62(2), pp.791–800.

Atkinson, D. et al., 1999. Automatic compensation of motion artifacts in MRI. *Magnetic Resonance in Medicine*, 41(1), pp.163–170.

Atkinson, D. et al., 1997. Automatic correction of motion artifacts in magnetic resonance images using an entropy focus criterion. *IEEE Transactions on Medical Imaging*, 16(6), pp.903–910. Available at: http://www.ncbi.nlm.nih.gov/pubmed/9533590%5Cnhttp://ieeexplore.ieee.org/lpdocs/e

pic03/wrapper.htm?arnumber=650886.

Atkinson, D. et al., 2004. Coil-based artifact reduction. *Magnetic Resonance in Medicine*, 52(4), pp.825–830. Available at: http://dx.doi.org/10.1002/mrm.20226.

Atkinson, D. & Hill, D.L.G., 2003. Reconstruction after rotational motion. *Magnetic Resonance in Medicine*, 49(1), pp.183–187.

Aukland, K., 1980. Methods for Measuring Renal Blood Flow: Total Flow and Regional Distribution. *Annual Review of Physiology*, 42(17), pp.543–555.

Babayeva, M. et al., 2015. Accuracy and Precision of Head Motion Information in Multi-Channel Free Induction Decay Navigators for Magnetic Resonance Imaging. *IEEE Transactions on Medical Imaging*, 34(9), pp.1879–1889.

Batchelor, P.G. et al., 2005. Matrix description of general motion correction applied to multishot images. *Magnetic Resonance in Medicine*, 54(5), pp.1273–1280.

de Bazelaire, C. et al., 2005. Arterial spin labeling blood flow magnetic resonance imaging for the characterization of metastatic renal cell carcinoma. *Academic Radiology*, 12(3), pp.347–357.

de Bazelaire, C. et al., 2008. Magnetic Resonance Imaging-Measured Blood Flow Change after Antiangiogenic Therapy with PTK787/ZK 222584~{C}orrelates with Clinical Outcome in Metastatic Renal Cell Carcinoma. *Clinical Cancer Research*, 14(17), pp.5548–5554.

de Bazelaire, C. et al., 2008. Magnetic Resonance Imaging - Measured Blood Flow Change after Antiangiogenic Therapy with PTK787/ZK 222584 Correlates with Clinical Outcome in Metastatic Renal Cell Carcinoma. *Clinical Cancer Research*, 14(17), pp.5548–5554.

de Bazelaire, C.M. et al., 2004. MR imaging relaxation times of abdominal and pelvic tissues measured in vivo at 3.0 T: Preliminary results. *Radiology*, 230(3), pp.652–659. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=14990831&retmode=ref&cmd=prlinks%5Cnpapers3://publication/doi/10.1148/radiol.2303021331.

Becker, A.S. & Rossi, C., 2017. Renal Arterial Spin Labeling Magnetic Resonance Imaging. *Nephron*, 135(1), pp.1–5. Available at: http://www.karger.com/?doi=10.1159/000450797.

Beeman, S.C. et al., 2011. Measuring glomerular number and size in perfused kidneys using MRI. *American Journal of Physiology*, 300(6), pp.F1454–F1457.

Beierwaltes, W.H. et al., 2013. Assessment of Renal Function; Clearance, the Renal

Microcirculation, Renal Blood Flow, and Metabolic Balance. *Comprehensive Physiology*, 3, pp.165–200. Available at: http://dx.doi.org/10.1002/cphy.c120008.

Bernstein, M.A., King, K.F. & Zhou, X.J., 2004. *Handbook of MRI Pulse Sequences*, Elsevier.

Berr, S.S. et al., 1999. Perfusion of the Kidney Using Extraslice Spin Tagging (EST) MRI. *J Magn Reson Imaging*, 18(184), pp.886–891.

Berr, S.S. & Mai, V.M., 1999. Extraslice spin tagging (EST) magnetic resonance imaging for the determination of perfusion. *Journal of Magnetic Resonance Imaging*, 9(1), pp.146–150.

Bertram, J.F. et al., 2011. Human nephron number: implications for health and disease. *Pediatric Nephrology*, 26(9), pp.1529–1533. Available at: http://dx.doi.org/10.1007/s00467-011-1843-8.

Bland, J.M. & Altman, D.G., 1986. Statistical Methods for Assessing Agreement Between Two Methods of Clinical Measurement. *Lancet*, 327, pp.307–310. Available at: http://www.sciencedirect.com/science/article/pii/S0140673686908378.

Bloch, F., 1946. Nuclear Induction. *Phys. Rev.*, 70(7–8), pp.460–474. Available at: http://link.aps.org/doi/10.1103/PhysRev.70.460.

Bracewell, R.N., 2000. *The Fourier Transform And Its Applications (3rd. Ed)*, McGraw-Hill.

Brau, A.C.S. & Brittain, J.H., 2006. Generalized self-navigated motion detection technique: Preliminary investigation in abdominal imaging. *Magnetic Resonance in Medicine*, 55(2), pp.263–270.

Breidthardt, T. et al., 2015. The pathophysiology of the chronic cardiorenal syndrome : a magnetic resonance imaging study. *Eur Radiol*, 25(6), pp.1684–1691.

Brown, R.W. et al., 2014. *Magnetic Resonance Imaging: Physical Principles and Sequence Design (2nd Ed.)*, Wiley-Blackwell.

Buxton, R.B. et al., 1998. A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magnetic Resonance in Medicine*, 40(3), pp.383–396.

Bydder, M. et al., 2003. SMASH navigators. *Magnetic Resonance in Medicine*, 49(3), pp.493–500. Available at: http://dx.doi.org/10.1002/mrm.10388.

Bydder, M., Larkman, D.J. & Hajnal, J. V, 2002. Detection and elimination of motion artifacts by regeneration of k-space. *Magnetic Resonance in Medicine*, 47(4), pp.677–686. Available at:

http://dx.doi.org/10.1002/mrm.10093.

Bydder, M., Larkman, D.J. & Hajnal, J. V, 2002. Generalized SMASH imaging. *Magnetic Resonance in Medicine*, 47(1), pp.160–170. Available at: http://dx.doi.org/10.1002/mrm.10044.

Cahoon, G., 2011. Techniques in Pediatric MRI – Tips for Imaging Children. *Paediatric*, pp.6–16.

Di Cataldo, S. et al., 2011. Motion artifact correction in ASL images: An improved automated procedure. *Proceedings - 2011 IEEE International Conference on Bioinformatics and Biomedicine, BIBM 2011*, (i), pp.410–413.

Chantler, C. et al., 1969. Glomerular filtration rate measurement in man by the single injection methods using 51Cr-EDTA. *Clin Sci*, 37(1), pp.169–180.

Chavhan, G.B., Babyn, P.S. & Vasanawala, S.S., 2013. Abdominal MR imaging in children: motion compensation, sequence optimization, and protocol organization. *Radiographics*, 33(3), pp.703–719. Available at: http://radiographics.rsnajnls.org/cgi/doi/10.1148/rg.333125027%5Cnpapers2://publication/doi/10.1148/rg.333125027.

Chen, Q. et al., 1997. STAR-HASTE: Perfusion imaging without magnetic susceptibility artifact. *Magnetic Resonance in Medicine*, 38(3), pp.404–408.

Cheng, J.Y. et al., 2015. Free-breathing pediatric MRI with nonrigid motion correction and acceleration. *Journal of Magnetic Resonance Imaging*, 42(2), pp.407–420.

Cheng, J.Y. et al., 2012. Nonrigid motion correction in 3D using autofocusing withlocalized linear translations. *Magnetic Resonance in Medicine*, 68(6), pp.1785–1797.

Chowdhury, A.H. et al., 2012. A Randomized, Controlled, Double-Blind Crossover Study on the Effects of 2-L Infusions of 0.9% Saline and Plasma-Lyte® 148 on Renal Blood Flow Velocity and Renal Cortical Tissue Perfusion in Healthy Volunteers. *Annals of Surgery*, 256(1), pp.18–24. Available at: http://content.wkhealth.com/linkback/openurl?sid=WKPTLP:landingpage&an=00000658-201207000-00005.

Collignon, A. et al., 1995. 3D Multi-Modality Medical Image Registration Using Feature Space Clustering. *First International Conference on Computer Vision, Virtual Reality and Robotics in Medicine*, pp.195–204. Available at: http://link.springer.com/10.1007/978-3-540-49197-2_22.

Collignon, A. et al., 1995. Automated multi-modality image registration based on information theory. *14th International Conference on Information Processing in Medical Imaging*, pp.263–274.

Conlin, C.C. et al., 2017. Renal plasma flow (RPF) measured with multiple-inversion-time arterial spin labeling (ASL) and tracer kinetic analysis: Validation against a dynamic contrast-enhancement method. *Magnetic Resonance Imaging*, 37, pp.51–55. Available at: http://dx.doi.org/10.1016/j.mri.2016.11.010.

Cutajar, M. et al., 2014. Comparison of ASL and DCE MRI for the non-invasive measurement of renal blood flow: Quantification and reproducibility. *European Radiology*, 24(6), pp.1300–1308. Available at: http://dx.doi.org/10.1007/s00330-014-3130-0.

Cutajar, M. et al., 2015. Renal blood flow using arterial spin labelling MRI and calculated filtration fraction in healthy adult kidney donors Pre-nephrectomy and post-nephrectomy. *European Radiology*, 25(8), pp.2390–2396. Available at: http://dx.doi.org/10.1007/s00330-015-3594-6.

Cutajar, M. et al., 2012. Repeatability of renal arterial spin labelling MRI in healthy subjects. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 25(2), pp.145–153. Available at: http://dx.doi.org/10.1007/s10334-011-0300-9.

Dai, W. et al., 2008. Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. *Magnetic Resonance in Medicine*, 60(6), pp.1488–1497. Available at: http://dx.doi.org/10.1002/mrm.21790.

Dai, W. et al., 2012. Reduced resolution transit delay prescan for quantitative continuous arterial spin labeling perfusion imaging. *Magnetic Resonance in Medicine*, 67(5), pp.1252–1265.

Daly, C., 2007. Is early chronic kidney disease an important risk factor for cardiovascular disease? A Background Paper prepared for the UK Consensus Conference on Early Chronic Kidney Disease. *Journal of Diabetes and its Complications*, pp.26–30.

Delanaye, P. et al., 2012. Normal reference values for glomerular filtration rate: What do we really know? *Nephrology Dialysis Transplantation*, 27(7), pp.2664–2672.

Deoni, S.C.L., Rutt, B.K. & Peters, T.M., 2003. Rapid combined T1 and T2 mapping using gradient recalled acquisition in the steady state. *Magnetic Resonance in Medicine*, 49(3), pp.515–526.

Detre, J.A. et al., 2012. Applications of arterial spin labeled MRI in the brain. *Journal of Magnetic*

Resonance Imaging, 35(5), pp.1026–1037. Available at: http://dx.doi.org/10.1002/jmri.23581.

Detre, J.A. et al., 1992. Perfusion imaging. *Magnetic Resonance in Medicine*, 23(1), pp.37–45. Available at: https://www.scopus.com/inward/record.uri?eid=2-s2.0-0026508892&partnerID=40&md5=0e42ba909fcc2ee3c09f122cd5d2e478.

Detre, J. a et al., 1994. Tissue specific perfusion imaging using arterial spin labeling. *NMR in biomedicine*, 7(1–2), pp.75–82.

Dixon, W.T. et al., 1991. Multiple inversion recovery reduces static tissue signal in angiograms. *Magnetic Resonance in Medicine*, 18(2), pp.257–268.

Dixon, W.T. et al., 1986. Projection angiograms of blood labeled by adiabatic fast passage. *Magnetic Resonance in Medicine*, 3(3), pp.454–462.

Do, H.P., Jao, T.R. & Nayak, K.S., 2014. Myocardial arterial spin labeling perfusion imaging with improved sensitivity. *Journal of Cardiovascular Magnetic Resonance*, 16(1).

Dong, J. et al., 2013. Quantitative assessment of acute kidney injury by noninvasive arterial spin labeling perfusion MRI: a pilot study. *Science China - Life Sciences*, 56(8), pp.745–50. Available at: http://www.ncbi.nlm.nih.gov/pubmed/23740361.

Dyverfeldt, P. et al., 2014. Reduction of motion artifacts in carotid MRI using free-induction decay navigators. *Journal of Magnetic Resonance Imaging*, 40(1), pp.214–220.

Eckardt, K.U. et al., 2013. Evolving importance of kidney disease: From subspecialty to global health burden. *The Lancet*, 382(9887), pp.158–169.

Edelman, R.R. et al., 1994. Qualitative mapping of cerebral blood flow and functional localization with echo-planar MR imaging and signal targeting with alternating radio frequency. *Radiology*, 192(2), pp.513–520.

Edelman, R.R. & Chen, Q., 1998. EPISTAR MRI: multislice mapping of cerebral blood flow. *Magnetic Resonance in Medicine*, 40(6), pp.800–805.

Ehman, R.L. & Felmlee, J.P., 1989. Adaptive technique for high-definition MR imaging of moving structures. *Radiology*, 173(1), pp.255–263. Available at: http://eutils.ncbi.nlm.nih.gov/entrez/eutils/elink.fcgi?dbfrom=pubmed&id=2781017&retmode=ref&cmd=prlinks.

Evans, R.G. et al., 2004. Mechanisms underlying the differential control of blood flow in the renal

medulla and cortex. *Journal of Hypertension*, 22(8), pp.1439–1451. Available at: http://dx.doi.org/10.1097/01.hjh.0000133744.85490.9d.

Fassett, R.G. et al., 2011. Biomarkers in chronic kidney disease: a review. *Kidney International*, 80(8), pp.806–821. Available at: http://dx.doi.org/10.1038/ki.2011.198.

Feinberg, D.A. et al., 1986. Halving MR imaging time by conjugation: demonstration at 3.5 kG. *Radiology*, 161(2), pp.527–531.

Feinberg, D.A. & Oshio, K., 1991. GRASE (Gradient and Spin-Echo) MR and Imaging: A New Fast Clinical Imaging Technique. *Radiology*, 181, pp.597–602.

Fenchel, M. et al., 2006. Perfusion MR imaging with FAIR true FISP spin labeling in patients with and without renal artery stenosis: initial experience. *Radiology*, 238(3), pp.1013–1021.

Filippi, M. & Rocca, M. a, 2007. Magnetization transfer magnetic resonance imaging of the brain, spinal cord, and optic nerve. *Neurotherapeutics*, 4(3), pp.401–413.

Fu, Z.W. et al., 1995. Orbital navigator echoes for motion measurements in magnetic resonance imaging. *Magnetic Resonance in Medicine*, 34(5), pp.746–753.

Gach, H.M., Nguyen, T. & Wang, Y., 2006. CASL Perfusion MRI of the Kidney using Real-time Tracking, Free Breathing Navigator. In *Proceedings of the 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine, Seattle, p. 2990*.

Galazzo, I.B. et al., 2014. Reducing blurring artifacts in 3D-GRASE ASL by integrating new acquisition and analysis strategies. In *Proc. Intl. Soc. Mag. Reson. Med.* p. 2704.

Gallichan, D. & Marques, J.P., 2017. Optimizing the acceleration and resolution of three-dimensional fat image navigators for high-resolution motion correction at 7T. *Magnetic Resonance in Medicine*, 77, pp.547–558. Available at: http://doi.wiley.com/10.1002/mrm.26127.

Gallichan, D., Marques, J.P. & Gruetter, R., 2016. Retrospective correction of involuntary microscopic head movement using highly accelerated fat image navigators (3D FatNavs) at 7T. *Magnetic Resonance in Medicine*, 75(3), pp.1030–1039.

Gardener, A.G. & Francis, S.T., 2010. Multislice perfusion of the kidneys using parallel imaging: Image acquisition and analysis strategies. *Magnetic Resonance in Medicine*, 63(6), pp.1627–1636. Available at: http://dx.doi.org/10.1002/mrm.22387.

Gerich, J.E. et al., 2001. Renal gluconeogenesis: Its importance in human glucose homeostasis.

Diabetes Care, 24, pp.382–391.

Gillis, K.A. et al., 2014. Inter-study reproducibility of arterial spin labelling magnetic resonance imaging for measurement of renal perfusion in healthy volunteers at 3 Tesla. *BMC Nephrology*, 15(1), p.23. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3909760&tool=pmcentrez& rendertype=abstract.

Gillis, K.A. et al., 2016. Non-Contrast Renal Magnetic Resonance Imaging to Assess Perfusion and Corticomedullary Differentiation in Health and Chronic Kidney Disease. *Nephron*, 133(3), pp.183–192.

Gillis, K.A., 2016. *Oxidative stress and its Haemodynamic Consequences in Chronic Kidney Disease*. PhD Thesis, University of Glasgow.

Golay, X., Hendrikse, J. & Lim, T.C.C., 2004. Perfusion imaging using arterial spin labeling. *Topics in magnetic resonance imaging : TMRI*, 15(1), pp.10–27. Available at: papers3://publication/uuid/3B1D5756-24CD-4DD4-903D-0E7FAD62542D.

Green, M.A. & Hutchins, G.D., 2011. Positron Emission Tomography (PET) Assessment of Renal Perfusion. *Seminars in Nephrology*, 31(3), pp.291–299. Available at: http://dx.doi.org/10.1016/j.semnephrol.2011.05.008.

Griswold, M.A. et al., 2002. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6), pp.1202–1210. Available at: http://dx.doi.org/10.1002/mrm.10171.

Griswold, M.A. et al., 1999. Resolution enhancement in single-shot imaging using simultaneous acquisition of spatial harmonics (SMASH). *Magnetic Resonance in Medicine*, 41(6), pp.1236–1245.

Günther, M., 2014. Perfusion imaging. *Journal of Magnetic Resonance Imaging*, 40(2), pp.269–279.

Günther, M., Oshio, K. & Feinberg, D.A., 2005. Single-shot 3D imaging techniques improve arterial spin labeling perfusion measurements. *Magnetic Resonance in Medicine*, 54(2), pp.491–498.

Haacke, E.M. & Patrick, J.L., 1986. Reducing motion artifacts in two-dimensional Fourier transform imaging. *Magn Reson Imaging*, 4(4), pp.359–376.

Haase, A. et al., 1986. FLASH imaging. Rapid NMR imaging using low flip-angle pulses. *Journal of Magnetic Resonance (1969)*, 67(2), pp.258–266.

Hall, J.E., 2016. *Guyton and Hall Textbook of Medical Physiology (13th Ed.)*, Elsevier.

Hallowell, L.M. et al., 2008. Reviewing the process of preparing children for MRI. *Pediatric Radiology*, 38(3), pp.271–279.

Hammon, M. et al., 2016. Reproducibility of Kidney Perfusion Measurements With Arterial Spin Labeling at 1.5 Tesla MRI Combined With Semiautomatic Segmentation for Differential Cortical and Medullary Assessment. *Medicine*, 95(11), p.e3083. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26986143%5Cnhttp://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4839924.

Han, P.K. et al., 2016. Whole-brain perfusion imaging with balanced steady-state free precession arterial spin labeling. *NMR in Biomedicine*, 29(3), pp.264–274.

Haufe, S.E., Riedmüller, K. & Haberkorn, U., 2006. Nuclear Medicine Procedures for the Diagnosis of Acute and Chronic Renal Failure. *Nephron Clin Pract*, 103(2), pp.c77--c84. Available at: http://dx.doi.org/10.1159/000091576.

Hawkes, D.J., Studholme, C. & Hill, D.L.G., 1999. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 32(1), pp.71–86. Available at: http://www.sciencedirect.com/science/article/B6V14-3W83V9V-R/1/44f708e8fdb91c491ffcc992762ef514%5Cnhttp://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6V14-3W83V9V-R&_user=18704&_rdoc=1&_fmt=&_orig=search&_sort=d&view=c&_acct=C000002018&_version=1&_ur.

Heilmann, M. et al., 2012. Quantification of glomerular number and size distribution in normal rat kidneys using magnetic resonance imaging. *Nephrology, dialysis, transplantation*, 27(1), pp.100–7. Available at: http://www.ncbi.nlm.nih.gov/pubmed/21642513.

Henkelman, R.M. et al., 1993. Quantitative interpretation of magnetization transfer. *Magn Reson Med*, 29(6), pp.759–766. Available at: http://onlinelibrary.wiley.com/doi/10.1002/mrm.1910290607/abstract?systemMessage=Wiley+Online+Library+will+be+unavailable+for+up+to+3+hours+on+Saturday+19th+March+2016+from++11%3A00-14%3A00+GMT+%2F+07%3A00-10%3A00+EDT+%2F+19%3A00-22%3A00+SGT+for+essential.

Henkelman, R.M., Stanisz, G.J. & Graham, S.J., 2001. Magnetization transfer in MRI: a review. *NMR in Biomedicine*, 14(2), pp.57–64. Available at: http://dx.doi.org/10.1002/nbm.683.

Hennig, J., Nauerth, A. & Friedburg, H., 1986. RARE Imaging: A Fast Imaging Method for Clinical MR. *Magnetic Resonance in Medicine*, 3, pp.823–833.

Herget-Rosenthal, S., 2011. Imaging Techniques in the Management of Chronic Kidney Disease: Current Developments and Future Perspectives. *Seminars in Nephrology*, 31(3), pp.283–290. Available at: http://dx.doi.org/10.1016/j.semnephrol.2011.05.011.

Hess, A.T. et al., 2012. Real-time motion and B0 correction for localized adiabatic selective refocusing (LASER) MRSI using echo planar imaging volumetric navigators. *NMR in Biomedicine*, 25(2), pp.347–358.

Heusch, P. et al., 2014. Functional evaluation of transplanted kidneys using arterial spin labeling MRI. *Journal of Magnetic Resonance Imaging*, 40(1), pp.84–89.

Hill, D.L. et al., 2001. Medical image registration. *Physics in Medicine and Biology*, 46(3), pp.R1–R45.

Hoerger, T.J. et al., 2015. The future burden of CKD in the United States: A simulation model for the CDC CKD initiative. *American Journal of Kidney Diseases*, 65(3), pp.403–411.

Holland, A.E., Goldfarb, J.W. & Edelman, R.R., 1998. Diaphragmatic and cardiac motion during suspended breathing: preliminary experience and implications for breath-hold MR imaging. *Radiology*, 209(2), pp.483–489.

Hu, P. et al., 2011. Motion correction using coil arrays (MOCCA) for free-breathing cardiac cine MRI. *Magnetic Resonance in Medicine*, 66(2), pp.467–475.

Huang, F. et al., 2010. Data Convolution and Combination Operation (COCOA) for motion ghost artifacts reduction. *Magnetic Resonance in Medicine*, 64(1), pp.157–166.

Huang, Y. et al., 2011. Measurement and comparison of T1 relaxation times in native and transplanted kidney cortex and medulla. *Journal of Magnetic Resonance Imaging*, 33(5), pp.1241–1247.

Hueper, K. et al., 2015. Functional MRI detects perfusion impairment in renal allografts with delayed graft function. *American Journal of Physiology*, 308(12), pp.F1444-51. Available at: http://www.ncbi.nlm.nih.gov/pubmed/25925250.

Ingle, R.R. et al., 2013. Nonrigid autofocus motion correction for coronary MR angiography with

a 3D cones trajectory. *Magnetic Resonance in Medicine*, 72(2), pp.347–361. Available at: http://dx.doi.org/10.1002/mrm.24924.

Inker, L.A. et al., 2012. Estimating Glomerular Filtration Rate from Serum Creatinine and Cystatin C. *N Engl J Med*, 367(1), pp.20–29. Available at: http://dx.doi.org/10.1056/NEJMoa1114248.

Inoue, T. et al., 2012. Is there no future for renal BOLD-MRI? *Kidney International*, 82(8), pp.934–934.

James, M.T., Hemmelgarn, B.R. & Tonelli, M., 2010. Early recognition and prevention of chronic kidney disease. *The Lancet*, 375(9722), pp.1296–1309. Available at: http://dx.doi.org/10.1016/S0140-6736(09)62004-3.

Jenkinson, M. et al., 2012. FSL. *NeuroImage*, 62(2), pp.782–790.

Johnson, P.M. et al., 2016. Retrospective 3D motion correction using spherical navigator echoes. *Magnetic Resonance Imaging*, 34(9), pp.1274–1282. Available at: http://dx.doi.org/10.1016/j.mri.2016.06.006.

Jones, R.A. et al., 2011. Magnetic resonance imaging evaluation of renal structure and function related to disease: Technical review of image acquisition, postprocessing, and mathematical modeling steps. *Journal of Magnetic Resonance Imaging*, 33(6), pp.1270–1283. Available at: http://dx.doi.org/10.1002/jmri.22335.

Kaewlai, R. & Abujudeh, H., 2012. Nephrogenic Systemic Fibrosis. *American Journal of Roentgenology*, 199(1), pp.W17--W23. Available at: http://dx.doi.org/10.2214/AJR.11.8144.

Kainz, A. et al., 2015. Prediction of prevalence of chronic kidney disease in diabetic patients in countries of the European Union up to 2025. *Nephrology, dialysis, transplantation.*, 30 Suppl 4(suppl 4), p.iv113-8. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26209733 [Accessed November 23, 2016].

Karger, N. et al., 2000. Quantitation of renal perfusion using arterial spin labeling with FAIR-UFLARE. *Magnetic Resonance Imaging*, 18(6), pp.641–647. Available at: http://www.sciencedirect.com/science/article/pii/S0730725X00001557.

KDIGO, C.W.G., 2013. KDIGO 2012 Clinical Practice Guideline for the Evaluation and Management of Chronic Kidney Disease. *Kidney International Supplements*, 3(1), pp.4–4. Available at: dx.doi/10.1038/kisup.2012.76.

Kerr, M. et al., 2012. Estimating the financial cost of chronic kidney disease to the NHS in England. *Nephrology Dialysis Transplantation*, 27(SUPPL. 3).

Kiefer, C. et al., 2009. A Feasibility Study on Model-based Evaluation of Kidney Perfusion Measured by Means of FAIR Prepared True-FISP Arterial Spin Labeling (ASL) on a 3-T MR Scanner. *Academic Radiology*, 16(1), pp.79–87. Available at: http://dx.doi.org/10.1016/j.acra.2008.04.024.

Kim, D.W. et al., 2017. Measurement of arterial transit time and renal blood flow using pseudocontinuous ASL MRI with multiple post-labeling delays: Feasibility, reproducibility, and variation. *Journal of Magnetic Resonance Imaging*, 46(3), pp.813–819. Available at: http://doi.wiley.com/10.1002/jmri.25634.

Kim, S.G., 1995. Quantification of Relative Cerebral Blood-Flow Change by Flow-Sensitive Alternating Inversion-Recovery (FAIR) Technique - Application to Functional Mapping. *Magnetic Resonance in Medicine*, 34(3), pp.293–301.

Klein, S. et al., 2009. Adaptive stochastic gradient descent optimisation for image registration. *International Journal of Computer Vision*, 81(3), pp.227–239.

Klein, S. et al., 2010. elastix: A Toolbox for Intensity-Based Medical Image Registration. *IEEE Transactions on Medical Imaging*, 29(1), pp.196–205.

Klein, S., Staring, M. & Pluim, J.P.W., 2007. Evaluation of optimization methods for nonrigid medical image registration using mutual information and B-splines. *IEEE T Image Processing*, 16(12), pp.2879–90.

Knox, F.G., Ritman, E.L. & Romero, J.C., 1984. Intrarenal Distribution of Blood Flow: Evolution of a New Approach to Measurement. *Kidney International*, 25(3), pp.473–479.

Kober, T. et al., 2011. Head motion detection using FID navigators. *Magnetic Resonance in Medicine*, 66(1), pp.135–143. Available at: http://dx.doi.org/10.1002/mrm.22797.

Kober, T., Gruetter, R. & Krueger, G., 2012. Prospective and retrospective motion correction in diffusion magnetic resonance imaging of the human brain. *NeuroImage*, 59(1), pp.389–398. Available at: http://dx.doi.org/10.1016/j.neuroimage.2011.07.004.

Van Der Kouwe, A.J.W., Benner, T. & Dale, A.M., 2006. Real-time rigid body motion correction and shimming using cloverleaf navigators. *Magnetic Resonance in Medicine*, 56(5), pp.1019–1032.

Kwong, K.K. et al., 1995. MR perfusion studies with T1-weighted echo planar imaging. *Magn Reson Med*, 34(6), pp.878–887.

Lanzman, R.S. et al., 2012. Arterial Spin-labeling MR Imaging of Renal Masses: Correlation with Histopathologic Findings. *Radiology*, 265(3), pp.799–808.

Lanzman, R.S. et al., 2010. Quantification of renal allograft perfusion using arterial spin labeling MRI: initial results. *European Radiology*, 20(6), pp.1485–1491. Available at: http://dx.doi.org/10.1007/s00330-009-1675-0.

Larkman, D.J., Atkinson, D. & Hajnal, J. V, 2004. Artifact Reduction Using Parallel Imaging Methods. *Top Magn Reson Imaging*, 15(4), pp.267–275.

Lauterbur, P.C., 1973. Image formation by induced local interactions. Examples employing nuclear magnetic resonance. *Nature*, 242, pp.190–191.

Lee, V.S. et al., 2007. Renal function measurements from MR renography and a simplified multicompartmental model. *American Journal of Physiology, Renal Physiology*, 292(5), pp.F1548--F1559. Available at: http://dx.doi.org/10.1152/ajprenal.00347.2006.

Lerman, L.O., Rodriguez-Porcel, M. & Romero, J.C., 1999. The development of x-ray imaging to study renal function. *Kidney International*, 55(2), pp.400–416.

Levey, A.S. et al., 2009. A new equation to estimate glomerular filtration rate. *Ann. Intern. Med.*, 150(9), pp.604–612.

Levin, A., Levin, A. & Stevens, P.E., 2011. Early detection of CKD: the benefits, limitations and effects on prognosis. *Nat. Rev. Nephrol*, 7(10), pp.446–45786. Available at: www.nature.com/nrneph.

Li, L.-P. et al., 2017. Evaluation of Renal Blood Flow in Chronic Kidney Disease Using Arterial Spin Labeling Perfusion MRI. *Kidney International Reports*, 2(1), pp.36–43. Available at: http://linkinghub.elsevier.com/retrieve/pii/S2468024916300754.

Liang, X. et al., 2014. A variable flip angle-based method for reducing blurring in 3D GRASE ASL. *Phys Med Biol*, 59(18), pp.5559–5573. Available at: http://iopscience.iop.org/0031-9155/59/18/5559/pdf/0031-9155_59_18_5559.pdf.

Lin, W. et al., 2007. Image metric-based correction (autofocusing) of motion artifacts in high-resolution trabecular bone imaging. *Journal of Magnetic Resonance Imaging*, 26(1), pp.191–197.

Liu, J. & Drangova, M., 2011. Rapid six-degree-of-freedom motion detection using prerotated baseline spherical navigator echoes. *Magnetic Resonance in Medicine*, 65(2), pp.506–514.

Loktyushin, A. et al., 2015. Blind multirigid retrospective motion correction of MR images. *Magnetic Resonance in Medicine*, 73(4), pp.1457–1468. Available at: http://dx.doi.org/10.1002/mrm.25266.

Loktyushin, A. et al., 2013. Blind retrospective motion correction of MR images. *Magnetic Resonance in Medicine*, 70(6), pp.1608–1618. Available at: http://dx.doi.org/10.1002/mrm.24615.

Lopez-Giacoman, S. & Madero, M., 2015. Biomarkers in chronic kidney disease, from kidney function to kidney damage. *World journal of nephrology*, 4(1), pp.57–73. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4317628&tool=pmcentrez&rendertype=abstract.

Luh, W.M. et al., 1999. QUIPSS II with thin-slice TI1 periodic saturation: A method for improving accuracy of quantitative perfusion imaging using pulsed arterial spin labeling. *Magnetic Resonance in Medicine*, 41(6), pp.1246–1254.

Lustig, M., Donoho, D. & Pauly, J.M., 2007. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6), pp.1182–1195.

Maclaren, J. et al., 2012. Measurement and Correction of Microscopic Head Motion during Magnetic Resonance Imaging of the Brain. *PLoS ONE*, 7(11), pp.3–11.

Maclaren, J. et al., 2013. Prospective motion correction in brain imaging: A review. *Magnetic Resonance in Medicine*, 69(3), pp.621–636.

MacLaren, J. et al., 2011. Combined prospective and retrospective motion correction to relax navigator requirements. *Magnetic Resonance in Medicine*, 65(6), pp.1724–1732.

Maclaren, J., Aksoy, M. & Bammer, R., 2015. Contact-free physiological monitoring using a markerless optical system. *Magnetic resonance in medicine*, 74(2), pp.571–577.

Maes, F., Vandermeulen, D. & Suetens, P., 1999. Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information. *Medical Image Analysis*, 3(4), pp.373–386.

Maintz, J.B. & Viergever, M.A., 1998. A survey of medical image registration. *Med Image Anal*, 2(1), pp.1–36.

Manduca, a et al., 2000. Autocorrection in MR imaging: adaptive motion correction without navigator echoes. *Radiology*, 215(3), pp.904–909.

Mansfield, P., 1977. Multiplanar Image Formation using NMR spin-echoes. *J. Phys. Chem. Solid. State. Phys*, 10, pp.L55–L58.

Martirosian, P. et al., 2004. FAIR true-FISP perfusion imaging of the kidneys. *Magnetic Resonance in Medicine*, 51(2), pp.353–361. Available at: http://dx.doi.org/10.1002/mrm.10709.

McDonald, G.S. & Leigh Jr., J.S., 1973. A New Method for Measuring Longitudinal Relaxation Times. *Journal of Magnetic Resonance*, (9), pp.358–362.

McGee, K.P. et al., 2000. Image metric-based correction (autocorrection) of motion effects: analysis of image metrics. *J Magn Reson Imaging*, 11(2), pp.174–181.

Meiboom, S. & Gill, D., 1958. Modified spin-echo method for measuring nuclear relaxation times. *Review of Scientific Instruments*, 29(8), pp.688–691.

Meyers, M.A., Charnsangavej, C. & Oliphant, M., 2011. *Meyers' Dynamic Radiology of the Abdomen - Normal and Pathologic Anatomy*, Springer.

Michaely, H.J. et al., 2012. Renal BOLD-MRI does not reflect renal function in chronic kidney disease. *Kidney International*, 81(7), pp.684–689. Available at: http://dx.doi.org/10.1038/ki.2011.455.

Michaely, H.J. et al., 2004. Renal disease: value of functional magnetic resonance imaging with flow and perfusion measurements. *Investigative radiology*, 39(11), pp.698–705. Available at: http://www.ncbi.nlm.nih.gov/pubmed/15486531.

Mugler, J.P., 1999. Improved three-dimensional GRASE imaging with the SORT phase-encoding strategy. *Journal of Magnetic Resonance Imaging*, 9(4), pp.604–612.

Mutsaerts, H.J.M.M. et al., 2015. Multi-vendor reliability of arterial spin labeling perfusion {MRI} using a near-identical sequence: Implications for multi-center studies. *NeuroImage*, 113, pp.143–152. Available at: http://dx.doi.org/10.1016/j.neuroimage.2015.03.043.

Nery, F. et al., 2016a. Automatic rejection of motion corrupted data in renal ASL. In *Proceedings of the ESMRMB 33nd Annual Meeting, Vienna, Austria; p. S362-S363*.

Nery, F. et al., 2017a. Impact of image registration on renal T1 mapping in children with chronic kidney disease. In *Proceedings of the ISMRM 25th Annual Meeting & Exhibition, Honolulu, HI, USA, 1306*.

Nery, F. et al., 2017b. Improved reproducibility of longitudinal renal ASL perfusion measurements in children with chronic kidney disease using retrospective motion correction. In *Proceedings of the ISMRM 25th Annual Meeting & Exhibition, Honolulu, HI, USA, 1887*.

Nery, F. et al., 2015a. Investigation and minimisation of the sources of image artefact in 3D-GRASE ASL of the kidney. In *Proceedings of the ESMRMB 32nd Annual Meeting, Edinburgh, Scotland; p. S370-S371*.

Nery, F. et al., 2015b. Optimisation of 3D GRASE ASL acquisition schemes for renal perfusion measurement. In *Proceedings of the British Chapter of the ISMRM Annual Meeting*. p. 26.

Nery, F. et al., 2016b. Single-shot 3D GRASE ASL for fast, motion-insensitive whole kidney perfusion mapping. In *Proceedings of the ESMRMB 33nd Annual Meeting, Vienna, Austria; p. S285-S286*.

Nguyen, T.D. et al., 2003. Direct monitoring of coronary artery motion with cardiac fat navigator echoes. *Magnetic Resonance in Medicine*, 50(2), pp.235–241.

Niendorf, T. et al., 2015. How bold is blood oxygenation level-dependent (BOLD) magnetic resonance imaging of the kidney? Opportunities, challenges and future directions. *Acta Physiologica*, 213(1), pp.19–38. Available at: http://dx.doi.org/10.1111/apha.12393.

Niles, D.J. et al., 2016. Longitudinal Assessment of Renal Perfusion and Oxygenation in Transplant Donor-Recipient Pairs Using Arterial Spin Labeling and Blood Oxygen Level-Dependent Magnetic Resonance Imaging. *Investigative Radiology*, 51(2), pp.113–120. Available at: http://content.wkhealth.com/linkback/openurl?sid=WKPTLP:landingpage&an=00004424-201602000-00006.

Norris, D.G., 1991. Ultrafast low-angle RARE: U-FLARE. *Magnetic Resonance in Medicine*, 17, pp.539–542.

O'Neill, W.C., 2000. Sonographic evaluation of renal failure. *Am. J. Kidney Dis.*, 35(6), pp.1021–1038. Available at: http://www.sciencedirect.com/science/article/pii/S0272638600700369.

Odille, F. et al., 2008. Generalized Reconstruction by Inversion of Coupled Systems (GRICS) applied to free-breathing MRI. *Magnetic Resonance in Medicine*, 60(1), pp.146–157. Available at: http://dx.doi.org/10.1002/mrm.21623.

Ogg, R.J., Kingsley, P.B. & Taylor, J.S., 1994. WET, a T1- and B1-Insensitive Water-Suppression Method for in Vivo Localized 1H NMR Spectroscopy. *Journal of Magnetic Resonance, Series B*, 104(1), pp.1–10.

Oliveira, F.P.M. & Tavares, J.M.R.S., 2014. Medical image registration: a review. *Computer Methods in Biomechanics and Biomedical Engineering*, 17(2), pp.73–93. Available at: http://www.ncbi.nlm.nih.gov/pubmed/22435355%5Cnhttp://www.tandfonline.com/doi/abs/10.1080/10255842.2012.670855.

Ooi, M.B. et al., 2013. Prospective motion correction using inductively coupled wireless RF coils. *Magnetic Resonance in Medicine*, 70(3), pp.639–647.

OpenStax, 2016. *Anatomy & Physiology*, OpenStax, available at https://openstax.org/details/anatomy-and-physiology. Available at: https://openstax.org/details/anatomy-and-physiology.

Oppelt, A. et al., 1986. FISP — a new fast MRI sequence. *Electromedica*, 54, pp.15–18. Available at: http://mri-q.com/uploads/3/2/7/4/3274160/oppelt_fisp_1986.pdf.

Oshio, K. & Feinberg, D. a, 1991. GRASE (Gradient- and spin-echo) imaging: a novel fast MRI technique. *Magnetic Resonance in Medicine*, 20, pp.344–349.

Ott, C. et al., 2013. Vascular and renal hemodynamic changes after renal denervation. *Clinical journal of the American Society of Nephrology*, 8(7), pp.1195–201. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3700691&tool=pmcentrez&rendertype=abstract.

Pallone, T.L., Edwards, A. & Mattson, D.L., 2012. Renal Medullary Circulation. *Comprehensive Physiology*, 2, pp.97–140. Available at: http://dx.doi.org/10.1002/cphy.c100036.

Park, S.-H., Wang, D.J.J. & Duong, T.Q., 2013. Balanced steady state free precession for arterial spin labeling MRI: Initial experience for blood flow mapping in human brain, retina, and kidney. *Magnetic Resonance Imaging*, 31(7), pp.1044–1050. Available at: http://dx.doi.org/10.1016/j.mri.2013.03.024.

Parkes, L.M. & Tofts, P.S., 2002. Improved accuracy of human cerebral blood perfusion measurements using arterial spin labeling: Accounting for capillary water permeability. *Magnetic Resonance in Medicine*, 48(1), pp.27–41.

Pennec, X., Cachier, P. & Ayache, N., 1999. Understanding the "Demon's Algorithm": 3D non-rigid registration by gradient descent. *International Conference on Medical Image*

Computing And Computer-Assisted Intervention (MICCAI), 1679, pp.597–605. Available at: http://apps.isiknowledge.com/InboundService.do?Func=Frame&product=WOS&action=r etrieve&SrcApp=Papers&UT=000171179200064&SID=X167jP544jG@o4LegiD&SrcAuth= mekentosj&Init=Yes&mode=FullRecord&customersID=mekentosj&DestFail=http://access .isiproducts.com/custom_i.

Petersen, E.T. et al., 2006. Non-invasive measurement of perfusion: a critical review of arterial spin labelling techniques. *British Journal of Radiology*, 79(944), pp.688–701. Available at: http://dx.doi.org/10.1259/bjr/67705974.

Pluim, J.P.W., Maintz, J.B. a & Viergever, M. a, 2003. Mutual information based registration of medical images: a survey. *IEEE Transactions on medical imaging*, XX(Y), pp.1–21.

Power, J.D. et al., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage*, 59(3), pp.2142–2154. Available at: http://dx.doi.org/10.1016/j.neuroimage.2011.10.018.

Prasad, P. V, Edelman, R.R. & Epstein, F.H., 1996. Noninvasive evaluation of intrarenal oxygenation with BOLD MRI. *Circulation*, 94(12), pp.3271–3275.

Press, W. et al., 2002. *Numerical Recipes: The Art of Scientific Computing (2nd. Ed.)*, Cambridge University Press.

Pruessmann, K.P. et al., 1999. SENSE: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5), pp.952–962.

Regan, M.C. et al., 1995. Regional renal blood flow in normal and disease states. *Urol. Res.*, 23(1), pp.1–10.

Ren, T. et al., 2016. Evaluation of renal allografts function early after transplantation using intravoxel incoherent motion and arterial spin labeling MRI. *Magnetic Resonance Imaging*, 34(7), pp.908–914. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0730725X16300340.

Ridgway, G., 2008. Rice/Rician distribution - Functions for Rice/Rician PDF: stats, samples, and distribution fitting. Available at: https://uk.mathworks.com/matlabcentral/fileexchange/14237-rice-rician-distribution [Accessed April 5, 2017].

Ritt, M. et al., 2010. Measurement of kidney perfusion by magnetic resonance imaging: Comparison of MRI with arterial spin labeling to para-aminohippuric acid plasma clearance

in male subjects with metabolic syndrome. *Nephrology Dialysis Transplantation*, 25(4), pp.1126–1133.

Robbins, H. & Monro, S., 1951. A Stochastic Approximation Method. *The Annals of Mathematical Statistics*, 22(3), pp.400–407.

Roberts, D.A. et al., 1995. Renal perfusion in humans: MR imaging with spin tagging of arterial water. *Radiology*, 196(1), pp.281–6. Available at: http://www.ncbi.nlm.nih.gov/pubmed/7784582.

Robson, M.D. et al., 2003. Magnetic Resonance : An Introduction to Ultrashort TE (UTE) Imaging. *Journal of Computer Assisted Tomography*, 27(6), pp.825–846.

Robson, P.M. et al., 2009. Strategies for reducing respiratory motion artifacts in renal perfusion imaging with arterial spin labeling. *Magnetic Resonance in Medicine*, 61(6), pp.1374–1387.

Robson, P.M. et al., 2016. Volumetric Arterial Spin-labeled Perfusion Imaging of the Kidneys with a Three-dimensional Fast Spin Echo Acquisition. *Academic Radiology*, 23(2), pp.144–154. Available at: http://dx.doi.org/10.1016/j.acra.2015.09.013.

Romero, M., Buxbaum, J.L. & Palmer, S.L., 2014. Magnetic resonance imaging of the gut: a primer for the luminal gastroenterologist. *The American journal of gastroenterology*, 109(4), p.497–509; quiz 510. Available at: http://www.ncbi.nlm.nih.gov/pubmed/24394750.

Rossi, C. et al., 2012. Histogram Analysis of Renal Arterial Spin Labeling Perfusion Data Reveals Differences Between Volunteers and Patients With Mild Chronic Kidney Disease. *Investigative Radiology*, 47(8), pp.490–496.

Rueckert et al., D., 1999. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Trans. Med. Imag.*, 18(8), pp.712–21. Available at: http://www.ncbi.nlm.nih.gov/pubmed/10534053.

Saladin, K.S., Sullivan, S.J. & Gan, C.A., 2016. *Human Anatomy (5th Ed.)*, McGraw-Hill Education.

Sandilands, E.A. et al., 2013. Measurement of renal function in patients with chronic kidney disease. *Br J Clin Pharmacol*, 76(4), pp.504–515. Available at: http://dx.doi.org/10.1111/bcp.12198.

Sarnak, M.J. et al., 2003. Kidney Disease as a Risk Factor for Development of Cardiovascular Disease: A Statement From the American Heart Association Councils on Kidney in

Cardiovascular Disease, High Blood Pressure Research, Clinical Cardiology, and Epidemiology and Prevention. *Circulation*, 108(17), pp.2154–2169.

Schewzow, K. et al., 2010. Image Registration in ASL-Perfusion Imaging of Kidney - Impact on Image Quality. *Proceedings of the ISMRM & ESMRMB 2010 Joint Annual Meeting & Exhibition, Stockholm, Sweden, 1780*.

Schmitt, F., Stehling, M.K. & Turner, R., 1998. *Echo-planar Imaging - Theory, Technique and Application*, Springer.

Schneider, A. et al., 2011. Bench-to-bedside review: contrast enhanced ultrasonography--a promising technique to assess renal perfusion in the ICU. *Critical care*, 15(3), p.157.

Schneider, A.G. et al., 2012. Renal perfusion evaluation with contrast-enhanced ultrasonography. *Nephrology Dialysis Transplantation*, 27(2), pp.674–681.

Schneider, A.G., Goodwin, M.D. & Bellomo, R., 2013. Measurement of kidney perfusion in critically ill patients. *Critical Care*, (17), p.220. Available at: http://link.springer.com/10.1007/978-3-642-35109-9_51%5Cnpapers3://publication/doi/10.1007/978-3-642-35109-9_51.

Schneider, M.P. et al., 2012. Reversibility of the effects of aliskiren in the renal versus systemic circulation. *Clinical Journal of the American Society of Nephrology*, 7(2), pp.258–264.

Schultz, C.L. et al., 1984. The effect of motion on two-dimensional Fourier transformation magnetic resonance images. *Radiology*, 152(1), pp.117–121.

Schwartz, G.J. et al., 2009. New Equations to Estimate GFR in Children with CKD. *Journal of the American Society of Nephrology*, 20(3), pp.629–637. Available at: http://dx.doi.org/10.1681/ASN.2008030287.

Seif, M. et al., 2015. Image registration for triggered and non-triggered DTI of the human kidney: Reduced variability of diffusion parameter estimation. *Journal of Magnetic Resonance Imaging*, 41(5), pp.1228–1235.

Semelka, R.C. et al., 1996. HASTE MR imaging: description of technique and preliminary results in the abdomen. *J Magn Reson Imaging*, 6(4), pp.698–699. Available at: http://www.ncbi.nlm.nih.gov/pubmed/8835965.

Shannon, C.E., 1948. A mathematical theory of communication. *The Bell System Technical Journal*, 27, pp.379–423. Available at: http://cm.bell-

labs.com/cm/ms/what/shannonday/shannon1948.pdf.

Shimizu, K. et al., 2017. Arterial Transit Time-corrected Renal Blood Flow Measurement with Pulsed Continuous Arterial Spin Labeling MR Imaging. *Magnetic Resonance in Medical Sciences*, 16(1), pp.38–44. Available at: https://www.jstage.jst.go.jp/article/mrms/16/1/16_mp.2015-0117/_article.

Shirzadi, Z. et al., 2015. Automated removal of spurious intermediate cerebral blood flow volumes improves image quality among older patients: A clinical arterial spin labeling investigation. *Journal of Magnetic Resonance Imaging*, 42(5), pp.1377–1385.

Silva, a C. et al., 1995. Multi-slice MRI of rat brain perfusion during amphetamine stimulation using arterial spin labeling. *Magnetic Resonance in Medicine*, 33(2), pp.209–214.

Siva, S. et al., 2013. An analysis of respiratory induced kidney motion on four-dimensional computed tomography and its implications for stereotactic kidney radiotherapy. *Radiation Oncology*, 8(1), p.248. Available at: http://ro-journal.biomedcentral.com/articles/10.1186/1748-717X-8-248.

Skare, S. et al., 2015. Properties of a 2D fat navigator for prospective image domain correction of nodding motion in brain MRI. *Magnetic Resonance in Medicine*, 73(3), pp.1110–1119.

Smyser, C.D. et al., 2010. Longitudinal analysis of neural network development in preterm infants. *Cerebral Cortex*, 20(12), pp.2852–2862.

Sodickson, D.K. & Manning, W.J., 1997. Simultaneous Acquisition of Spatial Harmonics (SMASH): Fast Imaging with Radiofrequency Coil Arrays. *Magnetic Resonance in Medicine*, 38, pp.591–603.

Song, R. et al., 2011. Evaluation of respiratory liver and kidney movements for MRI navigator gating. *Journal of Magnetic Resonance Imaging*, 33(1), pp.143–148.

Song, R., Loeffler, R.B. & Hillenbrand, C.M., 2010. Improved renal perfusion measurement with a dual navigator-gated Q2TIPS fair technique. *Magnetic Resonance in Medicine*, 64(5), pp.1352–1359.

Sotiras, A., Davatzikos, C. & Paragios, N., 2012. *Deformable medical image registration: A survey*,

Sourbron, S.P. et al., 2008. MRI-measurement of perfusion and glomerular filtration in the human kidney with a separable compartment model. *Investigative Radiology*, 43(1), pp.40–48.

Stenvinkel, P., 2010. Chronic kidney disease: A public health priority and harbinger of premature cardiovascular disease. *Journal of Internal Medicine*, 268(5), pp.456–467.

Stevens, L.A. & Levey, A.S., 2009. Measured GFR as a Confirmatory Test for Estimated GFR. *Journal of the American Society of Nephrology*, 20(11), pp.2305–2313. Available at: http://dx.doi.org/10.1681/ASN.2009020171.

Studholme, C., Hill, D.L.G. & Hawkes, D.J., 1995. Multiresolution voxel similarity measures for MR-PET registration. In *Information Processing in Medical Imaging (IPMI '95)*. pp. 287–298.

Studler, U. et al., 2010. Impact of motion on T1 mapping acquired with inversion recovery fast spin echo and rapid spoiled gradient recalled-echo pulse sequences for delayed gadolinium-enhanced MRI of cartilage (dGEMRIC) in volunteers. *Journal of Magnetic Resonance Imaging*, 32(2), pp.394–398.

Sugimori, H. et al., 2013. Evaluation of renal blood flow using multi-phase echo-planar magnetic resonance imaging and signal targeting with alternating radiofrequency (EPISTAR) in 3-T magnetic resonance imaging. *Radiological Physics and Technology*, 6(1), pp.86–91.

Taal, M.W. et al., 2012. *Brenner & Rector's The Kidney (9th Edition)*, Saunders, Elsevier.

Tan, H. et al., 2009. A fast, effective filtering method for improving clinical pulsed arterial spin labeling {MRI}. *J. Magn. Reson. Imaging*, 29(5), pp.1134–1139. Available at: http://dx.doi.org/10.1002/jmri.21721.

Tan, H., Koktzoglou, I. & Prasad, P. V, 2014. Renal perfusion imaging with two-dimensional navigator gated arterial spin labeling. *Magnetic Resonance in Medicine*, 71(2), pp.570–579. Available at: http://dx.doi.org/10.1002/mrm.24692.

Tanenbaum, A.B. et al., 2015. A method for reducing the effects of motion contamination in arterial spin labeling magnetic resonance imaging. *Journal of Cerebral Blood Flow & Metabolism*, 35(10), pp.1697–702. Available at: http://www.ncbi.nlm.nih.gov/pubmed/26036937.

Telischak, N.A., Detre, J.A. & Zaharchuk, G., 2015. Arterial spin labeling MRI: Clinical applications in the brain. *J. Magn. Reson. Imaging*, 41(5), pp.1165–1180. Available at: http://dx.doi.org/10.1002/jmri.24751.

Thesen, S. et al., 2000. Prospective Acquisition Correction for head motion with image-based tracking for real-time fMRI. *Magnetic Resonance in Medicine*, 44(3), pp.457–465.

Thirion, J.P., 1998. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical image analysis*, 2(3), pp.243–260.

Thoeny, H.C. & De Keyzer, F., 2011. Diffusion-weighted MR imaging of native and transplanted kidneys. *Radiology*, 259(1), pp.25–38.

Thomas, D.L. et al., 2016. Reduction of motion artefacts in multi-shot 3D GRASE Arterial Spin Labelling using Autofocus. In *Proceedings of the ISMRM 24th Annual Meeting & Exhibition, Singapore, 2874*.

Tisdall, M.D. et al., 2016. Prospective motion correction with volumetric navigators (vNavs) reduces the bias and variance in brain morphometry induced by subject motion. *NeuroImage*, 127, pp.11–22. Available at: http://dx.doi.org/10.1016/j.neuroimage.2015.11.054.

Tisdall, M.D. et al., 2012. Volumetric navigators for prospective motion correction and selective reacquisition in neuroanatomical MRI. *Magnetic Resonance in Medicine*, 68(2), pp.389–399.

Tofts, P.S. et al., 2012. Precise measurement of renal filtration and vascular parameters using a two-compartment model for dynamic contrast-enhanced MRI of the kidney gives realistic normal values. *European Radiology*, 22(6), pp.1320–1330. Available at: http://dx.doi.org/10.1007/s00330-012-2382-9.

Tortora, G.J. & Derrickson, B., 2014. *Principles of Anatomy & Physiology (14th Ed.)*,

Tyler, D.J. et al., 2007. Magnetic resonance imaging with ultrashort TE (UTE) PULSE sequences: Technical considerations. *Journal of Magnetic Resonance Imaging*, 25(2), pp.279–289.

Viola, P. & Wells, W.M.I., 1997. Alignment by maximization of mutual information. *Proceedings of IEEE International Conference on Computer Vision*, 24(2), pp.16–23.

Vold, R. et al., 1968. Measurement of spin relaxation in complex systems. *J Chem Phys*, 48, pp.3831–3832.

Wan, L. et al., 2008. An assessment of the accuracy of renal blood flow estimation by Doppler ultrasound. *Intensive Care Medicine*, 34(8), pp.1503–1510. Available at: http://dx.doi.org/10.1007/s00134-008-1106-8.

Wang, J. et al., 2012. Hemodynamic Effects of Furosemide on Renal Perfusion as Evaluated by ASL-MRI. *Academic Radiology*, 19(10), pp.1194–1200. Available at:

http://dx.doi.org/10.1016/j.acra.2012.04.021.

Wang, J.J. et al., 1998. Perfusion quantitation in transplanted rat kidney by MRI with arterial spin labeling. *Kidney International*, 53(6), pp.1783–1791.

Wang, Y. et al., 2011. Incidence of nephrogenic systemic fibrosis after adoption of restrictive gadolinium-based contrast agent guidelines. *Radiology*, 260(1), pp.105–111.

Wang, Z. et al., 2008. Empirical optimization of ASL data analysis using an ASL data processing toolbox: ASLtbx. *Magnetic Resonance Imaging*, 26(2), pp.261–269.

Warmuth, C. et al., 2007. Accuracy of blood flow values determined by arterial spin labeling: A validation study in isolated porcine kidneys. *J. Magn. Reson. Imaging*, 26(2), pp.353–358.

Warmuth, C. et al., 2005. Arterial Spin Labeling Perfusion Imaging of the Kidneys using Prospective Navigators and Slice Following. *Proceedings of the 13th ISMRM Annual Meeting, 1918*.

Wasung, M.E., Chawla, L.S. & Madero, M., 2015. Biomarkers of renal function, which and when? *Clinica Chimica Acta*, 438(1), pp.350–357. Available at: http://dx.doi.org/10.1016/j.cca.2014.08.039.

Waszak, M. et al., 2016. Prospective head motion correction using FID-guided on-demand image navigators. *Magnetic Resonance in Medicine*, 78(1), pp.193–203.

Welch, E.B. et al., 2002. Spherical navigator echoes for full 3-D rigid body motion measurement in MRI. *Magnetic Resonance in Medicine*, 47, pp.32–41.

Wentland, A.L. et al., 2009. Quantitative MR Measures of Intrarenal Perfusion in the Assessment of Transplanted Kidneys. *Academic Radiology*, 16(9), pp.1077–1085. Available at: http://dx.doi.org/10.1016/j.acra.2009.03.020.

White, M.J. et al., 2009. Motion artifact correction in free-breathing abdominal MRI using overlapping partial samples to recover image deformations. *Magnetic Resonance in Medicine*, 62(2), pp.440–449. Available at: http://dx.doi.org/10.1002/mrm.22017.

White, N. et al., 2010. PROMO: Real-time prospective motion correction in MRI using image-based tracking. *Magnetic Resonance in Medicine*, 63(1), pp.91–105.

Williams, D.S. et al., 1992. Magnetic resonance imaging of perfusion using spin inversion of arterial water. *Proceedings of the National Academy of Sciences*, 89(1), pp.212–6. Available at:

http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=48206&tool=pmcentrez&rendertype=abstract.

Winkelmann, R., Bornert, P. & Dossel, O., 2005. Ghost artifact removal using a parallel imaging approach. *Magnetic Resonance in Medicine*, 54(4), pp.1002–1009.

Winter, J.D., St. Lawrence, K.S. & Margaret Cheng, H.L., 2011. Quantification of renal perfusion: Comparison of arterial spin labeling and dynamic contrast-enhanced MRI. *Journal of Magnetic Resonance Imaging*, 34(3), pp.608–615.

Wolff, S.D. & Balaban, R.S., 1989. Magnetization transfer contrast (MTC) and tissue water proton relaxation in vivo. *Magnetic Resonance in Medicine*, 10(1), pp.135–144. Available at: http://www.ncbi.nlm.nih.gov/pubmed/2547135.

Wong, E.C., 2014. An introduction to ASL labeling techniques. *Journal of Magnetic Resonance Imaging*, 40(1), pp.1–10. Available at: http://dx.doi.org/10.1002/jmri.24565.

Wong, E.C. et al., 2006. Velocity-selective arterial spin labeling. *Magnetic Resonance in Medicine*, 55(6), pp.1334–1341. Available at: http://dx.doi.org/10.1002/mrm.20906.

Wong, E.C., Buxton, R.B. & Frank, L.R., 1998. A theoretical and experimental comparison of continuous and pulsed arterial spin labeling techniques for quantitative perfusion imaging. *Magnetic Resonance in Medicine*, 40(3), pp.348–55. Available at: http://www.ncbi.nlm.nih.gov/pubmed/9727936.

Wong, E.C., Buxton, R.B. & Frank, L.R., 1998. Quantitative imaging of perfusion using a single subtraction (QUIPSS and QUIPSS II). *Magnetic Resonance in Medicine*, 39(5), pp.702–708.

Wood, M.L. & Henkelman, R.M., 1984. MR image artifacts from periodic motion. *Medical Physics*, 12(2), pp.143–151.

Wu, W.-C. et al., 2011. Renal Perfusion 3-T MR Imaging: A Comparative Study of Arterial Spin Labeling and Dynamic Contrast-enhanced Techniques. *Radiology*, 261(3), pp.845–853.

Ye, F.Q. et al., 2000. Noise reduction in 3D perfusion imaging by attenuating the static signal in arterial spin tagging (ASSIST). *Magnetic Resonance in Medicine*, 44(1), pp.92–100.

Yuan, Q. et al., 2016. Intratumor Heterogeneity of Perfusion and Diffusion in Clear-Cell Renal Cell Carcinoma: Correlation With Tumor Cellularity. *Clinical Genitourinary Cancer*, 14(6), pp.e585–e594. Available at: http://dx.doi.org/10.1016/j.clgc.2016.04.007.

Zaharchuk, G. et al., 1999. Multislice perfusion and perfusion territory imaging in humans with

separate label and image coils. *Magnetic Resonance in Medicine*, 41(6), pp.1093–1098.

Zaitsev, M. et al., 2006. Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system. *NeuroImage*, 31(3), pp.1038–1050.

Zaitsev, M., Maclaren, J. & Herbst, M., 2015. Motion artifacts in MRI: A complex problem with many partial solutions. *Journal of Magnetic Resonance Imaging*, 42(4), pp.887–901.

Zhang, J.L. et al., 2014. New magnetic resonance imaging methods in nephrology. *Kidney International*, 85(4), pp.768–778. Available at: http://www.sciencedirect.com/science/article/pii/S0085253815562765.

Zhang, W. et al., 1995. NMR measurement of perfusion using arterial spin labeling without saturation of macromolecular spins. *Magn Reson Med*, 33(3), pp.370–376.

Zhang, Y. et al., 2016. Tumor Vascularity in Renal Masses: Correlation of Arterial Spin-Labeled and Dynamic Contrast-Enhanced Magnetic Resonance Imaging Assessments. *Clinical Genitourinary Cancer*, 14(1), pp.e25–e36. Available at: http://dx.doi.org/10.1016/j.clgc.2015.08.007.

Zimmer, F. et al., 2013. Quantitative Renal Perfusion Measurements in a Rat Model of Acute Kidney Injury at 3T: Testing Inter- and Intramethodical Significance of ASL and DCE-MRI. *PLoS ONE*, 8(1).

Zitová, B. & Flusser, J., 2003. Image registration methods: A survey. *Image and Vision Computing*, 21(11), pp.977–1000.

Zun, Z., Shankaranarayanan, A. & Zaharchuk, G., 2014. Pseudocontinuous arterial spin labeling with prospective motion correction (PCASL-PROMO). *Magnetic Resonance in Medicine*, 72(4), pp.1049–1056.