Signal sampling and processing in magnetic resonance applications

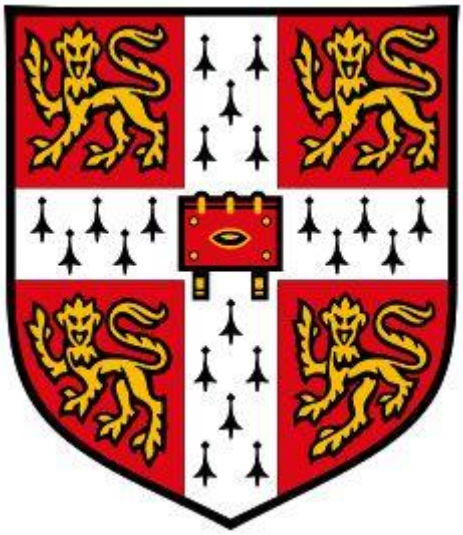

Andi Reci

Trinity College

This dissertation is submitted for the degree of Doctor of Philosophy

September 2018

Department of Chemical Engineering and Biotechnology

University of Cambridge

Acknowledgements

I would like to thank Professor Lynn Gladden for her guidance throughout the course of the degree, for allowing me the right amount of freedom in the projects that I pursued and for her encouragement and faith during tough research times. I am also indebted to Dr Andy Sederman for his day-to-day support in the lab; I have been lucky to draw from his vast knowledge and experience in the magnetic resonance field.

I would like to thank several researchers with whom I have collaborated, whose work is acknowledged accordingly in the text. My work would not have been complete without their input. Dr Daan de Kort, Dr Mohamed Ainte, Dr Chris Robertson and Tian Li of the Magnetic Resonance Research Centre, University of Cambridge are thanked for the help with experiments; Dr Martin Benning of the Department of Applied Mathematics and Theoretical Physics is thanked for his help with inverse problems.

A special acknowledgement goes to all members of the Magnetic Resonance Research Centre, University of Cambridge. Because of the nature of my projects, I have had research interactions with most of them. In particular, I would like to thank Dr Mick Mantle for his expertise in the magnetic resonance field.

Dr Andy York of Johnson Matthey plc is thanked for providing a monolith and Lee Pratt of the Department of Chemical Engineering and Biotechnology is thanked for the help in building the bubble injection rig.

For financial support, I am grateful to the Gates Cambridge Scholarship programme. The Gates Cambridge Scholarship programme was also the main financial contributor, alongside Trinity College and the Department of Chemical Engineering and Biotechnology, for my participation at the MRPM14 conference in Florida, US. Trinity College has been my home for the last 7 years and I thank them for the overseas bursary during my undergraduate years. In particular, I am grateful to my college tutor, Professor Hamish Low, and tutor secretary, Janice Chambers, for putting up with my requests.

My family has been a major moral support throughout the course of the PhD. Thank you for everything.

Declaration

This thesis is the original work of the author, it contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. Neither this work, nor any part thereof, has ever been submitted for any other degree. The research described herein was performed at the Magnetic Resonance Research Centre, in the Department of Chemical Engineering and Biotechnology, University of Cambridge, between October 2015 and September 2018. This thesis contains not more than 65,000 words.

Andi Reci

Abstract

In this thesis, signal sampling and processing techniques are developed for magnetic resonance applications, to improve the estimation of magnetic resonance parameters and to reduce experimental acquisition times.

Two processing techniques are developed for Nuclear Magnetic Resonance (NMR) relaxation and diffusion experiments. L_1 regularization is recommended for extracting parameter distributions which are *a priori* known to be composed of sparse features. L_1 regularization is shown to be stable at signal-to-noise ratios < 20 and capable of resolving relaxation time constants and diffusion coefficients which differ by as little as 10%, such as in relaxation and diffusion studies of hexane/dodecane in porous media. Modified Total Generalized Variation (MTGV) regularization is recommended for extracting parameter distributions for which there is no prior knowledge of whether they are composed of sparse or smooth features. MTGV regularization is shown to perform better than conventional processing techniques and L_1 regularization over a range of simulated distributions.

A method for optimising sampling patterns for relaxation and diffusion experiments, based on the Cramér-Rao Lower Bound theory, is presented. The method is validated against pulsed field gradient NMR diffusion data of two experimental systems. In the first experimental system, the sampling pattern is optimised for the most accurate estimation of the lognormal distribution parameters of an emulsion droplet size distribution of toluene in water. In the second experimental system, the sampling pattern is optimised for the most accurate estimation of the bi-exponential model parameters of a binary mixture of methane/ethane adsorbed in a zeolite. The proposed method predicts an uncertainty in estimating the model parameters which is < 10% different from the uncertainty estimated from the experimental data sampled using the same sampling pattern.

Signal sampling and processing techniques are subsequently combined to reduce experimental acquisition times, which opens opportunities for studying unsteady systems over a long acquisition time and investigating fast-changing phenomena. A 32-fold decrease in the experimental acquisition time is achieved in extracting 3D spatially resolved spin-spin relaxation maps. This is expected to be useful in investigating porous media systems. Three-component velocity maps on a 2D image, acquired every 4 ms, are used to capture, for the first time, the hydrodynamics of a bubble burst event. The experimental data are used to validate the predictions of numerical works.

Published work

First author

A. Reci, A. J. Sederman and L. F. Gladden, Obtaining sparse distributions in 2D inverse problems, J. Magn. Reson. 281 (2017), 188-198.

A. Reci, A. J. Sederman and L. F. Gladden, Retaining both discrete and smooth features in 1D and 2D NMR relaxation and diffusion experiments, J. Magn. Reson. 284 (2017), 39-47.

A. Reci, A. J. Sederman and L. F. Gladden, Optimising magnetic resonance sampling patterns for parametric characterisation, J. Magn. Reson. 294 (2018), 35-43.

A. Reci, A. J. Sederman and L. F. Gladden, Experimental evidence of velocity profile inversion in developing laminar flow using magnetic resonance velocimetry, J. Fluid. Mech. 851 (2018), 545-557.

A. Reci, D. W. de Kort, A. J. Sederman and L. F. Gladden, Accelerating the estimation of 3D spatially resolved T_2 distributions, J. Magn. Reson. 296 (2018), 93-102.

A. Reci, M. I. Ainte, A. J. Sederman, M. D. Mantle and L. F. Gladden, Optimising sampling patterns for bi-exponentially decaying signals, Magn. Reson. Imaging 56 (2009), 14-18.

Co-author

D. W. de Kort, A. Reci, N. P. Ramskill, M. Appel, H. de Jong, M. D. Mantle, A. J. Sederman and L. F. Gladden, Acquisition of spatially-resolved displacement propagators using compressed sensing APGSTE-RARE MRI, J. Magn. Reson. 295 (2018), 45-56.

V. Corona, M. Benning, M. J. Erhardt, L. F. Gladden, R. Mair, A. Reci, A. J. Sederman, S. Reichelt and C.-B. Schöenlieb, Enhancing joint reconstruction and segmentation with non-convex Bregman iteration (2018) arXiv:1807.01660.

Table of contents

Chapter 1

Introduction

In science and engineering applications, experimental observations need to be related to physical parameters through a mathematical model; an illustration of the relationship is given in Fig. 1.1. If the physical parameters and the mathematical model are known, computing the expected experimental observations is known as a forward problem. If the physical parameters and the experimental observations are known, finding the underlying mathematical model is known as a model identification problem. If the mathematical model and the experimental observations are known, estimating the physical parameters is known as an inverse problem. Inverse problems occur in areas as diverse as geophysics [1], astronomy [2], combinatorial chemistry [3], systems biology [4], various branches of physics [5] and tomographic techniques [6]. In chemical engineering applications, inverse problems occur in chemical reaction kinetics [7], heat transfer [8], mass transfer [9] and adsorption [10] among other areas.

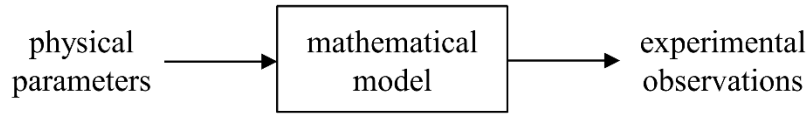

Fig 1.1. Schematic of the relationship between the experimental observations, the mathematical model and the physical parameters in most science and engineering problems. The unknown factor out of the three defines the type of problem: forward, model identification or inverse.

Inverse problems are challenging for several reasons [11]. For some inverse problems, no value of the physical parameters exactly fit the experimental observations; this is most commonly caused by measurement noise. The main design decision for this type of inverse problem is the choice of the appropriate regression technique. For some other inverse problems, there are infinite values of the physical parameters that exactly fit the experimental observations; this is common in applications where many parameters are estimated from very few experimental observations. The main design decision for this type of inverse problems is the choice of the prior physical information needed to discriminate

between the different solutions. For some other inverse problems, the estimation of the physical parameters may be unstable to small changes in the experimental observations; this is commonly caused by the mathematical model that relates the physical parameters to the experimental observations or the sampling scheme used for the acquired observations. The main design decision for this type of inverse problem is the choice of the appropriate stabilization technique. Inverse problems with an infinite number of solutions or unstable solutions are often classified as ill-conditioned inverse problems.

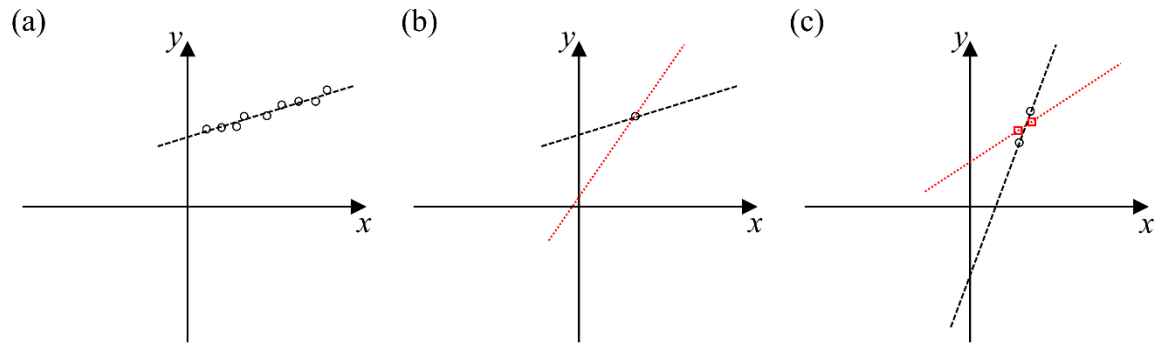

Fig. 1.2. Illustration of different cases of a straight line fit between variables *x* and *y*. (a) A case when no exact fit line exists. (b) A case when an infinite number of exact fit lines exist. (c) A case when a single exact fit line exists, but the fit is unstable; '□' marker observations correspond to slightly different observations to '○' marker observations.

The challenges discussed are now illustrated with respect to an example. One of the simplest and most common inverse problems is the estimation of the gradient and the intercept of a straight line fit to a series of experimental observations. When a lot of experimental observations exist (Fig. 1.2(a)), there is no gradient and intercept that fits exactly all observations. Therefore, a decision must be made on whether to minimize the sum of the absolute residuals, squared residuals or some other objective function. When only a single experimental observation exists (Fig. 1.2(b)), there are an infinite number of gradients and intercepts that fit exactly this observation. The range of solutions can be narrowed down by having some prior physical insight of the problem; for example, there might be some physical reason to believe that the line should pass through the origin. When exactly two data points are acquired (Fig. 1.2(c)), there is a single pair of gradient and intercept which fit the observations exactly. However, the gradient and the intercept can vary widely for small changes in the acquired observations. A method to stabilize this problem could be, for example, to sample the two points further away from each other.

This work focusses on ill-conditioned inverse problems in Nuclear Magnetic Resonance (NMR) and Magnetic Resonance Imaging (MRI). Although the main application of NMR and MRI remains in the medical field, this work focuses on the application of these techniques in engineering areas, particularly in chemical process and reaction engineering, where they have found increasing use; the reader is referred to the reviews of Gladden *et al*. [12-16] for a comprehensive literature survey.

1.1. Ill-conditioned inverse problems in NMR

In NMR applications, ill-conditioned inverse problems occur in the estimation of spin-lattice relaxation time constant, T_1, distributions, spin-spin relaxation time constant, T_2, distributions and diffusion coefficient, D, distributions. T_1 and T_2 are two fundamental chemically specific NMR parameters that relate to the rate of transfer of energy between magnetic spins and the lattice and between neighbouring magnetic spins, respectively. The diffusion coefficient is a molecular property that describes the variance in the displacement of a molecule per unit time during Brownian motion. T_1 and T_2 distributions are routinely used in investigating the permeability, pore size distribution and pore structure of rocks [17-24] and elucidating the composition of crude oils [25-27] in the oil and gas industry; extracting soil water content in soil engineering [28-30]; and the study of cement microporous structure in cement science [31, 32]. D distributions are widely used in extracting size distributions of polymer chain lengths, reverse micelle systems, emulsion droplet systems, biomass components and hydrocarbon chain lengths [33-39].

T_1, T_2 and D distributions can also be obtained jointly through 2D NMR correlation experiments, such as T_1-T_2 [40-42] and D-T_2 [43-45] experiments. 2D joint distributions give better chemical resolution and convey information more compactly, e. g. the T_1/T_2 ratio obtained from a T_1-T_2 map is commonly used as a measure of adsorbate-adsorbent strength in catalysis. Other types of 2D NMR correlation experiments, such as T_2-T_2 [46-48] and D-D [49-51] experiments are commonly used in studying exchange processes between different environments.

The extraction of the 1D and 2D distributions of T_1, T_2 and D is an ill-conditioned inverse problem mainly because the NMR signal is related to the distributions through a Laplace transform; the numerical inversion of the Laplace transform is unstable. The instability

can be improved by choosing appropriately the sampling scheme and the prior knowledge input to the reconstruction process.

1.2. Ill-conditioned inverse problems in MRI

The magnitude of an MRI image details the spatial distribution of, most commonly, the ^{1}H nuclear density. The acquired signal is related to the MRI image through a Fourier transform. The extraction of the MRI image from a signal sampled at an appropriate frequency is not an ill-conditioned inverse problem; an exact, robust and fast numerical inversion of the Fourier transform exists [52]. However, to make the acquisition time shorter, under sampling of the signal is often needed. The reconstruction of the MRI image from an under sampled signal is an ill-conditioned inverse problem because there are infinite solutions to the problem. Reconstructions of the magnitude of the MRI image from under sampled signals have been used in studying absorption processes within cereal-based wafer materials [53], core flooding and displacement processes in rocks [54-56], rising bubbles in fluidised beds [57] and extracting chemical composition maps [58].

The phase of an MRI image can be encoded with the velocity of the fluid being studied, a technique known as MR velocimetry [59-62]. The reconstruction of the phase of the MRI image from an under sampled signal is also an ill-conditioned inverse problem. In science and engineering applications, reconstructions of the phase of the MRI image from under sampled signals have been used in studying multiphase flow [63-66] and flow in microfluidics systems [67, 68].

The stability of the reconstruction of the magnitude or the phase of the MRI images from under sampled signals can be improved by the appropriate choice of the prior knowledge input to the reconstruction process.

1.3. Ill-conditioned inverse problems in joint NMR-MRI applications

An area of research which is attracting increasing interest, and which leads to ill-conditioned inverse problems, is obtaining spatially resolved distributions of NMR

parameters. These experiments encode in each pixel of the MRI image the local NMR parameter distribution. For example, spatially resolved T_1 and T_2 distributions are used in determining the local pore size distribution, permeability and displacement processes in rocks [69-72]. Spatially resolved D distributions have found application in studying the local structure of polymers and local emulsion droplet size distributions [73-75]. Spatially resolved propagators, which measure local displacement distributions, have been used in monitoring dissolution processes in rocks [76]. The spatial resolution of 2D joint NMR parameter distributions is less common because of the prohibitively long acquisition time. However, spatially resolved T_1-T_2 and D-T_2 distributions of porous media have been reported [77-79].

The spatially resolved distribution of NMR parameters is typically related to the acquired signal through a combination of a Laplace and a Fourier transform. The presence of the Laplace transform and the fact that the acquired signal is typically highly under sampled in the frequency domain (to reduce the acquisition time) makes the extraction of the spatially resolved distribution of NMR parameters an ill-conditioned inverse problem.

1.4. Scope of work

The objective of this work is the development of experimental acquisition sampling patterns and reconstruction techniques for particular ill-conditioned inverse problems in NMR and MRI, to make the estimation of parameter distributions more accurate or the experimental acquisition time shorter. The thesis is structured as follows.

Chapter 2 introduces the basic NMR and MRI concepts used throughout the work, while Chapter 3 outlines the mathematical methods on which the developed techniques in this work are based.

Novel reconstruction techniques for the extraction of relaxation time constant and diffusion coefficient distributions, developed in this work, are presented in Chapters 4 and 5. Chapter 4 describes a method for reconstructing distributions which are *a priori* known to be sparse, while Chapter 5 describes a reconstruction method for retaining both discrete and smooth features in the distributions.

An approach to designing optimal sampling patterns for NMR relaxation and diffusion experiments is described in Chapters 6 and 7. Chapter 6 illustrates this approach with

respect to the extraction of the lognormal distribution parameters of an emulsion droplet size distribution, while Chapter 7 illustrates the approach with respect to the resolution of the components of a binary gaseous mixture.

Chapter 8 illustrates how the appropriate combination of an acquisition sampling pattern and a reconstruction technique can lead to quantitative 3D spatially resolved T_2 distributions from highly under sampled signals.

In Chapter 9, highly under sampled MR velocimetry signals and a recently proposed reconstruction technique are used to report the first experimental measurement of the hydrodynamics of a bubble burst event.

Chapter 10 is somewhat unrelated to the main theme of this work. MR velocimetry is used to shed light on a discrepancy between the predictions of numerical and analytical methods about the development of the velocity profile at the entrance of a pipe during laminar flow.

The conclusions and the proposed future work are presented in Chapter 11.

References

[1] A. G. Ramm, Some inverse scattering problems of geophysics, Inverse Problems 1 (1985), 133-172.

[2] I. J. D. Craig and J. C. Brown, Inverse Problems in Astronomy, CRC Press, 1986.

[3] X. Li, Z. Li and L. Wang, The inverse problems for some topological indices in combinatorial chemistry, J. Comput. Biol. 20 (2003), 47-55.

[4] H. W. Engl, C. Flamm, P. Kügler, J. Lu, S. Müller and P. Schuster, Inverse problems in systems biology, Inverse Problems 25 (2009), 123014.

[5] R. G. Newton, Inverse problems in physics, SIAM Rev. 12 (1970), 346-356.

[6] S. R. Arridge and J. C. Schotland, Optical tomography: forward and inverse problems, Inverse Problems 25 (2009), 123010.

[7] K. H. Ebert, P. Deuflhard and W. Jäger, Modelling of Chemical Reaction Systems, Springer, 1981.

[8] G. Stolz, Numerical solutions to an inverse problem of heat conduction for simple shapes, J. Heat Transfer 82 (1960), 20-25.

[9] L. B. Dantas, H. R. B. Orlande and R. M. Cotta, An inverse problem of parameter estimation for heat and mass transfer in capillary porous media, Int. J. Heat Mass Transfer 46 (2003), 1587-1598.

[10] J. D. Wilson, Statistical approach to the solution of first-kind integral equations arising in the study of materials and their properties, J. Mater. Sci. 27 (1982), 3911-3924.

[11] R. C. Aster, B. Borchers and C. Thurber, Parameter Estimation and Inverse Problems, 2nd ed., Academic Press, 2012.

[12] L. F. Gladden and P. Alexander, Applications of nuclear magnetic resonance imaging in process engineering, Meas. Sci. Technol. 7 (1996), 423-435.

[13] L. F. Gladden, Applications of *in situ* magnetic resonance techniques in chemical reaction engineering, Top. Catal. 8 (1999), 87-95.

[14]L. F. Gladden, Magnetic resonance: Ongoing and future role in chemical engineering research, AIChE J. 49 (2003), 1-9.

[15]M. D. Mantle and A. J. Sederman, Dynamic MRI in chemical process and reaction engineering, Prog. Nucl. Magn. Reson. Spec. 43 (2003), 3-60.

[16]L. F. Gladden, Magnetic resonance in reaction engineering: beyond spectroscopy, Curr. Opin. Chem. Eng. 2 (2013), 331-337.

[17]K. R. Brownstein and C. E. Tarr, Importance of classical diffusion in NMR studies of water in biological cells, Phys. Rev. A19 (1979), 2446-2453.

[18]W. E. Kenyon, P. I. Day, C. Straley and J. F. Willemsen, A three-part study of NMR longitudinal relaxation properties of water-saturated sandstones, SPE Formation Evaluation 3 (1988), 622-636.

[19]R. L. Kleinberg and M. A. Horsfield, Transverse relaxation processes in porous sedimentary rock, J. Magn. Reson. 88 (1990), 9-19.

[20]G. C. Borgia, G. Brighenti, P. Fantazzini, G. D. Fanti and E. Mesini, Specific surface and fluid transport in sandstones through NMR studies, SPE Formation Evaluation 7 (1992), 206-210.

[21]L. L. Latour, R. L. Kleinberg and A. Sezginer, Nuclear Magnetic Resonance properties of rocks at elevated temperature, J. Colloid Interface Sci. 150 (1992), 535-548.

[22]R. L. Kleinberg, C. Straley, W. E. Kenyon, R. Akkurt and S. A. Farooqui, Nuclear Magnetic Resonance of rocks, T_1 vs T_2, SPE Paper 26470 (1993).

[23]R. L. Kleinberg, W. E. Kenyon and P. P. Mitra, Mechanism of NMR relaxation of fluids in rock, J. Magn. Reson. 108 (1994), 206-214.

[24]P. J. Barrie, Characterization of porous media using NMR methods, Annu. Rep. NMR Spectrosc. 41 (2000), 265-278.

[25]R. Freedman and N. Heaton, Fluid characterization using Nuclear Magnetic Resonance logging, Petrophysics 45 (2004), 241-250.

[26]D. E. Freed, Dependence on chain length of NMR relaxation times in mixtures of alkanes, J. Chem. Phys. 126 (2007), 174502.

[27] J. J. Chen, M. Hürlimann, J. Paulsen, D. Freed, S. Mandal and Y.-Q. Song, Dispersion of T_1 and T_2 Nuclear Magnetic Resonance relaxation in crude oils, ChemPhysChem 15 (2014), 2676-2681.

[28] J. Votrubová, M. Šanda, M. Cislerová, M. H. G. Amin and L. D. Hall, The relationships between MR parameters and the content of water in packed samples of two soils, Geoderma, 95 (2000), 267-282.

[29] L. R. Stingaciu, A. Pohlmeier, P. Blümler, L. Weihermüller, D. van Dusschoten, S. Stapf and H. Vereecken, Characterization of unsaturated porous media by high-field and low-field NMR relaxometry, Water Resour. Res. 45 (2009), W08412.

[30] A. Pohlmeier, S. Haber-Pohlmeier and S. Stapf, A Fast Field Cycling Nuclear Magnetic Resonance relaxometry study of natural soils, Vadose Zone J. 8 (2009), 735-742.

[31] A. Plassais, M. P. Pomiès, N. Legueux, P. Boch, J.-P. Korb, D. Petit and F. Barberon, Micropore size analysis by NMR in hydrated cement, Magn. Reson. Imaging 21 (2002), 369-371.

[32] J.-P. Korb, Nuclear magnetic relaxation of liquids in porous media, New Journal of Physics 13 (2011), 035016.

[33] A. Chen, D. Wu and C. S. Johnson, Determination of molecular weight distributions for polymers using diffusion-ordered NMR, J. Am. Chem. Soc. 117 (1995), 7965-7970.

[34] A. Jerschow and N. Müller, Diffusion-separated Nuclear Magnetic Resonance spectroscopy of polymer mixtures, Macromolecules 31 (1998), 6573-6578.

[35] S. J. Law and M. M. Britton, Sizing of reverse micelles in microemulsions using NMR measurements of diffusion, Langmuir 28 (2012), 11699-11706.

[36] K. J. Packer and C. Rees, Pulsed NMR studies of restricted diffusion. I. Droplet size distributions in emulsions, J. Colloid Interface Sci. 40 (1972), 206-218.

[37] P. T. Callaghan, K. W. Jolley and R. S. Humphrey, Diffusion of fat and water in cheese as studied by pulsed field gradient Nuclear Magnetic Resonance, J. Colloid Interface Sci. 93 (1983), 521-529.

[38]M. C. Miller, A. Klyosov, D. Platt and K.H. Mayo, Using pulse field gradient NMR diffusion measurements to define molecular size distributions in glycan preparations, Carbohydr. Res. 344 (2009), 1205-1212.

[39]D. E. Freed, L. Burcaw and Y.-Q. Song, Scaling laws for diffusion coefficients in mixtures of alkanes, Phys. Rev. Lett. 94 (2005), 067602.

[40]M. D. Hürlimann, L. Burcaw and Y.-Q. Song, Quantitative characterization of food products by two-dimensional D-T_2 and T_1-T_2 distribution functions in a static gradient, J. Colloid Interface Sci. 297 (2006), 303-311.

[41]D. Weber, J. Mitchell, J. McGregor and L. F. Gladden, Comparing strengths of surface interactions for reactants and solvents in porous catalysts using two-dimensional NMR relaxation correlations, J. Phys. Chem. C 113 (2009), 6610-6615.

[42]C. D'Agostino, J. Mitchell, M. D. Mantle and L. F. Gladden, Interpretation of NMR relaxation as a tool for characterising the adsorption strength of liquids inside porous materials, Chem. Eur. J. 20 (2014), 13009-13015.

[43]M. D. Hürlimann and L. Venkataramanan, Quantitative measurement of two-dimensional distribution functions of diffusion and relaxation in grossly inhomogeneous fields, J. Magn. Reson. 157 (2002), 31-42.

[44]M. D. Hürlimann, L. Venkataramanan and C. Flaum, The diffusion-spin relaxation time distribution function as an experimental probe to characterize fluid mixtures in porous media, J. Chem. Phys. 117 (2002), 10223.

[45]J.-P. Korb, N. Vorapalawut, B. Nicot and R. G. Bryant, Relation and correlation between NMR relaxation times, diffusion coefficients, and viscosity of heavy crude oils, J. Phys. Chem. C 119 (2015), 24439-24446.

[46]J. Lee, C. Labadie, C. S. Springer Jr. and G. S. Harbison, Two-dimensional inverse Laplace transform NMR: Altered relaxation times allow detection of exchange correlation, J. Am. Chem. Soc. 115 (1993), 7761-7764.

[47]L. Monteilhet, J.-P. Korb, J. Mitchell and P. J. McDonald, Observation of exchange of micropore water in cement pastes by two-dimensional T_2-T_2 nuclear magnetic resonance relaxometry, Phys. Rev. E 74 (2006), 061404.

[48] Y. Song, L. Zielinski and S. Ryu, Two-dimensional NMR of diffusion systems, Phys. Rev. Lett. 100 (2008), 248002.

[49] P. T. Callaghan and I. Furó, Diffusion-diffusion correlation and exchange as a signature for local order and dynamics, J. Chem. Phys. 120 (2004), 4032-4038.

[50] P. L. Hubbard, K. M. McGrath and P. T. Callaghan, A study of anisotropic water self-diffusion and defects in lamellar mesophase, Langmuir 21 (2004), 4340-4346.

[51] Y. Qiao, P. Galvosas, T. Adalsteinsson, M. Schönhoff and P. T. Callaghan, Diffusion exchange NMR spectroscopic study of dextran exchange through polyelectrolyte multilayer capsules, J. Chem. Phys. 122 (2005), 214912.

[52] J. W. Cooley and J. W. Tukey, An algorithm for the machine calculation of complex Fourier series, Math. Comput. 19 (1965), 297-301.

[53] P. Parasoglou, D. Malioutov, A. J. Sederman, J. Rasburn, H. Powell, L. F. Gladden, A. Blake and M. L. Johns, Quantitative single point imaging with compressed sensing, J. Magn. Reson. 201 (2009), 72-80.

[54] D. Xiao and B. J. Balcom, *k-t* Acceleration in pure phase encode MRI to monitor dynamic flooding processes in rock core plugs, J. Magn. Reson. 243 (2014), 114-121.

[55] N. P. Ramskill, I. Bush, A. J. Sederman, M. D. Mantle, M. Benning, B. C. Anger, M. Appel and L. F. Gladden, Fast imaging of laboratory core floods using 3D compressed sensing RARE MRI, J. Magn. Reson. 270 (2016), 187-197.

[56] N. P. Ramskill, A. J. Sederman, M. D. Mantle, M. Appel, H. de Jong and L. F. Gladden, In situ chemically-selective monitoring of multiphase displacement processes in a carbonate rock using 3D Magnetic Resonance Imaging, Transp. Porous Med. 121 (2018), 15-35.

[57] H. T. Fabich, A.J. Sederman and D. J. Holland, Development of ultrafast UTE imaging for granular systems, J. Magn. Reson. 273 (2016), 113-123.

[58] E. von Harbou, H. T. Fabich, M. Benning, A. B. Tayler, A. J. Sederman, L. F. Gladden and D. J. Holland, Quantitative mapping of chemical compositions with MRI using compressed sensing, J. Magn. Reson. 261(2015), 27-37.

[59] P. T. Callaghan, Translational Dynamics and Magnetic Resonance, Oxford University Press, 2011.

[60] E. Fukushima, Nuclear magnetic resonance as a tool to study flow, Annu. Rev. Fluid Mech. 31 (1999), 95-123.

[61] C. J. Elkins and M. T. Alley, Magnetic resonance velocimetry: applications of magnetic resonance imaging in the measurement of fluid motion, Exp. Fluids 43 (2007), 823-858.

[62] L. F. Gladden and A. J. Sederman, Recent advances in flow MRI, J. Magn. Reson. 229 (2013), 2-11.

[63] D. J. Holland, D. M. Malioutov, A. Blake, A. J. Sederman and L. F. Gladden, Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing, J. Magn. Reson. 203 (2010), 236-246.

[64] M. Benning, L. F. Gladden, D. Holland, C.-B. Schönlieb and T. Valkonen, Phase reconstruction from velocity-encoded MRI measurements – A survey of sparsity-promoting variational approaches, J. Magn. Reson. 238 (2014), 26-43.

[65] A. B. Tayler, D. J. Holland, A. J. Sederman and L. F. Gladden, Exploring the origins of turbulence in multiphase flow using compressed sensing MRI, Phys. Rev. Lett. 108 (2012), 264505.

[66] A. B. Tayler, M. Benning, A. J. Sederman, D. J. Holland and L. F. Gladden, Ultrafast magnetic-resonance-imaging velocimetry of liquid-liquid systems: Overcoming chemical-shift artefacts using compressed sensing, Phys. Rev. E 89 (2014), 063009.

[67] J. L. Paulsen, V. S. Bajaj and A. Pines, Compressed sensing of remotely detected MRI velocimetry in microfluidics, J. Magn. Reson. 205 (2010), 196-201.

[68] T. Z. Teisseyre, J. L. Paulsen, V. S. Bajaj, N. W. Halpern-Manners and A. Pines, Compressive sampling with prior information in remotely detected MRI of microfluidic devices, J. Magn. Reson. 216 (2012), 13-20.

[69] J. H. Strange and J. B. W. Webber, Spatially resolved pore size distributions by NMR, Meas. Sci. Technol. 8 (1997), 555-561.

[70] M. Li, D. Xiao, L. Romero-Zerón, and B. J. Balcom, Monitoring oil displacement processes with *k-t* accelerated spin echo SPI, Magn. Reson. Chem. 54 (2016), 197-204.

[71] M. Li, S. Vashaee, L. Romero-Zeró, F. Marica and B. J. Balcom, A magnetic resonance study of low salinity waterflooding for enhanced oil recovery, Energy Fuels 31 (2017). 10802-10811.

[72] I. Shikhov, M. N. d'Eurydice, J.-Y. Arns and C. H. Arns, An experimental and numerical study of relative permeability estimates using spatially resolved T_1-z NMR, Transp. Porous Med. 118 (2017), 225-250.

[73] F. Weigand, D. E. Demco, B. Blümich and H. W. Spiess, Spatially resolved NMR spin diffusion in solid polymers, J. Magn. Reson. A 120 (1996), 190-200.

[74] B. Blümich, P. Blümler, L. Gasper, A. Guthausen, V. Göbbels, S. Laukemper-Ostendorf, K. Unseld and G. Zimmer, Spatially resolved NMR in polymer science, Macromol. Symp. 141 (1999), 83-93.

[75] K. G. Hollingsworth and M. L. Johns, Spatially resolved droplet sizing using inverse Abel transforms, J. Magn. Reson. 176 (2005), 71-78.

[76] A. A. Colbourne, A. J. Sederman, M. D. Mantle and L. F. Gladden, Accelerating flow propagator measurements for the investigation of reactive transport in porous media, J. Magn. Reson. 272 (2016), 68-72.

[77] Y. Zhang and B. Blümich, Spatially resolved D-T_2 correlation NMR of porous media, J. Magn. Reson. 242 (2014), 41-48.

[78] S. Vashaee. B. Newling, B. MacMillan, F. Marica, M. Li and B. J. Balcom, Local diffusion and diffusion-T_2 distribution measurements in porous media, J. Magn. Reson. 278 (2017), 104-112.

[79] S. Vashaee, M. Li, B. Newling, B. MacMillan, F. Marica, H. T. Kwak, J. Gao, A. M. Al-harbi and B. J. Balcom, Local T_1-T_2 distribution measurements in porous media, J. Magn. Reson. 287 (2018), 113-122.

Chapter 2

NMR and MRI theory

The idea that certain atomic nuclei absorb and emit radiation in the radio frequency range of the electromagnetic spectrum when placed in a magnetic field originates in the early work of Purcell [1] and Bloch [2]. The phenomenon, known as Nuclear Magnetic Resonance (NMR), has since found applications as an analytical tool in countless areas. Of the many different NMR techniques, the ability to obtain spatially resolved nuclear information from a sample, known as Magnetic Resonance Imaging (MRI) [3, 4], has been particularly beneficial to mankind, through its use in medical applications. Perhaps less known are the applications of NMR an MRI to science and engineering problems, where the scope of the present work lies.

A comprehensive treatment of NMR and MRI fundamentals and methods can be found in the classic texts of Callaghan [5,6], Levitt [7] and Keeler [8]. This chapter gives only a brief introduction to the techniques used in the present work.

2.1. Zeeman interaction

Sub-atomic particles of the nucleus of an atom (protons, neutrons) have a property known as spin, which is a form of angular momentum, but which does not arise from the rotational motion of these particles. Some nuclei, most notably ^{1}H and ^{13}C, have a net overall spin number, I, which results in a net magnetic moment or dipole and $2I + 1$ energy levels being available to the nucleus. These energy levels are degenerate when no external magnetic field is applied, but their energies split when an external magnetic field is applied, due to the Zeeman interaction. For a ^{1}H nucleus, which is the nucleus used in the present work, $I = ½$ and the energy split between the two energy levels formed when an external magnetic field of magnetic field, $\boldsymbol{B}_0$, is applied is given by:

$$\Delta E = \frac{h}{2\pi}\gamma B_0 \ . \qquad (2.1)$$

In Eq. (2.1), h is the Planck's constant, γ = 4257 Hz G^{-1} is the gyromagnetic ratio of ^1H, and B_0 is the magnetic field strength. The population of these energy levels follows a Boltzmann distribution with temperature at thermal equilibrium; the population of the lower energy level is slightly higher than the population of the higher energy level. Although the magnetic dipole of any nucleus can be in any direction, nuclei in the lower energy level have on average a net magnetic dipole along the direction of $\boldsymbol{B}_0$; nuclei in the higher energy level have on average a net magnetic dipole in the direction opposite to $\boldsymbol{B}_0$. The slight difference in the population between the energy levels leads to a net macroscopic magnetization vector, $\boldsymbol{M}$, along the direction of $\boldsymbol{B}_0$. An illustration of these ideas is shown in Fig. 2.1.

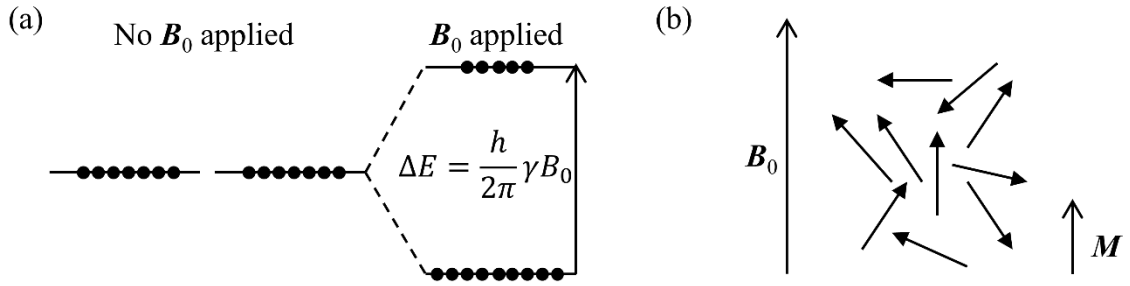

Fig. 2.1. (a) Illustration of the splitting of energy levels available to a nucleus with I = ½ when an external magnetic field, $\boldsymbol{B}_0$, is applied. (b) Illustration of the formation of a net magnetization vector, $\boldsymbol{M}$, when the magnetic dipoles of nuclei (the randomly oriented arrows in the figure) are placed under an external magnetic field, $\boldsymbol{B}_0$.

The external magnetic field, $\boldsymbol{B}_0$, exerts a torque on the individual magnetic dipoles, which results in a precession of the magnetic dipoles about $\boldsymbol{B}_0$. Macroscopically, this leads to a precession of the magnetization vector, $\boldsymbol{M}$, about $\boldsymbol{B}_0$ with a Larmor frequency w_0 which turns out to be in the radio frequency (r.f.) range of the electromagnetic spectrum:

$$w_0 = \gamma B_0 \ . \qquad (2.2)$$

Even for B_0 ~ 10 T, which is at the higher end of the magnetic field strengths used, the difference in populations between the energy levels is so small that the magnetization vector magnitude is orders of magnitude smaller than B_0, which means that the magnetization vector is not detectable in practice. To make the magnetization vector detectable, an oscillating magnetic field, $\boldsymbol{B}_1$, is applied perpendicular to $\boldsymbol{B}_0$, with a magnitude such that $B_1 << B_0$ and a frequency w_0; this is referred to as applying an r.f. pulse. The r.f. pulse perturbs the system from thermal equilibrium through a resonant

interaction with the magnetic dipole of the nuclei; it is this interaction which makes $\boldsymbol{M}$ detectable, as described in the following sections.

2.2. Vector model

The interaction of the overall external magnetic field, $\boldsymbol{B} = \boldsymbol{B}_0 + \boldsymbol{B}_1$, with the magnetization vector is ultimately described by quantum mechanics. The Bloch vector model will instead be used in this work as it offers an intuitive understanding. However, it is noted that the vector model breaks down when spins interact, at which point a quantum mechanical treatment is needed.

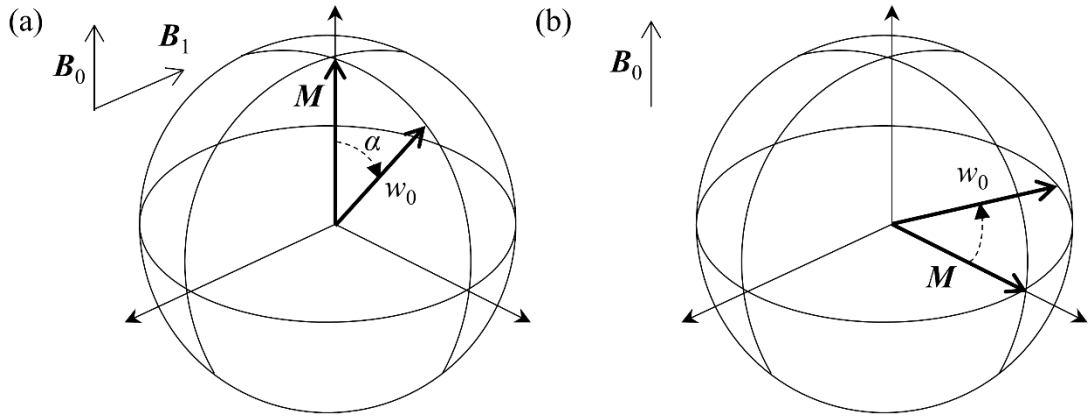

Fig. 2.2. (a) Illustration of the rotation of the magnetization vector, $\boldsymbol{M}$, during the application of the r.f. pulse, in the rotating frame of reference. (b) An illustration of the rotation of $\boldsymbol{M}$ about $\boldsymbol{B}_0$ in the lab frame of reference, after the r.f. pulse is applied with a flip angle of 90°.

Under the influence of $\boldsymbol{B}$, $\boldsymbol{M}$ evolves with time according to the Bloch equation:

$$\frac{d\boldsymbol{M}}{dt} = \gamma \boldsymbol{M} \times \boldsymbol{B}. \qquad (2.3a)$$

This leads to $\boldsymbol{M}$ rotating simultaneously about $\boldsymbol{B}_0$ and $\boldsymbol{B}_1$ with a frequency w_0. This rotation is difficult to visualize, but a helpful aid is to work in a rotating frame of reference (as opposed to the lab frame of reference), which rotates about $\boldsymbol{B}_0$ with a frequency of w_0; in this rotating frame of reference, $\boldsymbol{M}$ rotates about $\boldsymbol{B}_1$ with a frequency w_0. For example, assume that the r.f. pulse applied is given by $\boldsymbol{B}_1 = B_1 \cos(2\pi w_0 t)\, \boldsymbol{j}$, where $\boldsymbol{j}$ is a unit vector in the y-axis and that $\boldsymbol{B}_0$ is applied in the z-direction. In the rotating frame of reference, the solution to Eq. (2.3a) becomes:

$$M_z = M\cos(2\pi w_0 t)\ , \tag{2.3b}$$

$$M_x = M\sin(2\pi w_0 t)\ , \tag{2.3c}$$

$$M_y = 0\ , \tag{2.3d}$$

where M_x, M_y and M_z are the time-varying components of $\boldsymbol{M}$. When the r.f. pulse is applied for a time of t_p, $\boldsymbol{M}$ is rotated through a flip angle $\alpha = 2\pi w_0 t_p$ about $\boldsymbol{B}_1$. This is illustrated in Fig. 2.2(a). The most common flip angles used are 90° and 180°. If a flip angle of 90° is used, Eq. (2.3a) can be subsequently solved using $\boldsymbol{B}_1 = \boldsymbol{0}$ in the lab frame of reference to get:

$$M_z = 0\ , \tag{2.3e}$$

$$M_x = M\sin(2\pi w_0 t + \varphi_1)\ , \tag{2.3f}$$

$$M_y = M\cos(2\pi w_0 t + \varphi_1)\ , \tag{2.3g}$$

for some phase angle φ_1. Therefore, $\boldsymbol{M}$ rotates in the lab frame of reference with a frequency w_0 about $\boldsymbol{B}_0$. The magnetization vector now has components in the perpendicular directions to $\boldsymbol{B}_0$, which makes it detectable. An illustration is given in Fig. 2.2(b).

2.3. Relaxation

When the r.f. pulse is applied, the thermal equilibrium is disturbed. Once the r.f. pulse is removed, the system moves towards thermal equilibrium, which is for the magnetization vector $\boldsymbol{M}$ to be aligned with $\boldsymbol{B}_0$. The return to equilibrium is described by two relaxation processes: spin-lattice and spin-spin relaxation. By convention, it is assumed that $\boldsymbol{B}_0$ is aligned with the *z*-direction.

Spin-lattice relaxation

Spin-lattice relaxation describes how the component of $\boldsymbol{M}$ in the direction of $\boldsymbol{B}_0$ (also known as longitudinal magnetization), M_z, recovers from a magnitude of $M_{z\alpha}$ when the r.f. pulse has just been applied (which depends on the flip angle, α) to the maximum value, M, at equilibrium. This process is driven by the random thermal motion of molecules, which causes local fields to oscillate. When the oscillation frequency equals the Larmor frequency, magnetic poles of individual spins are rotated, with the net effect being a

gradual alignment of the magnetic dipoles with $\boldsymbol{B}_0$. The process is typically described by a first order rate model, such that:

$$M_z(t) = M_{z\alpha} + (M - M_{z\alpha})(1 - e^{-t/T_1}) \,, \tag{2.4}$$

where T_1 is commonly referred to as the longitudinal relaxation time constant.

Spin-spin relaxation

The process by which the components of $\boldsymbol{M}$ in the direction perpendicular to $\boldsymbol{B}_0$ (also known as the transverse magnetization vector), for instance M_x, decrease from a magnitude of $M_{x\alpha}$ when the r.f. pulse has just been applied to zero at equilibrium is governed by two mechanisms. The first mechanism is spin-lattice relaxation, as described above. The second mechanism comes from neighbouring spins precessing at slightly different Larmor frequencies, corresponding to the local field strength, which causes dephasing of the spins (spin-spin relaxation). This mechanism is irreversible. The net effect of these two mechanisms is for M_x to decrease exponentially with a transverse relaxation time constant $T_2 < T_1$:

$$M_x(t) = M_{x\alpha} e^{-\frac{t}{T_2}} \,. \tag{2.5}$$

The homogeneity of the externally applied magnetic field, $\boldsymbol{B}_0$, is rarely ideal. This results in additional dephasing of the spins, leading to an effective relaxation time constant $T_2^* < T_2$. However, this additional dephasing is reversible through the application of an appropriate r.f. pulse sequence, as is discussed in Section 2.7.

2.4. Magnetic field gradients

Commercial magnets have the capability to apply a magnetic field gradient, $\boldsymbol{G}$, in the direction of the external magnetic field $\boldsymbol{B}_0$. The local magnetic field then depends on the position, $\boldsymbol{r}$, from the centre of the magnetic field gradients. As a result, the local Larmor frequency is given by:

$$w(\boldsymbol{r}) = \gamma(B_0 + \boldsymbol{r} \cdot \boldsymbol{G}) \,. \tag{2.6}$$

Magnetic field gradients, together with r.f. pulses, form the basic building blocks of an NMR experiment, which is known as a pulse sequence. Magnetic field gradients are used

in encoding spatial information and are therefore the main tool behind diffusion, imaging and velocimetry experiments.

2.5. Signal detection

Typically, two receiver coils are placed with their axis perpendicular to the direction of $\boldsymbol{B}_0$ and to each other. As $\boldsymbol{M}$ rotates about $\boldsymbol{B}_0$ and simultaneously relaxes back to equilibrium, voltages are induced in these coils due to Faraday's Law, which are translated into a NMR signal. The NMR signal is complex, where the real and imaginary parts hold information from the two receiver coils, separately.

The simplest NMR pulse sequence consists of a 90° r.f. pulse followed by signal acquisition, as illustrated in Fig. 2.3. The noise-free NMR signal acquired, known as the Free Induction Decay (FID), is given by:

$$\hat{y}(t) = e^{-t/T_2} \int_{-\infty}^{+\infty} \rho(w) e^{-2\pi w t i} \mathrm{d}w, \qquad (2.7)$$

where $\rho(w)\mathrm{d}w$ is proportional to the number of nuclei with a Larmor frequency in the range w to $w + \mathrm{d}w$ (spin density). It is assumed that a single T_2 characterizes the whole system. If T_2 is very large (no relaxation effects), $\hat{y}(t)$ and $\rho(w)$ are a Fourier pair. Therefore, $\rho(w)$ can be obtained through a Fourier transform of $\hat{y}(t)$; peaks in $\rho(w)$ reveal the Larmor frequencies present in the sample and the relative population of the nuclei having that Larmor frequency. If T_2 is relatively small, this is manifested in peak broadening of $\rho(w)$, giving each peak a Lorentzian shape. This is illustrated in Fig. 2.3.

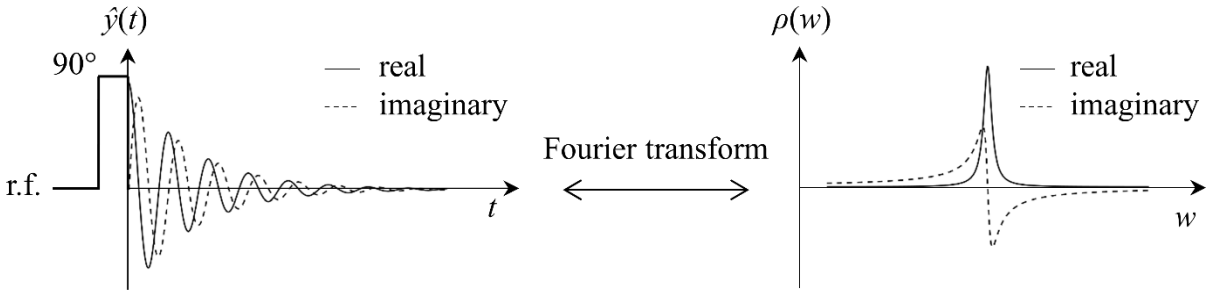

Fig. 2.3. Illustration of the simplest pulse sequence composed of a 90° r.f. pulse, followed by signal acquisition. Through Fourier transformation of the signal, the spectrum of Larmor frequencies is obtained (in this case illustrated by a single Lorentzian peak).

2.6. Chemical shift

Not all the nuclei are in the same chemical environment; The most common example is nuclei in different functional groups. The local magnetic field experienced by nuclei in different chemical environments is different. This causes the nuclei to have different Larmor frequencies and it is manifested as multiple peaks in the Fourier transform of the NMR signal (also referred to as the NMR spectrum). NMR spectra are dependent on B_0. As a result, they are typically normalized to give a chemical shift, δ (ppm):

$$\delta = 10^6 \left(\frac{w - w_{\mathrm{ref}}}{w_{\mathrm{ref}}}\right), \tag{2.8}$$

where w_{ref} is the Larmor frequency of a reference chemical (most commonly tetramethylsilane for ^{1}H). Further peak splitting occurs from the magnetic interaction of nuclei in close proximity or connected through covalent bonds, known as *J*-coupling. The detailed peak splitting leads to complex NMR spectra which are chemical fingerprints. This is the most commonly used NMR method to elucidate chemical structure and composition. The present work is not focused on NMR spectral analysis. However, chemical shift information will be given where necessary.

2.7. Spin echo

The homogeneity of the external magnetic field, $\boldsymbol{B}_0$, is rarely perfect. Corrections can be made through the application of shimming magnetic field gradients, but some inhomogeneity remains. This is particularly marked when materials with different magnetic susceptibilities are present in the sample, such as air and water, or when traces of paramagnetic material exist, such as in porous media applications. Therefore, the NMR signal decays with a transverse relaxation time $T_2^* < T_2$. This can be problematic for several reasons: it leads to a lower signal-to-noise ratio; only short pulse sequences are viable as the signal decays very fast; the T_2^* decay is more difficult to interpret than the T_2 decay, which is more meaningful because of its chemical and environmental specificity.

The spin echo [9] is one of the basic building blocks of NMR pulse sequences as it retrieves the T_2 decay from the T_2^* decay by refocusing the dephasing that occurs due to $\boldsymbol{B}_0$ inhomogeneity. The spin echo pulse sequence is composed of a 90° r.f. pulse followed by

a 180° r.f. pulse after an echo time τ. During the echo time, the spins precessing at different Larmor frequencies due to $\boldsymbol{B}_0$ inhomogeneity diphase, leading to a drop in the NMR signal. After the application of the 180° r.f. pulse, the spins' direction is inverted and they start rephasing. After an equivalent time of τ, the spins are completely in phase (before starting to dephase again), and this is known as a spin echo. The magnitude of the signal at the spin echo is then related to the magnitude of the signal just after the application of the 90° r.f. pulse through a T_2 decay. An illustration of the formation of the spin echo is given in Fig. 2.4.

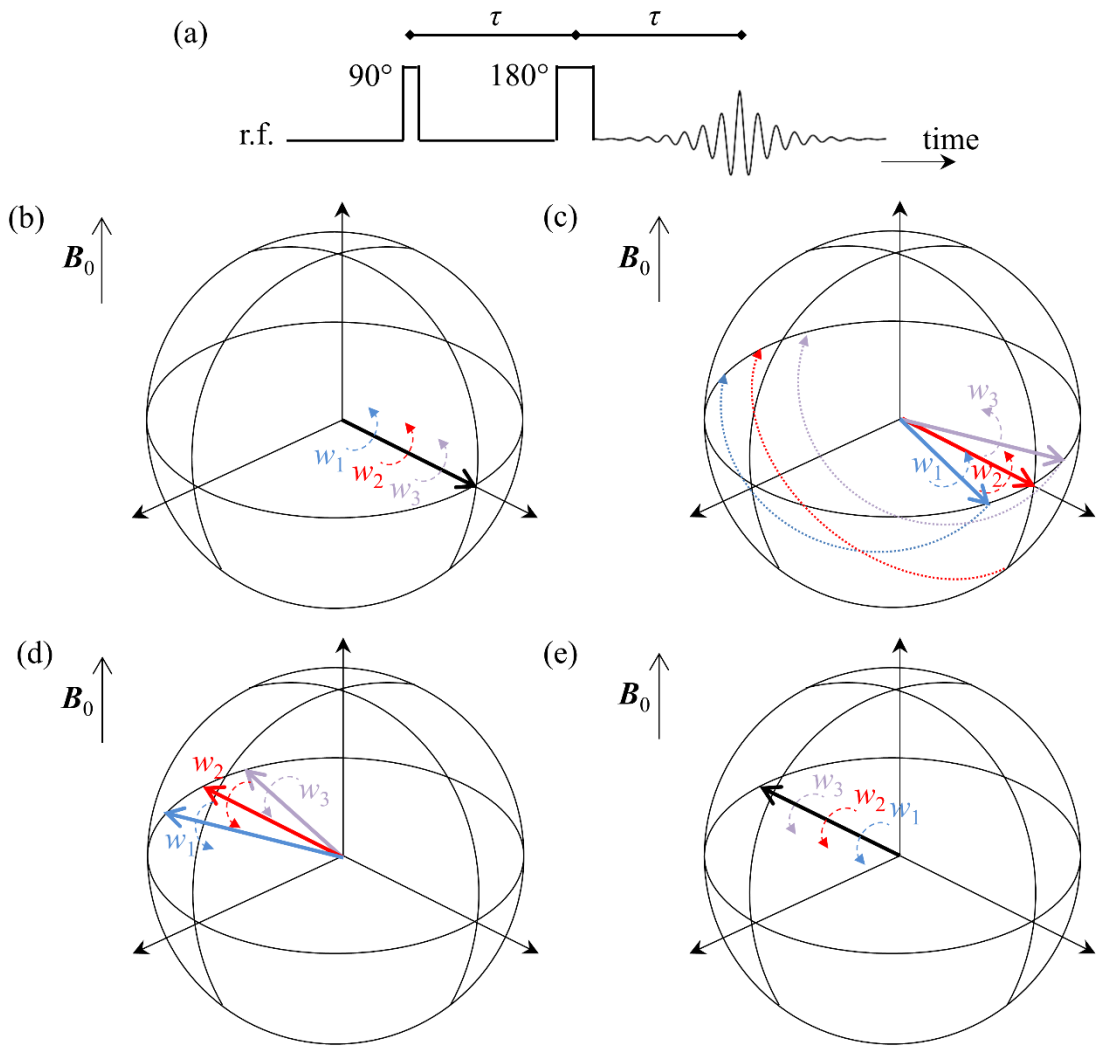

Fig. 2.4. (a) The spin echo pulse sequence. (b) Just after the application of the 90° r.f. pulse, all spins are in phase and they start precessing at slightly different Larmor frequencies. Three representative spins with $w_1 \neq w_2 \neq w_3 \neq w_1$ are followed. (c) After an echo time τ, spins have dephased. When the 180° r.f. pulse is applied, spins are inverted, a process illustrated by the dotted lines. (d) Position of spins just after the application of the 180° r.f. pulse. (e) Spins refocus after an equivalent time of τ; a spin echo is formed. Spins start dephasing again after this point.

2.8. Stimulated echo

In certain applications, such as when a chemical species interacts strongly with a surface, $T_2 << T_1$, which causes a fast decaying of the signal. A common technique to overcome this limitation is to store the transverse magnetization in the $\boldsymbol{B}_0$ direction for a time t_{store} by the application of a 90° r.f. pulse at a time τ after an initial 90° r.f. pulse, as shown in the pulse sequence in Fig. 2.5. During t_{store}, the stored magnetization decays with a time constant T_1, which is a slower decay. In order to recall the magnetization back into the transverse direction, another 90° r.f. pulse is applied and an echo is formed after a time τ. This is known as a stimulated echo sequence and is used in some of the diffusion experiments in the present work. It is noted that other echoes may also form because the second 90° r.f. pulse is, in practice, imperfect and some remnant transverse magnetization persists. These echoes are undesirable and are formed whenever a train of imperfect r.f. pulses is applied [10]. The undesirable echoes are suppressed using spoiler magnetic field gradients, as shown in Fig. 2.5, or phase cycling techniques. These are very important areas in NMR experimental design, but they will not be the focus of the present work.

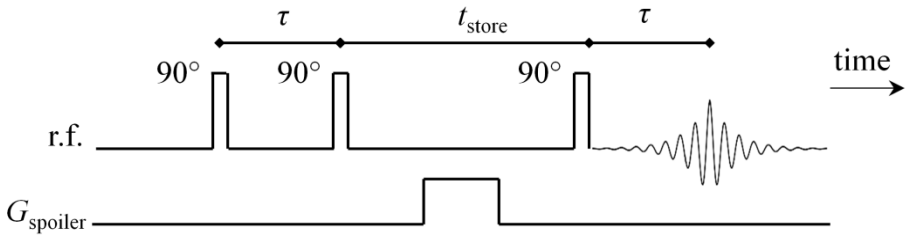

Fig. 2.5. Schematic of the stimulated echo pulse sequence.

2.9. T_1 estimation

The most commonly used pulse sequence to measure the relaxation time constant, T_1, is the Inversion Recovery (IR) pulse sequence, shown in Fig. 2.6. An initial 180° r.f. pulse flips the magnetization vector in the opposite direction. The magnetization vector has no component in the transverse direction and, therefore, decays with a relaxation time constant, T_1. To make the magnetization vector measurable, the magnetization vector is rotated into the transverse plane after a time t_1 using a 90° r.f. pulse, after which the NMR signal is acquired. The process is repeated for several values of t_1, which gives a number

of NMR signals. These are individually Fourier transformed and the area under the NMR spectrum, which will confusingly be known as the NMR signal, is plotted as function of t_1.

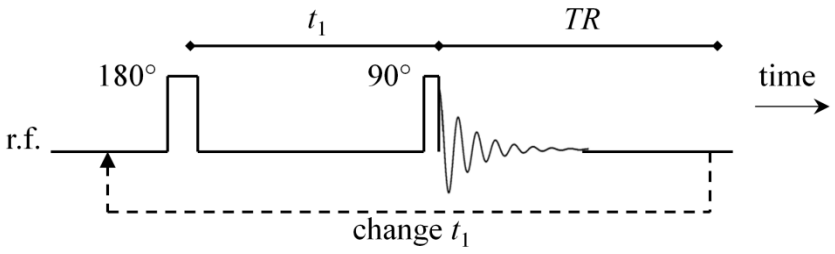

Fig. 2.6. Schematic of the Inversion Recovery (IR) pulse sequence. *TR* refers to the recovery time which is needed for the longitudinal magnetization vector to return to its equilibrium condition before starting the next loop of the experiment.

For a system characterised by a single T_1, the noise-free NMR signal, is given by:

$$\hat{y}(t_{1,i}) = A\left(1 - 2e^{-t_{1,i}/T_1}\right), \tag{2.9}$$

where $A = \hat{y}(t_1 \to \infty)$ is a scaling factor. From the expression in Eq. (2.9), T_1 can be retrieved through a simple fitting from any computational package. If the system is characterised by a distribution of T_1, $u(T_1)$, where $u(T_1)\mathrm{d}T_1$ is the contribution to the signal from species with a relaxation time constant in the range T_1 to $T_1+\mathrm{d}T_1$, the noise-free NMR signal is given by:

$$\hat{y}(t_{1,i}) = \int_0^\infty u(T_1)\left(1 - 2e^{-t_{1,i}/T_1}\right)\mathrm{d}T_1 \ . \tag{2.10}$$

The distribution $u(T_1)$ needs to be obtained numerically. As a result, Eq. (2.10) is discretized by seeking a solution only for several relaxation times spaced in some pattern between a minimum and maximum value, where the relaxation times are believed to lie within:

$$\hat{y}(t_{1,i}) = \sum_j u(T_{1,j})\left(1 - 2e^{-t_{1,i}/T_{1,j}}\right). \tag{2.11}$$

Eq. (2.11) can be written more compactly in matrix form as:

$$\hat{\boldsymbol{y}} = \boldsymbol{K}\,\boldsymbol{u}\ , \tag{2.12}$$

where $\hat{\boldsymbol{y}}$ is the signal column vector such that $\hat{y}_i = \hat{y}(t_{1,i})$, $\boldsymbol{K}$ is a kernel matrix such that $K_{ij} = 1 - 2\exp(-t_{1,i} / T_{1,j})$, and $\boldsymbol{u}$ is the distribution column vector such that $u_j = u(T_{1,j})$. So

far, it has been assumed that the NMR signal is not corrupted by noise. However, noise is unavoidable and, therefore, the noisy acquired signal is given by:

$$\boldsymbol{y} = \boldsymbol{K}\,\boldsymbol{u} + \boldsymbol{e}\ , \tag{2.13}$$

where $\boldsymbol{e}$ is an unknown noise column vector, such that e_i is the noise related to y_i. $\boldsymbol{u}$ needs to be estimated from a knowledge of $\boldsymbol{y}$ and $\boldsymbol{K}$, which is much of the scope of the present work.

2.10. T_2 estimation

The most commonly used pulse sequence in measuring the transverse relaxation time constant, T_2, is the Carr-Purcell Meiboom-Gill (CPMG) [11, 12] pulse sequence, illustrated in Fig. 2.7(a). The CPMG pulse sequence is composed of a train of equidistant 180° r.f. pulses following an initial 90° r.f. pulse. After each 180° r.f. pulse, a spin echo forms, but only the middle point of the echo is acquired. For a system characterised by a single T_2, the noise-free acquired signal follows an exponential decay:

$$\hat{y}\left(t_{2,i}\right) = Ae^{-t_{2,i}/T_2}\ , \tag{2.14}$$

where $t_{2,i} = 2i\tau$, i is the echo number; $A = \hat{y}(t_2 \rightarrow 0)$ is a scaling factor, and τ is the echo time. For a system characterised by a distribution of T_2, $u(T_2)$, a similar argument to the discussion in Section 2.9 gives the noise-free NMR signal in a discrete setting as:

$$\hat{y}\left(t_{2,i}\right) = \sum_j u\left(T_{2,j}\right)e^{-t_{2,i}/T_{2,j}}\ . \tag{2.15}$$

In the presence of noise, the acquired signal is given in a compact matrix form as:

$$\boldsymbol{y} = \boldsymbol{K}\,\boldsymbol{u} + \boldsymbol{e}\ , \tag{2.16}$$

where $\boldsymbol{y}$ is the signal column vector such that $y_i = y(t_{2,i})$, $\boldsymbol{K}$ is a kernel matrix such that $K_{ij} = \exp(-t_{2,i} / T_{2,j})$, $\boldsymbol{u}$ is the distribution column vector such that $u_j = u(T_{2,j})$ and $\boldsymbol{e}$ is an unknown noise column vector, such that e_i is the noise related to y_i. It is noted that Eq. (2.16) is the same as Eq. (2.13), which means that the same numerical methods can be used to obtain $\boldsymbol{u}$ from a knowledge of $\boldsymbol{y}$ and $\boldsymbol{K}$ for both T_1 and T_2 experiments.

In some of the experiments in the present work, a variation of the CPMG sequence, known as Periodic Refocusing of J Evolution by Coherence Transfer (PROJECT) [13] is used. This consists in the addition of a 90° r.f. pulse between every other 180° r.f. pulse

(illustrated in Fig. 2.7(b)). The PROJECT pulse sequence offers an improvement over the CPMG pulse sequence in cases where the frequency difference between spins coupled through *J*-coupling is of the order of $1/\tau$ which would lead to signal modulation if the CPMG sequence was used.

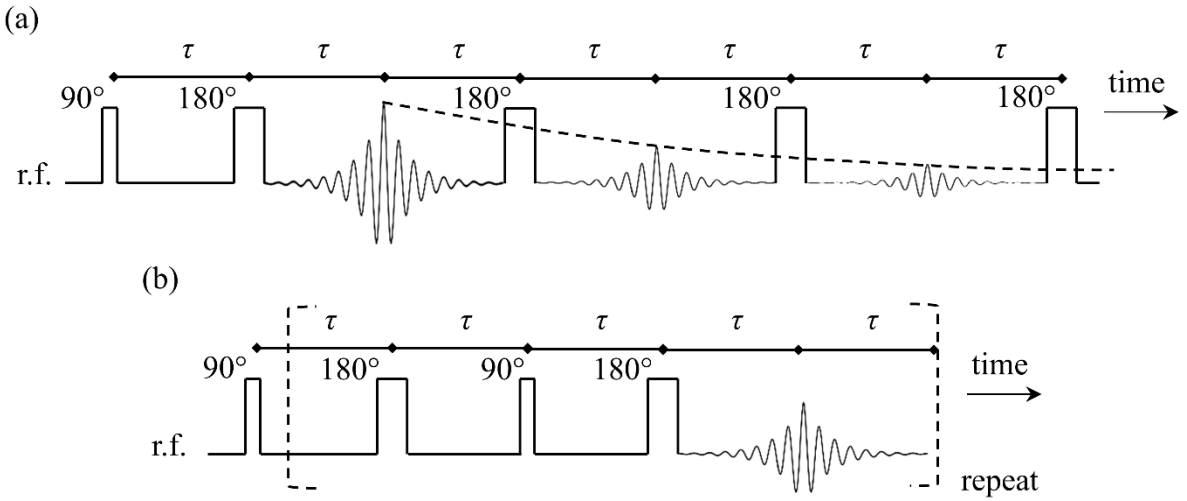

Fig. 2.7. (a) Schematic of the CPMG pulse sequence. Only the middle point of each echo is acquired. (b) Schematic of the PROJECT pulse sequence. Similarly, only the middle point of each echo is acquired.

2.11. Diffusion coefficient estimation

Apart from the chemical shift and the relaxation time constants, another property which can be measured using NMR is the self-diffusion coefficient, *D*, the foundations to which go back to the work of Carr and Purcell [11]. The most common pulse sequence used to estimate *D* is based on a spin echo and the application of two pulses of magnetic field gradients, known as pulsed field gradients (PFG), either side of the 180° r.f. pulse, as shown in Fig. 2.8(a). The first gradient pulse imparts on the spins located at different spatial locations a different Larmor frequency. Spins precess at this imposed frequency and dephase over the time in which the pulsed gradient is applied, δ. The spin direction is reversed after the application of the 180° r.f. pulse. At the end of the application of the second pulsed field gradient, spins would completely rephase if the spins did not change location during the time, Δ, between the application of the two pulsed field gradients, as illustrated in Fig. 2.8(b). The spin echo formed would then be similar to the spin echo that would form without the PFG application. However, spins are in continuous random thermal motion. As they follow a Brownian motion in the time Δ, the spins experience a

different Larmor frequency when the second pulsed field gradient is applied. As a result, the spins do not completely rephase at the spin echo, leading to a loss in the NMR signal, as illustrated in Fig. 2.8(c). The signal loss is related to the root mean square distance travelled by the spins, which is in turn related to the diffusion coefficient, D. The experiment is repeated for several gradient magnitude values, g_i.

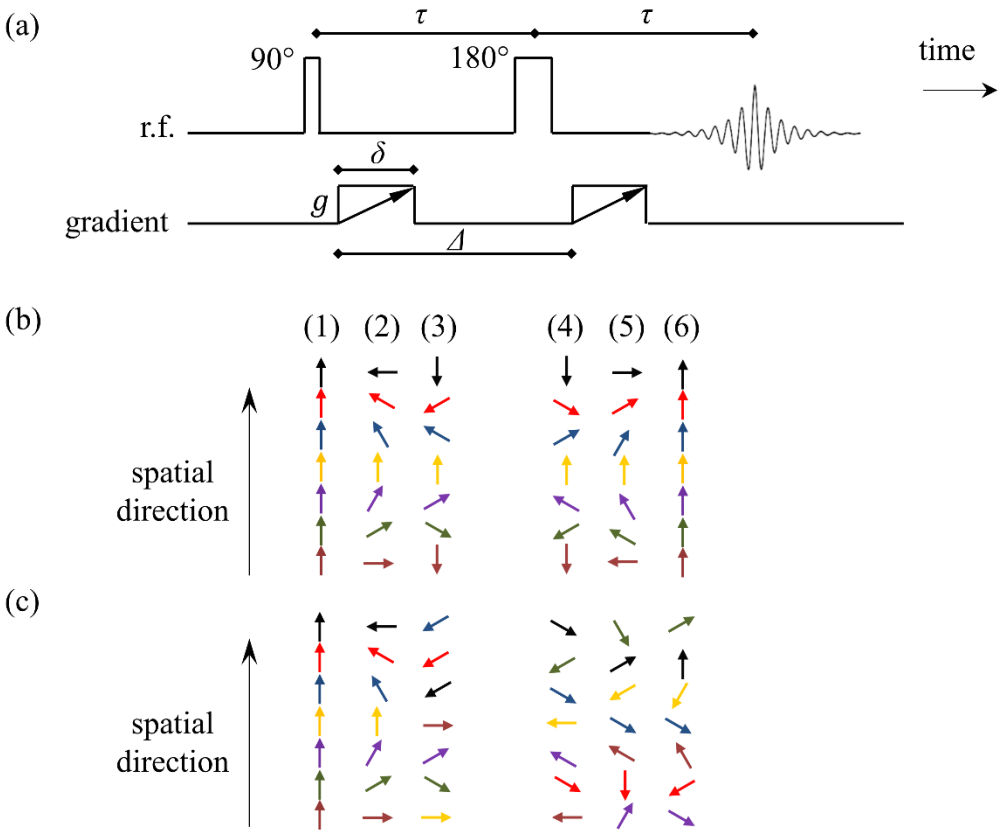

Fig. 2.8. (a) Schematic of the PFG NMR diffusion pulse sequence based on a spin echo. The magnitude of the pulsed field gradients, g, is changed during repeats of the pulse sequence. (b) Illustration of the behaviour of spins in different spatial locations during different stages of the pulse sequence, under the assumption of no spin motion: (1) Just after the application of the 90° r.f. pulse; (2) Some time during the application of the first pulsed gradient; (3) After the first pulsed gradient is switched off; (4) Just after the application of the 180° r.f. pulse; (5) Some time during the application of the second pulsed gradient; (6) After the second pulsed gradient is switched off. (c) Same as (b) but for the case when spins are allowed to diffuse randomly.

For a system characterised by a single D, the noise-free signal attenuation is given by the Stejskal-Tanner equation [14]:

$$\hat{y}(g_i) = Ae^{-\gamma^2 g_i{}^2 \delta^2 D\left(\Delta - \delta/3\right)} \quad . \tag{2.17}$$

For a system characterised by a distribution of D, $u(D)$, a similar argument to the discussion in section 2.9 gives the noise-free NMR signal in a discrete setting as:

$$\hat{y}(g_i) = \sum_j u(D_j) e^{-\gamma^2 g_i^2 \delta^2 D_j (\Delta - \delta/3)} \,. \qquad (2.18)$$

In the presence of noise, the acquired signal is given in a compact matrix form as:

$$\boldsymbol{y} = \boldsymbol{K}\,\boldsymbol{u} + \boldsymbol{e} \,, \qquad (2.19)$$

where $\boldsymbol{y}$ is the signal column vector such that $y_i = y(g_i)$, $\boldsymbol{K}$ is a kernel matrix such that $K_{ij} = \exp(-\gamma^2 g_i^2 \delta^2 D_j (\Delta - \delta / 3))$, $\boldsymbol{u}$ is the distribution column vector such that $u_j = u(D_j)$, and $\boldsymbol{e}$ is an unknown noise column vector, such that e_i is the noise related to y_i. It is noted that Eq. (2.19) is the same as Eqs. (2.13) and (2.16), which means that the same numerical methods can be used to obtain $\boldsymbol{u}$ from a knowledge of $\boldsymbol{y}$ and $\boldsymbol{K}$ for T_1, T_2 and D experiments.

In some of the experiments of the present work, a PFG NMR diffusion pulse sequence based on the stimulated echo [15] is used, instead of the spin echo. Eqs. (2.17-2.19) still hold for this type of experiment.

2.12. Joint parameter estimation

The pulse sequences described so far are used in obtaining the distribution of single parameters, either T_1, T_2 or D. Acquiring joint distributions of any two of these parameters could be beneficial for a number of reasons: a better chemical resolution can be achieved for species which are very similar in terms of one of the parameters, exchange processes between chemical environments can be studied, and correlations between different parameters are presented more compactly. The pulse sequences which acquire joint distribution functions of these parameters are known as 2D NMR correlation experiments. The pulse sequences for 2D NMR correlation experiments are constructed by simply stacking in series the pulse sequences used in estimating the individual parameters. Only an example will be given in this section; other pulse sequences are constructed and analysed similarly. Fig. 2.9 illustrates the pulse sequence for a T_1-T_2 experiment. The IR pulse sequence is combined with the CPMG pulse sequence and the noise-free NMR signal acquired is given by:

$$\hat{Y}(t_{1,i}, t_{2,j}) = \sum_k \sum_l U(T_{1,k}, T_{2,l}) e^{-t_{2,j}/T_{2,l}} (1 - 2e^{-t_{1,i}/T_{1,k}}) \,, \qquad (2.20)$$

where $U(T_1,T_2)$ is the joint distribution function and $t_{2,j} = 2j\tau$. In the presence of noise, the NMR signal can be written compactly in a matrix form as:

$$\boldsymbol{Y} = \boldsymbol{K}_1 \, \boldsymbol{U} \, \boldsymbol{K}_2{}^{\mathrm{T}} + \boldsymbol{E} \, , \tag{2.21}$$

where $\boldsymbol{Y}$ is the signal matrix such that $Y_{ij} = Y(t_{1,i}, t_{2,j})$, $\boldsymbol{K}_1$ is a kernel matrix such that $K_{1,ik} = 1 - 2 \exp(-t_{1,i} / T_{1,k})$, $\boldsymbol{K}_2$ is a kernel matrix such that $K_{2,jl} = \exp(-t_{2,j} / T_{2,l})$, $\boldsymbol{U}$ is the distribution matrix such that $U_{kl} = U(T_{1,k}, T_{2,l})$,and $\boldsymbol{E}$ is an unknown noise matrix, such that E_{ij} is the noise related to Y_{ij}. Eq. (2.21) can be rewritten in terms of vectorised matrices $\boldsymbol{Y}_{\mathrm{v}}$, $\boldsymbol{U}_{\mathrm{v}}$ and $\boldsymbol{E}_{\mathrm{v}}$, formed by stacking the columns of $\boldsymbol{Y}$, $\boldsymbol{U}$ and $\boldsymbol{E}$, respectively, as:

$$\boldsymbol{Y}_{\mathrm{v}} = \boldsymbol{K} \, \boldsymbol{U}_{\mathrm{v}} + \boldsymbol{E}_{\mathrm{v}} \, , \tag{2.22}$$

where $\boldsymbol{K} = \boldsymbol{K}_2 \otimes \boldsymbol{K}_1$ ($\otimes$ stands for the Kronecker product). It is noted that Eq. (2.22) is the same as Eqs. (2.13), (2.16) and (2.19), which means that the same numerical methods can be used to obtain $\boldsymbol{U}$ from a knowledge of $\boldsymbol{Y}$ and $\boldsymbol{K}$ for both single and joint parameter estimation problems.

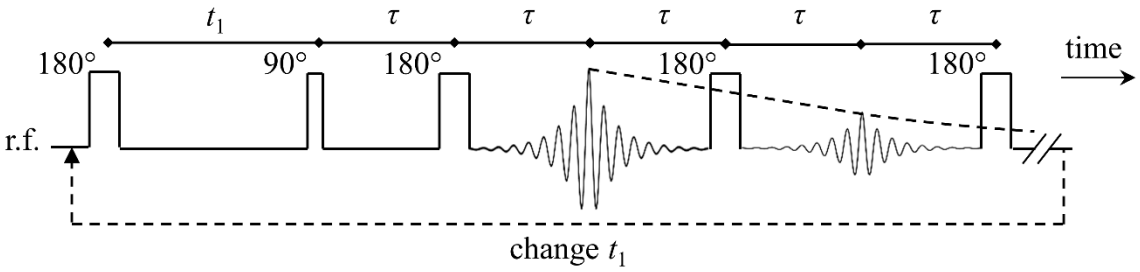

Fig. 2.9. The pulse sequence used to estimate the joint T_1-T_2 distribution. An IR and a CPMG pulse sequence are stacked in series. Only the middle point in each echo is acquired.

2.13. Principles of MRI

The main aim in MRI is to obtain a spatially resolved nuclear spin density, $\rho(\boldsymbol{r})$, referred to simply as the image. This is achieved through the use of time-varying magnetic field gradients, $\boldsymbol{G}(t)$, which cause the local Larmor frequency, $w(\boldsymbol{r})$, to change according to Eq. (2.6). Under the influence of such a gradient, the noise-free complex NMR signal acquired is given by:

$$\hat{y}(t) = \iiint \rho(\boldsymbol{r}) e^{i\phi(\boldsymbol{r},t)} \, \mathrm{d}\boldsymbol{r} \, , \tag{2.23}$$

where $\phi(\boldsymbol{r},t)$ is the spatial and time-varying phase which evolves according to:

$$\phi(\boldsymbol{r},t) = \int_0^t w(\boldsymbol{r},t) dt = \gamma \, \boldsymbol{r} \cdot \int_0^t \boldsymbol{G}(t) \mathrm{d}t \, . \tag{2.24}$$

It is assumed that the signal has been heterodyned, a process by which the reference oscillation frequency of γB_0 is filtered out, and that relaxation effects are negligible. With the definition of an inverse space variable, $\boldsymbol{k}$:

$$\boldsymbol{k} = (2\pi)^{-1}\gamma \int_0^t \boldsymbol{G}(t)\mathrm{d}t \ , \tag{2.25}$$

the noise-free NMR signal acquired is given by:

$$\hat{y}(\boldsymbol{k}) = \iiint \rho(\boldsymbol{r}) e^{i2\pi \boldsymbol{r}\cdot\boldsymbol{k}} \, \mathrm{d}\boldsymbol{r} \ . \tag{2.26}$$

Eq. (2.26) shows that $\hat{y}(\boldsymbol{k})$ (known as *k*-space) and $\rho(\boldsymbol{r})$ are a Fourier pair. Therefore, the image is obtained through Fourier transformation of the *k*-space.

The NMR signal is acquired digitally and therefore the *k*-space is stored as a matrix, with the sizes of each dimension being typically a power of 2; this is done to exploit the fast computation of the image through the Fast Fourier Transform (FFT) method [16]. In the discrete setting and under the presence of noise, the *k*-space matrix $\boldsymbol{Y}$, vectorised as $\boldsymbol{Y}_\mathrm{v}$, is given by:

$$\boldsymbol{Y}_\mathrm{v} = \boldsymbol{\mathcal{F}}\, \boldsymbol{U}_\mathrm{v} + \boldsymbol{E}_\mathrm{v} \ , \tag{2.27}$$

where $\boldsymbol{\mathcal{F}}$ is the discrete Fourier transform operator, $\boldsymbol{U}_\mathrm{v}$ is the vectorised image and $\boldsymbol{E}_\mathrm{v}$ is a vector of noise elements. It is noted that Eq. (2.27) is similar to Eq. (2.13), (2.16), (2.19) and (2.22), which means that NMR and MRI signals can be cast in the same matrix equation format.

2.14. *k*-space manipulation

To reconstruct the MRI image, the *k*-space matrix data needs to be acquired. The way this is done is often referred to as *k*-space traversal, a geometric interpretation of which was first introduced by Ljunggren [17]. Although different *k*-space traversal techniques lead to the same matrix being completed, each has its own advantages and disadvantages in terms of, among other reasons, speed of acquisition, artefacts introduced, signal-to-noise considerations and hardware requirements. For illustration purposes, only 2D imaging is considered in the following sections.

One of the simplest traversal techniques is the spin-warp method [18], illustrated in Fig. 2.10(a). Initially, a constant magnetic field gradient is applied in a spatial direction,

referred to as the phase direction and, by the end of its application, spins in that direction will have different phases. This is interpreted as moving along the k_{phase} direction in *k*-space. A constant magnetic field gradient is then applied in the perpendicular direction, known as the read direction and the signal is acquired during this period. This is interpreted as moving along the k_{read} direction in *k*-space and acquiring only regular points on a Cartesian grid. The process is then repeated for different strengths of the magnetic field in the phase direction until the whole *k*-space is acquired.

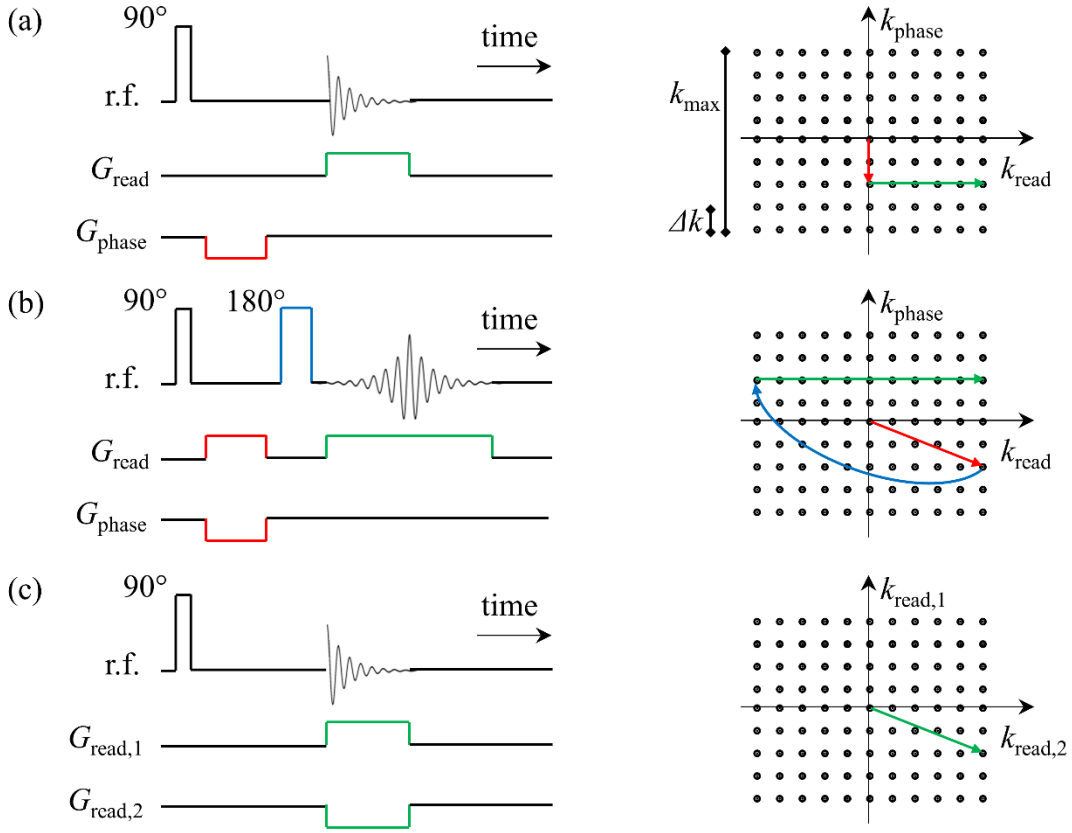

Fig. 2.10. (a) Illustration of a spin-warp imaging pulse sequence (left) and the corresponding *k*-space traversal (right). (b) Illustration of a spin echo imaging pulse sequence (left) and the corresponding *k*-space traversal (right). (c) Illustration of a radial imaging pulse sequence (left) and the corresponding *k*-space traversal (right).

A different traversal technique, which leads to a better signal-to-noise ratio and which leads to less artefacts, uses spin echoes and is illustrated in Fig. 2.10(b). The application of a 180° r.f. pulse is interpreted as moving in the opposite direction with respect to the centre of *k*-space. It is noted that the phase and read gradients can be switched on at the same time, leading to a diagonal movement in *k*-space.

The *k*-space data can also be acquired in a non-Cartesian way. Fig. 2.10(c) shows an example where the NMR data is acquired while *k*-space is traversed radially, by keeping

the magnetic field gradients in both perpendicular directions switched on during acquisition (both directions are read directions). While non-Cartesian sampling generally leads to a more efficient *k*-space coverage, the data needs to be regridded in a regular Cartesian plane before Fourier transformation; this process leads to artefacts being introduced in the MRI image.

If the *k*-space data is acquired in a regular Cartesian way, the geometric dimensions of *k*-space, illustrated in Fig. 2.10(a), are related to the geometric dimensions of the MRI image through the inverse relationships:

$$\Delta k = \frac{1}{FOV}\,, \tag{2.28a}$$

$$k_{\max} = \frac{1}{\Delta x}\,, \tag{2.28b}$$

where *FOV* refers to the field-of-view and Δx is the resolution of the MRI image.

2.15. Slice selection

The r.f. pulses that have been used so far are known as hard pulses; the shape of the pulse is typically rectangular and the duration is short (~10 μs). As a result, the bandwidth of frequencies that are excited, Δw, which is inversely proportional to the pulse duration, is so large such that all spins in the sample are excited and the NMR signal has a contribution from all of them. In 2D imaging, it is useful to selectively excite and acquire NMR signal only from a slice of thickness Δz in the sample; this is done using soft pulses.

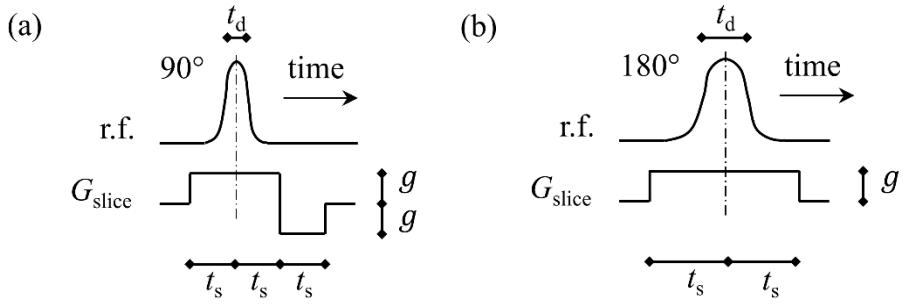

Fig. 2.11. (a) Illustration of soft 90° Gaussian r.f. pulse. The reversal of the slice selective gradient is done to refocus spins within the excited slice which have dephased during the application of the pulse. (b) Illustration of soft 180° Gaussian r.f. pulse. There is no need to reverse the slice selective gradient because the 180° pulse self-refocuses the dephased spins.

A plethora of different soft pulses exist and the choice is primarily made based on the shape of the excited slice required. If the shape of the excited slice is relatively unimportant, Gaussian pulses in the presence of a slice selective magnetic field gradient are commonly used, as they are simple to implement. In this work, Gaussian pulses are used for both soft 90° and 180° r.f. pulses, as illustrated in Fig. 2.11. The slice thickness excited by the soft r.f. pulse is given by:

$$\Delta z = {}^{1}/_{\gamma t_{\mathrm{d}} g}\,, \tag{2.29}$$

where t_{d} (s) represents the Full-Width-Half-Maximum value of the Gaussian pulse.

2.16. Ultrafast imaging techniques

One of the limitations of the imaging techniques described in Fig. 2.10 is the long acquisition time. For each 90° r.f. pulse application, a single line in *k*-space is acquired, after which a recovery time ~5T_1 is allowed before the next r.f. pulse application, to recover the magnetization vector. Ultrafast imaging techniques aim to sample the *k*-space more efficiently, reducing the total acquisition time. Different ultrafast imaging techniques exist [19, 20], but the discussion will be limited to the Rapid Acquisition with Relaxation Enhancement (RARE) method [21] and spiral imaging [22], which are used in this work.

The RARE pulse sequence and the corresponding *k*-space traversal route are illustrated in Fig. 2.12(a). The basic idea is that for each 90° r.f. pulse application, several lines of *k*-space are acquired (RARE factor). The RARE factor is chosen such that the time needed for the acquisition of the *k*-space lines during an application of a 90° r.f. pulse is << T_2. RARE imaging has found most application in medical T_2 contrast imaging where the RARE factor is chosen such that spatial features with different T_2 values can be differentiated. The main disadvantage of RARE imaging is the use of a train of slice-selective 180° r.f. pulses, which, if not perfect, can lead to complex stimulated echo patterns [23, 24].

The spiral imaging pulse sequence and the corresponding *k*-space traversal route are illustrated in Fig. 2.12(b). The whole *k*-space is sampled with a single application of a 90° r.f. excitation pulse, which leads to a very short acquisition time. The centre of *k*-space is acquired at the start of the sequence which gives a high signal-to-noise ratio. One of the disadvantages of spiral imaging is that the data is acquired in a non-Cartesian way.

Therefore, a Non-Uniform Fast Fourier Transform (NUFFT) technique is needed to reconstruct the MRI image. Since no 180° r.f. pulses are used, the signal decay with time is controlled by T_2^*, rather than T_2. As a result, the magnetic field needs to be relatively homogeneous in order for T_2^* to be much larger than the acquisition time. Spiral imaging has found application in Magnetic Resonance velocimetry (Section 2.18) because of its stability to flow artefacts [25] and the relatively low phase accrual related to the points of *k*-space close to the centre [26], where most of the energy of *k*-space is concentrated.

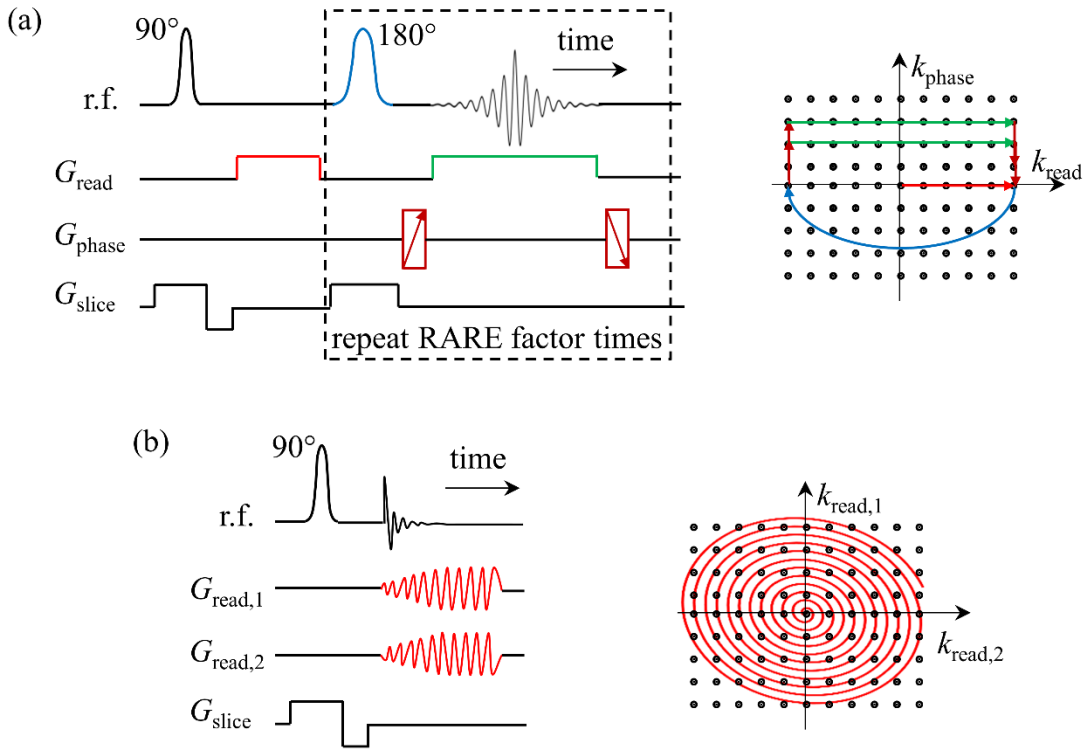

Fig. 2.12. (a) Illustration of the RARE pulse sequence and the corresponding *k*-space traversal. (b) Illustration of the spiral imaging pulse sequence and the corresponding *k*-space traversal.

2.17. Gradient mapping

The use of large gradient strengths and fast gradient switching in ultrafast imaging techniques, such as spiral imaging, can lead to strong eddy currents being generated, which in turn lead to gradient waveform non-ideality. The actual gradient experienced by the sample may be significantly different from the design gradient, which leads to poor MRI image reconstructions. MRI image reconstructions can be improved by measuring the actual gradient waveform and using this waveform, rather than the design gradient

waveform, in the reconstruction process. Although different methods exist [27, 28], the discussion will be limited to the method of Duyn *et al*. [29], as modified by Tayler [26].

The gradient mapping pulse sequence is illustrated in Fig. 2.13. To measure a gradient waveform in the read direction, $G_{\text{read}}(t)$, a small volume element positioned at a distance d_{read} in the read direction from the centre of the magnet is selected (this is done by using an offset for the frequency of the 90° r.f. pulse). The gradient waveform to be measured is then applied and an NMR signal is acquired at the same time. The experiment is repeated with no gradient waveform applied, to account for off-resonance effects. The difference in the phase of the two NMR signals, $\Delta\phi(t)$, is then given by:

$$\Delta\phi(t) = \int_0^t \gamma G_{\text{read}}(t) d_{\text{read}} \mathrm{d}t = \gamma d_{\text{read}} k_{\text{read}}(t)\,, \tag{2.30}$$

where $k_{\text{read}}(t)$ is the component of the inverse space variable in the read direction. The actual gradient waveform is then obtained by numerically differentiating $k_{\text{read}}(t)$.

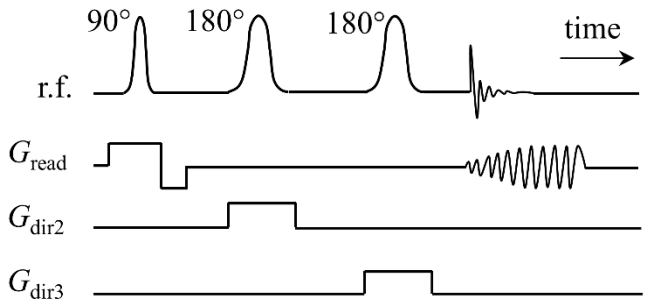

Fig. 2.13. Illustration of the pulse sequence used in this work to map the gradient waveform. The two perpendicular directions to the read direction are referred to as 'dir2' and 'dir3'. The gradient being mapped in the figure is the one corresponding to spiral imaging, but it could be any other gradient waveform.

2.18. Magnetic Resonance velocimetry

Magnetic Resonance (MR) velocimetry is a technique to measure coherent motion. The technique works by encoding the velocity of the fluid in either the magnitude (Time-of-Flight method [30]) or the phase (phase encoding method) of the MR image. The discussion will be limited to the velocity phase encoding method as it is by far the more robust and it is the method used in this work. Unlike conventional velocity measuring techniques of Laser Doppler Velocimetry or Anemometry, Hot Wire Anemometry and Pitot tubes, MR velocimetry is able to measure velocities at different spatial locations

simultaneously and is routinely used to measure all three components of the flow field in 2D slice images, which can be acquired in any direction, as well as 3D volume images. MR velocimetry is non-invasive in the sense that no measuring devices or tracer particles are inserted in the flow, as is the case with the most prominent technique of Particle Image Velocimetry or with Doppler Ultrasound Velocimetry. MR velocimetry can be used to study optically opaque systems and gives chemically-selective information. The main disadvantages of MR velocimetry include the need to apply strong magnetic fields and the inability to study ferromagnetic systems. The measurable velocities are limited on the lower end by diffusion and noise (typically ~ 10 μm s^{-1}) and on the higher end by liquid leaving the imaging region.

The principle behind velocity phase encoding is to apply a time varying magnetic field gradient, $\boldsymbol{G}_{\text{vel}}(t)$, in the direction of the velocity to be measured before the imaging part of the pulse sequence. The imaging part of the pulse sequence can be any of the imaging pulses sequences discussed in Sections 2.14 and 2.16. The accrued phase before the imaging starts is then given by:

$$\phi(t) = \int_0^t \gamma \boldsymbol{G}_{\text{vel}}(t) \cdot \boldsymbol{r}(t) \mathrm{d}t = \gamma \left(\boldsymbol{r}_0 \cdot \int_0^t \boldsymbol{G}_{\text{vel}}(t) \mathrm{d}t + \frac{\mathrm{d}\boldsymbol{r}}{\mathrm{d}t} \cdot \int_0^t t \boldsymbol{G}_{\text{vel}}(t) \mathrm{d}t + \cdots \right), \qquad (2.31)$$

where a Taylor expansion about $t = 0$ was performed and $\boldsymbol{r}_0$ is the position of spins at $t = 0$. The first term (zeroth moment of $\boldsymbol{G}_{\text{vel}}(t)$) corresponds to the phase accrual due to the initial position of spins. The second term (first moment of $\boldsymbol{G}_{\text{vel}}(t)$) corresponds to the phase accrual due to the velocity of spins. The higher order terms contain the second and higher moments of $\boldsymbol{G}_{\text{vel}}(t)$. The time varying velocity encoding gradient is typically composed of a pair of bipolar gradient pulses of magnitude g, duration δ and distance Δ, as illustrated in Fig. 2.14(a). This velocity encoding gradient has a zeroth moment equal to zero and negligible second and higher moments [31]. The first moment is easily calculated and the phase accrual at the end of the application of $\boldsymbol{G}_{\text{vel}}(t)$ is given by:

$$\phi = \gamma g \delta \Delta v, \qquad (2.32)$$

where v is the velocity in the direction of the velocity encoding gradient. Some phase accrual may occur during slice selection, the imaging part of the pulse sequence or due to generated eddy currents; a reference image with no velocity encoding gradients is typically acquired to take into account these effects. A reference image with the flow switched off

is also typically acquired to account for flow effects on the imaging part of the pulse sequence.

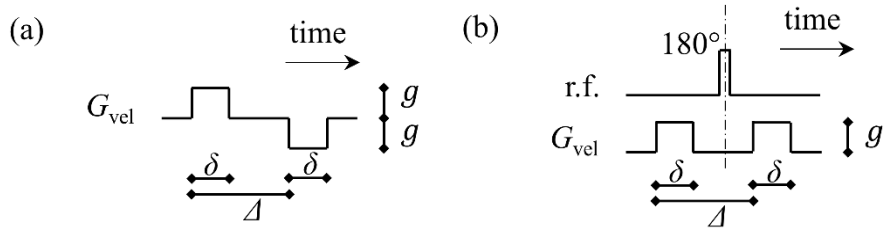

Fig. 2.14. Illustration of the velocity encoding magnetic field gradient (a) with no 180° r.f. pulse and, (b) with a 180° r.f. pulse.

An alternative choice for the velocity encoding gradient is to have the two gradient pulses with the same polarity, but with a 180° r.f. pulse applied in between, as illustrated in Fig. 2.14(b); the effect of the 180° r.f. pulse is to invert the phase accrued due to the first gradient pulse.

2.19. Flow compensation

For temporally changing flow conditions, acquiring a reference image with no velocity encoding gradient may not be appropriate as it would be acquired at different flow conditions to the velocity image acquired with the velocity encoding gradient on. This can be improved by designing magnetic field gradients which have a zeroth moment and first moment equal to zero; these are called flow compensated magnetic field gradients. An example of a flow compensated slice selective gradient for a 90° r.f. pulse, which is used in this work, is shown in Fig. 2.15. For other flow compensated slice selective or non-slice selective magnetic field gradients, the reader is referred to the work of Pope and Yao [32].

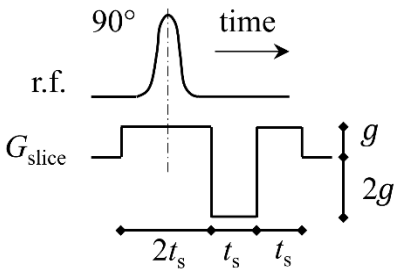

Fig. 2.15. Illustration of a flow compensated slice selective magnetic field gradient for a 90° r.f. pulse.

References

[1] E. M. Purcell, H. C. Torrey and R. V. Pound, Resonance absorption by nuclear magnetic moments in a solid, Phys. Rev. 69 (1946), 37-38.

[2] F. Bloch, Nuclear induction, Phys. Rev. 70 (1946), 460-474.

[3] P. C. Lauterbur, Image formation by induced local interactions: Examples employing nuclear magnetic resonance, Nature 242 (1973), 190-191.

[4] P. Mansfield and P. K. Grannell, NMR 'diffraction' in solids, J. Phys. C. 6 (1973), L422-L426.

[5] P. T. Callaghan, Principles of Nuclear Magnetic Resonance Microscopy, Oxford University Press, 1991.

[6] P. T. Callaghan, Translational Dynamics & Magnetic Resonance, Oxford University Press, 2011.

[7] M. H. Levitt, Spin Dynamics: Basics of Nuclear Magnetic Resonance, WILEY, 2001.

[8] J. Keeler, Understamding NMR Spectroscopy, 2nd ed., WILEY, 2010.

[9] E. L. Hahn, Spin echoes, Phys. Rev. 80 (1950), 580-594.

[10] J. Hennig, Echoes – how to generate, recognize, use or avoid them in MR-imaging sequences, Conc. Magn. Reson. 3 (1991), 125-143.

[11] H. Y. Carr and E. M. Purcell, Effects of diffusion on free precession in nuclear magnetic resonance experiments, Phys. Rev. 94 (1954), 630-638.

[12] S. Meiboom and D. Gill, Modified spin-echo method for measuring nuclear relaxation times, Rev. Sci. Instrum. 29 (1958), 688-691.

[13] J. Aguilar, M. Nilsson, G. Bodenhausen and G. A. Morris, Spin echo NMR spectra without J modulation, Chem. Commun. 48 (2012), 811-813.

[14] E. O. Stejskal and J. E. Tanner, Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient, J. Chem. Phys. 42 (1965), 288-292.

[15] J. E. Tanner, Use of the stimulated echo in NMR diffusion studies, J. Chem. Phys. 52 (1970), 2523-2526.

[16]J. W. Cooley and J. W. Tukey, An algorithm for the machine calculation of complex Fourier series, Math. Comput. 19 (1965), 297-301.

[17]S. Ljunggren, A simple graphical representation of Fourier-based imaging methods, J. Magn. Reson. 54 (1983), 338-343.

[18]W. A. Edelstein, J. M. S. Hutchinson, G. Johnson and T. W. Redpath, Spin warp NMR imaging and applications to human whole-body imaging, Phys. Med. Biol. 25 (1980), 751-756.

[19]A. Haase, Snapshot FLASH MRI. Application to T_1, T_2, and chemical-shift imaging, Magn. Reson. Med. 13 (1990), 77-89.

[20]P. Mansfield, Multi-planar image formation using NMR spin echoes, J. Phys. C: Solid State Phys. 10 (1977), L55-L58.

[21]J. Hennig, A. Naureth and H. Friedburg, RARE imaging: A fast imaging method for clinical MR, Magn. Reson. Med. 3 (1986), 823-833.

[22]C. B. Ahn, J. H. Kim and Z. H. Cho, High-speed spiral-scan echo planar NMR imaging – I, IEEE Trans. Med. Imaging MI-5 (1986), 2-7.

[23]G. P. Liney, A. J. Knowles, D. J. Manton, L. W. Turnbull, S. J. Blackband and A. Horsman, Comparison of conventional single echo and multi-echo sequences with a fast spin-echo sequence for quantitative T_2 mapping: Application to the prostate, J. Magn. Reson. Imaging 6 (1996), 603-607.

[24]M. S. Sussman, L. Vidarsson, J. M. Pauly and H. M. Cheng, A technique for rapid single-echo spin-echo T_2 mapping, Magn. Reson. Med. 64 (2010), 536-545.

[25]C. H. Meyer, B. S. Hu, D. G. Nishimura and A. Macovski, Fast spiral coronary artery imaging, Magn. Reson. Med. 28 (1992), 202-213.

[26]A. B. Tayler, Experimental characterisation of bubbly flow using MRI, PhD Thesis, 2011, University of Cambridge.

[27]N. D. Zanche, C. Barmet, J. A. Nordmeyer-Massner and K. P. Pruessmann, NMR probes for measuring magnetic fields and field dynamics in MR systems, Magn. Reson. Imaging 60 (2008), 176-186.

[28]H. Han, R. P. MacGregor and B. J. Balcom, Pure phase encode magnetic field gradient monitor, J. Magn. Reson. 201 (1009), 212-217.

[29]J. H. Duyn, Y. Frank, J. A. Frank and J.W. van der Veen, Simple correction method for *k*-space trajectory deviations in MRI, J. Magn. Reson. 132 (1998), 150-153.

[30]J. R. Singer, Blood flow rates by Nuclear Magnetic Resonance measurements, Science 130 (1959), 1652-1653.

[31]A. J. Sederman, M. D. Mantle, C. Buckley and L. F. Gladden, MRI technique for measurement of velocity vectors, acceleration and autocorrelation functions in turbulent flow, J. Magn. Reson. 166 (2004), 182-189.

[32]J. M. Pope and S. Yao, Quantitative NMR imaging of flow, Concepts Magn. Reson. 5 (1993), 281-302.

Chapter 3

Mathematical techniques

This chapter introduces the mathematical techniques used throughout the thesis. Conventional Inverse Laplace Transform (ILT) methods, which are used in the extraction of 1D and 2D distributions of relaxation time constants and diffusion coefficients, are presented in Section 3.1. These techniques are used as benchmarks for the new ILT methods developed in Chapters 4 and 5. The idea of Compressed Sensing for MRI acquisitions and a reconstruction method, the Primal-Dual Hybrid Gradient Method (PDHGM), are then introduced in Sections 3.2 and 3.3, respectively. The PDHGM algorithm is used to solve some of the Compressed Sensing related problems posed in Chapters 4, 5, 8 and 9. The Cramér-Rao Lower Bound (CRLB) theory, which is a statistical method for estimating the uncertainty in the extraction of parameters from noisy experimental data, is discussed in Section 3.4. The CRLB theory forms the basis of the sampling pattern optimisation method developed in Chapters 6 and 7, and is also used in Chapter 4.

3.1. Inverse Laplace Transform methods

It was shown in Sections 2.9-2.12 that the vectorised NMR signal, $\boldsymbol{Y}_{\mathrm{v}}$, from either a 1D T_1, T_2, D experiment, or a 2D NMR correlation experiment encoded with the joint distribution of any two of these parameters, can be written in the form:

$$\boldsymbol{Y}_{\mathrm{v}} = \boldsymbol{K}\,\boldsymbol{U}_{\mathrm{v}} + \boldsymbol{E}_{\mathrm{v}}\,, \tag{3.1}$$

where $\boldsymbol{K}$ is a kernel function which depends on the type of experiment, $\boldsymbol{U}_{\mathrm{v}}$ is the vectorised distribution that needs to be extracted, and $\boldsymbol{E}_{\mathrm{v}}$ represents the noise vector that corrupts the signal. Methods for extracting $\boldsymbol{U}_{\mathrm{v}}$ from a knowledge of $\boldsymbol{Y}_{\mathrm{v}}$ and $\boldsymbol{K}$ are known collectively as Inverse Laplace Transform (ILT) methods; the reason for the name is related to the kernel function being composed of exponentially decaying functions.

The naïve solution:

$$\boldsymbol{U}_{\mathrm{v}} = \boldsymbol{K}^{-1}\,\boldsymbol{Y}_{\mathrm{v}}\ , \tag{3.2}$$

is not suitable for three major reasons. The inverse of the kernel function may not exist as $\boldsymbol{K}$ is not necessarily a square matrix; it typically has more columns than rows which means that there are an infinite number of solutions for $\boldsymbol{U}_{\mathrm{v}}$. Even if $\boldsymbol{K}$ is designed to be a square matrix, its condition number, $\kappa(\boldsymbol{K})$, defined as the ratio of its largest to the smallest singular value [1], is so large that small changes in the acquired signal $\boldsymbol{Y}_{\mathrm{v}}$ (due to the random nature of noise) can lead to very large changes in the reconstructed $\boldsymbol{U}_{\mathrm{v}}$. The large condition number is caused by the constituent exponentially decaying functions of $\boldsymbol{K}$ being very similar to each other. The fact that $\boldsymbol{K}$ may not be a square matrix or that it has a large condition number makes the extraction of $\boldsymbol{U}_{\mathrm{v}}$ an ill-conditioned inverse problem. Physically, $\boldsymbol{U}_{\mathrm{v}}$ should not have any negative entries because it represents amounts of material. However, there is nothing constraining $\boldsymbol{U}_{\mathrm{v}}$ not to have negative entries in the solution according to Eq. (3.2).

3.1.1. Non-Negative Least Squares method

The non-negativity constraint was first addressed with the Non-Negative Least Squares (NNLS) method developed by Lawson and Hanson [2], who used an active set method to estimate $\boldsymbol{U}_{\mathrm{v}}$ from the following minimization problem:

$$\boldsymbol{U}_{\mathrm{v}} = \arg\min_{\boldsymbol{U}_{\mathrm{v}} \geq \boldsymbol{0}} \frac{1}{2} \|\boldsymbol{K}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2\ , \tag{3.3}$$

where the inequality is element-wise. The L_p-norm of a vector $\boldsymbol{a}$, $\|\boldsymbol{a}\|_p$, is defined as:

$$\|\boldsymbol{a}\|_p = \left(\sum_{i=1}^{n} |a_i|^p\right)^{1/p}\ . \tag{3.4}$$

Faster algorithms for imposing the non-negativity constraint have been developed [3], but they are all based in solving the minimization problem in Eq. (3.3). In this work, the routine `lsqnonneg` in MATLAB [4] was used to implement the NNLS method.

Although the NNLS method has been very successful, its main disadvantage is that it does not address the ill-conditioning of the inverse problem. As a result, solutions are unstable to noise and the instability increases with the noise level.

3.1.2. Tikhonov regularization

Tikhonov regularization [5] is the most commonly used method to address the ill-conditioning issue of the inverse problem, although other methods exist [6]. In

Tikhonov regularization, the distribution $\boldsymbol{U}_{\mathrm{v}}$ is estimated from the following minimization problem:

$$\boldsymbol{U}_{\mathrm{v}} = \arg\min_{\boldsymbol{U}\geq\boldsymbol{0}} \left(\frac{1}{2}\|\boldsymbol{K}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \alpha\|\boldsymbol{R}\,\boldsymbol{U}_{\mathrm{v}}\|_2^2\right), \tag{3.5}$$

where α is a regularization parameter and $\boldsymbol{R}$ is a regularization matrix. The first term is referred to as the fidelity term, while the second term is referred to as the penalty term. There are different explanations as to why Tikhonov regularization improves the reconstructions. The simplest explanation is that it narrows down the range of potential solutions of $\boldsymbol{U}_{\mathrm{v}}$, by searching for solutions that have a minimal penalty term. Another explanation comes from the solution to Eq. (3.5), neglecting the non-negativity constraint:

$$\boldsymbol{U}_{\mathrm{v}} = (\boldsymbol{K}^{\mathrm{T}}\boldsymbol{K} + \alpha\boldsymbol{R}^{\mathrm{T}}\boldsymbol{R})^{-1}\boldsymbol{K}^{\mathrm{T}}\boldsymbol{Y}_{\mathrm{v}}\,. \tag{3.6}$$

The matrix being inverted in Eq. (3.6) is generally of a much lower condition number than the matrix that needs to be inverted for the solution of Eq. (3.3) which, when neglecting the non-negativity constraint, is:

$$\boldsymbol{U}_{\mathrm{v}} = (\boldsymbol{K}^{\mathrm{T}}\boldsymbol{K})^{-1}\boldsymbol{K}^{\mathrm{T}}\boldsymbol{Y}_{\mathrm{v}}\,. \tag{3.7}$$

This renders Tikhonov regularization reconstructions more stable to noise than NNLS reconstructions. Other explanations can be given in terms of a singular value decomposition [7] or overfitting argument [8].

For 1D NMR experiments, the regularization matrix, $\boldsymbol{R}$, is either taken to be the identity matrix or the second order derivative matrix; the second order derivative matrix is a tridiagonal matrix with entries 1, -2, 1 on diagonals of number -1, 0, 1. The choice of the second order derivative matrix, also known as the Phillips-Twomey method [9, 10], constrains $\boldsymbol{U}_{\mathrm{v}}$ to be smooth. For 2D NMR correlation experiments, $\boldsymbol{R}$ is almost exclusively chosen to be the identity matrix, because of the difficulty of defining the second order derivative in two dimensions.

The regularization parameter, α, controls the amount of regularization imposed on the reconstructed distribution. A larger value of α leads to smaller values of the penalty term in Eq. (3.5) (hence, smoother distribution) and larger values of the fidelity term (hence, worse fit to the experimental data). This is known as the bias-variance trade-off. There is, therefore, a need to rationalize the choice of α. The Generalized Cross Validation (GCV) technique [11] is used in this work to choose α, although other techniques such as the

L-curve [12] and the Morozov Discrepancy Principle [13] exist. The choice was made mainly on the statistical basis of GCV. GCV is a leave-one-out cross-validation technique: one data point is used as a validation set while the other data points are used as the training set. The process is repeated for all possible choices of the validation set. The average test error over all possible validation sets is then referred to as the GCV score. A GCV score for a range of values of α, GCV(α), is calculated and the best regularization parameter is chosen to be the one that minimizes the GCV(α). An analytical solution exists for GCV(α) [11]:

$$\mathrm{GCV}(\alpha) = \frac{\|\boldsymbol{K}\,\boldsymbol{U}_{\mathrm{V}}(\alpha) - \boldsymbol{Y}_{\mathrm{V}}\|_2^2}{(n - df(\alpha))^2}\,, \tag{3.8}$$

where $\boldsymbol{U}_{\mathrm{V}}(\alpha)$ is the reconstructed distribution for the choice of α, n is the size of $\boldsymbol{U}_{\mathrm{V}}$ and df(α) is referred to as the number of fitted parameters, or degrees of freedom, given by:

$$\mathrm{df}(\alpha) = \boldsymbol{K}(\boldsymbol{K}^{\mathrm{T}}\boldsymbol{K} + \alpha\,\boldsymbol{R}^{\mathrm{T}}\boldsymbol{R})^{-1}\boldsymbol{K}^{\mathrm{T}}\,. \tag{3.9}$$

GCV is a model selection method, similar to the Akaike Information Criterion [14], Bayesian Information Criterion [15] or C_p method [16], in the sense that it invokes a parsimony argument; as the regularization parameter α changes, the number fitted parameters df(α) varies, and model selection methods decide on the optimal α based on a trade-off between fidelity to the data and df(α). Although GCV has been the most commonly used model selection method, the other methods have also been used in the inversion of NMR data [17]. Practically, all these methods give very similar results, but GCV is easier to calculate as it does not require a knowledge of the noise level in the data.

Several algorithms have been developed over the years to numerically solve Eq. (3.5) [18-23]. In this work, Eq. (5) is cast as a quadratic programming problem [24] and the routine `quadprog` in MATLAB [4] is subsequently applied.

3.1.3. Data compression

For 2D NMR correlation experiments, the data size can be very large, particularly when one of the distributions being extracted is the T_2 relaxation time constant. As a result, the numerical computation of Eqs. (3.3) and (3.5) can be time and memory consuming. It is, therefore, common to perform a pre-compression of the NMR data before feeding it to the appropriate ILT algorithm. The compression technique used in this work is the one developed by Venkataramanan *et al*. [21]. As discussed in Section 2.12, the acquired NMR data for a 2D NMR correlation experiment can be written in a matrix form:

$$\boldsymbol{Y} = \boldsymbol{K}_1 \, \boldsymbol{U} \, \boldsymbol{K}_2{}^{\mathrm{T}} + \boldsymbol{E} \, . \tag{3.10}$$

Let the singular value decomposition of $\boldsymbol{K}_1$ and $\boldsymbol{K}_2$ be $\boldsymbol{U}_1\boldsymbol{\Sigma}_1\boldsymbol{V}_1{}^{\mathrm{T}}$ and $\boldsymbol{U}_2\boldsymbol{\Sigma}_2\boldsymbol{V}_2{}^{\mathrm{T}}$, where $\boldsymbol{U}_1$, $\boldsymbol{U}_2$, $\boldsymbol{V}_1$ and $\boldsymbol{V}_2$ are orthogonal matrices of respective sizes $p \times s_1$, $q \times s_2$, $n \times s_1$ and $m \times s_2$, while $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_2$ are diagonal square matrices of respective sizes $s_1 \times s_1$ and $s_2 \times s_2$ with sorted values on the diagonals (known as singular values). Because of the large condition number of $\boldsymbol{K}_1$ and $\boldsymbol{K}_2$, the singular values span orders of magnitude and there are only a few important singular values. The compression is based on the truncation of the singular values such that only the most important ones are kept; the choice of how many are kept is typically made based on the noise level of the NMR signal. Let the truncated diagonal matrices be $\boldsymbol{\Sigma}_{1\mathrm{t}}$ and $\boldsymbol{\Sigma}_{2\mathrm{t}}$, of respective sizes $s_{1\mathrm{t}} \times s_{1\mathrm{t}}$ and $s_{2\mathrm{t}} \times s_{2\mathrm{t}}$. The corresponding truncated orthogonal matrices $\boldsymbol{U}_{1\mathrm{t}}$, $\boldsymbol{U}_{2\mathrm{t}}$, $\boldsymbol{V}_{1\mathrm{t}}$ and $\boldsymbol{V}_{2\mathrm{t}}$ are of respective sizes $p \times s_{1\mathrm{t}}$, $q \times s_{2\mathrm{t}}$, $n \times s_{1\mathrm{t}}$ and $m \times s_{2\mathrm{t}}$. Pre-multiplying by $\boldsymbol{U}_{1\mathrm{t}}{}^{\mathrm{T}}$ and post-multiplying by $\boldsymbol{U}_{2\mathrm{t}}$ both sides of Eq. (3.10), the following expression is obtained:

$$\boldsymbol{Y}_{\mathrm{t}} = \boldsymbol{K}_{1t} \, \boldsymbol{U} \, \boldsymbol{K}_{2t}{}^{\mathrm{T}} + \boldsymbol{E}_{\mathrm{t}} \, , \tag{3.11}$$

where $\boldsymbol{Y}_{\mathrm{t}} = \boldsymbol{U}_{1\mathrm{t}}{}^{\mathrm{T}}\boldsymbol{Y}\,\boldsymbol{U}_{2\mathrm{t}}$, $\boldsymbol{K}_{1\mathrm{t}} = \boldsymbol{\Sigma}_{1\mathrm{t}}\,\boldsymbol{V}_{1\mathrm{t}}{}^{\mathrm{T}}$, $\boldsymbol{K}_{2\mathrm{t}} = \boldsymbol{\Sigma}_{2\mathrm{t}}\,\boldsymbol{V}_{2\mathrm{t}}{}^{\mathrm{T}}$ and $\boldsymbol{E}_{\mathrm{t}} = \boldsymbol{U}_{1\mathrm{t}}{}^{\mathrm{T}}\boldsymbol{E}\,\boldsymbol{U}_{2\mathrm{t}}$. Eq. (3.11) is of the same form as Eq. (3.10) but the size of matrices involved is much smaller. Therefore, by applying the above compression method to the NMR signal and the kernel matrices, the reconstruction of distributions becomes less time and memory demanding.

Apart from making the reconstruction of distributions less time and memory demanding, the data compression technique described has also the ‘side-effect’ of making the inverse problem less ill-conditioned; the condition number of $\boldsymbol{K}_{1\mathrm{t}}$ and $\boldsymbol{K}_{2\mathrm{t}}$ is smaller than the condition number of $\boldsymbol{K}_1$ and $\boldsymbol{K}_2$.

3.2. Reconstructions from incomplete data

Classical sampling theory [25, 26] states that if a signal is bandlimited in the frequency domain, it can be completely recovered from frequency samples only if the sampling frequency is above a certain limit. This sets a lower limit of samples that need to be acquired in the frequency domain to fully reconstruct a signal in the time domain. This barrier has been broken by developments in Compressed Sensing [27, 28]. Compressed Sensing states that, to a high probability, it is possible to reconstruct the signal from far fewer frequency samples than the classical sampling theory requires, as long as sampling

is performed randomly and the signal being reconstructed is sparse in some domain. Although reconstruction techniques from very few samples have been around for a long time [29, 30], Compressed Sensing has created a mathematical framework for assigning probabilities for the success of such reconstruction techniques. Since its formulation, Compressed Sensing has found applications in areas as diverse as astronomy [31], single-pixel cameras [32] and tomography [33].

In MRI applications, an under sampled dataset $\boldsymbol{Y}$, vectorised as $\boldsymbol{Y}_{\mathrm{v}}$, corresponds to a partially acquired *k*-space, which can be written as:

$$\boldsymbol{Y}_{\mathrm{v}} = \boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_{\mathrm{v}} + \boldsymbol{E}_{\mathrm{v}}\,, \tag{3.12}$$

where $\boldsymbol{S}$ is the subsampling matrix, $\boldsymbol{\mathcal{F}}$ is the Fourier operator, $\boldsymbol{U}_{\mathrm{v}}$ is the vectorised MRI image and $\boldsymbol{E}_{\mathrm{v}}$ represents the noise vector that corrupts the signal. Two major considerations need to be taken when applying Compressed Sensing to MRI applications: incoherence (therefore, the subsampling scheme) and the sparsifying domain (therefore, the reconstruction technique).

3.2.1. Incoherence and the subsampling scheme

The subsampling of *k*-space needs to be performed randomly to reduce coherent artefacts in the reconstructed MRI image. For measurements not corrupted by noise, this can be achieved by having a subsampling matrix, $\boldsymbol{S}$, which is incoherent [34]. An incoherent matrix is constructed by having the absolute inner product of any of its columns as low as possible. In the presence of measurement noise, this is generalised to $\boldsymbol{S}$ having a Restricted Isometry Property [27], which for most practical applications is relaxed to $\boldsymbol{S}$ being incoherent.

In MRI applications, the most important information is concentrated close to the centre of *k*-space. This prior knowledge forces researches to sample the centre of *k*-space more densely. However, this is in contradiction to the incoherence principle and it embodies the trade-off that exists between the sampling scheme being random and being biased towards the centre of *k*-space.

In practice, the most commonly used method is to draw random samples with higher probability of selection close to the centre of *k*-space, according to a specified probability density function [35, 36]. Other methods, such as Poisson disc sampling [37] and Diffusion Limited Aggregation [38] are based on analogies with natural systems and have the advantage of not needing to specify a probability density function. Recently, sampling

schemes which make use of some prior knowledge about the sample being imaged [39, 40] (typically obtained through some other imaging method) are drawing a lot of interest. All the mentioned sampling techniques tend not to sample the centre of *k*-space densely enough. Fully sampling a central region of *k*-space and then applying any of the mentioned techniques for the remaining part of *k*-space has been shown to give better results [41].

To ensure to a high probability the exact reconstruction of an MRI image of size n, which has only s non-zero entries in some domain, the number of *k*-space points that need to be acquired goes as $s \ln(n / s) \ll n$. As a result, a sparser MRI image in some domain can be reconstructed from a higher level of under sampling of *k*-space.

3.2.2. Sparsifying domain and the reconstruction techniques

The reconstruction of the MRI image, $\boldsymbol{U}_{\mathrm{v}}$, from under sampled *k*-space data, $\boldsymbol{Y}_{\mathrm{v}}$, is an ill-conditioned inverse problem; there are an infinite solutions of $\boldsymbol{U}_{\mathrm{v}}$ that fit the under sampled $\boldsymbol{Y}_{\mathrm{v}}$ exactly. A naïve method for reconstructing the MRI image is:

$$\boldsymbol{U}_{\mathrm{v}} = \boldsymbol{\mathcal{F}}^{-1}\boldsymbol{S}^{\mathrm{T}}\boldsymbol{Y}_{\mathrm{v}} \,. \tag{3.13}$$

This is known as a Zero-Filled reconstruction; it is equivalent to filling the missing points in *k*-space by zeros and performing an inverse Fourier transform. Zero-Filled reconstructions are suboptimal; they suffer from under sampling and noise artefacts. Better reconstructions are obtained by including in the reconstruction process the prior information that the MRI image is sparse in some domain, $\boldsymbol{\psi}$. One way to do this is to solve a constrained optimisation problem:

$$\boldsymbol{U}_{\mathrm{v}} = \arg\min_{\boldsymbol{U}_{\mathrm{v}}} \|\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 \quad \text{s.t.} \quad \|\boldsymbol{\psi}\,\boldsymbol{U}_{\mathrm{v}}\|_0 \leq \epsilon, \tag{3.14}$$

where the L_0-norm is a limiting case of the L_p-norm defined in Eq. (3.4); it is equal to the cardinality, or the number of non-zero elements, of the vector it is applied to. ε is a parameter to be set, which determines the amount of sparsity to be enforced on the transform of $\boldsymbol{U}_{\mathrm{v}}$. More commonly, Eq. (3.14) is written as an unconstrained regularization problem:

$$\boldsymbol{U}_{\mathrm{v}} = \arg\min_{\boldsymbol{U}_{\mathrm{v}}} \left(\frac{1}{2}\|\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \alpha\|\boldsymbol{\psi}\,\boldsymbol{U}_{\mathrm{v}}\|_0\right), \tag{3.15}$$

which is its Lagrangian form. The regularization parameter, α, depends on the ε parameter. Unfortunately, the numerical solution of the minimization problem in Eq. (3.15) is an NP

hard problem which is computationally intractable. Therefore, a convex optimization approximation is more commonly used:

$$\boldsymbol{U}_{\mathrm{v}} = \arg\min_{\boldsymbol{U}_{\mathrm{v}}} \left(\frac{1}{2} \| \boldsymbol{S}\, \boldsymbol{\mathcal{F}}\, \boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}} \|_2^2 + \alpha \| \boldsymbol{\psi}\, \boldsymbol{U}_{\mathrm{v}} \|_1 \right). \tag{3.16}$$

The most commonly used sparsifying domain, $\boldsymbol{\psi}$, is Total Variation (TV) [42], although Wavelet [43] and Total Generalized Variation [44] have also been used. TV regularization is based on the observation that MRI images typically contain large areas of constant intensity, followed by sharp drops in intensity at the edges. As a result, pixel-wise differences of the MRI image in both directions (discretized version of the first order derivative) contain very few non-zero elements.

3.3. Primal-Dual Hybrid Gradient Method

As discussed in Sections 3.1 and 3.2, many NMR and MRI reconstruction problems are regularization problems. These are convex optimization problems, for the solution of which there exist a plethora of different algorithms [45]. One family of algorithms that tackles these problems is proximal algorithms [46,47]. In these algorithms, a large convex optimization problem is split into small convex optimization problems and a proximal operator function is applied to all these small problems. A proximal operator function, which is described in more detail in Section 3.3.2, effectively moves a search point towards a domain and minimum of a function in a similar way that Newton's gradient descent method does.

In recent years, a class of the family of proximal algorithms, primal-dual algorithms, has gained much interest [48-51]. The main advantage of the primal-dual algorithms is the efficient splitting of the large problem, and consequently the lower computational cot involved. In this work, the Primal-Dual Hybrid Gradient Method (PDHGM) developed by Chambolle and Pock [50], which has been very successful in imaging applications, is used. The PDHGM method is an efficient algorithm for solving problems of the form:

$$(\boldsymbol{u}, \boldsymbol{v}) = \min_{\boldsymbol{u}} \max_{\boldsymbol{v}} f(\boldsymbol{u}) - g(\boldsymbol{v}) + \boldsymbol{v}^{\mathrm{T}} \boldsymbol{R}\, \boldsymbol{u}\,, \tag{3.17}$$

where $\boldsymbol{u}$ and $\boldsymbol{v}$ are vectors referred to as the primal and the dual variable, respectively, $\boldsymbol{R}$ is a matrix, and f and g are convex functions. The PDHGM is also applicable to the case when the primal and dual variables are matrices, rather than vectors; this case is not

discussed in the following sections. Practically, the only difference is that the third term in Eq. (3.17) would be substituted by the inner product of $\boldsymbol{v}$ and $\boldsymbol{R}\,\boldsymbol{u}$, $\langle \boldsymbol{v}, \boldsymbol{R}\,\boldsymbol{u} \rangle$.

A rigorous mathematical treatment, including a convergence proof is found in the work of Chambolle and Pock [50]. The derivation of the algorithm given in Section 3.3.1 is much simplified and is based on the treatment of Benning [52]. Section 3.3.2 introduces the proximal operator which is the main building block of algorithm. Three examples of the application of the proximal operator are given in Section 3.3.3. Some practicalities regarding the implementation of the PDHGM are given in Section 3.3.4.

3.3.1. Iteration scheme

Let the derivative of $f(\boldsymbol{u})$ with respect to $\boldsymbol{u}$ be $p(\boldsymbol{u})$ and the derivative of $g(\boldsymbol{v})$ with respect to $\boldsymbol{v}$ be $q(\boldsymbol{v})$. Then, the optimum of the problem in Eq. (3.17) is achieved when:

$$p(\boldsymbol{u}) + \boldsymbol{R}^{\mathrm{T}}\boldsymbol{v} = \boldsymbol{0}\,, \tag{3.18a}$$

$$q(\boldsymbol{v}) - \boldsymbol{R}\,\boldsymbol{u} = \boldsymbol{0}\,. \tag{3.18b}$$

A potential iteration scheme for solving Eqs. 3.18(a) and 3.18(b) is as follows:

$$\begin{bmatrix} p\big(\boldsymbol{u}^{(k+1)}\big) + \boldsymbol{R}^{\mathrm{T}}\boldsymbol{v}^{(k+1)} \\ q\big(\boldsymbol{v}^{(k+1)}\big) - \boldsymbol{R}\,\boldsymbol{u}^{(k+1)} \end{bmatrix} + \boldsymbol{T}\begin{bmatrix} \boldsymbol{u}^{(k+1)} - \boldsymbol{u}^{(k)} \\ \boldsymbol{v}^{(k+1)} - \boldsymbol{v}^{(k)} \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix}, \tag{3.19}$$

where k refers to the iteration step and $\boldsymbol{T}$ is a matrix to be chosen. It is seen that as the iterations converge, i.e. $\boldsymbol{u}^{(k+1)} \rightarrow \boldsymbol{u}^{(k)}$ and $\boldsymbol{v}^{(k+1)} \rightarrow \boldsymbol{v}^{(k)}$, Eqs. (3.18a) and (3.18b) are obeyed and, therefore, an optimum to Eq. (3.17) is obtained. Chambolle and Pock [50] made the following choice for $\boldsymbol{T}$:

$$\boldsymbol{T} = \begin{bmatrix} 1/\tau\,\boldsymbol{I} & -\boldsymbol{R}^{\mathrm{T}} \\ \boldsymbol{R} & 1/\sigma\,\boldsymbol{I} \end{bmatrix}, \tag{3.20}$$

where τ and σ are algorithmic parameters and $\boldsymbol{I}$ is the identity matrix. This choice of $\boldsymbol{T}$ makes the system of equations in Eq. (3.19) explicit in terms of $\boldsymbol{u}^{(k+1)}$ and $\boldsymbol{v}^{(k+1)}$:

$$\boldsymbol{u}^{(k+1)} + \tau\, p\big(\boldsymbol{u}^{(k+1)}\big) = \boldsymbol{u}^{(k)} - \tau\,\boldsymbol{R}^{\mathrm{T}}\boldsymbol{v}^{(k)}, \tag{3.21a}$$

$$\boldsymbol{v}^{(k+1)} + \sigma\, q\big(\boldsymbol{v}^{(k+1)}\big) = \boldsymbol{v}^{(k)} + \sigma\,\boldsymbol{R}\big(2\boldsymbol{u}^{(k+1)} - \boldsymbol{u}^{(k)}\big)\,. \tag{3.21b}$$

3.3.2. Proximal operator

It is common to write Eq. (3.21) in operator format:

$$(\boldsymbol{I} + \tau\,\partial f^{\mathrm{T}})\big(\boldsymbol{u}^{(k+1)}\big) = \boldsymbol{u}^{(k)} - \tau\,\boldsymbol{R}^{\mathrm{T}}\boldsymbol{v}^{(k)}, \tag{3.22a}$$

$$(\boldsymbol{I} + \sigma\, \partial g^{\mathrm{T}})\left(\boldsymbol{v}^{(k+1)}\right) = \boldsymbol{v}^{(k)} + \sigma\, \boldsymbol{R}\left(2\boldsymbol{u}^{(k+1)} - \boldsymbol{u}^{(k)}\right) . \tag{3.22b}$$

Solving for $\boldsymbol{u}^{(k+1)}$ and $\boldsymbol{v}^{(k+1)}$:

$$\boldsymbol{u}^{(k+1)} = (\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-\mathbf{1}}\left(\boldsymbol{u}^{(k)} - \tau\, \boldsymbol{R}^{\mathrm{T}}\boldsymbol{v}^{(k)}\right) , \tag{3.23a}$$

$$\boldsymbol{v}^{(k+1)} = (\boldsymbol{I} + \sigma\, \partial g^{\mathrm{T}})^{-\mathbf{1}}\left(\boldsymbol{v}^{(k)} + \sigma\, \boldsymbol{R}\left(2\boldsymbol{u}^{(k+1)} - \boldsymbol{u}^{(k)}\right)\right) . \tag{3.23b}$$

The operator $(\boldsymbol{I} + \partial f^{\mathrm{T}})^{-1}$ is known as a proximal operator. Therefore, each iteration step in the PDHGM algorithm consists in the application of two proximal operators for the updating of the primal and dual variables. The proximal operator is computed with the help of the following observation:

$$(\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-\mathbf{1}}(\boldsymbol{a}) = \arg\min_{\boldsymbol{u}}\left(\frac{1}{2}\|\boldsymbol{u} - \boldsymbol{a}\|_2^2 + \tau\, f(\boldsymbol{u})\right) . \tag{3.24}$$

To prove Eq. (3.24), it is noted that the minimum of the r.h.s. expression is achieved when its derivative with respect to $\boldsymbol{u}$ is equal to zero, which gives:

$$\boldsymbol{a} = \boldsymbol{u} + \frac{\partial f^{\mathrm{T}}(\boldsymbol{u})}{\partial \boldsymbol{u}} . \tag{3.25}$$

When Eq. (3.25) is written in an operator format, it gives the l.h.s. in Eq. (3.24).

As is seen in Eq. (3.24), the proximal operator applied to a vector $\boldsymbol{a}$, returns a vector that is a compromise between being as close as possible to $\boldsymbol{a}$ while minimizing a cost function, f, with the relative weight determined by τ.

3.3.3. Examples

Three examples of the proximal operator for common cost functions that are encountered in NMR and MRI applications are now presented.

Example 1

The simplest cost function is a set function:

$$f(\boldsymbol{u}) = I_C(\boldsymbol{u}) = \begin{cases} 0 & \boldsymbol{u} \in C \\ +\infty & \boldsymbol{u} \notin C \end{cases} , \tag{3.26}$$

where C is a set. Since the cost is infinite if $\boldsymbol{u}$ is not in the set, the solution to the proximal operator in Eq. (3.24) is the projection of $\boldsymbol{a}$ into the set C.

Example 2

Another simple cost function encountered in MRI applications is:

$$f(\boldsymbol{u}) = \chi_{\alpha}(\boldsymbol{u}) = \begin{cases} 0 & \|\boldsymbol{u}\|_{\infty} \leq \alpha \\ +\infty & \|\boldsymbol{u}\|_{\infty} > \alpha \end{cases}, \tag{3.27}$$

where the L_{∞}-norm is defined in Eq. (3.4); it represents the largest absolute value entry of $\boldsymbol{u}$. The function $\chi_{\alpha}(\boldsymbol{u})$ is known as the indicator function. The proximal operator of $\chi_{\alpha}(\boldsymbol{u})$:

$$(\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-\mathbf{1}}(\boldsymbol{a}) = \arg\min_{\boldsymbol{u}} \left(\frac{1}{2} \|\boldsymbol{u} - \boldsymbol{a}\|_2^2 + \tau\, \chi_{\alpha}(\boldsymbol{u}) \right), \tag{3.28}$$

is best computed by considering each element of $\boldsymbol{a}$ independently, because the minimization problem in Eq. (3.28) is separable. If an element of $\boldsymbol{a}$ is $\leq \alpha$, then the corresponding element in $\boldsymbol{u}$ must be equal to that entry of $\boldsymbol{a}$, because in this way the proximal function equals 0 and this is the smallest value it can reach. If an element of $\boldsymbol{a}$ is $> \alpha$, the corresponding element in $\boldsymbol{u}$ cannot be greater than α, because the indicator function blows up. The best value the element in $\boldsymbol{u}$ can get is α, because this minimizes the first term while keeping the indicator function at 0. Therefore, the proximal operator is given by:

$$(\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-\mathbf{1}}(\boldsymbol{a}) = \frac{\boldsymbol{a}}{\max(1, \boldsymbol{a}/\alpha)}, \tag{3.29}$$

where all operations are element-wise.

Example 3

A common function encountered in under sampled MRI data acquisitions, as discussed in Section 3.2, is:

$$f(\boldsymbol{u}) = \frac{1}{2} \|\boldsymbol{S}\, \boldsymbol{\mathcal{F}}\, \boldsymbol{u} - \boldsymbol{y}\|_2^2, \tag{3.30}$$

where $\boldsymbol{S}$ is a subsampling matrix and $\boldsymbol{\mathcal{F}}$ is the Fourier operator. By definition, the proximal operator of such a function is given by:

$$(\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-\mathbf{1}}(\boldsymbol{a}) = \arg\min_{\boldsymbol{u}} \left(\frac{1}{2} \|\boldsymbol{u} - \boldsymbol{a}\|_2^2 + \frac{\tau}{2} \|\boldsymbol{S}\, \boldsymbol{\mathcal{F}}\, \boldsymbol{u} - \boldsymbol{y}\|_2^2 \right). \tag{3.31}$$

Differentiating the r.h.s. with respect to $\boldsymbol{u}$ and setting it to zero:

$$\boldsymbol{u} - \boldsymbol{a} + \tau\, \boldsymbol{\mathcal{F}}^{-1} \boldsymbol{S}^{\mathrm{T}} \boldsymbol{S}\, \boldsymbol{\mathcal{F}}\, \boldsymbol{y} - \boldsymbol{\mathcal{F}}^{-1} \boldsymbol{S}^{\mathrm{T}} \boldsymbol{y} = \boldsymbol{0}. \tag{3.32}$$

Set $\underline{\boldsymbol{u}} = \boldsymbol{\mathcal{F}}\, \boldsymbol{u}$. After some rearrangement, Eq. (3.32) becomes:

$$(\boldsymbol{\mathcal{F}}^{-1} + \tau\, \boldsymbol{\mathcal{F}}^{-1} \boldsymbol{S}^{\mathrm{T}} \boldsymbol{S}) \underline{\boldsymbol{u}} = \boldsymbol{a} + \boldsymbol{\mathcal{F}}^{-1} \boldsymbol{S}^{\mathrm{T}} \boldsymbol{y}. \tag{3.33}$$

Taking the Fourier transform of both sides and rearranging, the proximal operator becomes:

$$(\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-1}(\boldsymbol{a}) = \boldsymbol{\mathcal{F}}^{-1}\underline{\boldsymbol{u}} = \boldsymbol{\mathcal{F}}^{-1}\left(\frac{\boldsymbol{\mathcal{F}}\,\mathrm{a}+\boldsymbol{S}^{\mathrm{T}}\boldsymbol{y}}{\mathrm{diag}(\boldsymbol{I}+\tau\boldsymbol{S}^{\mathrm{T}}\boldsymbol{S})}\right), \tag{3.34}$$

where diag($\boldsymbol{I}$ + τ $\boldsymbol{S}^{\mathrm{T}}\boldsymbol{S}$) extracts the main diagonal of the matrix in brackets and the division in Eq. (3.34) is element-wise.

3.3.4. Implementation

The PDHGM algorithm has two parameters, τ and σ, which control the convergence, speed and stability of the algorithm. It was proven by Chambolle and Pock [50] that convergence is guaranteed if:

$$\tau\,\sigma\,\|\boldsymbol{R}\|_{2,2}^{2} \leq 1\,. \tag{3.35}$$

The $L_{p,q}$ norm of the matrix $\boldsymbol{R}$, $||\boldsymbol{R}||_{p,q}$, is defined as:

$$\|\boldsymbol{R}\|_{p,q} = \left(\sum_j \left(\sum_i |R_{ij}|^{p}\right)^{q/p}\right)^{1/q}. \tag{3.36}$$

There is freedom to choose τ and σ, if Eq. (3.35) is obeyed, and the choice depends on the problem being studied.

Different stopping criteria could be used for the PDHGM algorithm; the criterion used in this work was:

$$\frac{\|\boldsymbol{u}^{(k+1)}-\boldsymbol{u}^{(k)}\|_2}{\|\boldsymbol{u}^{(k)}\|_2} < \mathrm{TOL}\,, \tag{3.37}$$

where the tolerance level, TOL, is set heuristically.

3.4. Cramér-Rao Lower Bound theory

The Cramér-Rao Lower Bound (CRLB) theory [53] is a statistical method which determines the lower limit for the uncertainty related to the estimate of the parameters of a model extracted from noisy experimental data. It is widely used as a metric for comparing the performance of different parameter estimation techniques in areas as diverse as optical tomography [54], radar measurements [55] and X-ray imaging [56].

Let $\boldsymbol{y}$ be a series of noisy experimental data points and $\hat{\boldsymbol{y}}$ be its noise-free model. Let $\boldsymbol{y}$ be described by a probability density function f that depends on parameters $\boldsymbol{\theta}$. If $\underline{\boldsymbol{\theta}}$ is an

unbiased estimator of $\boldsymbol{\theta}$, the uncertainty related to the estimator of each individual parameter has a lower limit:

$$\mathrm{var}\left(\underline{\theta}_l\right) \geq (\boldsymbol{F}^{-1})_{ll} \,, \tag{3.38}$$

where $\boldsymbol{F}$ is the Fisher Information Matrix [57], defined as:

$$F_{l_1 l_2} = \mathrm{E}\left(S_{l_1}(\boldsymbol{y})\, S_{l_2}(\boldsymbol{y})\right), \tag{3.40}$$

where E refers to the expected value and $S_l(\boldsymbol{y})$ is a score statistic defined as:

$$S_l(\boldsymbol{y}) = \frac{\partial}{\partial \theta_l} \ln f(\boldsymbol{y}|\boldsymbol{\theta}), \tag{3.41}$$

where $\ln f(\boldsymbol{y}|\boldsymbol{\theta})$ is known as the likelihood function. A proof is given below for the case when f depends on only one parameter θ; the proof for the multi-parameter case follows a similar path.

Proof

From the definition of an unbiased estimator:

$$\mathrm{E}\left(\underline{\theta}\right) = \theta, \tag{3.42}$$

which is alternatively written as:

$$\int \underline{\theta} f(\boldsymbol{y}|\theta) d\boldsymbol{y} = \theta. \tag{3.43}$$

Differentiating both sides of Eq. (3.43) with respect to θ and applying the chain rule, the following expression is obtained:

$$\int \underline{\theta} f(\boldsymbol{y}|\theta) \left(\frac{\partial}{\partial \theta} \log f(\boldsymbol{y}|\theta)\right) d\boldsymbol{y} = 1. \tag{3.44}$$

From the definition of the score statistics in Eq. (3.41) and the definition of expectation of a random variable, the following expression is obtained:

$$\mathrm{E}\left(\underline{\theta} S(\boldsymbol{y})\right) = 1. \tag{3.45}$$

On the other hand:

$$\mathrm{E}\left(S(\boldsymbol{y})\right) = \int \left(\frac{\partial}{\partial \theta} \log f(\boldsymbol{y}|\theta)\right) f(\boldsymbol{y}|\theta) d\boldsymbol{y} = \int \left(\frac{\partial}{\partial \theta} f(\boldsymbol{y}|\theta)\right) d\boldsymbol{y} = 0 \,, \tag{3.46}$$

from which it follows that the covariance between $\underline{\theta}$ and $S(\boldsymbol{y})$ is:

$$\mathrm{cov}\left(\underline{\theta}, S(\boldsymbol{y})\right) = \mathrm{E}\left(\underline{\theta} S(\boldsymbol{y})\right) - \mathrm{E}\left(\underline{\theta}\right)\mathrm{E}\left(S(\boldsymbol{y})\right) = 1 \,. \tag{3.47}$$

Since the correlation coefficient between any two random variables is less than 1:

$$\mathrm{var}(\underline{\theta})\mathrm{var}(S(\boldsymbol{y})) \geq \mathrm{cov}^2\left(\underline{\theta}, S(\boldsymbol{y})\right), \tag{3.48}$$

from which it follows that:

$$\mathrm{var}(\underline{\theta}) \geq 1/\mathrm{var}(S(\boldsymbol{y})) = 1/\mathrm{E}(S^2(\boldsymbol{y})), \tag{3.49}$$

which proves the CRLB theory for the one parameter case.

The calculation of the Fisher Information Matrix, $\boldsymbol{F}$, as described by Eqs. (3.40) and (3.41), is not straightforward. However, its calculation can be simplified for the case when the noise associated with the measured data points are independently and normally distributed with a constant standard deviation, σ. Let the noise related to the data points be $\boldsymbol{e}$ and, therefore, $\hat{\boldsymbol{y}} = \boldsymbol{y} - \boldsymbol{e}$. The likelihood function is then given by:

$$\ln f(\boldsymbol{y}|\boldsymbol{\theta}) = n \ln\left(\frac{1}{\sqrt{2\pi}\sigma}\right) - \frac{1}{2\sigma^2}(\boldsymbol{y} - \hat{\boldsymbol{y}})^{\mathrm{T}}(\boldsymbol{y} - \hat{\boldsymbol{y}}). \tag{3.50}$$

The score statistic in Eq. (3.41) becomes:

$$S_l(\boldsymbol{y}) = \frac{1}{\sigma^2}(\hat{\boldsymbol{y}} - \boldsymbol{y})^{\mathrm{T}} \frac{\partial \hat{\boldsymbol{y}}}{\partial \theta_l} = -\frac{1}{\sigma^2}\boldsymbol{e}^{\mathrm{T}} \frac{\partial \hat{\boldsymbol{y}}}{\partial \theta_l}. \tag{3.51}$$

Using the fact that the noise on different data points are independently distributed, the Fisher Information Matrix in Eq. (3.40) becomes:

$$F_{l_1 l_2} = \frac{1}{\sigma^4}\left(\sum_i \frac{\partial \hat{y}_i}{\partial \theta_{l_1}} \frac{\partial \hat{y}_i}{\partial \theta_{l_2}} \mathrm{E}(e_i^2)\right) = \frac{1}{\sigma^2}\left(\sum_i \frac{\partial \hat{y}_i}{\partial \theta_{l_1}} \frac{\partial \hat{y}_i}{\partial \theta_{l_2}}\right), \tag{3.52}$$

The calculation of the Fisher information matrix through Eq. (3.52) is trivial if the model relating $\hat{\boldsymbol{y}}$ to $\boldsymbol{\theta}$ is known.

References

[1] G. Strang, Introduction to Linear Algebra, 4th Ed., Cambridge University Press, 2009.

[2] C. L. Lawson and R. J. Hanson, Solving Least Squares Problems, Prentice-Hall, 1974.

[3] R. Bro and S. de Jong, A fast non-negativity-constrained least squares algorithm, J. Chemom. 11 (1997), 393-401.

[4] MATLAB Release 2016b, The Mathworks, Inc., Natick, Massachusetts, United States.

[5] A. N. Tikhonov and V. Y. Arsenin, Solutions of Ill-Posed Problems, V. H. Winston and Sons, 1977.

[6] É. Chouzenoux, S. Moussaoui, J. Idier and F. Mariette, Efficient maximum entropy reconstruction of nuclear magnetic resonance T_1-T_2 spectra, IEEE Trans. Sign. Process. 58 (2010), 6040-6051.

[7] G. H. Golub and C. F. van Loan, Matrix Computations, 3rd Ed., The Johns Hopkins University Press, 1996.

[8] T. Hastie, R. Tibshirani and J. Friedman, The Elements of Statistical Learning, 2nd Ed., Springer, 2008.

[9] B. L. Phillips, A technique for the numerical solution of certain integral equations of the first kind, J. ACM 9 (1962), 84-97.

[10] S. Twomey, On the numerical solution of Fredholm integral equations of the first kind by the inversion of the linear system produced by quadrature, J. ACM 10 (1963), 97-101.

[11] P. C. Hansen, Rank-Deficient and Discrete Ill-Posed Problems, SIAM, 1997.

[12] J. P. Butler, J. A. Reeds and S. V. Dawson, Estimating solutions of first kind integral equations with nonnegative constraints and optimal smoothing, SIAM J. Numer. Anal. 18 (1981), 381-397.

[13] G. H. Golub, M. Heath and G, Wahba, Generalized cross-validation as a method for choosing a good ridge parameter, Technometrics 21 (1979), 215-223.

[14] H. Akaike, A new look at the statistical model identification, IEEE Trans. Autom. Control 19 (1974), 716-723.

[15]G. Schwarz, Estimating the dimension of a model, Ann. Statist. 6 (1978), 461-464.

[16]C. L. Mallows, Somme comments on C_p, Technometrics, 15 (1973), 661-675.

[17]P. Babak, S. Kryuchkov and A. Kantzas, Parsimony and goodness-of-fit in multi-dimensional NMR inversion, J. Magn. Reson. 274 (2017), 46-56.

[18]S. J. Provencher, A constrained regularization method for inverting data represented by linear algebraic or integral equations, Comput. Phys. Commun. 27 (1982), 213-227.

[19]S. J. Provencher, CONTIN: A general purpose constrained regularization program for inverting noisy linear algebraic and integral equations, Comput. Phys. Commun. 27 (1982), 229-242.

[20]G. C. Borgia, R. J. S. Brown and P. Fantazzini, Uniform-penalty inversion of multiexponential decay data, J. Magn. Reson. 122 (1998), 65-77.

[21]L. Venkataramanan, Y.-Q. Song and M. D. Hürlimann, Solving Fredholm integrals of the first kind with tensor product structure in 2 and 2.5 dimensions, IEEE Trans. Signal Process. 50 (2002), 1017-1026.

[22]P. D. Teal and C. Eccles, Adaptive truncation of matrix decompositions and efficient estimation of NMR relaxation distributions, Inverse Problems 31 (2015), 045010.

[23]V. Bortolotti, R. J. S. Brown, P. Fantazini, G. Landi and F. Zama, Uniform penalty inversion of two-dimensional NMR relaxation data, Inverse Problems 33 (2016), 015003.

[24]J. Mitchell, T. C. Chandrasekera, and L. F. Gladden, Numerical estimation of relaxation and diffusion distributions in two dimensions, Prog. Nucl. Magn. Reson. Spectrosc. 62 (2012), 34-50.

[25]H. Nyquist, Certain topics in telegraph transmission theory, Trans. A. I. E. E. 47 (1928), 617-644.

[26]C. E. Shannon, Communication in the presence of noise, Proc. I. R. E. 37 (1949), 10-21.

[27]E. J. Candés, J. Romberg and T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, IEEE Trans. Inf. Theory 52 (2006), 489-509.

[28]D. L. Donoho, Compressed Sensing, IEEE Trans. Inf. Theory, 52 (2006), 1289-1306.

[29]J. A. Högbom, Aperture synthesis with a non-regular distribution of interferometer baselines, Astron. Astrophys. Suppl. Ser. 15 (1974), 417-426.

[30]S. F. Gull and G. J. Daniell, Image reconstruction from incomplete and noisy data, Nature 272 (1978), 686-690.

[31]J. Bobin, J.-L. Starck and R. Ottensamer, Compressed Sensing in astronomy, IEEE J. Sel. Top. Signal Process. 2 (2008), 718-726.

[32]F. Maghalhães, F. M. Arújo, M. V. Correira, M. Abolbashari and F. Farahi, Active illumination single-pixel camera based on compressive sensing, Appl. Optics 4 (2011), 405-414.

[33]R. Leary, Z. Saghi, P. A. Midgley and D. J. Holland, Compressed sensing electron tomography, Ultramicroscopy 131 (2013), 70-91.

[34]M. F. Duarte and Y. C. Eldar, Structured compressed sensing: From theory to applications, IEEE Trans. Signal Process. 59 (2011), 4053-4085.

[35]M. Lustig, D. Donoho and J. M. Pauly, Sparse MRI: The application of Compressed Sensing for rapid MR imaging, Magn. Reson. Med. 58 (2007), 1182-1195.

[36]M. Benning, L. Gladden, D. Holland, C.-B. Schönlieb and T. Valkonen, Phase reconstruction from velocity-encoded MRI measurements – A survey of sparsity - promoting variational approaches, J. Magn. Reson. 238 (2014), 26-43.

[37]N. Gdaniec, H. Eggers, P. Börnert, M. Doneva and A. Mertins, Robust abdominal imaging with incomplete breath-holds, Magn. Reson. Med. 71 (2014), 1733-1742.

[38]K.-V. Nguyen, J.-R. Li, G. Radecki and L. Ciobanu, DLA based compressed sensing for high resolution MR microscopy of neuronal tissue, J. Magn. Reson. 259 (2015), 186-191.

[39]F. Knoll, C. Clason, C. Diwoky and R. Stollberger, Adapted random sampling patterns for accelerated MRI, MAGMA 24 (2011), 43-50.

[40]D. Xiao and B. J. Balcom, Restricted *k*-space sampling in pure phase encode MRI of rock core plugs, J. Magn. Reson. 231 (2013), 126-132.

[41]B. Adcock, A. C. Hansen, C. Poon and B. Roman, Breaking the coherence barrier: A new theory for compressed sensing, arXiv:1302.0561v4 (2014).

[42]L. I. Rudin, S. Osher and E. Fatemi, Nonlinear total variation based noise removal algorithms, Physica D 60 (1992), 259-268.

[43]M. Guerquin-Kern, M. Haberlin, K. Pruessmann and M. Unser, A fast wavelet-based reconstruction method for magnetic resonance imaging, IEEE Trans. Med. Imaging 30 (2011), 1649-1660.

[44]K. Bredies, K. Kunisch and T. Pock, Total generalized variation, SIAM J. Imaging Sci. 3 (2010), 492-526.

[45]S. Boyd and L. Vandenberghe, Convex Optimization, Cambridge University Press, 2004.

[46]P. L. Combettes and J.-C. Pesquet, Proximal splitting methods in signal processing, (2010) arXiv:0912.3522v4.

[47]N. Parikh and S. Boyd, Proximal algorithms, Foundations and Trends in Optimization, 1 (2013), 123-231.

[48]A. Beck and M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, SIAM J. Imaging Sci. 2 (2009), 183-202.

[49]E. Esser, X. Zhang and T. F. Chan, A general framework for a class of first order primal-dual algorithms for convex optimization in imaging science, SIAM J. Imaging Sci. 3 (2010), 1015-1046.

[50]A. Chambolle and T. Pock, A first-order primal-dual algorithm for convex problems with applications to imaging, J. Math. Imaging Vis. 40 (2011), 120-145.

[51]T. Goldstein and S. Osher, The split Bregman method for l_1-regularized problems, SIAM J. Imaging Sci. 2 (2013), 323-343.

[52]M. Benning, Inverse problems, Lecture notes, DAMTP University of Cambridge, 2015.

[53]H. Cramér, Mathematical Methods of Statistics, Princeton University Press, 1963.

[54]V. Pera, D. H. Brooks and M. Niedre, On the use of the Cramér-Rao lower bound for diffuse optical imaging system design, J. Biomed. Opt. 19 (2014), 025002.

[55]H. Godrich, A. M. Haimovich and R. S. Blum, Target localization accuracy gain in MIMO radar-based systems, IEEE Trans. Inf. Theory 56 (2010), 2783-2803.

[56] A. S. Wang and N. J. Pelc, Sufficient statistics as a generalization of binning in spectral X-ray imaging, IEEE Trans. Med. Imaging 30 (2011), 84-93.

[57] R. A. Fisher, The Design of Experiments, Oliver and Boyd, 1953.

Chapter 4

Obtaining sparse distributions in 2D NMR correlation experiments

4.1. Introduction

2D NMR correlation experiments are used in estimating joint distribution functions of relaxation time constants, T_1, T_2 and the diffusion coefficient, D. It was discussed in Section 2.12 that for 2D NMR correlation experiments, the acquired NMR signal is related to the joint distribution function through a Laplace transform. The conventional ILT methods of NNLS and Tikhonov regularization, which are commonly used in reconstructing the joint distribution function, were discussed in Section 3.1. The NNLS method suffers from instability to noise at high noise levels; slight variations in the signal (caused by the random nature of noise) can lead to vastly different distributions being reconstructed. Tikhonov regularization improves the stability to noise, but its reconstructions are often plagued by over smoothing [1-3]. The inherent assumption of Tikhonov regularization, imposed through the penalty term in the reconstruction process, is that the distribution is smooth. This is generally a good assumption in applications where the system under investigation is highly heterogeneous, such as in the study of rocks [4] or complex chemical systems [5]. However, in applications where the dataset under investigation is composed of sparse features, such as when studying the relaxation times or diffusion coefficients characterising simple liquid mixtures, Tikhonov regularization can lead to over smoothing. Over smoothing becomes an issue when the sparse features are very close to each other, in which case Tikhonov regularization may not be able to distinguish between the different features; Tikhonov regularization is typically unable to resolve relaxation time constants or diffusion coefficients which are a factor of < 3 apart [1]. Hence, Tikhonov regularization suffers from a poor spectral resolution.

The development of Compressed Sensing, outlined in Section 3.2, has introduced a new form of regularization which is based on an L_1 minimization problem. The principle aim

of Compressed Sensing is the reduction in the number of samples needed to reconstruct a distribution which is sparse in some domain. However, the form of regularization that it uses, which from now on will be referred to as L_1 regularization, also promotes the sparsity of the distribution in a particular domain [6]. For distributions which are expected to be sparse, the enforcement of sparsity is synonymous with an increase in the spectral resolution of the reconstructed distribution. In addition to the potential for higher spectral resolution, an L_1 regularized problem is also expected to be more stable to noise than the non-regularized version. Indeed, L_1 regularization has been used in denoising sparse signals [7]. The potential of L_1 regularization to increase spectral resolution, while being stable to high levels of noise, has only recently begun to be realised in NMR applications [8-11]. Indeed, in the recent work of Benjamini and Basser [10, 11], L_1 regularization was used as a tool to reduce the required amount of data required for 2D NMR relaxation-diffusion reconstructions but did not specifically focus on the capability of L_1 regularization in reconstructing sparse distributions and its stability to noise.

In this chapter, L_1 regularization is introduced for processing 2D NMR correlation experiments for distributions which are *a priori* known to be sparse, with the aim of improving the spectral resolution, as compared to Tikhonov regularization, while still being more stable to noise compared to the NNLS method. The PDHGM algorithm described in Section 3.3 is used to numerically solve the L_1 regularized problem. A guide to the implementation of the algorithm is given, including the choice of the regularization parameter and the assignment of error estimates. The proposed regularization method and algorithm are demonstrated by applying it to T_1-T_2 and D-T_2 experimental data and the results are compared against NNLS and Tikhonov regularization results. Samples, in which the separation of components is of increasing difficulty, are investigated; the most difficult being mixtures of hexane and dodecane imbibed within porous silica beads which remain immersed within the bulk liquid mixture, thereby forming a system in which the discrimination of both chemical species existing in the bulk liquid state and within the porous medium is required. In this system, hexane and dodecane cannot be resolved spectroscopically neither in the bulk phase nor imbibed within the porous silica beads. Further, the similarity of the relaxation time constants and the diffusion coefficients of hexane and dodecane make it impossible to resolve these two components using the NNLS or Tikhonov regularization methods. Resolving the individual diffusion coefficients and relaxation time constants in both the bulk liquid phase and within the porous medium, as

well as quantifying the relative amount of each component in both phases, is invaluable in understanding the behaviour of multiple chemical components inside catalysts [12] and rocks [13] *in situ*.

This chapter is structured as follows. The L_1 regularization method and the proposed algorithm are outlined in Section 4.2. In Section 4.3 the materials and methods used are described, and a practical guide to implementing the proposed algorithm is suggested. The experimental results and discussion are presented in Section 4.4.

4.2. Mathematical treatment

As discussed in Section 2.12, the vectorised NMR signal acquired from a 2D NMR correlation experiment, $\boldsymbol{Y}_\mathrm{v}$, can be written in terms of a kernel matrix, $\boldsymbol{K}$, the vectorised joint distribution, $\boldsymbol{U}_\mathrm{v}$, and the vectorised noise matrix, $\boldsymbol{E}_\mathrm{v}$, as:

$$\boldsymbol{Y}_\mathrm{v} = \boldsymbol{K}\,\boldsymbol{U}_\mathrm{v} + \boldsymbol{E}_\mathrm{v}\ . \tag{4.1}$$

The motivation for this chapter is to develop a reconstruction technique which can distinguish close features in sparse distributions, while being stable to high levels of noise. To obtain a sparse distribution $\boldsymbol{U}_\mathrm{v}$, it is proposed to solve the following minimization problem:

$$\boldsymbol{U}_\mathrm{v} = \arg\ \min_{\boldsymbol{U}_\mathrm{v}\geq\boldsymbol{0}} \left(\frac{1}{2}\|\boldsymbol{K}\,\boldsymbol{U}_\mathrm{v} - \boldsymbol{Y}_\mathrm{v}\|_2^2 + \alpha\|\boldsymbol{U}_\mathrm{v}\|_1\right), \tag{4.2}$$

where the L_1-norm in the penalty term is defined in Section 3.1.1. The problem in Eq. (4.2) will be known as an L_1 regularization problem. The fact that Eq. (4.2) is regularized will give more stability to noise, as compared to the NNLS algorithm; the use of the L_1 penalty term is also expected to give a sparser solution than the use of the L_2 penalty term in Tikhonov regularization.

The minimization problem in Eq. (4.2) is numerically challenging because of the non-differentiability of the L_1-norm. To solve Eq. (4.2), it is first noted that:

$$\alpha\|\boldsymbol{U}_\mathrm{v}\|_1 = \max_{\boldsymbol{v}}(\boldsymbol{v}^\mathrm{T}\boldsymbol{U}_\mathrm{v}) \qquad \text{s.t. } \max_i v_i \leq \alpha\ , \tag{4.3}$$

where $\boldsymbol{v}$ is a dual variable. With the aid of Eq. (4.3), the minimization problem in Eq. (4.2) is written as:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{v}) = \arg \min_{\boldsymbol{U}_{\mathrm{v}} \geq \mathbf{0}} \max_{\boldsymbol{v}} \left(\frac{1}{2} \|\boldsymbol{K}\, \boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \boldsymbol{v}^{\mathrm{T}} \boldsymbol{U}_{\mathrm{v}} - \chi_{\alpha}(\boldsymbol{v}) \right), \tag{4.4}$$

where $\chi_{\alpha}(\boldsymbol{v})$ is an indicator function defined as:

$$\chi_{\alpha}(\boldsymbol{v}) = \begin{cases} 0 & \|\boldsymbol{v}\|_{\infty} \leq \alpha \\ +\infty & \|\boldsymbol{v}\|_{\infty} > \alpha \end{cases} . \tag{4.5}$$

Eq. (4.4) is in the form of the primal-dual problem discussed in Section 3.3:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{v}) = \min_{\boldsymbol{U}_{\mathrm{v}}} \max_{\boldsymbol{v}} f(\boldsymbol{U}_{\mathrm{v}}) - g(\boldsymbol{v}) + \boldsymbol{V}^{\mathrm{T}} \boldsymbol{U}_{\mathrm{v}} , \tag{4.6}$$

where $f(\boldsymbol{U}_{\mathrm{v}}) = ½||\boldsymbol{K}\, \boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}||_2^2$ and $g(\boldsymbol{v}) = \chi_{\alpha}(\boldsymbol{v})$. The PDHGM iteration scheme, outlined in Section 3.3, is then used to numerically solve the primal-dual problem in Eq. (4.4):

$$\boldsymbol{U}_{\mathrm{v}}{}^{(k+1)} = (\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-1} \left(\boldsymbol{U}_{\mathrm{v}}{}^{(k)} - \tau \boldsymbol{v}^{(k)} \right) , \tag{4.7a}$$

$$\boldsymbol{v}^{(k+1)} = (\boldsymbol{I} + \sigma\, \partial g^{\mathrm{T}})^{-1} \left(\boldsymbol{v}^{(k)} + \sigma \left(2\boldsymbol{U}_{\mathrm{v}}{}^{(k+1)} - \boldsymbol{U}_{\mathrm{v}}{}^{(k)} \right) \right) , \tag{4.7b}$$

where $(\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-1}$ and $(\boldsymbol{I} + \sigma\, \partial g^{\mathrm{T}})^{-1}$ are known as proximal operators; the proximal operator was defined in Section 3.3.2. The algorithm parameters τ and σ control the convergence speed and stability of the algorithm. The proximal operator for the function $g(\boldsymbol{v})$ was computed in Section 3.3.3. The computation of the proximal operator for $f(\boldsymbol{u})$ follows the same method as in Section 3.3.3; it is easily shown that the proximal operator of $f(\boldsymbol{u})$, when applied to a vector $\boldsymbol{a}$, is given by:

$$(\boldsymbol{I} + \tau\, \partial f^{\mathrm{T}})^{-1}(\boldsymbol{a}) = (\boldsymbol{I} + \tau\, \boldsymbol{K}^{\mathrm{T}} \boldsymbol{K})^{-1} (\boldsymbol{a} + \tau\, \boldsymbol{K}^{\mathrm{T}} \boldsymbol{Y}_{\mathrm{v}}) . \tag{4.8}$$

A pseudocode for the numerical solution of Eq. (4.2) is given in Table 4.1.

Table 4.1. Pseudocode to numerically solve Eq. (4.2), based on the PDHGM method.

Step 1.	Choose algorithm parameters τ, σ and the regularization parameter, α.
Step 2.	Set the convergence tolerance, TOL.
Step 3.	Calculate $\boldsymbol{C} = (\boldsymbol{I} + \tau\, \boldsymbol{K}^{\mathrm{T}} \boldsymbol{K})^{-1}$.
Step 4.	Initialize $\boldsymbol{v}^{(0)} = \underline{\boldsymbol{v}}^{(0)} = \mathbf{0}$, $\boldsymbol{U}_{\mathrm{v}}{}^{(0)} = \underline{\boldsymbol{U}}_{\mathrm{v}}{}^{(0)} \neq \mathbf{0}$.
Step 5.	Initialize count number, $k = 0$ and convergence tracker, $\varepsilon^{(0)} = 1$.
Step 6.	**while** $\varepsilon^{(k)} >$ TOL **do**
	a. $\underline{\boldsymbol{U}}_{\mathrm{v}}{}^{(k+1)} \leftarrow \boldsymbol{U}_{\mathrm{v}}{}^{(k)} - \tau\, \boldsymbol{v}^{(k)}$
	b. $\boldsymbol{U}_{\mathrm{v}}{}^{(k+1)} \leftarrow \boldsymbol{C}\, (\underline{\boldsymbol{U}}_{\mathrm{v}}{}^{(k+1)} + \tau\, \boldsymbol{K}^{\mathrm{T}} \boldsymbol{Y}_{\mathrm{v}})$
	c. $\boldsymbol{U}_{\mathrm{v}}{}^{(k+1)} \leftarrow \max(0, \boldsymbol{U}_{\mathrm{v}}{}^{(k+1)})$

All operations in this line are element-wise.

d. $\underline{\boldsymbol{v}}^{(k+1)} \leftarrow \boldsymbol{v}^{(k)} + \sigma\,(2\boldsymbol{U}_{\mathrm{v}}^{(k+1)} - \boldsymbol{U}_{\mathrm{v}}^{(k)})$

e. $\boldsymbol{v}^{(k+1)} \leftarrow \underline{\boldsymbol{v}}^{(k+1)}/\max(1, \underline{\boldsymbol{v}}^{(k+1)}/\,\alpha)$

All operations in this line are element-wise.

f. $\varepsilon^{(k+1)} \leftarrow ||\boldsymbol{U}_{\mathrm{v}}^{(k+1)} - \boldsymbol{U}_{\mathrm{v}}^{(k)}||_2 \,/\, ||\boldsymbol{U}_{\mathrm{v}}^{(k)}||_2$

g. $k \leftarrow k + 1$

end while

4.3. Materials and methods

4.3.1. Materials

NMR measurements are conducted on several liquid samples in bulk and imbibed within a porous medium. The liquids used are: water deionised using PURELAB®, n-hexane (Sigma Aldrich, ≥ 99.7% purity) and n-dodecane (Sigma Aldrich, ≥ 99.8% purity). The porous medium used is meso-porous silica beads of approximately spherical shape, of diameter 3-4 mm and with an approximately unimodal pore size of 10 nm, provided by Johnson Matthey plc. Gadolinium (III) chloride hexahydrate ($GdCl3.6H_2O$) (Alfa Aesar, ≥ 99.8% REO purity) is used to reduce the relaxation time constant of water. An electronic scale (Precisa 205 A), with a precision of 0.1 mg, is used for gravimetric measurements. NMR concentration measurements are validated against Gas Chromatography (GC) measurements (Agilent Technologies 7890A). The calibration standard error on the GC concentration measurements is 20 ng μl^{-1} on a 1 μl sample. The GC measurements were performed with the help of the PhD student Tian Li of the Magnetic Resonance Research Centre, University of Cambridge.

Four pairs of samples are prepared for testing the different reconstruction techniques. These correspond to samples in which resolution of the chemicals is of increasing difficulty.

(i) Two solutions of $GdCl_3.6H_2O$ in water and of concentrations of 0.25 mM and 0.36 mM are prepared. From these solutions, two samples were prepared: sample A1 is simply the 0.25 mM solution; sample A2 is composed of the two solutions, physically separated, and with a volume ratio 1:1.

(ii) Two samples composed of bulk binary liquid mixtures of hexane and dodecane are prepared, with respective mass fractions of dodecane of 0.267 ± 0.003 (sample B1) and 0.532 ± 0.003 (sample B2).

(iii) Two samples composed of binary mixtures of hexane and dodecane imbibed within silica beads are prepared as follows. A binary bulk mixture of known mass fraction is initially prepared, and the silica beads are immersed in the mixture. After 24 h, the beads are taken out and the draining liquid removed. The composition of the liquid imbibed within the beads is calculated by measuring the composition of the remaining bulk liquid using GC and performing a mass balance. The dodecane mass fractions of the liquid imbibed within the beads is calculated to be 0.20 ± 0.02 (sample C1) and 0.45 ± 0.05 (sample C2), respectively.

(iv) Two samples composed of mixtures of hexane and dodecane residing both inside the silica beads (intra particle) and between the silica beads (inter particle) are prepared by immersing the silica beads into each of two bulk binary mixtures of hexane and dodecane; the two bulk liquid mixtures being characterised by dodecane mass fractions of 0.246 ± 0.003 (sample E1) and 0.516 ± 0.003 (sample E2), respectively. The NMR measurements are then performed on the immersed beads.

4.3.2. NMR acquisitions

All experiments are conducted on a Bruker DMX300 magnet, operating at a resonant frequency of 300.13 MHz for ^{1}H observation. The maximum magnetic field gradient amplitude available on this magnet is 1176 G cm^{-1} and the r.f. coil is of 5 mm diameter. The experiments are performed at a temperature of 20 ± 1 °C.

1D NMR relaxation and diffusion experiments are performed on the individual chemicals used and on pure single-component hexane and dodecane imbibed within the silica beads, as follows. The CPMG and PROJECT pulse sequences, outlined in Section 2.10, are used to measure T_2; the IR pulse sequence, outlined in Section 2.9, is used to measure T_1; the Pulsed Gradient Stimulated Echo (PGSTE) pulse sequence, described in Sections 2.8 and 2.11, is used to measure D. The echo time in the CPMG and PROJECT sequences is 1 ms and the maximum number of echoes is such that the maximum time observed in the T_2 decay is ~ $5T_2$. For the IR pulse sequence, the maximum time observed in the T_1 decay is ~ $5T_1$. The diffusion time in the PGSTE sequence is Δ = 5 ms, the gradient pulse duration is δ = 1 ms and the maximum gradient strength is such that an attenuation of the signal to

~ 0.5% of the maximum value is achieved in the Stejskal-Tanner decay. The results from these 1D experiments are summarized in Table 4.2. The implementation of the PROJECT pulse sequence was done with the help of the PhD student Chris Robertson of the Magnetic Resonance Research Centre, University of Cambridge. The data from the 1D experiments are used as benchmarks for the 2D NMR correlation experiments, which are described below.

Table 4.2. The relaxation time constants, T_1, T_2 and diffusion coefficient, D, of the chemicals used in the experiments, obtained using 1D NMR relaxation and diffusion experiments. The quoted uncertainty is the 95% confidence interval of the fitting procedure.

Material	T_1 / s	T_2 / s	D ($\times 10^{-9}$) / $m^2\ s^{-1}$
0.25 mM $GdCl_3.6H_2O$ solution	0.150±0.005[a]	0.120±0.004[b]	2.00±0.03[d]
0.36 mM $GdCl_3.6H_2O$ solution	0.080±0.004[a]	0.080±0.003[b]	2.00±0.04[d]
pure bulk hexane	1.90±0.01[a]	1.60±0.01[c]	3.80±0.05[d]
pure bulk dodecane	1.00±0.02[a]	0.88±0.03[c]	0.76±0.02[d]
pure hexane imbibed in silica beads	1.30±0.06[a]	0.60±0.05[c]	2.50±0.06[d]
pure dodecane imbibed in silica beads	0.74±0.03[a]	0.48±0.02[c]	0.43±0.02[d]

[a]obtained from an IR pulse sequence.

[b]obtained from a CPMG pulse sequence.

[c]obtained from a PROJECT pulse sequence.

[d]obtained from a PGSTE pulse sequence.

T_1-T_2 and D-T_2 correlation experiments are performed on the samples described in Section 4.3.1. The pulse sequence used for the T_1-T_2 experiments of the water solutions is the IR-CPMG pulse sequence, outlined in Section 2.12. The pulse sequence used for the D-T_2 experiments of the water solutions is a PGSTE-CPMG sequence, described in Section 2.12. For the mixtures of hexane and dodecane, the CPMG sequence is substituted with the PROJECT sequence. The NMR acquisition parameters for the 2D NMR correlation experiments are kept the same as the acquisition parameters used for the 1D NMR experiments.

The number of data points used for the IR and PGSTE sequences embedded in the 2D NMR correlation experiments of T_1-T_2 and D-T_2 is 16, apart from the experiment to characterise liquid mixtures present in both the inter and intra particle space, i.e. samples E1 and E2, in which case 32 data points are used. The data points for the IR, PGSTE,

CPMG and PROJECT sequences are spaced equidistantly for all experiments, apart from the experiments on samples E1 and E2. In the case of samples E1 and E2, four components are expected in the distribution maps. In these systems, it is expected that the two intra particle relaxation time constants and diffusion coefficients are similar to each other, but relatively different to the two inter particle values. In turn, it is expected that the two inter particle measurements of the relaxation time constants and diffusion coefficients to be similar. For this case, the equidistant sampling scheme did not offer enough resolution. In particular, in the case in which 4 diffusion coefficients need to be identified, where $D_1 > D_2 \gg D_3 > D_4$, the equidistant sampling scheme was capable of resolving the two pairs (D_1, D_2) and (D_3, D_4) but was not capable of resolving within the pairs, i.e. D_1 from D_2, and D_3 from D_4. To get a better resolution of all four diffusion coefficients, the following sampling density of points, $\rho(b)$, is used:

$$\rho(b) \propto (e^{-bD_4} - e^{-bD_3}) + (e^{-bD_2} - e^{-bD_1}) , \tag{4.9}$$

where $b = \gamma^2 g^2 \delta^2 (\Delta - \delta/3)$. This sampling scheme effectively allows denser sampling in the regions where there is a larger difference between the signals acquired from components of similar diffusion coefficients. A more systematic approach to optimising sampling schemes is outlined in Chapters 6 and 7.

4.3.3. Implementation of the L_1 regularization algorithm

Before any of the NNLS, Tikhonov regularization and L_1 regularization methods are applied, the data are compressed according to the method outlined in Section 3.1.3. All distributions are reconstructed on a 32 × 32 grid, with equidistantly spaced points. This makes the quantification of the amounts of each component straightforward. Simulations showed that the relative performance of the NNLS method, Tikhonov and L_1 regularization with equidistantly spaced points in the distribution is the same as in the case when logarithmically spaced points are used, which is more common in the literature.

As discussed in Section 3.3.4, the convergence of the algorithm in Table 4.1 is guaranteed if the following condition is met:

$$\tau\sigma \leq 1 . \tag{4.10}$$

However, the particular choice of τ and σ is heuristic. A smaller τ increases the stability while reducing the convergence speed of the algorithm. A good compromise between stability and convergence speed is found when $\tau = 0.1$ and $\sigma = 10$. The best values of τ and

σ depend slightly on the scaling of the signal. To avoid this, it is best to scale the NMR signal to a maximum of 1, a technique which is followed in this study. A typical number of iterations needed for the algorithm to converge is ~10,000. The time in which this convergence is achieved for a 32 × 32 distribution map, for example a T_1-T_2 map, depends on the processing speed of the computer; with a 2.0 GHz CPU and 16 GB RAM, 10,000 iterations take ~17 s. The initialization of the algorithm has no impact on the reconstructed distribution because the function being minimized is convex. Practically, it is also observed that the convergence speed is unaffected.

4.3.4. Choice of the regularization parameter

The regularization parameter, α, in the L_1 regularization problem is chosen using a variant of GCV, which for the case of Tikhonov regularization was described in Section 3.1.2. An analytical expression for the GCV score does not exist in the case of L_1 regularization, in contrast to Tikhonov regularization. However, by writing Eq. (4.2) as:

$$\boldsymbol{U}_{\mathrm{v}} = \arg\min_{\boldsymbol{U}_{\mathrm{v}} \geq \boldsymbol{0}} \left(\frac{1}{2} \|\boldsymbol{K}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \alpha \sum_i \frac{U_{\mathrm{v},i}^2}{|U_{\mathrm{v},i}|} \right), \tag{4.11}$$

which looks similar to Tikhonov regularization in Section 3.1.2, it is possible to approximate the degrees of freedom of an L_1 regularization problem by [14, 15]:

$$\mathrm{df}(\alpha) = \mathrm{tr}(\boldsymbol{K}(\boldsymbol{K}^{\mathrm{T}}\boldsymbol{K} + \alpha\,\boldsymbol{Z}^{-})^{-1}\boldsymbol{K}^{\mathrm{T}}), \tag{4.12}$$

where $\boldsymbol{Z}^{-}$ is the generalized inverse of $\boldsymbol{Z}$, with $\boldsymbol{Z}$ defined as $\mathrm{diag}(|U_{\mathrm{v},i}|)$. This value of the degrees of freedom is used to calculate the GCV score, defined in Section 3.1.2. The best value of α is again found at the minimum of the GCV score. More accurate estimations of the degrees of freedom in the L_1 regularization problem have been recently devised [16-18]. However, they are more difficult to estimate because of the requirement for some prior knowledge of the distribution being estimated.

Other methods of choosing the regularization parameter for the L_1 regularization problem were considered; in particular, the adaptation of the L-curve method [19] in Tikhonov regularization to the L_1 regularization problem, known as the Pareto curve [20], and the Morozov Discrepancy Principle [21]. However, simulations showed that the Pareto curve is unsuitable for T_1-T_2 and D-T_2 experiments because there was no clear maximum in the curvature of the curve, while the Morozov discrepancy principle was very sensitive to the estimate of the noise level in the acquired signal. Therefore, these two methods are not used in this chapter in choosing the regularization parameter.

4.3.5. Estimate of uncertainties

The uncertainty in the magnitude of each peak in the distribution is estimated using an adaptation of the CRLB theory, which was outlined in Section 3.4. In overview, the CRLB theory states that the uncertainty in the estimation of each element in the distribution, $U_{\mathrm{v},i}$, has a lower bound, related to the Fisher Information Matrix, $\boldsymbol{F}$, by the expression:

$$\mathrm{var}(U_{\mathrm{v},i}) \geq (\boldsymbol{F}^{-1})_{ii} \,. \tag{4.13}$$

Given the model in Eq. (4.1) and assuming Gaussian noise of variance σ^2, it can be shown that the Fisher Information Matrix simplifies to $\boldsymbol{F} = 1/\sigma^2\, \boldsymbol{K}^{\mathrm{T}}\boldsymbol{K}$. However, since $\boldsymbol{K}$ is typically characterised by a large condition number, estimating the inverse of $\boldsymbol{F}$ in Eq. (4.13) is unstable. Therefore, it is impractical to assign uncertainty estimates for all the elements of $\boldsymbol{U}_{\mathrm{v}}$. However, it is possible to assign uncertainty estimates for the non-zero elements of $\boldsymbol{U}_{\mathrm{v}}$ (the peaks of the distribution) using the following method. If it is assumed that the distribution has k discrete peaks whose indices are *a priori* known to be in a set Ω (such that if $i \in \Omega$, $U_{\mathrm{v},i} > 0$ and if $i \notin \Omega$, $U_{\mathrm{v},i} = 0$), then $\boldsymbol{F} = 1/\sigma^2\, \boldsymbol{K}_\Omega{}^{\mathrm{T}}\boldsymbol{K}_\Omega$, where $\boldsymbol{K}_\Omega$ is the kernel matrix $\boldsymbol{K}$ with only the columns whose indices are in Ω. If the distribution has only a few peaks, the inverse of $\boldsymbol{F}$ is well-posed. The estimated uncertainties using this approach, termed the oracle estimates [22], are low because of the assumption of having prior knowledge of the position of the peaks. To correct for this assumption, it has been proved [23-25] that the true estimate of the uncertainty of any of the parameters is related to the oracle estimate, with a high probability, by only a constant factor:

$$\mathrm{var}(U_{\mathrm{v},i}) \approx c\, k \ln(n)\, \sigma^2 \left(\left(\boldsymbol{K}_\Omega{}^{\mathrm{T}} \boldsymbol{K}_\Omega \right)^{-1} \right)_{ii} , \tag{4.14}$$

where c is a constant which depends on the application and the scaling of the signal and, n is the size of the distribution. For T_1-T_2 and D-T_2 experiments, with the NMR signal being scaled to a maximum of 1, simulations over a range of distributions, distribution sizes and sparsity levels show that $c = 0.25$ gives reliable estimates for the uncertainties.

4.4. Results and discussion

Initially, experimental results on a single component sample are used to investigate the stability to noise of the L_1 regularization method, as compared to the NNLS method and Tikhonov regularization. Then, experimental results of bulk binary liquid mixtures and

binary liquid mixtures imbibed within the porous beads system are used to compare the spectral resolution of L_1 regularization with the NNLS method and Tikhonov regularization. Finally, the L_1 regularization method is used to process 2D NMR correlation experiments of mixtures of hexane and dodecane in different physical environments (i.e., inter and intra particle space) where the NNLS and Tikhonov regularization methods fail to distinguish between the different components.

4.4.1. Stability to noise

To test the stability of the different reconstruction techniques to noise, experiments are performed on a single component sample and increasing levels of Gaussian noise are added. T_1-T_2 and D-T_2 experiments are performed on sample A1, the 0.25 mM $GdCl_3.6H_2O$ solution. From Table 4.2, the distributions are expected to have one peak at $T_1 = 0.150 \pm 0.005$ s, $T_2 = 0.120 \pm 0.004$ s and $D = (2.00 \pm 0.03) \times 10^{-9}$ m^2 s^{-1}. Random Gaussian noise of increasing standard deviation is added to the 2D NMR signal, such that a range of signal-to-noise ratios (SNR) from 5000 to 5 is studied. The 2D distributions are then reconstructed by applying the NNLS method, Tikhonov regularization and L_1 regularization.

For SNR > 20, no difference is seen between the reconstructions from the different methods. For this reason, they are not shown here. For SNR < 20, the NNLS method fails to reconstruct the distribution accurately, incorrectly predicting more than one peak in the distribution. In contrast, Tikhonov regularization and L_1 regularization correctly predict the distributions at all SNR studied. Fig. 4.1 shows the reconstructed T_1-T_2 and D-T_2 distributions from the NNLS method, Tikhonov regularization and L_1 regularization at SNR = 20. The NNLS method incorrectly predicts two peaks in both T_1-T_2 and D-T_2 distributions, while Tikhonov regularization and L_1 regularization correctly predict the position of the single peak. Although Tikhonov regularization is observed to give broader distributions than L_1 regularization, the reconstructed distributions from Tikhonov and L_1 regularization are very similar.

It is concluded that L_1 regularization is more stable to noise than the NNLS method and comparable in stability to Tikhonov regularization. The stability to noise offered by the L_1 regularization method can be used to reduce the extent of signal averaging that is needed to achieve sufficient SNR for subsequent analysis. Since SNR $\propto B_0^{7/4}$ [26], where B_0 is the

permanent magnetic field strength, this method could be useful in extending measurement capabilities at low magnetic fields.

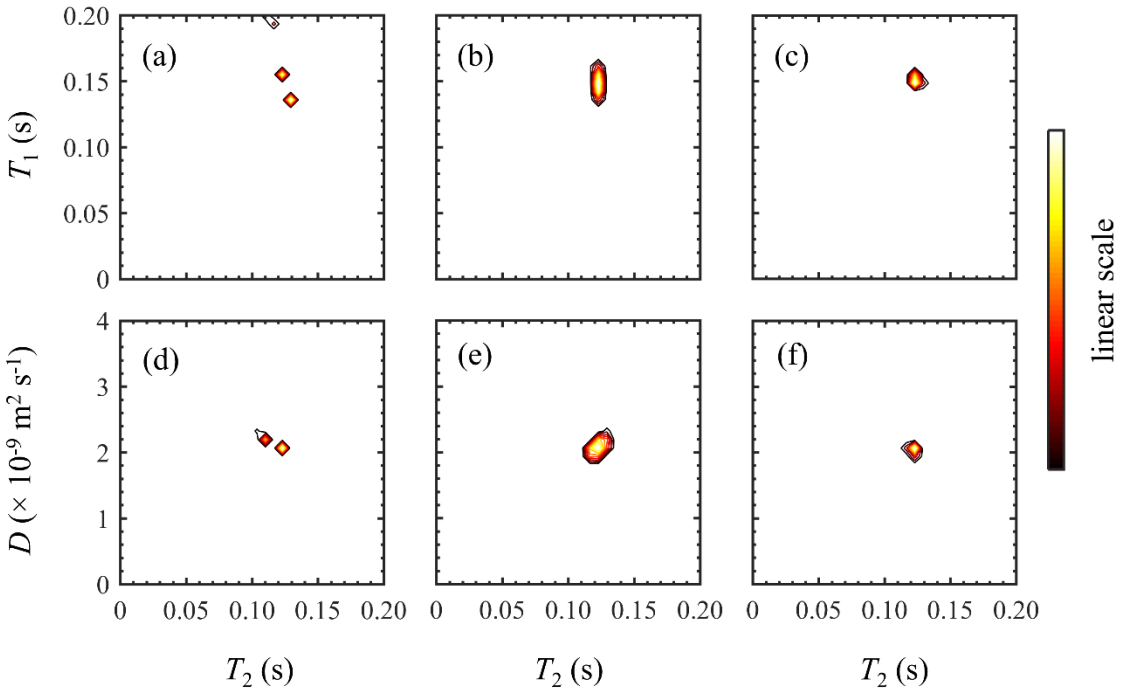

Fig. 4.1. (a, b, c) T_1-T_2 and (d, e, f) D-T_2 distributions of sample A1, a 0.25 mM $GdCl_3.6H_2O$ deionised water solution, reconstructed for SNR = 20 using (a, d) the NNLS method, (b, e) Tikhonov regularization, and (c, f) L_1 regularization.

4.4.2. Spectral resolution

To compare the spectral resolution of the different reconstruction techniques, experiments are performed on a two-component sample, in which the resolution of the signal from the individual components is expected to be challenging using Tikhonov regularization. T_1-T_2 and D-T_2 experiments are performed on sample A2, composed of two physically separate $GdCl_3.6H_2O$ solutions of concentrations 0.25 and 0.36 mM and volume ratio 1:1. From Table 4.2, it is expected that the distribution maps will show two peaks: one at $T_1 = 0.150 \pm 0.005$ s, $T_2 = 0.120 \pm 0.004$ s and $D = (2.00 \pm 0.03) \times 10^{-9}$ m^2 s^{-1} and the other at $T_1 = 0.080 \pm 0.004$ s, $T_2 = 0.080 \pm 0.003$ s and $D = (2.00 \pm 0.04) \times 10^{-9}$ m^2 s^{-1}.

Fig. 4.2 shows the reconstructed T_1-T_2 and D-T_2 distributions obtained using the NNLS method, Tikhonov regularization and L_1 regularization. It is observed that Tikhonov regularization cannot distinguish between the two solutions; the peaks associated with each of the two solutions are included in a single broad peak. This is to be expected since

the resolution of Tikhonov regularization is typically limited to a factor of 3 in the relaxation time constants or the diffusion coefficients [1]. In contrast, L_1 regularization and the NNLS method correctly predict two peaks. The position of these peaks is in excellent agreement with those given in Table 4.2 for each of the component liquids. In addition, the ratio of the amounts of each solution present, measured by the ratio of the areas of each peak in the distributions is 1.0 (to 2 s. f.), for both T_1-T_2 and D-T_2 distributions reconstructed from L_1 regularization and the NNLS method. This is consistent with the 1:1 volume ratio that the sample is prepared at. The SNR in the acquired T_1-T_2 and D-T_2 signals is ~3000 and 1900, respectively, which explains the good reconstruction from the NNLS method.

These results show that L_1 regularization has the potential to be used to resolve features in distributions which were previously unable to be resolved by conventional methods, such as Tikhonov regularization.

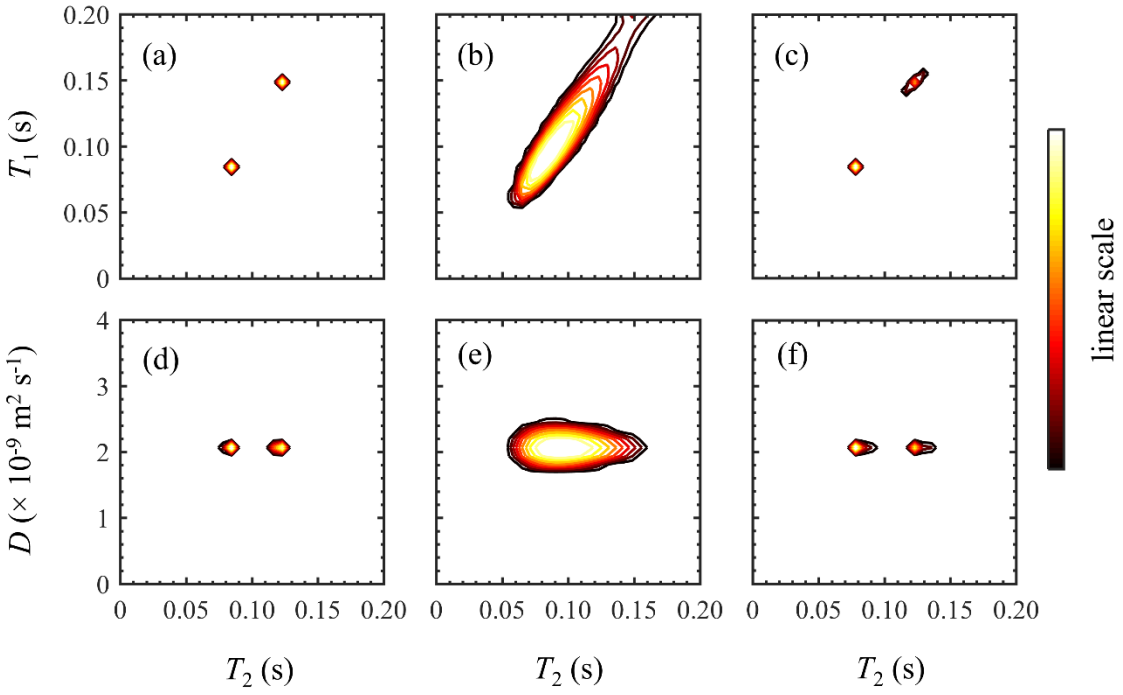

Fig. 4.2. (a, b, c) T_1-T_2 and (d, e, f) D-T_2 distributions of sample A2, composed of physically separate 0.25 mM and 0.36 mM $GdCl_3.6H_2O$ deionised water solutions, reconstructed using (a, d) the NNLS method, (b, e) Tikhonov regularization, and (c, f) L_1 regularization. The SNR in the NMR signal is ~3000 for the T_1-T_2 experiment and ~1900 for the D-T_2 experiment. The component with the larger T_2 is the 0.25 mM solution.

4.4.3. Resolving individual components of mixtures: hexane and dodecane in bulk liquid mixtures and within porous media

In this section, the different reconstruction techniques are applied to samples for which both the NNLS and Tikhonov regularization fail to predict the distributions, while L_1 regularization succeeds. T_1-T_2 and D-T_2 experiments are performed on bulk liquid mixtures of hexane and dodecane (samples B1 and B2) and mixtures of hexane and dodecane imbibed within silica beads (samples C1 and C2). Samples at different mass fractions of dodecane are included to demonstrate the quantitative nature of the method at different concentrations. It is well known that the relaxation time constants and diffusion coefficients of alkanes in a mixture become more similar relative to the pure single component values [27, 28]; the relaxation time constants and diffusion coefficient of an alkane in a mixture depend not only on the chain length of the alkane, but also on the mean chain length of the mixture. Therefore, it is expected that the peaks in the T_1-T_2 and D-T_2 distributions corresponding to hexane and dodecane are closer to one another than the values in Table 4.2, making it more difficult to resolve these chemicals.

Fig. 4.3 shows the T_1-T_2 distributions of the bulk mixtures of hexane and dodecane and of the mixtures of hexane and dodecane imbibed within silica beads, reconstructed using L_1 regularization. Figs. 4.3(a) and 4.3(b) show the T_1-T_2 distributions of two different bulk mixtures, with corresponding dodecane mass fraction of 0.267 ± 0.003 (sample B1) and 0.532 ± 0.003 (sample B2). Two peaks are clearly observed in both distributions: hexane at $T_1 = 1.5 \pm 0.1$ s, $T_2 = 1.4 \pm 0.1$ s and dodecane at $T_1 = 1.4 \pm 0.1$ s, $T_2 = 1.2 \pm 0.1$ s. Figures 4.3(c) and 4.3(d) show the T_1-T_2 distributions of two different mixtures of hexane and dodecane imbibed within silica beads, with corresponding dodecane mass fractions of 0.20 ± 0.02 (sample C1) and 0.45 ± 0.05 (sample C2). Two peaks are clearly identified in these distributions: intra particle hexane at $T_1 = 1.0 \pm 0.1$ s, $T_2 = 0.4 \pm 0.04$ s and intra particle dodecane at $T_2 = 0.10 \pm 0.02$ s. The relaxation time constant T_1 of intra particle dodecane is $T_1 = 0.94 \pm 0.04$ s for sample C1 and $T_1 = 0.78 \pm 0.04$ s for sample C2. The relaxation time constants T_2 of the intra particle liquids are significantly shorter than the corresponding values in the bulk liquid mixture due to the additional dephasing of nuclear spins caused by the magnetic heterogeneity in the porous medium and the interaction with the pore surface [29].

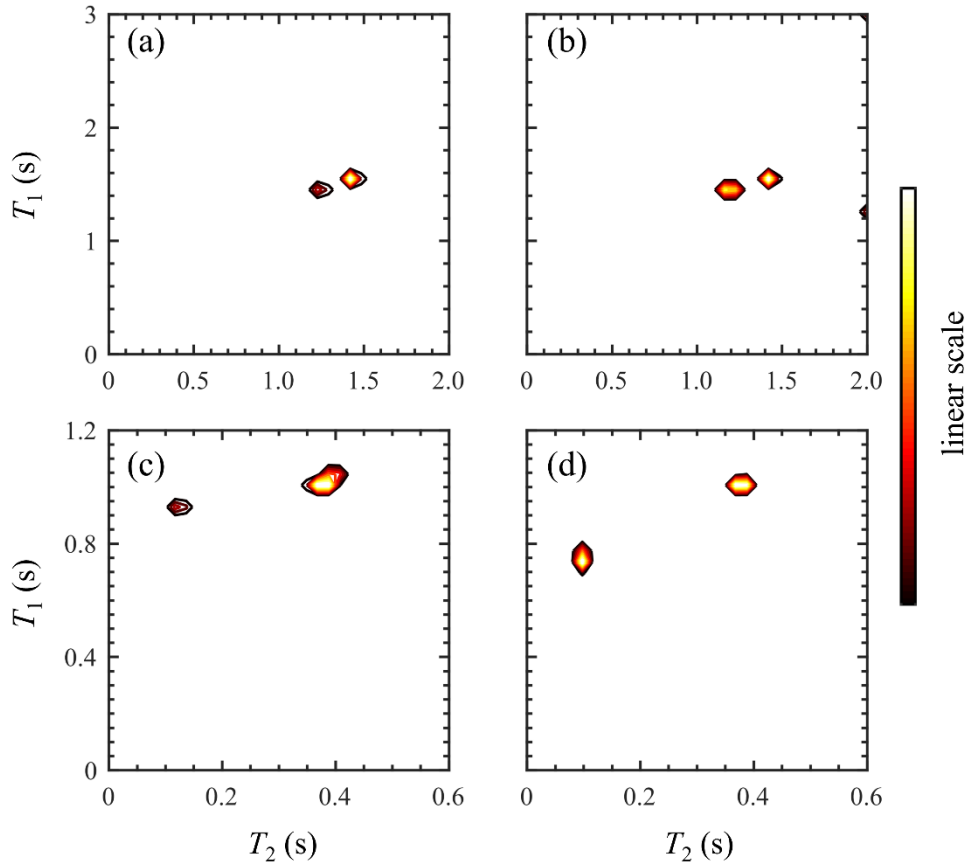

Fig. 4.3. T_1-T_2 distributions of mixtures of hexane and dodecane (a, b) in bulk and (c, d) imbibed within silica beads, reconstructed using the L_1 regularization method. The samples shown are (a) sample B1, (b) sample B2, (c) sample C1 and (d) sample C2. The component with the larger T_2 is hexane. The SNR in the NMR signal is ~2500 in all the experiments.

Fig. 4.4 shows the corresponding D-T_2 distributions for the same samples discussed in Fig. 4.3, reconstructed using L_1 regularization. Figs. 4.4(a) and 4.4(b) show the D-T_2 distributions of samples B1 and B2, while Figs. 4.4(c) and 4.4(d) show the D-T_2 distributions of samples C1 and C2. Hexane and dodecane are clearly identified in all these distributions. The relaxation time constants T_2 obtained using the D-T_2 maps differ by < 10% from the corresponding relaxation time constants obtained using the T_1-T_2 maps.

The results in Figs. 4.3 and 4.4 show that the L_1 regularization method can resolve components in T_1-T_2 and D-T_2 distribution maps even if their difference in relaxation constants is as low as 10%.

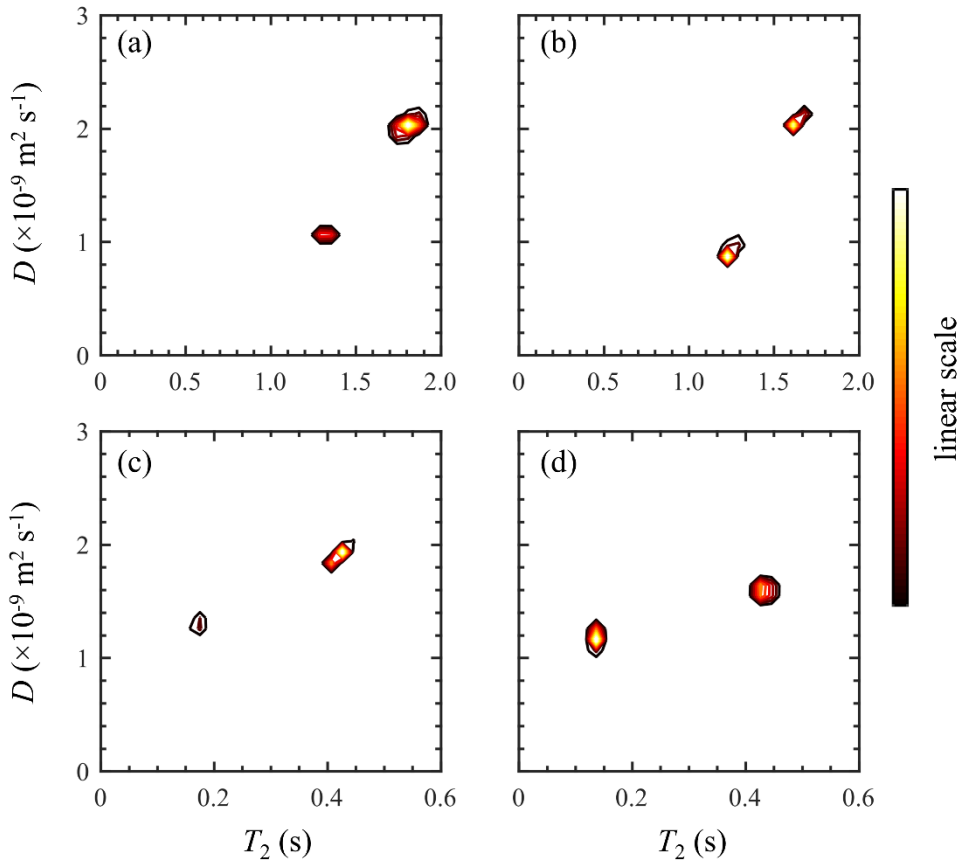

Fig. 4.4. D-T_2 distributions of mixtures of hexane and dodecane (a, b) in bulk and (c, d) imbibed within silica beads, reconstructed using the L_1 regularization method. The samples shown are (a) sample B1, (b) sample B2, (c) sample C1 and (d) sample C2. The component with the larger T_2 is hexane. The SNR in the NMR signal is ~2000 in all the experiments.

Both the NNLS method and Tikhonov regularization failed to reconstruct the T_1-T_2 and D-T_2 distributions correctly. One example of this is shown for sample C2 in Fig. 4.5. Sample C2 is chosen as it represents one of the cases in which, relative to the other samples, it should be easier for the NNLS and Tikhonov regularization methods to predict the correct distributions because the components have a difference in the expected T_2 of a factor of 4. Fig. 4.5(a) shows the D-T_2 distribution of sample C2, reconstructed using the NNLS method. Three peaks are predicted from the method, when only two are expected. The prediction of extra peaks is a typical artefact of the NNLS method. Fig. 4.5(b) shows the D-T_2 distribution of sample C2, reconstructed using the Tikhonov regularization method. The method predicts a single broad peak (extending over the range of the two expected peaks). This is a typical artefact of the Tikhonov regularization method.

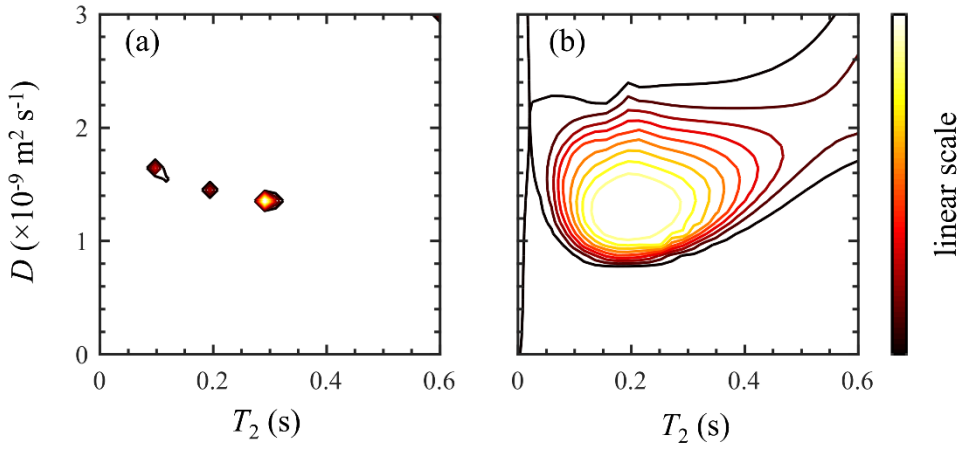

Fig. 4.5. D-T_2 distribution maps of sample C2, a mixture of hexane and dodecane imbibed within silica beads, reconstructed using (a) the NNLS and (b) Tikhonov regularization methods. The SNR in the NMR signal is ~2000.

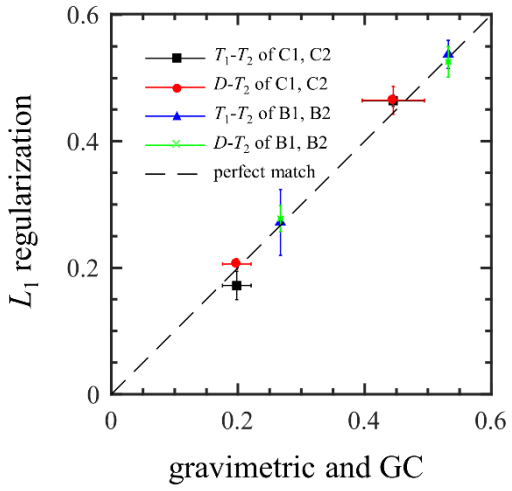

Fig. 4.6. Comparison of the dodecane mass fraction measured using L_1 regularization in the reconstruction of joint parameter distributions with the mass fraction measured using the gravimetric and GC methods described in Section 4.3.1.

A quantitative comparison of the mass fraction of dodecane measured from the T_1-T_2 and D-T_2 distributions in Figs. 4.3 and 4.4 against the mass fraction of dodecane measured using gravimetric and GC methods is shown in Fig. 4.6. The first observation is that the mass fractions of dodecane measured using the T_1-T_2 distributions has an average difference of 0.02 from the corresponding mass fractions of dodecane measured using the D-T_2 distributions. The second observation is that each of these measures has an average difference of 0.01 from GC and gravimetric measurements. The results presented in Fig. 4.6 show that, in addition to being able to resolve very close features, L_1 regularization

can be used quantitatively to obtain information on the amounts of each component present.

4.4.4. Resolving inter and intra particle hexane and dodecane compositions

In this section, L_1 regularization is applied to determine the diffusion coefficients, D, the relaxation time constants, T_2, and the chemical composition of hexane and dodecane mixtures imbibed within porous silica beads which remain immersed in the bulk liquid mixture. D-T_2 experiments are performed on samples E1 and E2. Four peaks are expected in the D-T_2 distributions, corresponding to intra particle dodecane, intra particle hexane, inter particle dodecane and inter particle hexane.

Fig. 4.7 shows the D-T_2 distributions of samples E1 and E2 reconstructed using L_1 regularization. Four components are clearly observed in each of the distributions. The values of the relaxation time constant, T_2, and the diffusion coefficient, D, extracted from these distributions for each of the components are given in Table 4.3. Although there is less than 10% difference between the relaxation time constants and diffusion coefficients of the inter particle or the intra particle components, L_1 regularization is capable of resolving the different components.

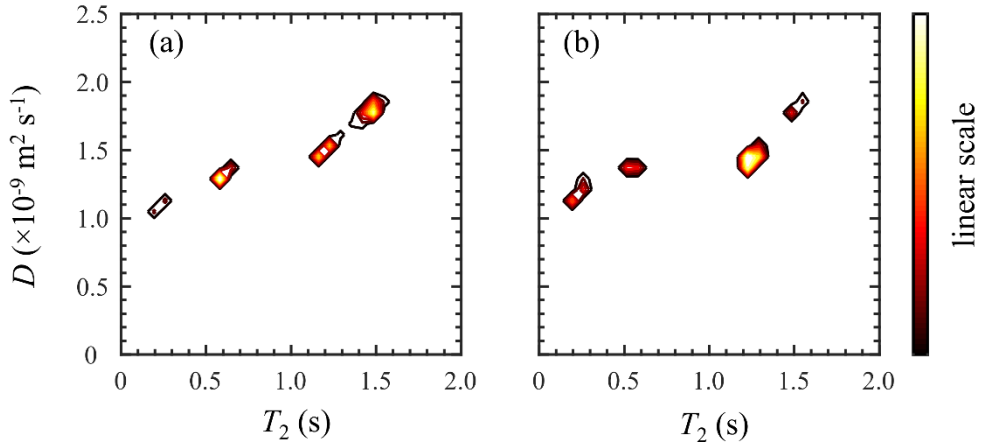

Fig. 4.7. D-T_2 distribution of (a) sample E1 and (b) sample E2, which are liquid mixtures of hexane and dodecane in silica beads with inter and intra particle liquid present. The maps are reconstructed using L_1 regularization. In an order of increasing T_2, the peaks correspond to intra particle dodecane, intra particle hexane, inter particle dodecane and inter particle hexane. The SNR in the NMR signal is ~1500 in both experiments.

The D-T_2 distributions are now used to estimate the total mass fraction of dodecane in each of the samples. For sample E1, the total dodecane mass fraction is estimated to be 0.22 ± 0.04; for sample E2, it is estimated to be 0.53 ± 0.02. These mass fractions differ

by less than 0.03 from the measurements using the gravimetric technique, demonstrating that L_1 regularization can be used quantitatively.

Table 4.3. The relaxation time constants, T_2, and the diffusion coefficients, D, of the components in samples E1 and E2, extracted from a D-T_2 correlation experiment processed with L_1 regularization. The D-T_2 distributions are shown in Fig. 4.7.

component	sample E1		sample E2	
	T_2 / s	D ($\times 10^{-9}$) / m^2 s^{-1}	T_2 / s	D ($\times 10^{-9}$) / m^2 s^{-1}
inter particle hexane	1.4±0.1	1.7±0.2	1.5±0.1	1.8±0.1
inter particle dodecane	1.2±0.1	1.5±0.1	1.3±0.1	1.5±0.1
intra particle hexane	0.6±0.1	1.3±0.1	0.6±0.1	1.4±0.1
intra particle dodecane	0.2±0.1	1.1±0.1	0.2±0.1	1.2±0.1

More importantly, the maps in Fig. 4.7 can be used to obtain the inter and intra particle composition of the liquid. The dodecane mass fraction for the inter particle liquid in samples E1 and E2 is 0.30 ± 0.02 and 0.59 ± 0.02, respectively. The dodecane mass fraction for the intra particle liquid in samples E1 and E2 is 0.19 ± 0.03 and 0.43 ± 0.02, respectively. The ability to obtain the difference in the concentration between the inter and intra particle liquid is invaluable in, for example, identifying mass transfer limitations in heterogeneous catalysts.

4.4.5. Potential and limitations

It was shown in Sections 4.4.1-4.4.4 that the proposed L_1 regularization method is able to reconstruct sparse T_1-T_2 and D-T_2 distribution, even if the components of the maps differ in relaxation time constants or diffusion coefficients by less than 10% or if SNR < 20. These are considerable improvements over the capabilities of the NNLS and Tikhonov regularization methods. Although L_1 regularization was applied to T_1-T_2 and D-T_2 experiments, it can be applied in a similar way to other 2D NMR correlation experiments, such as T_2-T_2 or D-D experiments. The only modification needed is the kernel function, $\boldsymbol{K}$, which is specific to a given application. Of particular interest could be the application to D-δ, or Diffusion Ordered Spectroscopy (DOSY) [30] experiments, where δ stands for the chemical shift.

L_1 regularization is recommended only for cases in which there is a prior knowledge that the distribution is sparse. If the true distribution is smooth, the algorithm can give

inaccurate reconstructions. When applying the L_1 regularization method, the typical indication that the true distribution might be smooth is if the reconstructed distribution shows a cluster of peaks within a narrow region.

L_1 regularization application is directly applicable to 1D relaxation and diffusion experimental data; the only modification needed is the kernel function, $\boldsymbol{K}$. However, the capability of L_1 regularization to resolve close sparse features is expected to be better when applied to 2D rather than 1D experimental data because higher dimensional data are sparser. In Chapter 5, L_1 regularization is compared to Tikhonov regularization and a new proposed regularization technique over a range of discrete and smooth, 1D and 2D simulated distributions.

4.5. Conclusions

L_1 regularization was proposed for the reconstruction of sparse distributions from 2D NMR correlation experiments. A robust algorithm was introduced to numerically solve the L_1 regularized problem and methods were described for objectively choosing the amount of regularization needed and assigning error estimates to the reconstructed distributions. L_1 regularization was applied to T_1-T_2 and D-T_2 correlation experiments. It was shown that L_1 regularization reconstructions are stable even at SNR < 20, while the conventional NNLS method fails at these high values of noise. In addition, it was shown that using L_1 regularization, components with relaxation time constants and diffusion coefficients differing by as little as 10% can be resolved; the conventional Tikhonov regularization method fails to resolve relaxation time constants and diffusion coefficients differing by a factor < 3. The increased resolution capability obtained from L_1 regularization was used to extract inter and intra particle liquid compositions of a binary mixture of hexane and dodecane. This measurement capability will be valuable in extending our ability to characterise the composition and surface interaction properties of multi-component mixtures in porous media.

References

[1] P. T. Callaghan, Translational Dynamics and Magnetic Resonance, Oxford University Press, 2011.

[2] K. P. Whitall and A. L. MacKay, Quantitative interpretation of NMR relaxation data, J. Magn. Reson. 84 (1989), 134-152.

[3] Y.-Q. Song, L. Venkataramanan and L. Burcaw, Determining the resolution of Laplace inversion spectrum, J. Chem. Phys. 122 (2005), 104104.

[4] M. D. Hürlimann and L. Venkataramanan, Quantitative measurement of two-dimensional distribution functions of diffusion and relaxation in grossly inhomogeneous fields, J. Magn. Reson. 157 (2002), 31-42.

[5] L. Huabing, X. Lizhi, G. Baoxin, Z. Zongfu, Z. Fangrong, D. Feng, Y. Huijun, V. Anferov and S. Anferova, Heavy oil component characterization with multi-dimensional unilateral NMR, Pet. Sci. 10 (2013), 402-407.

[6] R. G. Baraniuk, Compressive sensing, IEEE Signal Process. Mag. July (2007), 118-124.

[7] L. I. Rudin, S. Osher and E. Fatemi, Nonlinear total variation based noise removal algorithms, Physica D 60 (1992), 259-268.

[8] M. Urbańczyk, D. Bernin, W. Koźmiński and K. Kazimierczuk, Iterative thresholding algorithm for multiexponential decay applied to PGSE NMR data, Anal. Chem. 85 (2013), 1828-1833.

[9] X. Zhou, G. Su, L. Wang, S. Nie and X. Ge, The inversion of 2D NMR relaxometry data using L_1 regularization, J. Magn. Reson. 275 (2017), 46-54.

[10] D. Benjamini and P. J. Basser, Use of marginal distributions constrained optimization (MADCO) for accelerated 2D MRI relaxometry and diffusometry, J. Magn. Reson. 271 (2016), 40-45.

[11] D. Benjamini and P. J. Basser, Towards clinically feasible relaxation-diffusion MRI using MADCO, Micropor. Mesorpor. Mat. 269 (2018), 93-96.

[12] E. L. Cussler, Diffusion: Mass Transfer in Fluid Systems, 3rd ed., Cambridge University Press, 2009.

[13]M. D. Hürlimann, L. Venkataramanan and C. Flaum, The diffusion-spin relaxation time distribution function as an experimental probe to characterize fluid mixtures in porous media, J. Chem. Phys. 117 (2002), 10223.

[14]R. J. Gray, Flexible methods for analyzing survival data using splines, with applications to breast cancer prognosis, J. Am. Statist. Assoc. 87 (1992), 942-951.

[15]R. Tibshirani, Regression shrinkage and selection via the Lasso, J. R. Statist. Soc. B 58 (1996), 267-288.

[16]H. Zou, T. Hastie and R. Tibshirani, On the “degrees of freedom” of the Lasso, Ann. Statist. 35 (2007), 2173-2192.

[17]R. J. Tibshirani and J. Taylor, Degrees of freedom in Lasso problems, Ann. Statist. 40 (2012), 1198-1232.

[18]C. Dossal, M. Kachour, M. J. Fadili, G. Peyré and C. Chesnau, The degrees of freedom of the Lasso for general design matrix, Stat. Sinica 23 (2013), 809-828.

[19]P. C. Hansen, Analysis of discrete ill-posed problems by means of the L-curve, SIAM Rev. 34 (1992), 561-580.

[20]J. N. Tehrani, A. McEwan, C. Jin and A. van Schaik, L_1 regularization method in electrical impedance tomography by using the L_1-curve (Pareto frontier curve), Appl. Math. Model. 36 (2012), 1095-1105.

[21]V. A. Morozov, Methods for Solving Incorrectly Posed Problems, Springer-Verlag, 1984.

[22]B. Babadi, N. Kalouptsidis and V. Tarokh, Asymptotic achievability of the Cramér-Rao bound for noisy compressive sampling, IEEE Trans. Signal Process. 57 (2009), 1233-1236.

[23]P. J. Bickel, Y. Ritov and A. B. Tsybakov, Simultaneous analysis of Lasso and Dantzig selector, Ann. Stat. 37 (2009), 1705-1732.

[24]Z. Ben-Haim and Y. C. Eldar, The Cramér-Rao bound for estimating a sparse parameter vector, IEEE Trans. Signal Process. 58 (2010), 3384-3389.

[25]Z. Ben-Haim, Y. C. Eldar and M. Elad, Coherence-based performance guarantees for estimating a sparse vector under random noise, IEEE Trans. Signal Process. 58 (2010), 5030-5043.

[26]P. T. Callaghan, Principles of Nuclear Magnetic Resonance, Oxford University Press, 1991.

[27]D. E. Freed, L. Burcaw and Y.-Q. Song, Scaling laws for diffusion coefficients in mixtures of alkanes, Phys. Rev. Lett. 94 (2005), 067602.

[28]D. E. Freed, Dependence on chain length of NMR relaxation times in mixtures of alkanes, J. Chem. Phys. 126 (2007), 174502.

[29]D. Weber, J. Mitchell, J. McGregor and L. F. Gladden, Comparing strengths of surface interactions for reactants and solvents in porous catalysts using two-dimensional NMR relaxation correlations, J. Phys. Chem. C 113 (2009), 6610-6615.

[30]K. F. Morris and C. S. Johnson, Resolution of discrete and continuous molecular size distributions by means of diffusion-ordered 2D NMR spectroscopy, J. Am. Chem. Soc. 115 (1993), 4291-4299.

Chapter 5

Retaining both sparse and smooth features in 1D and 2D NMR relaxation and diffusion experiments

5.1. Introduction

In Chapter 4, a reconstruction method, L_1 regularization, was introduced for obtaining 2D joint distributions of relaxation time constants and diffusion coefficients from 2D NMR correlation experiments. The method is also directly applicable to obtaining 1D distribution functions of relaxation time constants and diffusion coefficients from 1D NMR experiments. L_1 regularization assumes that the distribution is sparse and is, therefore, appropriate to use only when there is prior knowledge that the distribution is sparse. If there is prior knowledge that the distribution is smooth, the conventional Tikhonov regularization technique, outlined in Section 3.1.2, is more appropriate to use. Both methods of regularization, Tikhonov and L_1, depend on the availability of some prior knowledge of what the true distribution is. Therefore, choosing the most appropriate method to use depends on the prior knowledge. However, in many practical cases, such knowledge may not be available, or it is exactly what one is trying to find out. Therefore, imposing an incorrect prior knowledge can lead to choosing the incorrect method which gives an inaccurate reconstruction.

In this chapter, a new form of regularization is proposed, based on an adaptation of Total Generalized Variation (TGV) regularization [1], which will be termed Modified Total Generalized Variation (MTGV) regularization. TGV regularization has been successfully used in image denoising and deblurring, as well as MRI image reconstruction [2-5]. TGV regularization sits in a family of reconstruction techniques that aim to exploit sparsity of the reconstructed distribution in more than one domain, such as the elastic net regularization [6] or sparsity averaging regularization [7]. The regularization method

proposed offers a compromise in distinguishing between smooth and sparse features in the distribution. Therefore, it is no longer required to know *a priori* whether the distribution is expected to be smooth or sparse, since MTGV regularization retains both features. MTGV regularization is introduced for obtaining 1D distribution functions of relaxation time constants and diffusion coefficients from 1D NMR experiments and for obtaining 2D joint distributions of relaxation time constants and diffusion coefficients from 2D NMR correlation experiments. Algorithms are developed for numerically solving the 1D and 2D MTGV regularized problems and a practical guidance to implementing them is given.

The work in this chapter focuses on simulated 1D and 2D NMR signals used to obtain T_1 and T_2 distributions. These distributions can be obtained either by performing dedicated 1D experiments or by performing a T_1-T_2 experiment and projecting the 2D distribution map in each dimension. The T_1 and T_2 distributions reconstructed from MTGV regularization are compared to the results obtained from Tikhonov and L_1 regularization. Simulations are performed on a range of different input T_1 and T_2 distributions, differing with respect to the variance of the peaks in the distributions, and at three different SNR values of 2000, 200 and 20. A new metric is proposed to compare the reconstructed distributions from each method with the input distributions, which allows a quantitative comparison to be made between Tikhonov, L_1 and MTGV regularizations. The comparison of the different regularization techniques is performed with regards to simulated rather than experimental data because in experimental data the true distribution is not known; this renders the objective comparison between the different regularization techniques from experimental data impossible.

The chapter is structured as follows. In Section 5.2, MTGV regularization is introduced and an outline of the algorithms proposed to solve the 1D and 2D MTGV regularized problems is given. Section 5.3 describes the simulations performed. The results and discussion are presented in Section 5.4.

5.2. Mathematical treatment

The motivation for this chapter is to develop a reconstruction technique which is capable of retaining both sparse and smooth features in the distributions and is, therefore, less reliant than Tikhonov regularization and L_1 regularization on prior knowledge about the distribution. To obtain such distributions, it is proposed to perform MTGV regularization,

which is defined slightly differently, depending on whether it is applied to 1D NMR data sets (called 1D MTGV regularization) or to 2D NMR datasets (called 2D MTGV regularization).

5.2.1. 1D MTGV regularization

As discussed in Sections 2.9-2.11, the NMR signal acquired from a 1D NMR experiment, $\boldsymbol{y}$, can be written in terms of a kernel matrix, $\boldsymbol{K}$, the distribution, $\boldsymbol{u}$, and the noise vector, $\boldsymbol{e}$, as:

$$\boldsymbol{y} = \boldsymbol{K}\,\boldsymbol{u} + \boldsymbol{e}\ . \tag{5.1}$$

1D MTGV regularization consists of extracting the distribution $\boldsymbol{u}$ from the following minimization problem:

$$(\boldsymbol{u}, \boldsymbol{w}) = \arg\ \min_{\boldsymbol{u}\geq\boldsymbol{0},\boldsymbol{w}} \left(\tfrac{1}{2}\|\boldsymbol{K}\,\boldsymbol{u} - \boldsymbol{y}\|_2^2 + \alpha\|\boldsymbol{u} - \boldsymbol{w}\|_1 + \beta\|\boldsymbol{Q}\,\boldsymbol{w}\|_1\right), \tag{5.2}$$

where $\boldsymbol{w}$ is an auxiliary vector and $\boldsymbol{Q}$ performs the second order derivative of the vector it is applied to; $\boldsymbol{Q}$ is a tri-diagonal matrix with elements -1, 2 and -1 in diagonals -1, 0 and 1. It is seen that the penalty term in MTGV regularization is composed of two parts: the first one enforces sparse features in the reconstructed distribution, while the second term enforces smooth features, with the weight of these terms controlled by the regularization parameters α and β, respectively. Therefore, MTGV regularization offers a compromise in retaining sparse and smooth features in the reconstructed distribution.

The minimization problem in Eq. (5.2) is challenging because of the non-differentiability of the L_1-norm. Similar to the approach used in Chapter 4, Eq. (5.2) is rewritten as:

$$(\boldsymbol{u}, \boldsymbol{w}, \boldsymbol{v}, \boldsymbol{z}) = \arg\ \min_{\boldsymbol{u}\geq\boldsymbol{0},\boldsymbol{w}} \max_{\boldsymbol{v},\boldsymbol{z}} \left(\tfrac{1}{2}\|\boldsymbol{K}\,\boldsymbol{u} - \boldsymbol{y}\|_2^2 + \boldsymbol{v}^{\mathrm{T}}(\boldsymbol{u} - \boldsymbol{w}) - \chi_\alpha(\boldsymbol{v}) + \boldsymbol{z}^{\mathrm{T}}\boldsymbol{Q}\,\boldsymbol{w} - \chi_\beta(\boldsymbol{z})\right), \tag{5.3}$$

where $\boldsymbol{v}$ and $\boldsymbol{z}$ are dual variables, and $\chi_\alpha(\boldsymbol{v})$ is an indicator function defined as:

$$\chi_\alpha(\boldsymbol{v}) = \begin{cases} 0 & \|\boldsymbol{v}\|_\infty \leq \alpha \\ +\infty & \|\boldsymbol{v}\|_\infty > \alpha \end{cases}\ . \tag{5.4}$$

The approach of Chambolle and Pock [8] is then followed to define new variables: $\boldsymbol{u}'$, $\boldsymbol{v}'$ and $\boldsymbol{R}$ as:

$$\boldsymbol{u}' = \begin{bmatrix}\boldsymbol{u}\\ \boldsymbol{w}\end{bmatrix} \qquad \boldsymbol{v}' = \begin{bmatrix}\boldsymbol{v}\\ \boldsymbol{z}\end{bmatrix} \qquad \boldsymbol{R} = \begin{bmatrix}\boldsymbol{I} & -\boldsymbol{I}\\ \boldsymbol{0} & \boldsymbol{Q}\end{bmatrix}. \tag{5.5}$$

Using the new definitions, Eq. (5.3) is written in the form of the primal-dual problem defined in Section 3.3:

$$(\boldsymbol{u}', \boldsymbol{v}') = \min_{\boldsymbol{u}'} \max_{\boldsymbol{v}'} f(\boldsymbol{u}') - g(\boldsymbol{v}') + \boldsymbol{v}'^{\mathrm{T}} \boldsymbol{R}\, \boldsymbol{u}' \,, \tag{5.6}$$

where $f(\boldsymbol{u}') = \frac{1}{2}||\boldsymbol{K}\,\boldsymbol{u} - \boldsymbol{y}||_2^2$ and $g(\boldsymbol{v}') = \chi_\alpha(\boldsymbol{v}) + \chi_\beta(\boldsymbol{z})$. The PDHGM iteration scheme, outlined in Section 3.3, is then used to numerically solve the minimization problem. The proximal operator for $f(\boldsymbol{u}')$ was derived in Section 4.1, while the proximal operator of $g(\boldsymbol{v}')$ is obtained from the sum of the proximal operators of its constituent functions, each of which was derived in Section 3.3.1.

A pseudocode for the numerical solution of Eq. (5.2) is given in Table 5.1.

Table 5.1. Pseudocode to numerically solve Eq. (5.2), based on the PDHGM method.

Step 1. Choose algorithm parameters τ, σ and the regularization parameters, α, β.

Step 2. Set the convergence tolerance, TOL.

Step 3. Calculate $\boldsymbol{C} = (\boldsymbol{I} + \tau\, \boldsymbol{K}^{\mathrm{T}}\boldsymbol{K})^{-1}$.

Step 4. Initialize $\boldsymbol{v}^{(0)} = \underline{\boldsymbol{v}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{z}^{(0)} = \underline{\boldsymbol{z}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{u}^{(0)} = \underline{\boldsymbol{u}}^{(0)} \neq \boldsymbol{0}$, $\boldsymbol{w}^{(0)} \neq \boldsymbol{0}$.

Step 5. Initialize count number, $k = 0$ and convergence tracker, $\varepsilon^{(0)} = 1$.

Step 6. **while** $\varepsilon^{(k)} >$ TOL **do**

a. $\underline{\boldsymbol{u}}^{(k+1)} \leftarrow \boldsymbol{u}^{(k)} - \tau\, \boldsymbol{v}^{(k)}$

b. $\boldsymbol{u}^{(k+1)} \leftarrow \boldsymbol{C}\,(\underline{\boldsymbol{u}}^{(k+1)} + \tau\, \boldsymbol{K}^{\mathrm{T}}\boldsymbol{y})$

c. $\boldsymbol{u}^{(k+1)} \leftarrow \max(0, \boldsymbol{u}^{(k+1)})$

All operations in this line are element-wise.

d. $\boldsymbol{w}^{(k+1)} \leftarrow \boldsymbol{w}^{(k)} + \tau\,(\boldsymbol{v}^{(k)} - \boldsymbol{Q}^{\mathrm{T}}\boldsymbol{z}^{(k)})$

e. $\underline{\boldsymbol{v}}^{(k+1)} \leftarrow \boldsymbol{v}^{(k)} + \sigma\,(2\boldsymbol{u}^{(k+1)} - \boldsymbol{u}^{(k)} - 2\boldsymbol{w}^{(k+1)} + \boldsymbol{w}^{(k)})$

f. $\boldsymbol{v}^{(k+1)} \leftarrow \underline{\boldsymbol{v}}^{(k+1)}/\max(1, \underline{\boldsymbol{v}}^{(k+1)}/\alpha)$

All operations in this line are element-wise.

g. $\underline{\boldsymbol{z}}^{(k+1)} \leftarrow \boldsymbol{z}^{(k)} + \sigma\, \boldsymbol{Q}\,(2\boldsymbol{w}^{(k+1)} - \boldsymbol{w}^{(k)})$

h. $\boldsymbol{z}^{(k+1)} \leftarrow \underline{\boldsymbol{z}}^{(k+1)}/\max(1, \boldsymbol{z}^{(k+1)}/\beta)$

All operations in this line are element-wise

i. $\varepsilon^{(k+1)} \leftarrow ||\boldsymbol{u}^{(k+1)} - \boldsymbol{u}^{(k)}||_2 \,/\, ||\boldsymbol{u}^{(k)}||_2$

j. $k \leftarrow k + 1$

end while

5.2.2. 2D MTGV regularization

As discussed in Section 2.12, the vectorised NMR signal acquired from a 2D NMR correlation experiment, $\boldsymbol{Y}_{\mathrm{v}}$, can be written in terms of a kernel matrix, $\boldsymbol{K}$, the vectorised joint distribution, $\boldsymbol{U}_{\mathrm{v}}$, and the vectorised noise matrix, $\boldsymbol{E}_{\mathrm{v}}$, as:

$$\boldsymbol{Y}_{\mathrm{v}} = \boldsymbol{K}\,\boldsymbol{U}_{\mathrm{v}} + \boldsymbol{E}_{\mathrm{v}}\ . \tag{5.7}$$

2D MTGV regularization consists of obtaining the reconstructed distribution from the following minimization problem:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{w}) = \arg\min_{\boldsymbol{U}_{\mathrm{v}} \geq \boldsymbol{0}, \boldsymbol{w}} \left(\frac{1}{2} \|\boldsymbol{K}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \alpha \|\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{w}\|_1 + \beta \|[\boldsymbol{Q}\,\boldsymbol{w}]_3\|_{2,1} \right), \tag{5.8}$$

where $\boldsymbol{w}$ is an auxiliary variable and $\boldsymbol{Q}$ performs a symmetrical second order derivative on the vectorised matrix it is applied to; the construction of $\boldsymbol{Q}$ is outlined below. The $L_{2,1}$-norm of a matrix was defined in Section 3.3.4, while the operation $[\boldsymbol{Q}\,\boldsymbol{w}]_3$ reshapes the vector $\boldsymbol{Q}\,\boldsymbol{w}$ into a matrix with three columns. The interpretation of the two penalty terms is similar to the case of 1D MTGV regularization.

$\boldsymbol{Q}$ is constructed according to the work of Ono and Yamada [5]. Define $\boldsymbol{D}$ to be a bi-diagonal matrix with elements -1 and 1 in diagonals 0 and 1 (with Neumann boundary conditions) and of the same size as $\boldsymbol{U}$. $\boldsymbol{D}_x = \boldsymbol{I} \otimes \boldsymbol{D}$, where $\boldsymbol{I}$ is the identity matrix of the same size as $\boldsymbol{D}$, performs the first order derivative of a matrix in the vertical direction (when applied to its vectorised version); $\boldsymbol{D}_y = \boldsymbol{D} \otimes \boldsymbol{I}$ performs the first derivative of a matrix in the horizontal direction (when applied to its vectorised version). As a result, $\boldsymbol{D}_{xy} = [\boldsymbol{D}_x{}^{\mathrm{T}}\ \boldsymbol{D}_y{}^{\mathrm{T}}]^{\mathrm{T}}$ performs the first derivative of a matrix in both directions (when applied to its vectorised version). As a result, the matrix $\boldsymbol{Q} = \boldsymbol{H}\,\boldsymbol{D}_{xy}$, where:

$$\boldsymbol{H} = \begin{bmatrix} \boldsymbol{D}_x{}^{\mathrm{T}} & \boldsymbol{0} \\ \boldsymbol{D}_y{}^{\mathrm{T}} & \boldsymbol{D}_x{}^{\mathrm{T}} \\ \boldsymbol{0} & \boldsymbol{D}_y{}^{\mathrm{T}} \end{bmatrix}, \tag{5.9}$$

performs the symmetrical second order derivative of a matrix (when applied to its vectorised version).

The minimization problem in Eq. (5.8) is written as:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{w}, \boldsymbol{v}, \boldsymbol{z}) = \arg\min_{\boldsymbol{U}_{\mathrm{v}} \geq \boldsymbol{0}, \boldsymbol{w}} \max_{\boldsymbol{v}, \boldsymbol{z}} \left(\frac{1}{2} \|\boldsymbol{K}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \boldsymbol{v}^{\mathrm{T}}(\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{w}) - \chi_\alpha(\boldsymbol{v}) + \boldsymbol{z}^{\mathrm{T}}\boldsymbol{Q}\,\boldsymbol{w} - \chi'_\beta(\boldsymbol{z}) \right), \tag{5.10}$$

where $\chi_\alpha(\boldsymbol{v})$ is defined as in Eq. (5.4), while $\chi'_\beta(\boldsymbol{z})$ is defined as:

$$\chi'_\beta(\boldsymbol{z}) = \begin{cases} 0 & \|[\boldsymbol{z}]_3\|_{2,\infty} \leq \beta \\ +\infty & \|[\boldsymbol{z}]_3\|_{2,\infty} > \beta \end{cases} . \tag{5.11}$$

Eq. (5.10) is then written in terms of combined variables, defined in Eq. (5.5), as was done in Section 5.2.1. With the definition of the combined variables, Eq. (5.10) becomes a primal-dual problem, defined in Section 3.3, where $f(\boldsymbol{U}_\mathrm{v}') = ½||\boldsymbol{K}\,\boldsymbol{U}_\mathrm{v} - \boldsymbol{Y}_\mathrm{v}||_2^2$ and $g(\boldsymbol{v}') = \chi_\alpha(\boldsymbol{v}) + \chi'_\beta(\boldsymbol{z})$. The PDHGM iteration scheme, outlined in Section 3.3, is then used to numerically solve the minimization problem. The proximal operator for $f(\boldsymbol{U}_\mathrm{v}')$ was derived in Section 4.2, while the proximal operator of $g(\boldsymbol{v}')$ is obtained from the sum of the proximal operators of its constituent functions. The proximal operator of $\chi_\alpha(\boldsymbol{v})$ was derived in Section 3.3.1. The proximal operator of $\chi'_\beta(\boldsymbol{z})$ is derived in a similar way to the proximal operator of $\chi_\alpha(\boldsymbol{v})$, and the result can be found in the work of Chambolle and Pock [8].

A pseudocode for the numerical solution of Eq. (5.8) is given in Table 5.2.

Table 5.2. Pseudocode to numerically solve Eq. (5.8), based on the PDHGM method. In the pseudocode, sum($\boldsymbol{A}$, i) sums the column or row vectors of matrix $\boldsymbol{A}$ along the i-th dimension and, repmat($\boldsymbol{A}$, k, i) stacks k copies of $\boldsymbol{A}$ in the i-th dimension.

Step 1. Choose algorithm parameters τ, σ and the regularization parameters, α, β.

Step 2. Set the convergence tolerance, TOL.

Step 3. Calculate $\boldsymbol{C} = (\boldsymbol{I} + \tau\,\boldsymbol{K}^\mathrm{T}\boldsymbol{K})^{-1}$.

Step 4. Initialize $\boldsymbol{v}^{(0)} = \underline{\boldsymbol{v}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{z}^{(0)} = \underline{\boldsymbol{z}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{U}_\mathrm{v}^{(0)} = \underline{\boldsymbol{U}}_\mathrm{v}^{(0)} \neq \boldsymbol{0}$, $\boldsymbol{w}^{(0)} \neq \boldsymbol{0}$.

Step 5. Initialize count number, $k = 0$ and convergence tracker, $\varepsilon^{(0)} = 1$.

Step 6. **while** $\varepsilon^{(k)} >$ TOL **do**

a. $\underline{\boldsymbol{U}}_\mathrm{v}^{(k+1)} \leftarrow \boldsymbol{U}_\mathrm{v}^{(k)} - \tau\,\boldsymbol{v}^{(k)}$

b. $\boldsymbol{U}_\mathrm{v}^{(k+1)} \leftarrow \boldsymbol{C}\,(\underline{\boldsymbol{U}}_\mathrm{v}^{(k+1)} + \tau\,\boldsymbol{K}^\mathrm{T}\boldsymbol{Y}_\mathrm{v})$

c. $\boldsymbol{U}_\mathrm{v}^{(k+1)} \leftarrow \max(0, \boldsymbol{U}_\mathrm{v}^{(k+1)})$

All operations in this line are element-wise.

d. $\boldsymbol{w}^{(k+1)} \leftarrow \boldsymbol{w}^{(k)} + \tau\,(\boldsymbol{v}^{(k)} - \boldsymbol{Q}^\mathrm{T}\boldsymbol{z}^{(k)})$

e. $\underline{\boldsymbol{v}}^{(k+1)} \leftarrow \boldsymbol{v}^{(k)} + \sigma\,(2\boldsymbol{U}_\mathrm{v}^{(k+1)} - \boldsymbol{U}_\mathrm{v}^{(k)} - 2\boldsymbol{w}^{(k+1)} + \boldsymbol{w}^{(k)})$

f. $\boldsymbol{v}^{(k+1)} \leftarrow \underline{\boldsymbol{v}}^{(k+1)}/\max(1, \underline{\boldsymbol{v}}^{(k+1)}/\alpha)$

All operations in this line are element-wise.

g. $\underline{\boldsymbol{z}}^{(k+1)} \leftarrow \boldsymbol{z}^{(k)} + \sigma\,\boldsymbol{Q}\,(2\boldsymbol{w}^{(k+1)} - \boldsymbol{w}^{(k)})$

---

h. $\boldsymbol{x} \leftarrow \sqrt{\text{sum}([\underline{\boldsymbol{z}}^{(k+1)}]_3^2, 2)}$

All operations in this line are element-wise.

i. $\boldsymbol{z}^{(k+1)} \leftarrow \underline{\boldsymbol{z}}^{(k+1)}/\max(1, \text{repmat}(\boldsymbol{x}, 3, 1)/\beta)$

All operations in this line are element-wise.

j. $\varepsilon^{(k+1)} \leftarrow ||\boldsymbol{U}_v^{(k+1)} - \boldsymbol{U}_v^{(k)}||_2 / ||\boldsymbol{U}_v^{(k)}||_2$

k. $k \leftarrow k + 1$

end while

5.2.3. Implementation of the algorithm

The condition for the convergence of the MTGV algorithm is similar to the condition for the TGV algorithm to converge, as has been discussed by Valkonen *et al.* [9]. The condition depends primarily on the choice of the two algorithmic parameters, τ and σ, which are described in the pseudocode. Although there is a relationship that they must obey, the particular choice of τ and σ is largely heuristic. The choice of $\tau = \sigma = 0.1$ is found to offer a good compromise between convergence speed and stability in our particular application. The values of τ and σ will depend slightly on how the NMR signal is scaled, determined by the largest data point in the acquired signal. To avoid this, it is best to scale the signal to a maximum of 1. In the present work, the number of iterations required to arrive at a reasonable convergence is ~10,000. The time in which this convergence is achieved for a 1D 32 × 1 distribution with a 2.0 GHz Intel® Core™ i5-4590T CPU and 16.4 GB RAM, is ~6 s, while the speed required to achieve this convergence for a 2D 32 × 32 distribution is ~23 s.

5.3. Simulation settings

The performance of Tikhonov, L_1 and MTGV regularizations in the reconstruction of individual T_1 and T_2 distributions are compared to each other by performing the following numerical experiments:

(i) Tikhonov regularization applied to 1D NMR T_1 and T_2 data, referred to as 1D Tikhonov regularization.

(ii) L_1 regularization applied to 1D NMR T_1 and T_2 data, referred to as 1D L_1 regularization.

(iii) MTGV regularization applied to 1D NMR T_1 and T_2 data, referred to as 1D MTGV regularization.

(iv) Tikhonov regularization applied to 2D T_1-T_2 data, with the reconstructed map projected in each dimension, referred to as 2D Tikhonov regularization.

(v) L_1 regularization applied to 2D T_1-T_2 data, with the reconstructed map projected in each dimension, referred to as 2D L_1 regularization.

(vi) MTGV regularization applied to 2D T_1-T_2 data, with the reconstructed map projected in each dimension, referred to as 2D MTGV regularization.

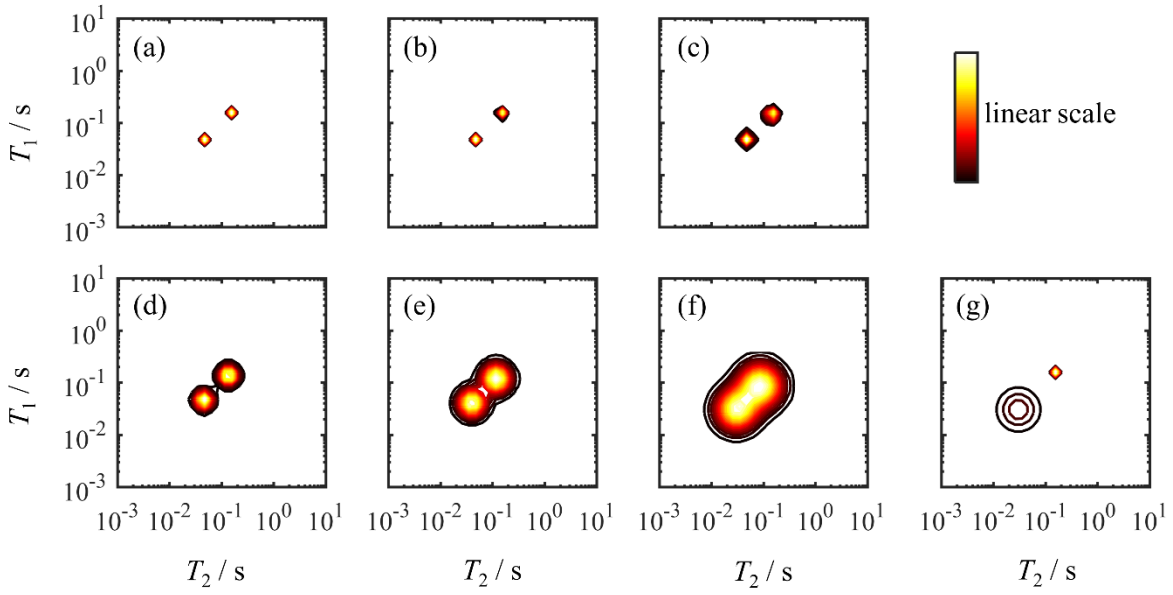

Fig. 5.1. T_1-T_2 input distributions used in simulations. The distributions in figures (a)-(g) are labelled accordingly A-G. All peaks are lognormally distributed, and the details are described in Table 5.3.

To investigate the performance of the different regularization methods, 7 simulated 2D T_1-T_2 distributions and their corresponding T_1 and T_2 projections are used. All these distributions are 32×32 logarithmically spaced maps that have two lognormally distributed peaks of the same magnitude and centred at $T_1 = T_2 = 0.05$ s and $T_1 = T_2 = 0.15$ s. The distributions, labelled A-G, are illustrated in Fig. 5.1. For distributions A-F, both peaks are of similar smoothness, with the level of smoothness increasing from distributions A to F. Distribution G is an example of a distribution where one of the peaks is discrete and the other is smooth. The lognormal distribution details for all the peaks are given in Table 5.3.

Noise-free NMR signals are simulated with 32 logarithmically spaced t_1 values and 256 linearly spaced t_2 steps from 0.001 s to 0.75 s, as outlined in Sections 2.9 and 2.10. Random Gaussian noise is then added to the signals such that 2D NMR experiments at SNR values of 2000, 200 and 20 are studied. For the respective 1D signals, the standard deviation of the noise is chosen according to the method described by Celik *et al*. [10]. Before processing, the 2D NMR data are compressed according to the method outlined in Section 3.1.3. The distribution maps are then reconstructed from the noisy simulated signals using the different regularization methods.

Table 5.3: Details of the peaks of distributions A-G shown in Fig. 5.1. Each distribution is composed of two peaks; the peak at the lower relaxation times is referred to as peak 1 while the peak at the higher relaxation times is referred to as peak 2. The distribution in each dimension, T_1 or T_2, is described by lognpdf(T_i,μ_1,σ_1) + λ lognpdf(T_i,μ_2,σ_2), for i = 1, 2, where μ, σ refer to the mean and standard deviation of each peak. Distributions A-F are designed to have the same area under each peak in a logarithmic scale. Distribution G is designed such that both peaks have the same maximum amplitude when projected in each dimension.

distribution	μ_1 (s)	σ_1 (s)	μ_2 (s)	σ_2 (s)	λ
A	0.050	0.003	0.150	0.010	10
B	0.050	0.005	0.150	0.016	10
C	0.050	0.008	0.150	0.025	10
D	0.050	0.013	0.150	0.040	10
E	0.050	0.020	0.150	0.063	10
F	0.050	0.032	0.150	0.100	10
G	0.050	0.032	0.150	0.010	2

To compare the reconstructed distributions with the true distributions, the following metric is proposed: If $U_{\mathrm{v,true},i}$ is the true distribution's i-th element and $U_{\mathrm{v,rec},i}$ the reconstructed distribution's i-th element, then a metric of how close the reconstructed distribution is to the true distribution is:

$$\eta = \sum_i \frac{\left(U_{\mathrm{v,rec},i} - U_{\mathrm{v,true},i}\right)^2}{\max\left(10^{-4}, U_{\mathrm{v,true},i}\right)} . \tag{5.12}$$

One reason for choosing such a metric instead of the more commonly used mean square error or peak-signal-to-noise ratio (PSNR) is that the latter do not give any structural information. In particular, any extra peaks in the reconstructed distribution, as compared to the true distribution are not heavily penalized. This issue is addressed in the proposed

metric by the division operation. A threshold of 10^{-4} was used for the division process, to account for distribution entries with $U_{\mathrm{v,true},i} \approx 0$. Simulations showed that reducing the threshold further made no difference to the calculation of η. The proposed metric is similar to the goodness of fit measure which is used in hypothesis testing in statistics [11]. The smaller the value of η, the closer the reconstructed map is to the true distribution. The perfect reconstruction has $\eta = 0$.

For Tikhonov and L_1 regularization, η is estimated for a range of the regularization parameter α ($10^{-4} – 10^{4}$) and the smallest value (corresponding to the best reconstruction) is recorded. For MTGV regularization, η is estimated for a range of the regularization parameters α ($10^{-3} – 10^{7}$ for 1D MTGV, $10^{1} – 10^{9}$ for 2D MTGV) and β ($10^{-10} – 10^{0}$ for 1D MTGV, $10^{-4} – 10^{4}$ for 2D MTGV) and the smallest value is recorded. A comparison of the lowest η from the different regularization methods is then a comparison of their performance, with the best method being the one with the smallest η.

5.4. Results and discussion

The reconstructed T_1 and T_2 distributions from the different regularization methods, Tikhonov, L_1 and MTGV are compared to each other using the metric η. Initially, the comparison is performed on distributions A-F, which contain two peaks that are of similar smoothness, with smoothness increasing from distribution A to F. Three values of SNR: 2000, 200 and 20 are investigated. The results are shown in Figs. 5.2-5.4, respectively. In each case, the smaller the value of η, the better the reconstruction of the distribution by that particular regularization method.

Comparing the performance of Tikhonov and L_1 regularization with MTGV regularization in processing the 1D NMR T_1 or T_2 data at SNR = 2000 in Fig. 5.2, it is clearly seen that MTGV regularization is always superior, apart from the case when the true distribution is very discrete (distribution A), in which case L_1 regularization is slightly more accurate. This is valid for both T_1 and T_2 distributions. Further, applying MTGV regularization to 2D NMR T_1-T_2 data is always superior to applying Tikhonov or L_1 regularization to 2D NMR T_1-T_2 data, apart from the case when the distribution is very discrete (distribution A), in which case L_1 regularization is again slightly more accurate.

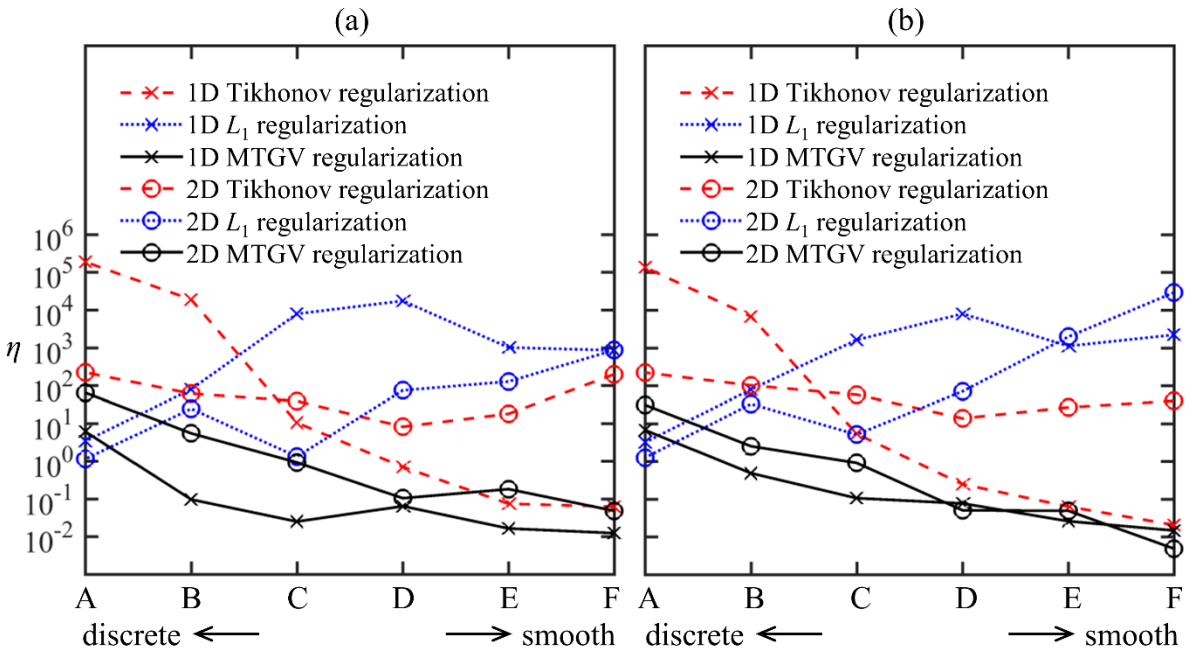

Fig. 5.2. Comparison of the performance of the different regularization techniques in reconstructing (a) the T_1 and (b) T_2 distributions for a range of smooth and discrete true distributions at SNR = 2000. Distributions A-F are shown in Fig. 5.1 and the metric η is defined in Eq. (5.12).

The same conclusions can be drawn for the reconstructions obtained at lower signal-to-noise ratios shown in Fig. 5.3 (SNR = 200) and Fig. 5.4 (SNR = 20). At all levels of noise, MTGV outperforms Tikhonov and L_1 regularization when processing 1D or 2D NMR data, apart from when the distribution is very discrete, in which case L_1 regularization is slightly more accurate.

The results in Figs. 5.2-5.4 also confirm the known strengths of Tikhonov and L_1 regularization; it is seen that when the true distribution is discrete (distributions A and B), L_1 regularization (whether performed on 1D or 2D data) performs better than Tikhonov regularization. When the true distribution is smooth (distributions C to F), Tikhonov regularization performs better than L_1 regularization. It is also seen that when Tikhonov regularization is applied to the T_1-T_2 discrete distributions A and B and the reconstructed maps are projected in each dimension, this gives better results than applying Tikhonov regularization to the individual 1D data. Therefore, 2D experiments give a better spectral resolution than 1D experiments, when processed with Tikhonov regularization. This is in agreement with the observations made by Celik *et al*. [10]. However, it is observed that when the true distribution is smooth (distributions C to F), this conclusion is no longer valid, and 1D experiments processed with Tikhonov regularization are better reconstructed

than 2D experiments processed with Tikhonov regularization and then projected in each dimension.

Considering the implementation of the MTGV regularizer in more detail, it was discussed in Section 5.2.1 that the regularization parameter β controls the amount of smoothness imposed in the reconstruction. A heuristic method for choosing the regularization parameter β is now introduced. In Figs. 5.2-5.4, the value of the metric η quoted for the MTGV regularization was that obtained by varying α and β independently and finding the minimum η obtained in this process. In addition to recording this minimum η, the values of α and β at which this minimum is obtained were also recorded. Fig. 5.5 shows these values of α and β for all the 1D and 2D MTGV regularization data displayed in Figs. 5.2-5.4. It is observed that the variation in the optimal α is much larger than the variation in the optimal β. As a result, to some approximation, β may be taken to be a constant, c. The value of c is expected to depend mainly on the type of distribution (discrete or smooth) and the noise level in the data. Fig. 5 shows that $c = 2\times10^{-5}$ for 1D MTGV regularization and $c = 0.3$ for 2D MTGV regularization are a good estimate for a range of distributions and noise levels. Similar values of c are expected when MTGV regularization is applied to other 1D and 2D NMR relaxation and diffusion data. The choice of the only remaining free regularization parameter, α, can be done in similar ways to the methods that exist for choosing the regularization parameter in Tikhonov regularization [12].

To investigate the effect that constraining β has on the performance of MTGV regularization, similar simulations to the ones described in Section 5.3 are performed on distributions A-F by constraining the regularization parameter β to $\beta = 2\times10^{-5}$ in 1D MTGV regularization and $\beta = 0.3$ in 2D MTGV regularization. An example of the comparison between the performance of Tikhonov, L_1 and MTGV regularization (with constrained β) at SNR = 20 is shown in Fig. 5.6 for both T_1 and T_2 distributions. A negligible difference is observed between the results in Fig. 5.6 and the results in Fig. 5.4, which compared the performance of Tikhonov, L_1 and MTGV regularization for an unconstrained β at SNR = 20. These results indicate that rendering MTGV regularization as a one-parameter regularization method by constraining β does not change any of the conclusions made in this chapter when β is considered a free parameter.

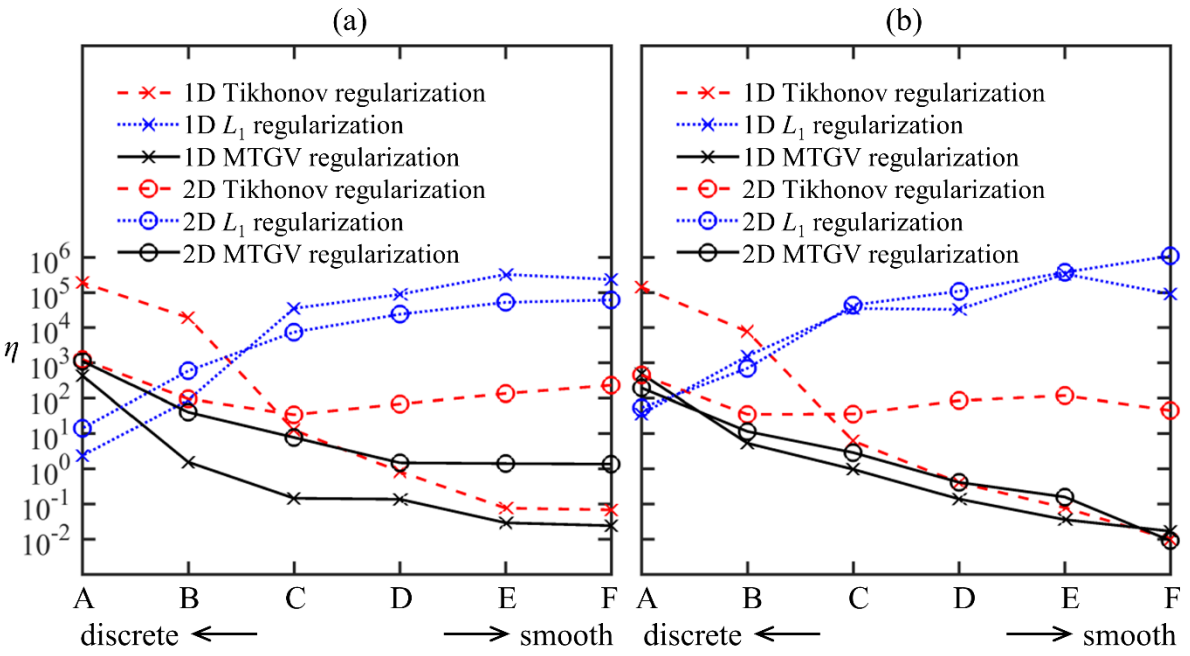

Fig. 5.3. Comparison of the performance of the different regularization techniques in reconstructing (a) the T_1 and (b) T_2 distributions for a range of smooth and discrete true distributions at SNR = 200. Distributions A-F are shown in Fig. 5.1 and the metric η is defined in Eq. (5.12).

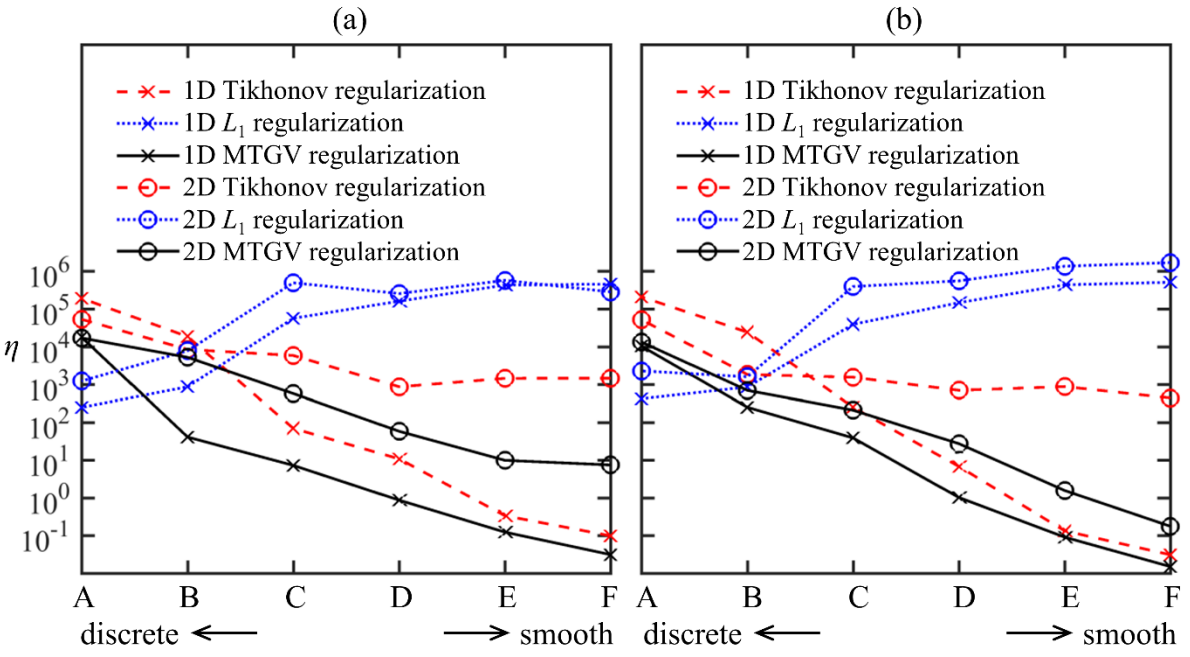

Fig. 5.4. Comparison of the performance of the different regularization techniques in reconstructing (a) the T_1 and (b) T_2 distributions for a range of smooth and discrete true distributions at SNR = 20. Distributions A-F are shown in Fig. 5.1 and the metric η is defined in Eq. (5.12).

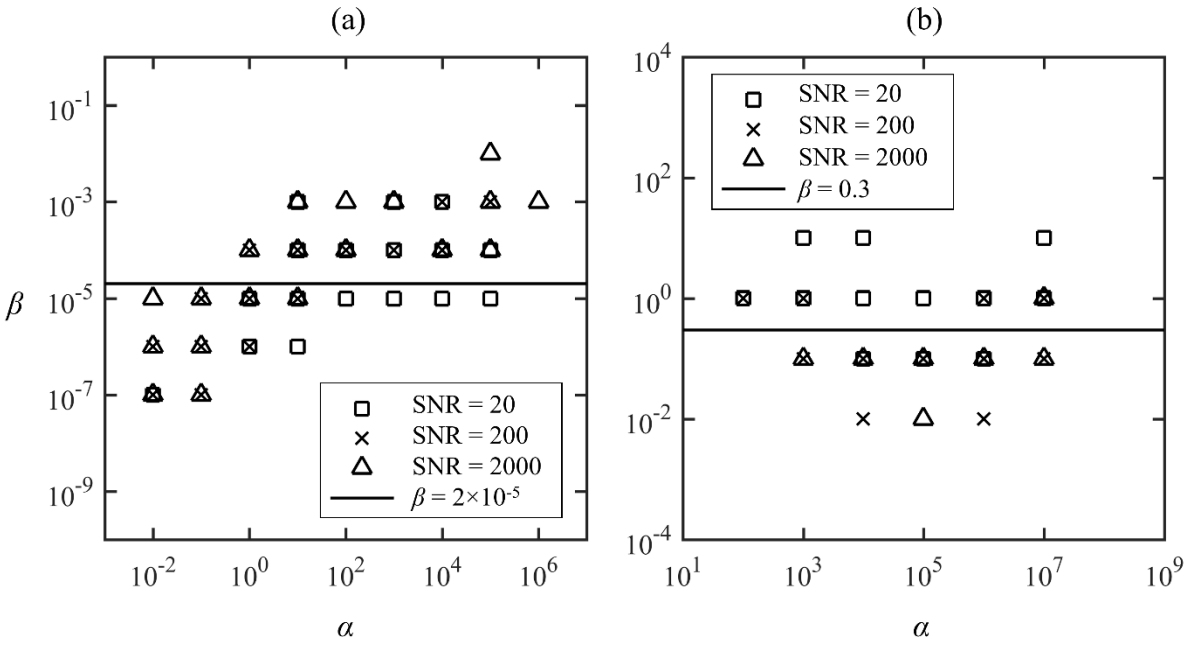

Fig. 5.5. Variation of the best α and β for (a) 1D MTGV regularization and (b) 2D MTGV regularization in the experiments described in Figs. 5.2-5.4. The line of best fits (a) $\beta = 2\times10^{-5}$ and (b) $\beta = 0.3$ are estimated by the best fitting of $\log_{10}\beta$ to a constant value. 108 points are plotted in both maps but some of them are not distinguishable because they overlap with other points.

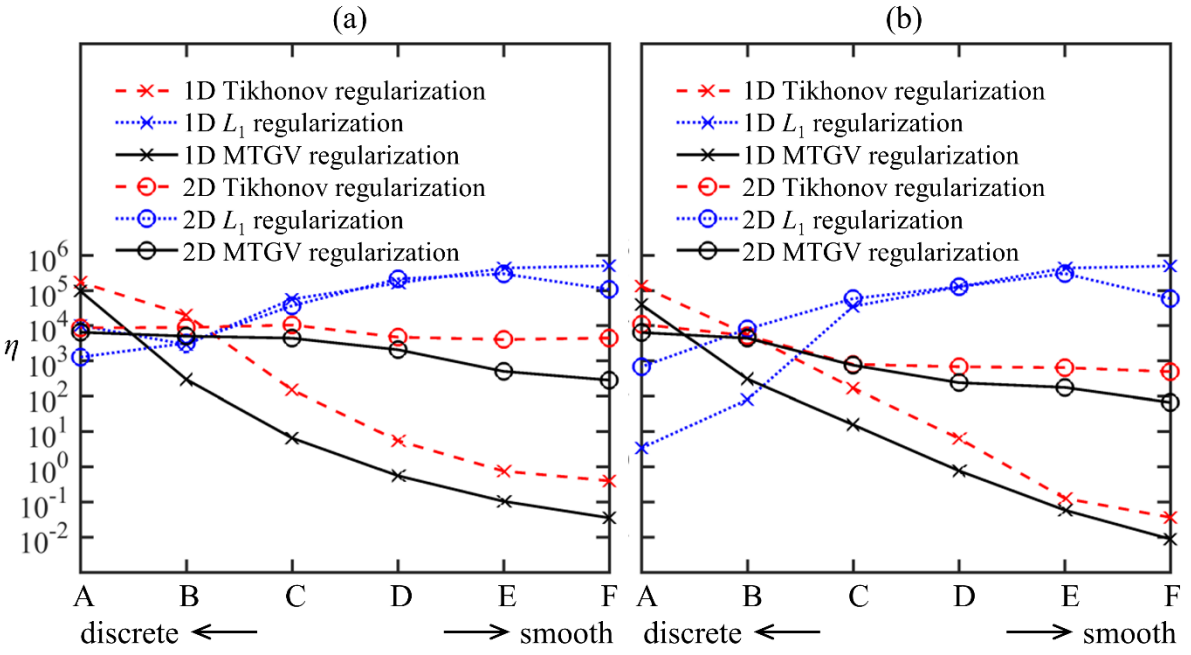

Fig. 5.6. Comparison of the performance of the different regularization techniques in reconstructing (a) the T_1 and (b) T_2 distributions for a range of smooth and discrete true distributions at SNR = 20. Distributions A-F are shown in Fig. 5.1 and the metric η is defined in Eq. (5.12). The value of β for MTGV regularization is constrained to (a) $\beta = 2\times10^{-5}$ for 1D MTGV and (b) $\beta = 0.3$ for 2D MTGV regularization.

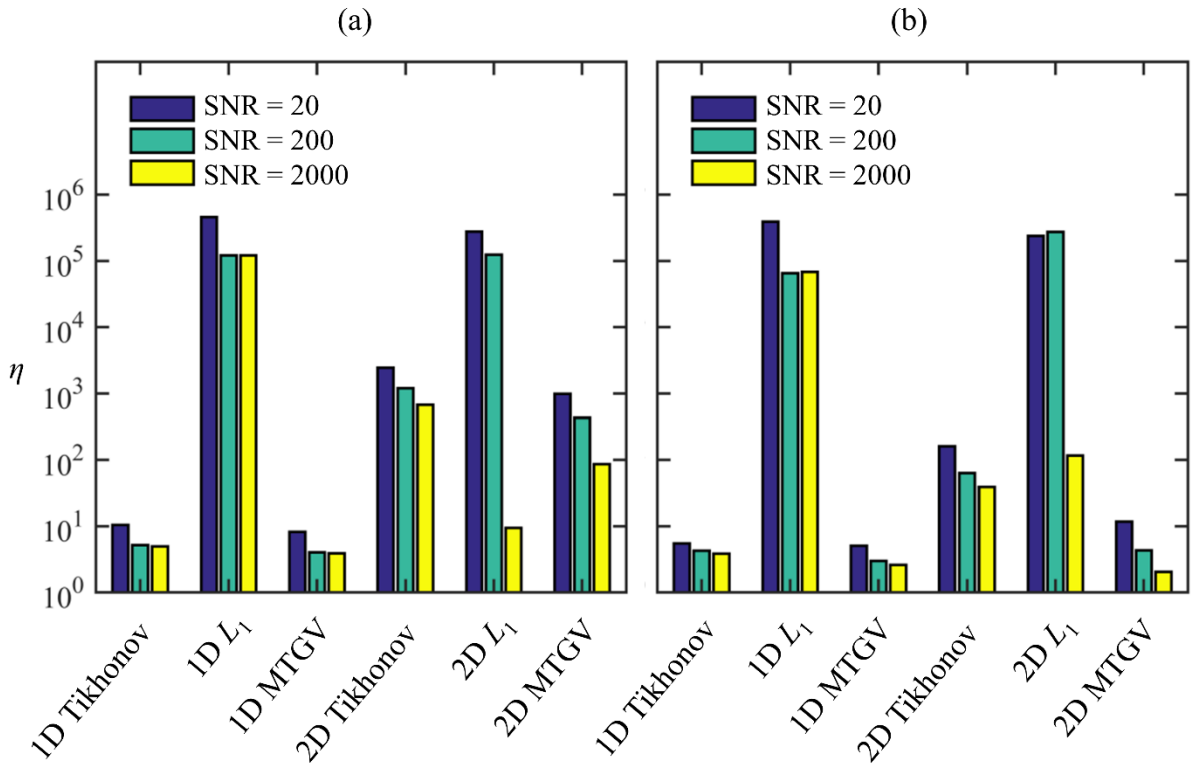

Fig. 5.7. Comparison of the performance of the different regularization techniques in reconstructing (a) the T_1 and (b) T_2 distributions for the true distribution G shown in Fig. 5.1 at three SNR: 20, 200 and 2000. The metric η is defined in Eq. (5.12). The value of β for MTGV regularization is constrained to (a) $\beta = 2\times10^{-5}$ for 1D MTGV and (b) $\beta = 0.3$ for 2D MTGV regularization.

The comparison between the different regularization techniques has so far been made on distributions A-F, which contain two peaks of similar smoothness, with the degree of smoothness increasing from A to F. An example of the comparison between Tikhonov, L_1 and MTGV regularization for a distribution which contains a mixture of discrete and smooth features is shown in Fig. 5.7. Fig. 5.7 compares the reconstructions of the T_1 and T_2 distributions of the true distribution G, illustrated in Fig. 5.1, at three different SNR: 20, 200 and 2000. MTGV regularization reconstructions are performed with a constrained regularization parameter β. It is clearly seen that 1D MTGV regularization always outperforms 1D Tikhonov and 1D L_1 regularizations and 2D MTGV regularization always outperforms 2D Tikhonov and 2D L_1 regularizations.

The results presented in this chapter suggest that MTGV regularization gives the best reconstructions of 1D and 2D NMR relaxation data, regardless of the type of distribution and the SNR in the NMR data. The only exception is the case when the true distribution

is known to only comprise of discrete peaks, in which case L_1 regularization gives slightly more accurate reconstructions.

Although the capabilities of MTGV regularization have been shown with respect to 1D and 2D NMR relaxation data, the methods can be easily adapted to 1D NMR diffusion data and 2D NMR correlation experiments which are diffusion encoded.

5.5. Conclusions

A new method of regularization of 1D and 2D NMR relaxation and diffusion data has been implemented, termed Modified Total Generalized Variation (MTGV) regularization. Unlike other methods of regularization of Tikhonov and L_1, MTGV regularization offers a compromise between preserving smooth and discrete features in the reconstructed distributions. This eliminates the requirement of knowing *a priori* what the distribution should look like before selecting the appropriate regularization technique to process the data. The improvements offered by MTGV regularization were demonstrated by applying it to simulated 1D T_1, 1D T_2 and 2D T_1-T_2 NMR data. MTGV regularization always outperforms Tikhonov and L_1 regularization, apart from the case when the true distribution is known to only comprise of discrete peaks, in which case L_1 regularization performs slightly better.

References

[1] K. Bredies, K. Kunisch and T. Pock, Total generalized variation, SIAM J. Imaging Sci. 3 (2010), 492-526.

[2] F. Knoll, K. Bredies, T. Pock and R. Stollberger, Second order total generalized variation (TGV) for MRI, Magn. Reson. Med. 65 (2011), 480-491.

[3] M. Benning, C. Brune, M. Burger and J. Müller, Higher-order TV methods-enhancement via Bergman iteration, J. Sci. Comput. 54 (2013), 269-310.

[4] M. Benning, L. F. Gladden, D. Holland, C. Schönlieb and T. Valkonen, Phase reconstruction from velocity-encoded MRI measurements - A survey of sparsity-promoting variational approaches, J. Magn. Reson. 238 (2014), 26-43.

[5] S. Ono and I. Yamada, Second-order total generalized variation constraint, IEEE International Conference on Acoustic, Speech and Signal Processing (ICASSP) (2014), 4938-4942.

[6] H. Zou and T. Hastie, Regularization and variable selection via the elastic net, J. Royal Stat. Soc. B 67 (2005), 301-320.

[7] R. E. Carrillo, J. D. McEwen, D. van de Ville, J.-P. Thiran and Y. Wiaux, Sparsity averaging for compressive imaging, IEEE Signal Proc. Lett. 20 (2013), 591-594.

[8] A. Chambolle and T. Pock, An introduction to continuous optimization for imaging, Acta Numerica, 25 (2016), 161-319.

[9] T. Valkonen, K. Bredies and F. Knoll, Total generalized variation in diffusion tensor imaging, SIAM J. Imaging Sci. 6 (2013), 487-525.

[10] H. Celik, M. Bouhrara, D. A. Reiter, K. W. Fishbein and R. G. Spencer, Stabilization of the inverse Laplace transform of multiexponential decay through introduction of a second dimension, J. Magn. Reson. 236 (2013), 134-139.

[11] D. C. Montgomery and G. C. Runger, Applied Statistics and Probability for Engineers, 6th Edition, Wiley, 2013.

[12] J. Mitchell, T. C. Chandrasekera and L. F. Gladden, Numerical estimation of relaxation and diffusion distributions in two dimensions, Prog. Nucl. Magn. Reson. Spectrosc. 62 (2012), 34-50.

Chapter 6

Optimising magnetic resonance sampling patterns for parametric characterisation

Chapters 4 and 5 showed how a good choice of the reconstruction technique improves the estimation of parameters from ill-conditioned inverse problems in NMR applications. In this chapter and Chapter 7, the focus shifts to the design of sampling patterns for ill-conditioned inverse problems in NMR applications.

6.1. Introduction

Sampling patterns have long been a central topic in experimental design [1]. A good design of sampling patterns leads to a more accurate estimation of the extracted properties, lower experimental acquisition time and better conditioning of the inverse problem. The design of sampling patterns is particularly critical in the application of techniques such as Compressed Sensing, discussed in Section 3.2, which rely in the extraction of parameters from few sampled data points. The design of optimal sampling strategies depends on the application and it has been approached differently for example in MRI [2], NMR spectroscopy [3], NMR relaxation time analysis [4, 5], electronic spectroscopy [6], X-ray ptychography [7] and Helium Atom Scattering [8].

In this chapter, a systematic, statistical method, based on the CRLB theory, outlined in Section 3.4, is introduced for the optimisation of sampling patterns for NMR diffusion and relaxation experiments. The method is illustrated with respect to the design of optimal sampling patterns for the most accurate estimate of parameters characterising a lognormal distribution. Lognormal distributions are ubiquitous in science and engineering, ranging from the description of the population distribution of organisms [9] to the size distribution of materials produced from particle processing techniques [10, 11]. In NMR applications,

lognormal distributions have been assumed to be a good approximation for polymer size distributions [12], molecular aggregate length distributions [13] and reverse micelle size distributions [14] obtained from diffusion experiments; T_1 distributions of heavy oils obtained from fast-field cycling [15] and inversion recovery experiments [16]; and T_2 distributions of heavy oils obtained from CPMG experiments [16, 17].

The proposed method is validated against NMR diffusion experiments of an emulsion of toluene in water. The accurate measurement of the emulsion droplet size distribution is important in the food, pharmaceutical and oil recovery industry, among other areas [18]. Since its development [19], the measurement of the emulsion droplet size distribution using NMR diffusion experiments has become an established characterisation technique [20-22]. The emulsion droplet size distribution obtained from NMR diffusion experiments is commonly approximated to a lognormal distribution [19, 23-25]; this is supported by population balance statistics between droplet breakage and coalescence during emulsification [26] and by experimental results from other characterisation techniques such as dynamic light scattering [27] and confocal scanning laser microscopy [28].

The complete optimisation of an NMR experiment would also require the optimisation of specific NMR acquisition parameters and reconstruction techniques, which are separate from the optimisation of the sampling pattern considered in this chapter. The optimisation of specific NMR acquisition parameters and reconstruction techniques for NMR diffusion experiments have been covered in detail elsewhere [22, 29, 30] and will not be addressed here.

The chapter is structured as follows. Section 6.2 introduces the theory behind obtaining the emulsion droplet size distribution from NMR diffusion experiments and the application of the CRLB theory to these experiments. The experimental setup and sampling patterns investigated are described in Section 6.3. The comparison between the predictions of the CRLB theory and NMR diffusion experimental results is presented in Section 6.4. The limitations, sensitivity and potential of the CRLB theory are also discussed in Section 6.4.

6.2. Mathematical treatment

Section 6.2.1 briefly reviews the theory behind the extraction of the emulsion droplet size distribution from NMR diffusion experiments. Section 6.2.2 outlines the application of the CRLB theory to the problem of optimising sampling patterns for the most accurate estimation of the lognormal distribution parameters of an emulsion droplet size distribution.

6.2.1. NMR diffusion of emulsion systems

The ideal NMR signal attenuation acquired from a diffusion experiment of an unconstrained chemical component, is described by the Stejskal-Tanner exponential decay, as discussed in Section 2.11. If the component is constrained to be in the dispersed phase of an emulsion and the diffusion time is such that the root mean square distance travelled due to Brownian motion is larger than the characteristic size of the droplet, the apparent diffusion coefficient of the component is smaller than the unconstrained diffusion coefficient, with the apparent value depending on the droplet size. Since the emulsion is characterized by a distribution of droplet sizes, the NMR signal attenuation is a multi-exponential decay, from which the droplet size distribution can be extracted. The noise-free NMR signal attenuation acquired from the dispersed phase of an emulsion with droplet number size distribution $f_0(a_j)$, has been calculated by Murday and Cotts [31]:

$$\hat{y}_i = A \sum_j {a_j}^3 f_0(a_j)\, R(g_i, a_j), \tag{6.1}$$

where a_j is the discretized list of droplet radii and A is a scaling factor. The standard notation for the number size distribution [32], $f_0(a_j)$, has been used. The factor ${a_j}^3$ could be combined with $f_0(a_j)$ to give the volume size distribution, $f_3(a_j)$, but we choose to focus on the number size distribution, for ease of interpretation. In Eq. (6.1):

$$R(g_i, a_j) = \exp\left(-2\gamma^2 {g_i}^2 \sum_{k=1}^{\infty} \frac{1}{\alpha_k^2\left(\alpha_k^2 a_j^2 - 2\right)} r(\alpha_k)\right), \tag{6.2}$$

where:

$$r(\alpha_k) = \frac{2\delta}{\alpha_k^2 D} + \frac{-2 - \exp\left(-\alpha_k^2 D(\Delta-\delta)\right) + 2\exp\left(-\alpha_k^2 D\Delta\right) + 2\exp\left(-\alpha_k^2 D\delta\right) - \exp\left(-\alpha_k^2 D(\Delta+\delta)\right)}{\left(\alpha_k^2 D\right)^2}. \tag{6.3}$$

In Eqs. 6.2 and 6.3, δ are Δ are the specific NMR diffusion acquisition parameters, defined in Section 2.11, D is the unconstrained diffusion coefficient of the dispersed phasel, g_i is

the list of magnetic field gradient values used, which also refer to the sampling pattern, and α_k are the solutions to the equation:

$$J_{3/2}(\alpha_k a_j) = \alpha_k a_j J_{5/2}(\alpha_k a_j) \,. \tag{6.4}$$

J_k is the Bessel function of the first kind and of order k. As discussed in Section 6.1, the droplet number size distribution, $f_0(a_j)$, is typically well approximated to a lognormal distribution:

$$f_0(a_j) = \frac{1}{a_j s_g \sqrt{2\pi}} \exp\left(-\frac{\left(\ln a_j - \bar{a}_{0,0}\right)^2}{2 s_g{}^2}\right), \tag{6.5}$$

where $\bar{a}_{0,0}$ is the geometric mean of the droplet size distribution and s_g is the geometric standard deviation of the droplet size distribution, following standard notation (the arithmetic mean of the droplet size distribution is $\bar{a}_{1,0}$ and the standard deviation of the droplet size distribution is s). If the droplet number size distribution is lognormal, the volume size distribution is also lognormal [33]. As a result, a similar analysis of what follows in the next sections can also be applied to the volume size distribution, $f_3(a_j)$.

6.2.2. Application of CRLB theory to emulsion systems

Combining Eqs. (6.1) and (6.5), the noise-free NMR diffusion signal, $\hat{y}_i$, is given by:

$$\hat{y}_i = A \sum_j \frac{a_j{}^2}{s_g \sqrt{2\pi}} \exp\left(-\frac{\left(\ln a_j - \bar{a}_{0,0}\right)^2}{2 s_g{}^2}\right) R(g_i, a_j), \tag{6.6}$$

and, therefore, depends on the lognormal distribution parameters $\bar{a}_{0,0}$ and s_g and the scaling factor, A. The accurate estimation of these parameters is the objective of the NMR experiment. In the presence of noise, the estimate of each of the parameters is uncertain. The uncertainty associated with the estimate of each of the parameters can be calculated from the CRLB theory, which was outlined in Section 3.4. The standard deviation associated with the estimation of any of the parameters, θ_l, has a lower limit:

$$\mathrm{std}(\theta_l) \geq \sqrt{(\boldsymbol{F}^{-1})_{ll}} \,, \tag{6.7}$$

where $\boldsymbol{F}$ is the Fisher Information Matrix. If the noise elements related to each acquired data point are independently normally distributed, with a constant variance, σ^2, it was shown in Section 3.4 that the Fisher Information can be constructed from:

$$F_{l_1 l_2} = \frac{1}{\sigma^2}\left(\sum_i \frac{\partial \hat{y}_i}{\partial \theta_{l_1}} \frac{\partial \hat{y}_i}{\partial \theta_{l_2}}\right). \tag{6.8}$$

Therefore, to construct $\boldsymbol{F}$, the partial derivatives of $\hat{y}_i$ with respect to A, $\bar{a}_{0,0}$ and s_g are needed, which are calculated from Eq. (6.6):

$$\frac{\partial \hat{y}_i}{\partial A} = \sum_j \frac{{a_j}^2}{s_g\sqrt{2\pi}} \exp\left(-\frac{\left(\ln a_j - \bar{a}_{0,0}\right)^2}{2{s_g}^2}\right) R(g_i, a_j) \,, \tag{6.9a}$$

$$\frac{\partial \hat{y}_i}{\partial \bar{a}_{0,0}} = A\sum_j \frac{{a_j}^2\left(\ln a_j - \bar{a}_{0,0}\right)}{{s_g}^3\sqrt{2\pi}} \exp\left(-\frac{\left(\ln a_j - \bar{a}_{0,0}\right)^2}{2{s_g}^2}\right) R(g_i, a_j) \,, \tag{6.9b}$$

$$\frac{\partial \hat{y}_i}{\partial s_g} = A\sum_j \frac{{a_j}^2\left[\left(\ln a_j - \bar{a}_{0,0}\right)^2 - {s_g}^2\right]}{{s_g}^4\sqrt{2\pi}} \exp\left(-\frac{\left(\ln a_j - \bar{a}_{0,0}\right)^2}{2{s_g}^2}\right) R(g_i, a_j) \,. \tag{6.9c}$$

The important result is that the uncertainty associated with the estimate of each parameter depends on A, $\bar{a}_{0,0}$, s_g, σ and most importantly on the sampling pattern, g_i. As a result, given the noise level and an experimental system (characterised by A, $\bar{a}_{0,0}$ and s_g), the error associated with the estimate of each parameter depends solely on the sampling pattern, g_i. Therefore, by optimizing the sampling pattern, it is possible to decrease to a minimum the uncertainty associated with the estimate of a given parameter, hence, making the estimate of that parameter as accurate as possible.

The optimal sampling pattern for the most accurate estimation of one of the parameters is not necessarily the optimal sampling pattern for the most accurate estimation of another parameter. Therefore, it is necessary to know what the most useful information for the particular application is, and to express it in terms of an objective function. This objective function could simply be the minimization of std(A), std($\bar{a}_{0,0}$) or std(s_g), or more complicated expressions such as different moments of the lognormal distribution [34]. The present method can cope with different objective functions. In this chapter, the choice is made to minimize the maximum error bar associated with each individual point in the droplet size number distribution, given by:

$$\eta = \max_j \left[\mathrm{std}\left(f_0(a_j)\right)\right] \approx \max_j \left[\left|\frac{\partial f_0(a_j)}{\partial \bar{a}_{0,0}}\mathrm{std}\left(\bar{a}_{0,0}\right)\right| + \left|\frac{\partial f_0(a_j)}{\partial s_g}\mathrm{std}\left(s_g\right)\right|\right] . \tag{6.10}$$

The derivation of Eq. (6.10) is outlined in the Appendix of this chapter. Using this definition, the best sampling pattern has the minimum objective function, η. The choice of the objective function is made purely to demonstrate the capabilities of the method, but other objective functions can be defined. It can be shown that std($\bar{a}_{0,0}$) $\propto \sigma A^{-1}$, std(s_g) $\propto \sigma A^{-1}$ and therefore $\eta \propto \sigma A^{-1}$. The ratio A/σ determines the signal-to-noise ratio (SNR). As a result, for a given SNR, the objective function, η, depends solely on the

lognormal distribution parameters $\bar{a}_{0,0}$ and s_g and the sampling pattern, g_i. In the following discussion, when stating σ, the work is referring to the noise level for a signal normalized to a maximum of 1.

6.3. Materials and methods

6.3.1. Experimental

An emulsion of toluene dispersed in water is prepared as follows. A non-ionic surfactant, Triton^{TM} X-100 (Sigma-Aldrich, laboratory grade), is first dissolved in deionised water. Toluene (Sigma-Aldrich, ≥ 99 % purity) is then mixed into the solution using a Janke & Kunkel RW-20 stirrer, equipped with a 2 cm radius stainless steel blade, driven at 150 rpm for 10 min. The emulsion formed is composed of 48.5 wt% water, 48.5 wt% toluene and 3 wt% surfactant. A small volume from the batch is transferred into a 5 mm NMR tube and a PFG NMR diffusion experiment is performed on this sample. The measurement time was 2.5 h. A second acquisition immediately after the first produced the same emulsion droplet size distribution, thereby confirming that the emulsion is stable during the course of the experiment.

Experiments are performed at 20 ± 0.5 °C on a Bruker DMX 300 spectrometer, operating at a resonant frequency of 300.13 MHz for ^{1}H nucleus observation. The maximum gradient amplitude available is 1176 G cm^{-1} and the r.f. coil has a diameter of 5 mm.

The unconstrained diffusion coefficient of toluene, D, is measured to be $(2.06 \pm 0.02) \times 10^{-9}$ m^2 s^{-1} using a stimulated echo PFG NMR pulse sequence outlined in Sections 2.8 and 2.11, with pulsed field gradient duration δ = 1.5 ms, diffusion time Δ = 200 ms and the magnetic field gradient amplitude, g_i, in the range 0.1 – 40 G cm^{-1} (16 equidistant points). This agrees with reported values at this temperature [35]. The same pulse sequence is then applied to the emulsion system, with the following acquisition parameters: δ = 5 ms, Δ = 400 ms, g_i = 4.8 – 300 G cm^{-1}. 252 gradient strength values, g_i, are used for two reasons: to provide an accurate estimate of the lognormal distribution parameters of the emulsion droplet size distribution, and to be used for the selection of different sampling patterns, as described in Section 6.3.2.

The attenuation of the aromatic peak intensity only is followed to estimate the droplet size distribution. The contribution to the aromatic peak intensity from the surfactant, as

compared to toluene, is small (calculated to be 0.9%) and is taken into account by subtracting it from the overall aromatic peak intensity at all the gradient values used. The signal attenuation from the aromatic peak of toluene in the PFG NMR diffusion experiment of the emulsion of toluene in water is shown in Fig. 6.1(a). A curvature is observed in the plot of the signal against the sampling pattern, g_i^2. This confirms that toluene experiences restricted diffusion resulting from being confined within droplets. The signal attenuation is fitted to a lognormal distribution for the droplet size distribution; the best fit lognormal distribution with $\bar{a}_{0,0} = 1.09$, $s_g = 0.38$ ($\bar{a}_{1,0} = 3.2$ μm, $s = 1.2$ μm) is shown in Fig. 6.1(b). The error associated with each individual point in the discretized lognormal distribution is negligible because of the relatively small σ (estimated to be 5.6×10^{-4}) and the large number of points in the decay (252).

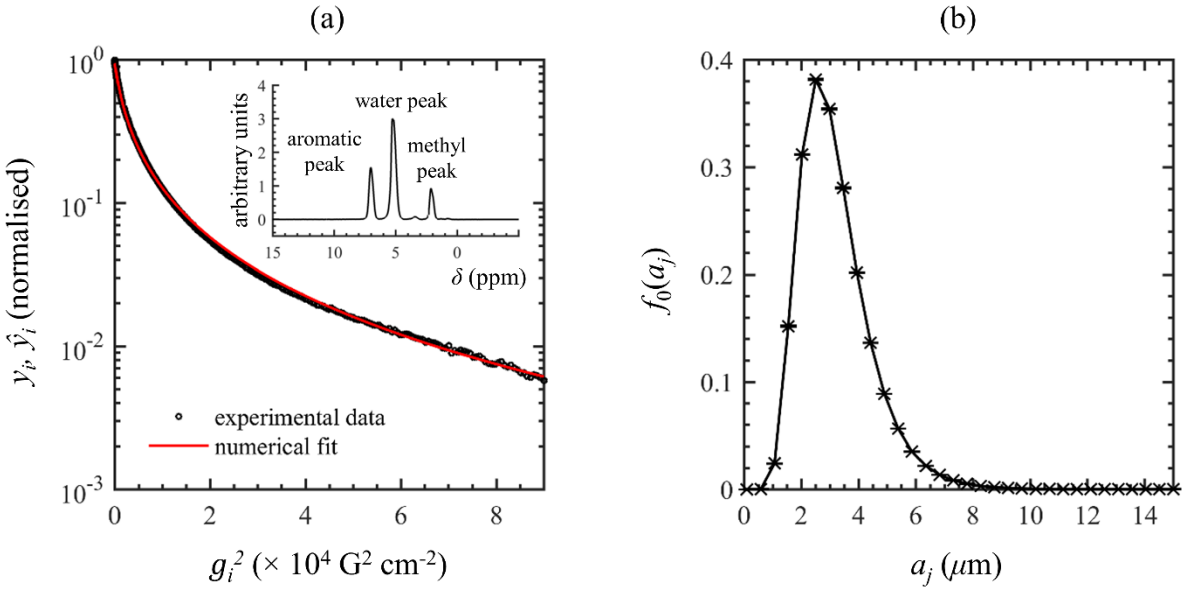

Fig. 6.1. (a) Experimental data for the attenuation of the aromatic peak of toluene (chemical shift $\delta \sim 7.1$ ppm with respect to tetramethylsilane, shown in the embedded figure) from the PFG NMR diffusion experiment of an emulsion of toluene in water. The numerical fit is estimated from the best fit lognormal distribution for the emulsion droplet size distribution. The estimated noise level from the fit is $\sigma = 5.6 \times 10^{-4}$. (b) Best fit lognormal distribution for the emulsion droplet size distribution with $\bar{a}_{0,0} = 1.09$, $s_g = 0.38$ ($\bar{a}_{1,0} = 3.2$ μm, $s = 1.2$ μm).

The emulsion droplet size distribution is reconstructed for a discrete list of droplet radii, a, composed of 32 points equidistantly spaced between 0.1-15 μm, as was used in the work of Hollingsworth and Johns [22]. Only the first 6 terms are kept from the solutions of the Bessel equation in Eq. (6.4); it was observed that keeping more terms made no difference to the droplet size distribution predicted.

6.3.2. Sampling patterns

A sampling pattern for the acquisition of the PFG NMR data consists of the choice of g_i, the list of pulsed field gradients at which the attenuation of the NMR signal is acquired. This involves choosing the number of pulsed field gradients used, n, and their respective values. n is typically a user-defined variable and larger values of n give more accurate estimate of parameters. In this chapter, n is kept fixed at 32 except when combining n and σ in a way which keeps either the experimental acquisition time or the parameter estimation accuracy fixed; these results are reported in Sections 6.4.3 and 6.4.4, respectively.

For a given number of points, n, the optimal choice of g_i would require setting the partial derivatives of the expression in Eq. (6.10) with respect to each individual g_i to zero and solving the set of resulting n simultaneous non-linear equations; this is computationally intractable. As a result, this chapter focuses on finding the optimal sampling pattern among classes of practical sampling patterns. Two classes of sampling patterns are considered: power law spacing of the linear version of sampled points and power law spacing of the logarithmic version of sampled points.

The power law spacing of the linear version of sampled points, which for simplicity will be referred to as a linear sampling pattern, takes the form:

$$g_i = g_1 + (g_n - g_1)\left(\frac{i-1}{n-1}\right)^r. \tag{6.11}$$

g_1 and g_n are the first and the last pulsed field gradients used and r determines the density distribution of the data points; $r > 1$ gives a higher density of points at the start of the decay, $r < 1$ gives a higher density of points at the end of the decay and $r = 1$ refers to equidistant pulsed field gradient values. Acquiring the point that has the largest amplitude is always beneficial because it contains the most information about the attenuation of the signal [36-38]. Therefore, g_1 was kept at the minimum pulsed field gradient used experimentally of 4.8 G cm^{-1}. As a result, the optimisation of the linear sampling pattern effectively consists of optimizing g_n and r.

The power law spacing of the logarithmic version of sampled points, which for simplicity will be referred to as a logarithmic sampling pattern, takes the form:

$$\log_{10} g_i = \log_{10} g_1 + (\log_{10} g_n - \log_{10} g_1)\left(\frac{i-1}{n-1}\right)^r, \tag{6.12}$$

where the parameters are defined similarly to the linear sampling pattern. Using the same assumptions as for the linear sampling scheme, optimisation of the logarithmic sampling pattern consists of optimizing g_n and r.

The comparison between the predictions of the CRLB theory and experimental results is performed as follows. Using a range of g_n from 20 to 300 G cm^{-1} and a range of r from 0.1 to 10, 10^4 linear sampling patterns and 10^4 logarithmic sampling patterns with $n = 32$ points are simulated using Eq. (6.11) and Eq. (6.12), respectively. For each of these sampling patterns, the CRLB theory is used to predict the objective function, η, defined in Eq. (6.10). This is compared to the objective function obtained from the experimental data which is calculated as follows. For each of the simulated sampling patterns, the closest experimental 32 g_i points (out of the total 252 g_i points acquired) are selected. The experimental attenuation data corresponding to these 32 points are then fitted to a lognormal distribution for the droplet size distribution and confidence intervals for $\bar{a}_{0,0}$ and s_g are extracted. std($\bar{a}_{0,0}$) and std(s_g) are then calculated from the confidence intervals using a z-test and these values are used to calculate the experimental objective function, η.

The difference between the simulated and the experimental sampling patterns is negligible. This ensures that the comparison between the predictions of the CRLB theory (which is performed using the simulated sampling pattern) and the experimental results (which uses the selected experimental points) is unaffected by mismatches in the corresponding sampling patterns used.

6.4. Results and discussion

In Section 6.4.1, the predictions of the CRLB theory are validated against the experimental data using the procedure outlined in Section 6.3.2. The limitations and sensitivity of the CRLB theory are subsequently discussed in Section 6.4.2. The CRLB theory is then used to optimise the sampling pattern for a fixed acquisition time experiment (Section 6.4.3) and to determine the potential for reduction in the acquisition time, while keeping the parameter estimation accuracy fixed (Section 6.4.4).

6.4.1. Validation of CRLB theory against experimental data

The comparison of the objective function, η, as predicted by the CRLB theory and as calculated from the experimental data is shown in Fig. 6.2. Small values of η indicate a good sampling pattern.

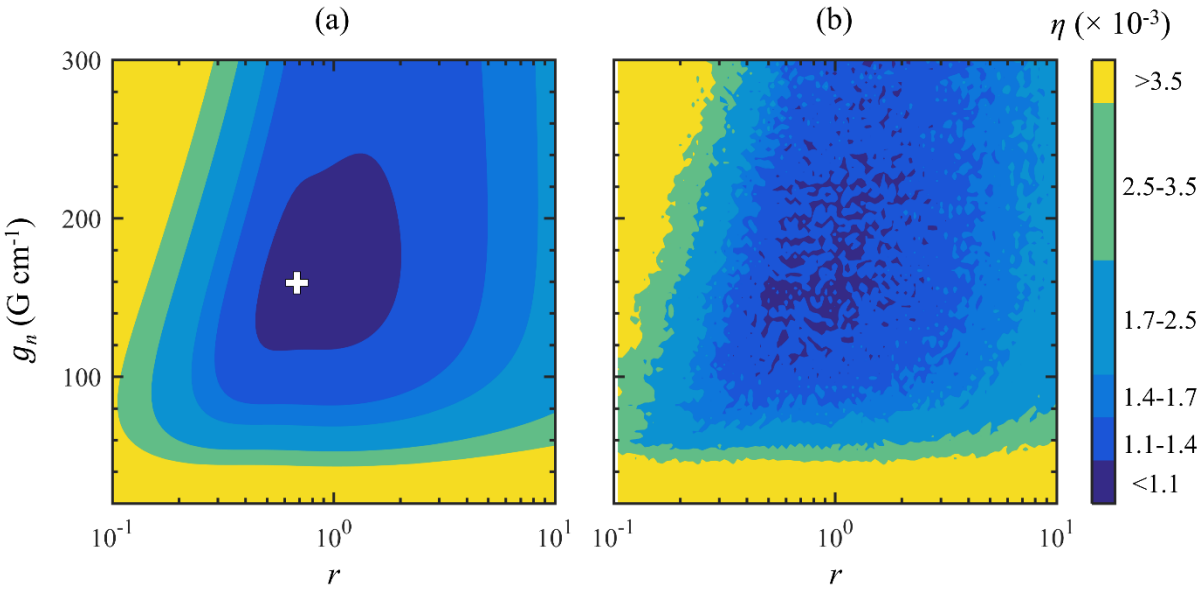

Fig. 6.2. (a) Contour map of the variation of the objective function, η, with respect to the parameters of the class of linear sampling patterns defined in Eq. (6.11) with n = 32, as predicted by the CRLB theory. (b) Contour map of the variation of the objective function, η, with respect to the parameters of the class of linear sampling patterns defined in Eq. (6.11) with n = 32, obtained from the experimental data shown in Fig. 6.1(a). The best linear sampling pattern in (a) is obtained for r = 0.67, g_n = 159 G cm^{-1}, and is identified as '+', at which point $\eta = 1.03 \times 10^{-3}$.

Fig. 6.2(a) shows the contour map of η as a function of linear sampling pattern parameters g_n and r, as predicted by the CRLB theory. Fig. 6.2(b) shows a similar contour map obtained from the experimental data. There is very good agreement between the CRLB and experimental contour maps; on average there is < 5% difference between the corresponding η values of the two maps, thereby validating the application of the CRLB theory to the optimization of sampling patterns. For a given value of r, for example r = 1 (corresponding to equidistant points), the value of the objective function is large at small g_n because late decay points are not sampled and, therefore, the uncertainty related to estimating the small droplet sizes (corresponding to small apparent diffusion coefficients) is large. The value of the objective function is also large at very large g_n because early decay points are not sampled properly and, therefore, the uncertainty related to estimating the large droplet sizes (corresponding to large apparent diffusion coefficients) is large. As

a result, for any given r, there is an optimum g_n which minimises the objective function. For a given g_n, for example g_n = 150 G cm^{-1}, the value of the objective function is large at small r because early decay points are sampled sparsely and, therefore, the uncertainty related to estimating the large droplet sizes is large. The value of the objective function is also large at very large r because late decay points are sampled sparsely and, therefore, the uncertainty related to estimating the small droplet sizes is large. As a result, for any given g_n, there is an optimum r which minimises the objective function. In summary, there is an optimum combination (g_n,r) at which the objective function obtains a minimum, which corresponds to the best linear sampling pattern. The best linear sampling pattern for 32 sampled points is predicted by the CRLB theory at r = 0.67 and g_n = 159 G cm^{-1}, at which point $\eta = 1.03 \times 10^{-3}$.

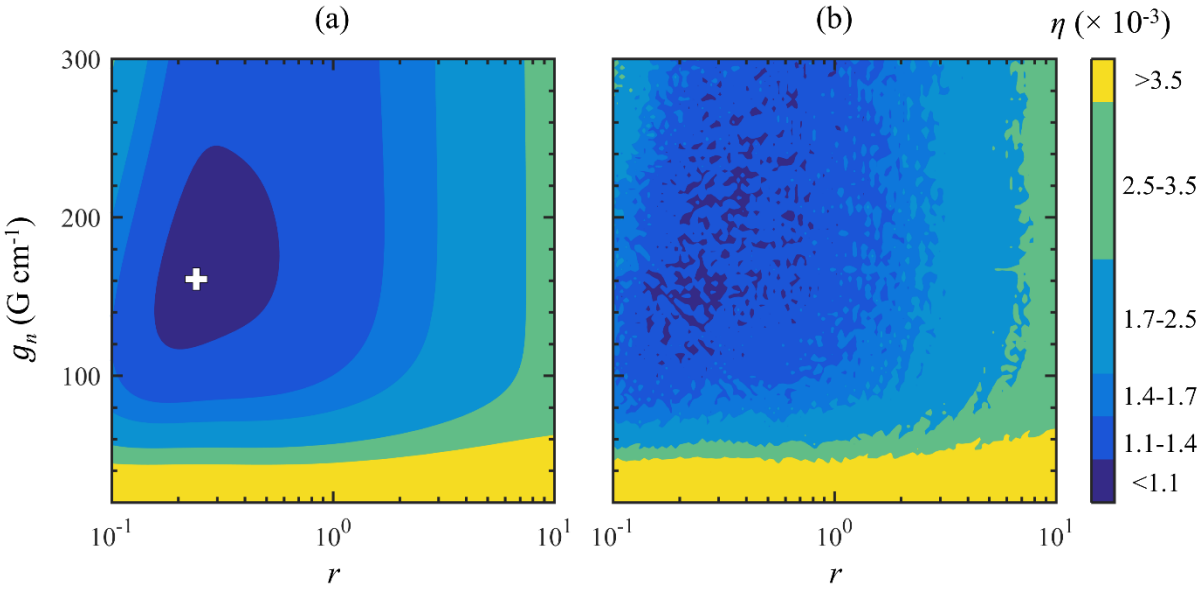

Fig. 6.3. (a) Contour map of the variation of the objective function, η, with respect to the parameters of the class of logarithmic sampling patterns defined in Eq. (6.12) with n = 32, as predicted by the CRLB theory. (b) Contour map of the variation of the objective function, η, with respect to the parameters of the class of logarithmic sampling patterns defined in Eq. (6.12) with n = 32, obtained from the experimental data shown in Fig. 6.1(a). The best logarithmic sampling pattern in (a) is obtained for r = 0.23, g_n = 161 G cm^{-1}, and is identified as +, at which point $\eta = 1.03 \times 10^{-3}$.

A similar comparison is made between the predictions of the CRLB theory and the experimental data for logarithmic sampling patterns of 32 points selected according to Eq. (6.12). The results are shown in Fig. 6.3. Again, very good agreement is observed between the contour maps; on average there is < 5% difference between the corresponding η values of the two maps, which further validates the application of the CRLB theory to

the optimization of sampling patterns. The best logarithmic sampling pattern for 32 sampled points is predicted by the CRLB theory at $r = 0.23$ and $g_n = 161$ G cm^{-1}, at which point $\eta = 1.03 \times 10^{-3}$.

It is noted that the minimum objective function value for linear sampling patterns is equal to the minimum objective function value for logarithmic sampling patterns ($\eta = 1.03 \times 10^{-3}$), although the two classes of sampling patterns are very different. This is an indication that this minimum objective function is a global minimum (with respect to all possible classes of sampling patterns). Given this observation, in the following sections, the discussion will be limited to linear sampling patterns. For completeness it is noted that, as discussed in Section 6.3.2, the proof that this minimum objective function is a global minimum for all sampling patterns would require the solution of an intractable computational problem.

6.4.2. Limitations and sensitivity

Discrepancies between the CRLB theory predictions and experimental data have been previously reported at large σ (low SNR) [36]. Considering the increasing interest in low-field measurements, where the signal-to-noise ratio may limit the extraction of parameters from the NMR signal, it is important to find the limit of applicability of the CRLB theory with respect to the SNR. The following discussion aims to find the limit of applicability of the CRLB theory to the present problem of estimating the lognormal distribution parameters of the emulsion droplet size distribution.

Monte Carlo type simulations [39] are performed as follows. The 32 experimental data points acquired at the optimal linear sampling pattern, reported in Section 6.4.1, are extracted. Random Gaussian noise is added to these acquired points, such that noise levels, σ, in the range 10^{-3}-5×10^{-1} are studied. For each noise level, 100 signals are simulated, differing only in the randomness of noise. Each of these signals is fitted to a lognormal distribution for the emulsion droplet size distribution. The values of $\bar{a}_{0,0}$ and s_g from each of these fittings are used to estimate std($\bar{a}_{0,0}$) and std(s_g). These standard deviations are used to calculate the objective function, η, according to Eq. (6.10). This is then compared to the CRLB predictions. The results are shown in Fig. 6.4(a). It is observed that there is very good agreement between the Monte Carlo simulations and the CRLB theory predictions up to $\sigma \sim 0.1$ (SNR ~ 10). This limit will be slightly different for systems

characterised by different $\bar{a}_{0,0}$ and s_g. However, a general conclusion can be made that the method is stable even at highly noisy data (i.e. SNR ~ 10).

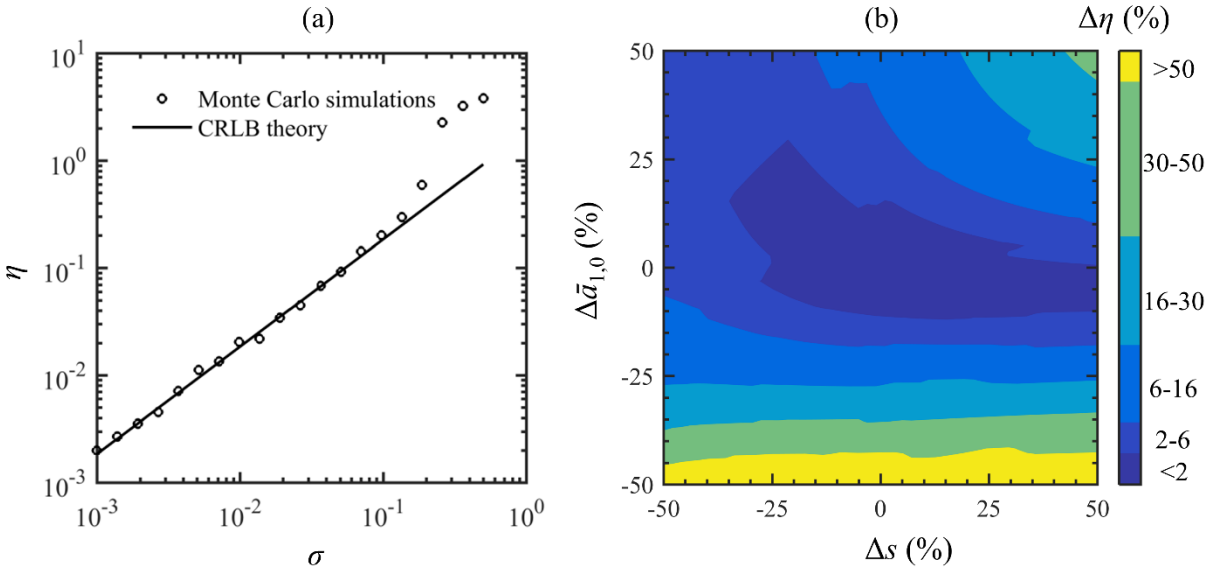

Fig. 6.4. (a) Variation of the objective function, η, with the noise level, σ, as calculated by Monte Carlo type simulations and as predicted by the CRLB theory. The sampling pattern used is the best linear sampling pattern for $n = 32$ with $r = 0.67$, $g_n = 159$ G cm^{-1}. (b) Contour map showing the percentage increase in the objective function, $\Delta\eta$, caused by the mis-design of the linear sampling pattern defined in Eq. (6.11) because of a percentage error in guessing the arithmetic mean, $\Delta\bar{a}_{1,0}$ (μm), and standard deviation, Δs (μm), of the emulsion droplet size distribution. The mis-designed sampling pattern used is the best linear obtained in Fig. 6.2(a) for $n = 32$ with $r = 0.67$, $g_n = 159$ G cm^{-1} and the noise level used in the CRLB predictions is $\sigma = 5.6 \times 10^{-4}$.

As discussed in Section 6.2.2, the application of the CRLB theory requires an initial guess of the lognormal distribution parameters, $\bar{a}_{0,0}$ and s_g. A valid question is: How accurate does this guess need to be? Since it is easier to work in terms of the arithmetic mean, $\bar{a}_{1,0}$ (μm), and standard deviation, s (μm), it is more sensible to ask about the accuracy with which $\bar{a}_{1,0}$ and s need to be *a priori* known. To answer this question, a sensitivity analysis is performed as follows. Assume that the optimal linear sampling pattern of 32 points found in Section 6.4.1 is designed on the basis that $\bar{a}_{1,0} = 3.2$ μm, $s = 1.2$ μm. Now, assume that this estimate for $\bar{a}_{1,0}$ and s is wrong, and the true value for the arithmetic mean and standard deviation of the lognormal distribution is actually $\bar{a}_{1,0t}$ and s_t. For these true values of $\bar{a}_{1,0t}$ and s_t, the CRLB theory can be used to find the minimum objective function, η_t, corresponding to the ideal linear sampling pattern, and the actual value of η evaluated for the actual sampling pattern used. The percentage difference of η from η_t is an indication of how sensitive the CRLB theory is to wrong initial guesses of the lognormal distribution

parameters. A contour map of the percentage difference of η from η_t, as a function of the percentage difference of $\bar{a}_{1,0}$ from $\bar{a}_{1,0t}$ and s from s_t is shown in Fig. 6.4(b). The key observation is that the CRLB theory is relatively insensitive to the guess of the standard deviation, s, as compared to the sensitivity to the arithmetic mean, $\bar{a}_{1,0}$. This is important as in general it is more difficult to guess s than $\bar{a}_{1,0}$. With regards to the sensitivity to the guess of $\bar{a}_{1,0}$, the CRLB theory is relatively insensitive around the true value, $\bar{a}_{1,0t}$, but becomes increasingly more sensitive the further away from the true value the guess is.

In this chapter, the prior knowledge about the lognormal distribution parameters was obtained by sampling a large number of points (252) at a low noise standard deviation, $\sigma = 5.6 \times 10^{-4}$. In practice, the acquisition of such a large number of points defeats the purpose of optimising the sampling pattern. It is therefore advised to sample a small number of points at a high noise standard deviation (through less signal averaging). A fit to these experimental data would provide a crude prior knowledge about the diffusion coefficients and the populations, which could then be used to optimise the sampling pattern using the proposed method.

The versatility of the CRLB theory to be adapted to different objectives is now considered with respect to two cases studies. In Section 6.4.3 the sampling scheme is optimised while keeping the acquisition time fixed, while in Section 6.4.4 the potential for reduction in the acquisition time is investigated while keeping the parameter estimation accuracy fixed.

6.4.3. Optimising the sampling scheme for a fixed acquisition time

The optimization of sampling patterns has so far been made for a fixed number of points, $n = 32$. As discussed in section 6.3.2, for a fixed noise level σ, acquiring more points will lead to a more accurate estimate for the lognormal distribution parameters. However, this comes at the cost of a longer acquisition time. If the acquisition time is limited by, for example, the stability of the emulsion under investigation, is it better to acquire fewer points with low σ (high SNR), or more points with higher σ (low SNR)?

The CRLB theory can be used to answer this question as follows. In a PFG NMR diffusion experiment, the acquisition time, t, is related to the number of points acquired, n, and the noise level, σ, by $t \propto n\ \sigma^{-2}$ [40]. Therefore, for a fixed acquisition time, $\sigma \propto n^{1/2}$. For a given n and the corresponding σ (for a fixed acquisition time), the CRLB theory can be used to determine the minimum objective function, η, for linear sampling patterns, as was done to produce Fig. 6.2(a). This is repeated for n in the range 3-32, and the plot of the

minimum η, as a function of n is shown in Fig. 6.5(a). At large n, the acquired data are noisy and therefore the best sampling pattern for that n has a relatively large minimum η. At small n, although the data are of a high signal-to-noise ratio, there are not many points to fit the lognormal distribution parameters to. As a result, the minimum η is again relatively large. The optimum linear sampling pattern is to acquire 7 points with linear sampling pattern parameters $r = 0.98$ (close to being equidistant points) and $g_n = 147$ G cm^{-1}, which is illustrated in Fig. 6.5(b).

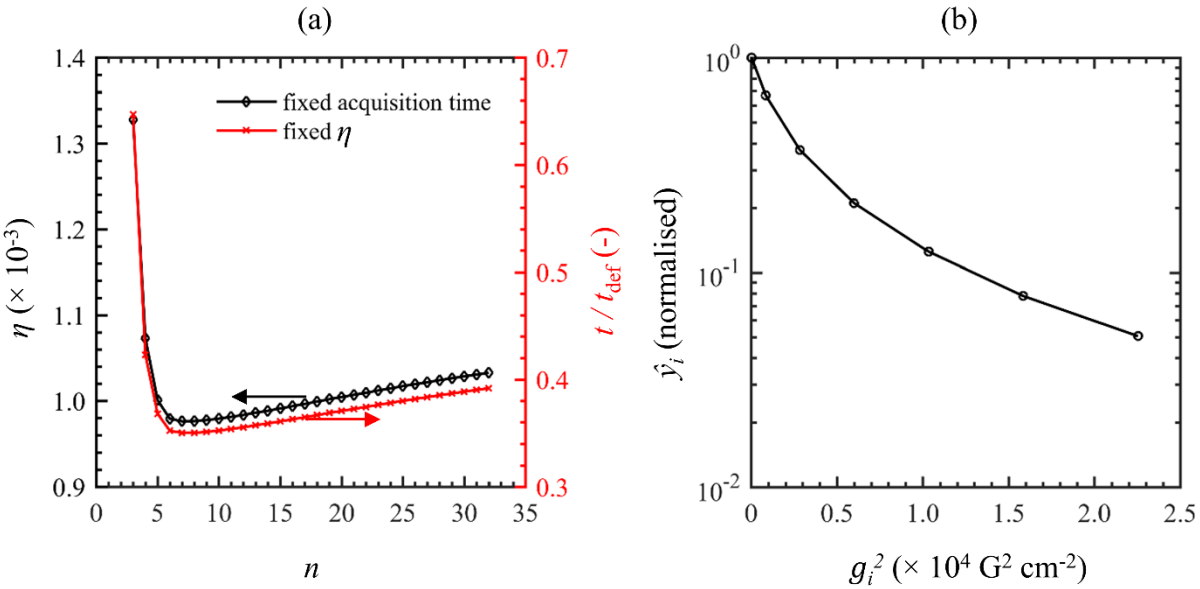

Fig. 6.5. (a) Variation of the minimum objective function, η, with the number of data points in a linear sampling pattern, n, calculated by the CRLB theory for the case of a fixed acquisition time experiment (◊). Variation of the reduction in acquisition time, t / t_{def}, with n calculated by the CRLB theory for the case of a fixed η experiment (×). The default acquisition time, t_{def}, was calculated for a default sampling pattern of 16 equidistant points and $g_n = 300$ G cm^{-1}. The lines are included to guide the eye. (b) Best linear sampling pattern for $n = 7$ with $r = 0.98$, $g_n = 147$ G cm^{-1}, which according to (a) is the optimal sampling pattern for both a fixed acquisition time experiment and a fixed η experiment. The line is included to guide the eye.

6.4.4. Reducing the acquisition time for a fixed parameter estimation accuracy

While the parameter estimation accuracy achieved using a default sampling pattern may be acceptable in an experiment, it is important to know whether the same accuracy can be achieved for a shorter acquisition time. This problem is investigated with respect to a hypothetical experiment, but the idea extends easily to other experiments. It is assumed that the system under the investigation is the emulsion used in this chapter, and that the default sampling pattern that was initially being used to study the emulsion by a PFG NMR

diffusion experiment was composed of 16 points equidistant from each other (linear sampling pattern with r = 1), with the largest gradient value being g_n = 300 G cm^{-1}. Using the CRLB theory, the objective function, η_{def}, that characterizes the parameter estimation accuracy of this default sampling pattern can be calculated after a reference noise level, σ_{ref}, is defined. As discussed in Section 6.4.3, the acquisition time for this sampling pattern is given by $t_{def} \propto 16\ \sigma_{ref}^{-2}$.

The reduction ratio in acquisition time, t / t_{def}, which can be achieved by better sampling, while keeping $\eta = \eta_{def}$ is calculated as follows. For a given n, the CRLB theory can be used to determine the minimum objective function, η, as a function of the noise level, σ, for linear sampling patterns, as was done to produce Fig. 6.2(a). The noise level is then tuned such that $\eta = \eta_{def}$. Using the tuned noise level, the acquisition time for the chosen sampling pattern is calculated as $t \propto n\ \sigma^{-2}$. The process is repeated for n in the range 3-32 and the variation of t / t_{def} with n is shown in Fig. 6.5(a). It is observed that the minimum acquisition time is achieved for a linear sampling pattern of n = 7 points with linear sampling pattern parameters r = 0.98 (close to being equidistant points) and g_n = 147 G cm^{-1}, which is illustrated in Fig. 6.5(b). The reduction ratio in acquisition time achieved is t / t_{def} = 0.35, corresponding to an acceleration factor of ~3.

The fact that the same optimal sampling pattern is achieved for both the process of keeping t fixed and minimizing η (discussed in Section 6.4.3) and the process of keeping η fixed and minimizing t is no coincidence. Indeed, it is shown in the appendix that $t_{\text{fixed }\eta} \propto \eta^2_{\text{fixed }t}$, which guarantees the same optimum n value where the minimum of either of these two functions is achieved.

6.5. Conclusions

A method was presented, based on the Cramér-Rao Lower Bound theory, for designing optimal sampling patterns for NMR experiments, illustrated with respect to the most accurate estimation of lognormal distribution parameters. The method was validated against experimental data of an NMR diffusion experiment of an emulsion of toluene in water, used to obtain the emulsion droplet size distribution. The difference between the predictions of the CRLB theory and experimental data was < 5% and it was shown that the CRLB theory is stable even at SNR ~ 10. Using the CRLB theory, it was demonstrated

that for a PFG NMR experiment which usually uses constant spacing between gradient increments, the acquisition time can be reduced by a typical factor of 3 while keeping the parameter estimation accuracy unchanged, through optimization of the sampling pattern. The presented approach can guide experimental design in wide range of experiments.

Appendix A: Derivation of Eq. (6.10)

Let x, z be two random variables, which may be correlated, and $h(x, z)$ be a function of these two random variables. Then, from the truncated Taylor series expansion [41]:

$$\mathrm{var}(h) \approx \left(\frac{\partial h}{\partial x}\right)^2 \mathrm{var}(x) + \left(\frac{\partial h}{\partial z}\right)^2 \mathrm{var}(z) + 2\frac{\partial h}{\partial x}\frac{\partial h}{\partial z}\mathrm{cov}(x,z). \quad \text{(A6.1)}$$

Since the correlation coefficient between two random variables is always ≤ 1, then:

$$\mathrm{cov}(x,z) \leq \mathrm{std}(x)\,\mathrm{std}(z). \quad \text{(A6.2)}$$

From the relationship between the variance and standard deviation, Eq. (A6.1) can then be written as:

$$\left(\mathrm{std}(h)\right)^2 \leq \left(\frac{\partial h}{\partial x}\right)^2 \left(\mathrm{std}(x)\right)^2 + \left(\frac{\partial h}{\partial z}\right)^2 \left(\mathrm{std}(z)\right)^2 + 2\left|\frac{\partial h}{\partial x}\right|\left|\frac{\partial h}{\partial z}\right| \mathrm{std}(x)\mathrm{std}(z) = \left(\left|\frac{\partial h}{\partial x}\right| \mathrm{std}(x) + \left|\frac{\partial h}{\partial z}\right| \mathrm{std}(z)\right)^2. \quad \text{(A6.3)}$$

It follows that a reasonable approximation for the lower limit of std(h) is:

$$\mathrm{std}(h) \approx \left|\frac{\partial h}{\partial x}\right| \mathrm{std}(x) + \left|\frac{\partial h}{\partial z}\right| \mathrm{std}(z). \quad \text{(A6.4)}$$

Appendix B: Proof that $t_{\text{fixed }\eta} \propto \eta^2_{\text{fixed }t}$

It is easily shown from Section 6.2.2 that for a given experimental system, the objective function, η, can be expressed as:

$$\eta \propto \sigma\, \eta_s, \quad \text{(B6.1)}$$

where η_s is a function that depends only on the sampling pattern (number of points and their spacing). Eq. (B6.1), together with the fact that $t \propto n\ \sigma^{-2}$ leads to the following observations:

$$t_{\text{fixed }\eta} \propto n\, {\eta_s}^2, \quad \text{(B6.2a)}$$

$$\eta^2_{\text{fixed t}} \propto n\,\eta_s{}^2, \qquad \text{(B6.2b)}$$

from which it follows that $t_{\text{fixed }\eta} \propto \eta^2_{\text{fixed }t}$.

References

[1] R. A. Fisher, Statistical Methods for Research Workers, 10th ed., Oliver and Boyd, 1948.

[2] M. Lustig, D. Donoho and J. M. Pauly, Sparse MRI: The application of Compressed Sensing for rapid MR imaging, Magn. Reson. Med. 58 (2007), 1182-1195.

[3] M. Mobli, M. W. Maciejewski, A. D. Schuyler, A. S. Stern and J. C. Hoch, Sparse sampling methods in multidimensional NMR, Phys. Chem. Chem. Phys. 14 (2012), 10835-10843.

[4] C. K. Anand, A. D. Bain and A. Sharma, Optimized sampling patterns for multi-dimensional T_2 experiments, J. Magn. Reson. 197 (2009), 63-70.

[5] Y.-Q. Song, Y. Tang, M. D. Hürlimann and D. G. Cory, Real-time optimization of nuclear magnetic resonance experiments, J. Magn. Reson. 289 (2018), 72-78.

[6] S. Roeding, N. Klimovich and T. Brixner, Optimizing sparse sampling for 2D electronic spectroscopy, J. Chem. Phys. 146 (2017), 084201.

[7] T. B. Edo, D. J. Batey, A. M. Maiden, C. Rau, U. Wagner, Z. D. Pešić, T. A. Waigh and J. M. Rodenburg, Sampling in X-ray ptychography, Phys. Rev. A 87 (2013), 053850.

[8] A. Jones, A. Tamtögl, I. Calvo-Almazán and A. Hansen, Continuous Compressed Sensing for surface dynamical processes with Helium Atom Scattering, Sci. Rep. 6 (2016), 27776.

[9] A. L. Koch, The logarithm in biology. 1. Mechanisms generating the log-normal distribution exactly, J. Theoret. Biol. 12 (1966), 276-290.

[10] E. Limpert, W. A. Stahel and M. Abbt, Log-normal distributions across the sciences: Keys and clues, BioScience 51 (2001), 341-352.

[11] L. B. Kiss, J. Söderlund, G. A. Niklasson and C. G. Granqvist, New approach to the origin of lognormal size distributions of nanoparticles, Nanotechnology 10 (1999), 25-28.

[12]N. H. Williamson, M. Nydén and M. Röding, The lognormal distribution models for estimating molecular weight distributions of polymers using PGSE NMR, J. Magn. Reson. 267 (2016), 54-62.

[13]A. J. Baldwin, S. J. Anthony-Cahill, T. P. J. Knowles, G. Lippens, J. Christodoulou, P. D. Barker and C. M. Dobson, Measurement of myloid fibril length distributions by inclusion of rotational motion in solution NMR diffusion measurements, Agnew. Chem. Int. Ed. 47 (2008), 3385-3387.

[14]S. J. Law and M. M. Britton, Sizing of reverse micelles in microemulsions using NMR measurements of diffusion, Langmuir 28 (2012), 11699-11706.

[15]J.-P. Korb, A. Louis-Joseph and L. Benamsili, Probing structure and dynamics of bulk and confined crude oils by multiscale NMR spectroscopy, diffusometry, and relaxometry, J. Chem. Phys. B 117 (2013), 7002-7014.

[16]L. Benamsili, J.-P. Korb, G. Hamon, A. Louis-Joseph, B. Bouyssiere, H. Zhou and R. G. Bryant, Multi-dimensional Nuclear Magnetic Resonance characterizations of dynamics and saturations of brine/crude oil/mud filtrate mixtures confined in rocks: The role of asphaltene, Energy Fuels 28 (2014), 1629-1640.

[17]Z. Yang and G. J. Hirasaki, NMR measurement of bitumen at different temperatures, J. Magn. Reson. 192 (2008), 280-293.

[18]P. Becher and M. N. Yudenfreund (Ed.), Emulsions, Lattices and Dispersions, Marcel Dekker, 1978.

[19]K. J. Packer and C. Rees, Pulsed NMR studies of restricted diffusion. I. Droplet size distributions in emulsions, J. Colloid Interface Sci. 40 (1972), 206-218.

[20]P. T. Callaghan, K. W. Jolley and R. S. Humphrey, Diffusion of fat and water in cheese as studied by pulsed field gradient Nuclear Magnetic Resonance, J. Colloid Interface Sci. 93 (1983), 521-529.

[21]J. C. van Den Enden, D. Waddington, H. van Aalst, C. G. Van Kralingen and K. J. Packer, Rapid determination of water droplet size distributions by PFG-NMR, J. Colloid Interface Sci. 140 (1990), 105-113.

[22]K. G. Hollingsworth and M. L. Johns, Measurement of emulsion droplet sizes using PFG NMR and regularization methods, J. Colloid Interface Sci. 258 (2003), 383-389.

[23]I. Fourel, J. P. Guillement and D. Le Botlan, Determination of water droplet size distributions by low resolution PFG-NMR, J. Colloid Interface Sci. 164 (1994), 48-53.

[24]G. J. W. Goudappel, J. P. M. van Duynhoven and M. M. Mooren, Measurement of oil droplet size distributions in food oil/water emulsions by time domain pulsed field gradient NMR, J. Colloid Interface Sci. 239 (2001), 535-542.

[25]M. L. Johns and L. F. Gladden, Sizing of emulsion droplets under flow using flow-compensating NMR-PFG techniques, J. Magn. Reson 154 (2002), 142-145.

[26]E. S. Rajagopal, Statistical theory of particle size distributions in emulsions and suspensions, Kolloid Z. 162 (1959), 85-92.

[27]J. C. Thomas, The determination of log normal particle size distributions by dynamic light scattering, J. Colloid Interface Sci. 117 (1987), 187-192.

[28]G. van Dalen, Determination of the water droplet size distribution of fat spreads using confocal scanning laser microscopy, J. Microsc. 208 (2002), 116-133.

[29]Y. Katz and U. Nevo, Quantification of pore size distribution using diffusion NMR: Experimental design and physical insights, J. Chem. Phys. 140 (2014), 164201.

[30]D. C. Alexander, A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features, Magn. Reson. Med. 60 (2008), 439-448.

[31]J. S. Murday and R. M. Cotts, Self-diffusion coefficient of liquid lithium, J. Chem. Phys. 48 (1968), 4938-4945.

[32]M. Alderliesten, Mean particle diameters. Part I: Evaluation of definition systems, Part. Part. Syst. Charact. 7 (1990), 233-241.

[33]T. Hatch and S. P. Choate, Statistical description of the size properties of non-uniform particulate substances, J. Franklin Inst. 207 (1929), 369-387.

[34]O. V. Petrov and S. Stapf, Parametrization of NMR relaxation curves in terms of logarithmic moments of the relaxation time distribution, J. Magn. Reson. 279 (2017), 29-38.

[35]D. E. O'Reilly and E. M. Peterson, Self-diffusion coefficients and rotational correlation times in polar liquids. III. Toluene, J. Chem. Phys. 56 (1972), 2262-2266.

[36]J. A. Jones, P. Hodgkinson, A. L. Barker and P. J. Hore, Optimal sampling strategies for the measurement of spin-spin relaxation times, J. Magn. Reson. Series B 113 (1996), 25-34.

[37]Y. Zhang, H. N. Yeung, M. O'Donnell and P. L. Carson, Determination of sample time for T_1 measurement, J. Magn. Reson. Imaging 8 (1998), 675-681.

[38]G. H. Weiss, R. K. Gupta, J. A. Ferretti and E. D. Becker, The choice of optimal parameters for measurement of spin-lattice relaxation times. I. Mathematical formulation, J. Magn. Reson. 37 (1980), 369-379.

[39]B. Efron, The Jackknife, the Bootstrap and Other Resampling Plans, SIAM, 1982.

[40]P. T. Callaghan, Principles of Nuclear Magnetic Resonance Microscopy, Oxford University Press, 1993.

[41]J. R. Taylor, An Introduction to Error Analysis, University Science Books, 1997.

Chapter 7

Optimising sampling patterns for bi-exponentially decaying signals

In Chapter 6, a systematic, statistical approach based on the CRLB theory was described for optimising sampling patterns for a wide range of NMR experiments; the findings of the CRLB theory were validated against PFG NMR diffusion experimental data of an emulsion system, characterised by a lognormal distribution of droplet sizes. In this chapter, the same method is applied to the problem of optimising the sampling pattern for the most accurate estimation of the parameters of a bi-exponential model.

7.1. Introduction

Bi-exponentially decaying signals are common in many areas of science and engineering, such as in photoluminescence [1, 2] and reaction kinetics [3, 4]. In NMR applications, bi-exponential models have been assumed to be a good approximation for T_2 and $T_{1\rho}$ decays, and diffusion attenuations [5-9]. The fitting of bi-exponential decays is a long-standing challenge and remains a subject of debate [10-14] because of the relatively large uncertainty related to the estimated parameters of the model.

The findings of the CRLB theory are initially validated against PFG NMR diffusion experimental data of a binary methane/ethane gaseous mixture adsorbed in a zeolite; zeolites are increasingly being used to separate gaseous hydrocarbon mixtures in the gas processing industry [15, 16]. The CRLB theory is subsequently used to determine the limit of resolution of the parameters in a bi-exponential model.

The chapter is structured as follows. Section 7.2 outlines the application of the CRLB theory to a bi-exponentially decaying signal. The experimental settings are described in Section 7.3. Section 7.4 presents the results and discussion.

7.2. Mathematical treatment

A noise-free, bi-exponentially decaying signal, $\hat{y}_i$, is described as:

$$\hat{y}_i = w_1 \exp(-b_i D_1) + w_2 \exp(-b_i D_2) \,, \tag{7.1}$$

where b_i represent the sampling pattern, w_1 and w_2 are the populations of the two components and, D_1 and D_2 represent the decay rate constants, where for convenience $D_1 < D_2$. The noise-free signal, $\hat{y}_i$, depends on 4 parameters, θ_l (l = 1, 2, 3, 4): w_1, w_2, D_1 and D_2. As in Chapter 6, it will be assumed that each signal data point is corrupted by noise elements which are random and independently normally distributed with zero mean and variance σ^2 (Gaussian noise). The noise variance depends on the scaling of the signal. In this chapter, the stated value of σ corresponds to a scaled signal of maximum 1, which corresponds to $w_1 + w_2 = 1$. In the presence of noise, the estimate of the bi-exponential model parameters becomes uncertain and there is an uncertainty associated with the estimate of each parameter. A common characteristic of this uncertainty is the standard deviation of any of the estimators, std(θ_l). A lower limit for std(θ_l) is given by the Cramér-Rao Lower Bound (CRLB) theory, which was outlined in Section 3.4:

$$\text{std}(\theta_l) \geq \sqrt{(\boldsymbol{F}^{-1})_{ll}} \,, \tag{7.2}$$

where $\boldsymbol{F}$ is the Fisher Information Matrix. The Fisher Information Matrix can be computed from:

$$F_{l_1 l_2} = \frac{1}{\sigma^2}\left(\sum_i \frac{\partial \hat{y}_i}{\partial \theta_{l_1}} \frac{\partial \hat{y}_i}{\partial \theta_{l_2}}\right) . \tag{7.3}$$

Therefore, to construct $\boldsymbol{F}$, the partial derivatives of $\hat{y}_i$ with respect to w_1, w_2, D_1 and D_2 are needed, which can be calculated from Eq. (7.1):

$$\frac{\partial \hat{y}_i}{\partial w_1} = \exp(-b_i D_1) \,, \tag{7.4a}$$

$$\frac{\partial \hat{y}_i}{\partial w_2} = \exp(-b_i D_2) \,, \tag{7.4b}$$

$$\frac{\partial \hat{y}_i}{\partial D_1} = -w_1 b_i \exp(-b_i D_1) \,, \tag{7.4c}$$

$$\frac{\partial \hat{y}_i}{\partial D_2} = -w_2 b_i \exp(-b_i D_2) \,. \tag{7.4d}$$

The important result is that the uncertainty associated with the estimate of each parameter

depends on σ, w_1, w_2, D_1, D_2 and most importantly on the sampling pattern, b_i. As a result, given the noise level and an experimental system (characterised by w_1, w_2, D_1 and D_2), the error associated with the estimate of each parameter depends solely on the sampling pattern, b_i. Therefore, by optimizing the sampling pattern, it is possible to decrease to a minimum the uncertainty associated with the estimate of a given parameter, hence, making the estimate of that parameter as accurate as possible.

The sampling pattern that minimises the uncertainty related to one of the parameters does not necessarily minimise the uncertainty related to the other parameters. Therefore, an objective function to be minimised needs to be defined. The presented method can be easily adapted to different objective functions. In this chapter, it is chosen to minimise:

$$\eta = \max_{l}[\mathrm{std}(\theta_l)/\theta_l] \times 100\% \,, \tag{7.5}$$

which practically means minimising the maximum percentage error of any of the estimated parameters.

7.3. Materials and methods

7.3.1. Experimental

A PFG NMR diffusion experiment is conducted on a binary gaseous mixture adsorbed on a porous material. The gases used are methane (BOC, >99% purity) and ethane (Air Liquide, >99.99%)). The porous material used is a microporous β-zeolite (BEA), supplied by Tosoh Corporation. The zeolite bed is of cylindrical shape, with a height of 10 mm and diameter of 5 mm. The zeolite has a mean crystallite size of 15 μm and Si:Al ratio of 124. The adsorption of the gases in the zeolite was carried out on a Schlenk vacuum line; further details are found elsewhere [17].

The PFG NMR experiment is performed on a Bruker DMX300 spectrometer, operating at a resonant frequency of 300.13 MHz for ^{1}H observation, at a temperature of 20 ± 0.5 °C. The maximum gradient amplitude available is 1176 G cm^{-1} and the r.f. coil has a diameter of 5 mm. The 13-interval bipolar gradient stimulated echo PFG NMR pulse sequence [18], illustrated in Fig. 7.1(a), is used for diffusion measurements. This pulse sequence was chosen because it reduces the effects of internal magnetic field gradients generated by porous media. The relevant NMR acquisition parameters used, annotated in Fig. 1(a), are

T = 100 ms, δ = 0.25 ms and τ = 0.46 ms. The echo time, τ, is minimised to increase the signal-to-noise ratio and to reduce relaxation weighting. The small remaining relaxation weighting does not affect the following analysis. 248 gradient strength values, g_i, spaced equidistantly in the range 1.5-50 G cm^{-1}, are used for two reasons: to provide an estimate for the diffusion coefficients and the populations of methane and ethane, and to be used for the selection of different sampling patterns, as described in Section 7.3.2.

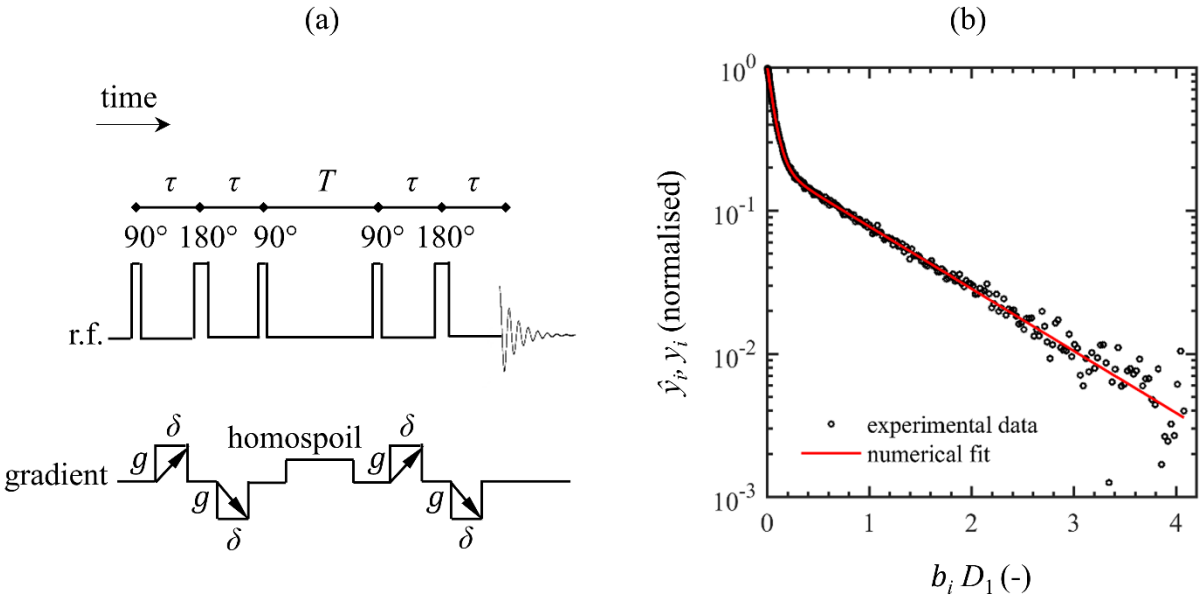

Fig. 7.1. (a) PFG pulse sequence used for diffusion measurements. (b) Experimental PFG NMR diffusion data of a binary gaseous mixture methane/ethane and the corresponding bi-exponential fit. The sampling points, b_i, have been non-dimensionalised by the diffusion coefficient of slowest diffusing component, ethane, $D_1 = 9.24 \times 10^{-9}$ m^2 s^{-1}.

The adsorbed methane and ethane require deconvolution to be resolved spectroscopically [17], which is not ideal. In this chapter, the attenuation of the convolved spectral peak, shown in Fig. 7.1(b), is investigated. The attenuation is well described by the bi-exponential model described in Eq. (7.1), with $b_i = \gamma^2 g_i^2 \delta^2 (4T + 6\tau - 2/3\delta)$ [18], where γ is the gyromagnetic ratio of ^{1}H nuclei. The estimated model parameters, which are used as ground truth for the following analysis, are: $w_1 = 0.199 \pm 0.001$ (ethane population), $D_1 = (9.24 \pm 0.07) \times 10^{-9}$ m^2 s^{-1} (ethane diffusion coefficient), $w_2 = 0.801 \pm 0.001$ (methane population) and, $D_2 = (1.42 \pm 0.01) \times 10^{-7}$ m^2 s^{-1} (methane diffusion coefficient). The quoted uncertainty corresponds to the standard deviation of the bi-exponential fit. The uncertainty related to each of the parameters is small because of the low σ (estimated to be 0.0032) and the large number of points in the decay.

The preparation of the samples and implementation of the pulse sequence were done with the help of the PhD student Mohammed Ainte of the Magnetic Resonance Research Centre, University of Cambridge.

7.3.2. Sampling patterns

A sampling pattern for the acquisition of the PFG NMR data consists of the choice of sampling points, b_i, at which the attenuation of the NMR signal is acquired. This involves choosing the number of data points, n, and their respective values. As discussed in Section 6.3.2, for a given number of points, n, the optimal choice of b_i would require setting the partial derivatives of the expression in Eq. (7.5) with respect to each individual b_i to zero and solving the set of resulting n simultaneous non-linear equations; this is computationally intractable. As a result, this chapter focuses on finding the optimal sampling pattern for the class of linear sampling patterns, defined as:

$$b_i = b_1 + (b_n - b_1)\left(\frac{i-1}{n-1}\right)^r. \tag{7.6}$$

b_1 and b_n are the first and last sampling points and r determines the density distribution of sampling points. The class of logarithmic sampling patterns, which was discussed in Chapter 6, was also investigated, but the optimised logarithmic sampling pattern had the same objective function value as the optimised linear sampling pattern studied. This is consistent with the results in Chapter 6 and it is the reason for focusing only on the class of linear sampling patterns.

As discussed in Section 6.3.2, sampling a larger number of points, n, leads to a more accurate estimate of the model parameters. Therefore, the comparison between the predictions of the CRLB theory and experimental results is performed with respect to sampling patterns of a particular number of sampling points, $n = 32$. Other values of n could have been chosen, but $n = 32$ is a commonly used experimental value. Using the smallest possible value for b_1 is supported by the results of other works, as discussed in Section 6.3.2. Therefore, the smallest experimentally used value, $b_1 = 4.9 \times 10^5$ m^{-2} s, is used when validating the predictions of the CRLB theory against experimental data. Having fixed n and b_1, optimising the linear sampling pattern consists of optimising b_n and r.

The comparison between the predictions of the CRLB theory and experimental results is performed as follows. Using a range of b_n from 5.4×10^7 to 4.8×10^8 m^{-2} s and a range of

r from 0.6 to 10, 10^4 linear sampling patterns are constructed using Eq. (7.6). For each of these constructed sampling patterns, the CRLB theory is used to predict the objective function, η, defined in Eq. (7.5). This is compared to the objective function obtained from the experimental data which is calculated as follows. The closest experimental 32 b_i points to the constructed sampling patterns (out of the total 248 experimental sampling points acquired) are selected. The experimental attenuation data corresponding to these 32 points is fitted to a bi-exponential decay. The standard deviations in the estimation of each parameter, std(θ_l), are then extracted and these values are used to calculate the experimental objective function, η.

7.4. Results and discussion

In Section 7.4.1, the predictions of the CRLB theory are validated against the experimental data using the procedure outlined in Section 7.3.2. The CRLB theory is then used in Section 7.4.2 to advise on which bi-exponential systems (characterised by w_1 and D_2/D_1) can be resolved and under which conditions they can be resolved (in terms of the noise standard deviation, σ, and the number of sampled points, n).

7.4.1. Validation of CRLB theory from experimental data

The comparison of the objective function, η, as predicted by the CRLB theory and as calculated from the experimental data is shown in Fig. 7.2. Small values of η indicate a good sampling pattern. Fig. 7.2(a) shows the contour map of η as a function of the linear sampling pattern parameters b_nD_1 and r, as predicted by the CRLB theory. Fig. 7.2(b) shows a similar contour map obtained from the experimental data. There is very good agreement between the CRLB and experimental contour maps; on average there is < 10% difference between the corresponding η values of the two maps, thereby validating the application of the CRLB theory to the optimization of sampling patterns for bi-exponentially decaying signals.

For a given value of r, for example r = 1 (corresponding to equidistant points), the value of the objective function is large at small b_nD_1 because late decay points are not sampled and, therefore, the uncertainty related to estimating the slower diffusing component (ethane) is large. The value of the objective function is also large at very large b_nD_1 because early decay points are not sampled properly and, therefore, the uncertainty related

to estimating the faster diffusing component (methane) is large. As a result, for any given r, there is an optimum b_nD_1 which minimises the objective function. For a given b_nD_1, for example b_nD_1 = 2.5, the value of the objective function is large at small r because early decay points are sampled sparsely and, therefore, the uncertainty related to estimating the faster diffusing component (methane) is large. The value of the objective function is also large at very large r because late decay points are sampled sparsely and, therefore, the uncertainty related to estimating the slower diffusing component (ethane) is large. As a result, for any given b_nD_1, there is an optimum r which minimises the objective function. In summary, there is an optimum combination (b_nD_1,r) at which the objective function obtains a minimum, which corresponds to the best linear sampling pattern. The best linear sampling pattern for 32 sampled points is predicted by the CRLB theory at r = 1.07 and b_nD_1 = 2.34, at which point η = 1.9%. Therefore, using the best linear sampling pattern, the maximum percentage error for any of the estimated parameters is 1.9%.

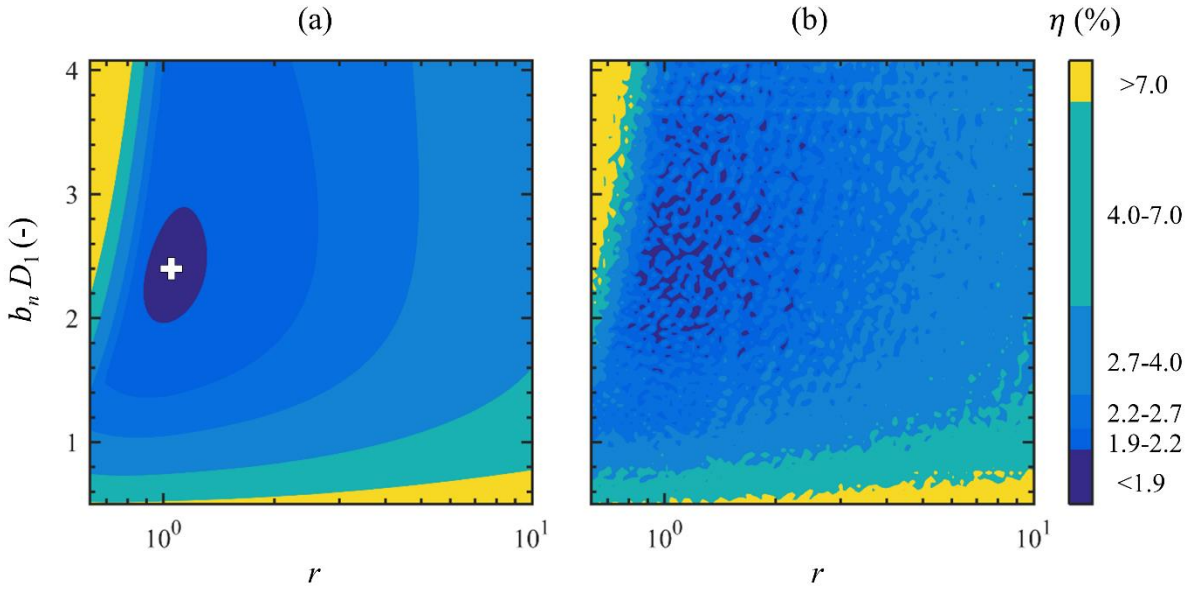

Fig. 7.2. (a) Contour map of the variation of the objective function, η, with respect to the parameters of the family of linear sampling patterns, defined in Eq. (7.6), with n = 32, as predicted by the CRLB theory. (b) Contour map of the variation of the objective function, η, with respect to the parameters of the class of linear sampling patterns obtained from the experimental data shown in Fig. 7.1(b). The best linear sampling pattern in (a) is obtained for r = 1.07, b_nD_1 = 2.34, and is identified as '+', at which point η = 1.9%.

For completeness, the objective function of a commonly used sampling pattern with 32 equidistantly sampled points (r = 1), with the largest sampled point being b_nD_1 = 5 (typical rule of thumb), is η = 3.1%. This value of η is >50% larger than the value achieved using

the optimised linear sampling pattern (η = 1.9%). The calculation was performed using the CRLB theory for the two-component system under investigation and assuming σ = 0.0032.

7.4.2. Resolution limit of bi-exponential decays

In Section 7.4.1, the CRLB theory predictions were validated against an experimental bi-exponential system characterised by D_2/D_1 = 15 (ratio of the largest to the smallest diffusion coefficients) and w_1 = 0.20 (population of slowest diffusing component). The aim of this section is to use the CRLB theory to investigate which bi-exponential systems (characterised by w_1 and D_2/D_1) can be resolved and under which conditions they can be resolved (in terms of the noise standard deviation, σ, and the number of sampled points, n).

This question is investigated as follows. For a given experimental system, characterized by w_1 and D_2/D_1, and a noise standard deviation, σ, a linear sampling pattern with a given number of points, n, and $b_1D_1 = 0$ is optimised in terms of r and b_nD_1. This is performed using a similar method to the method used to produce Fig. 7.2(a). For the optimised linear sampling pattern, the value of η is recorded. This is repeated for 100 values of w_1 in the range 0.0-1.0 (non-inclusive) and 100 values of D_2/D_1 in the range 1-1000 (non-inclusive at the lower limit). A contour map of the optimal η as a function of w_1 and D_2/D_1, for the chosen noise standard deviation, σ, and sampling number of points, n, is then constructed. At this point, the resolution limit needs to be defined; in this chapter the resolution limit is taken to be η < 30%; i.e., if the percentage error in the estimate of all the parameters is < 30 %, the bi-exponential decay is defined as being resolvable. It is noted that a different resolution limit definition is easily implemented. With the definition of the resolution limit, the contour map is divided into two regions: resolvable and non-resolvable, for the chosen values of σ and n. The whole procedure is performed for σ = 0.001, 0.005, 0.02, 0.05 and n = 16, 32, 64. The results are presented in Fig. 7.3.

The first observation from Fig. 7.3 is that the contour maps are not symmetric about w_1 = 0.5. The asymmetry suggests that it is easier to resolve bi-exponential decays when the slowest diffusing component has the largest population than when the slowest diffusing component has the smallest population. This is explained by the fact that the most information about the slow diffusing component is contained at the high b_i points. However, at high b_i points, the signal becomes comparable to the noise. Therefore, a high population of the slow diffusing component is needed if it is to be resolved. Another

qualitative observation from Fig. 7.3 is that to resolve components of similar diffusion coefficients, $D_2/D_1 \sim 1$, many points at a high signal-to-noise ratio need to be acquired.

The maps shown in Fig. 7.3 can be used quantitatively in many ways. For example, the methane/ethane gaseous mixture investigated in this chapter ($D_2/D_1 = 15$, $w_1 = 0.20$) was easily resolved ($\eta = 1.9\%$) using $n = 32$ points because the noise standard deviation was relatively low ($\sigma = 0.0032$). According to Fig. 7.3(b), the limit of resolution of these components for $n = 32$ points is at $\sigma \sim 0.05$. Therefore, the PFG NMR experimental acquisition time could be accelerated by a factor of ~ 4 through less signal averaging, while keeping the components still resolvable (assuming acquisition time $\propto \sqrt{(1/\sigma)}$ [19]). Accelerating the PFG NMR experimental acquisition time could be important if temporally changing phenomena are being investigated.

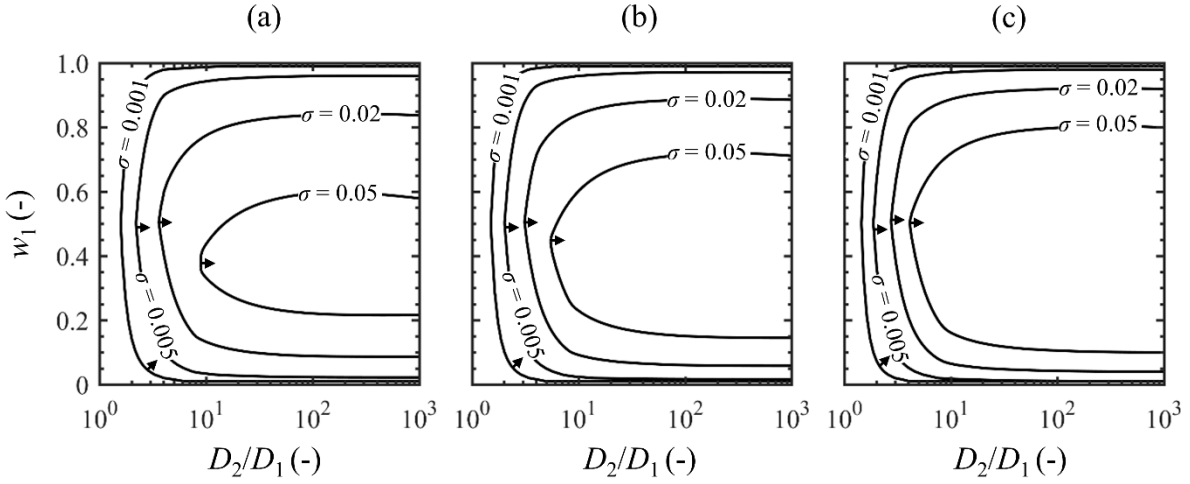

Fig. 7.3. Resolution limit of bi-exponential decays, defined as $\eta < 30\%$, for a range of experimental systems characterised by a range of w_1, D_2/D_1 and σ for (a) $n = 16$, (b) $n = 32$, and (c) $n = 64$ sampling points. The arrows in the map point towards the resolvable region.

7.5. Conclusions

The method proposed in Chapter 6 for optimising sampling patterns was applied to the problem of optimising sampling patterns for the most accurate estimation of the parameters of a bi-exponential model. The findings of the method were validated against PFG NMR diffusion data of a binary gaseous mixture of methane/ethane; < 10% difference was observed. The method was subsequently used to investigate which experimental systems that are modelled as bi-exponential decays can be resolved and under what conditions they can be resolved.

References

[1] T. Gokus, L. Cognet, J. G. Duque, M. Pasquali, A. Hartschuh and B. Lounis, Mono- and biexponential luminescence decays of individual single-walled carbon tubes, J. Phys. Chem. 114 (2010), 14025-14028.

[2] P. Coppens, J. Sokolow, E. Trzop, A. Makal and Y. Chen, On the biexponential decay of the photoluminescence of the two crystallographically-indepenedent molecules in crystals of [Cu(I)(phen)(PPh3)2][BF4], J. Phys. Chem. Lett. 4 (2013), 579-582.

[3] D. Harding, M. S. Ford, T. R. Walsh and S. R. Mackenzie, Dramatic size effects and evidence of structural isomers in the reactions of rhodium clusters, Rhn±, with nitrous oxide, Phys. Chem. Chem. Phys. 9 (2007), 2130-2136.

[4] P. Mondal and M. Meuwly, Solvent composition drives the rebinding kinetics of nitric oxide to microperoxidase, Sci. Rep. 8 (2018), 5281.

[5] T. Niendorf, R. M. Dijkhuzien, D. G. Norris, M. van Lookeren Campagne and K. Nicolay, Biexponential diffusion attenuation in various states of brain tissue: Implications for diffusion-weighted imaging, Magn. Reson. Med. 36 (1996), 847-857.

[6] R. Mulkern, H. P. Zengingonul, R. L. Robertson, P. Bogner, K. H. Zou, H. Gudbjartsson, C. R. G. Guttmann, D. Holtzman, W. Kyriakos, F. A. Jolesz and S. E. Maier, Multi-component apparent diffusion coefficients in human brain: Relationship to spin-lattice relaxation, Magn. Reson. Med. 44 (2000), 292-300.

[7] C. A. Clark and D. le Bihan, Water diffusion compartmentation and anisotropy at high *b* values in the human brain, Magn. Reson. Med. 44 (2000), 852-859.

[8] Z. Ababneh, H. Beloeil, C. B. Berde, G. Gambarota, S. E. Maier and R. V. Mulkern, Biexponential parametrization of diffusion and T_2 relaxation decay curves in a rat muscle edema model: Decay curve components and water compartments, Magn. Reson. Med. 54 (2005), 524-531.

[9] A. Sharafi, G. Chang and R. R. Regatte, Bi-component $T_{1\rho}$ and T_2 relaxation mapping of skeletal muscle in-vivo, Sci. Rep. 7 (2017), 14115.

[10] M. Nilsson, M. A. Connell, A. L. Davis and G. A. Morris, Biexponential fitting of diffusion-ordered NMR data: Practicalities and limitations, Anal. Chem. 78 (2006), 3040-3045.

[11] M. S. Shazeeb and C. H. Sotak, Limitations in biexponential fitting of NMR inversion-recovery curves, J. Magn. Reson. 276 (2017), 14-21.

[12] R. V. Mulkern, M. Balasubramanian and S. E. Maier, On the perils of multi exponential fitting of diffusion MR data. J. Magn. Reson. Imaging 45 (2016), 1545-1547.

[13] G. Pagès, A. Bonny, V. Gilard and M. Malet-Martino, Pulsed field gradient NMR with sigmoid shape gradient sampling to produce more detailed diffusion ordered spectroscopy maps of real complex mixtures: Examples with medicine analysis, Anal. Chem. 88 (2016), 3304-3309.

[14] R. Bai, D. Benjamini, J. Cheng and P. J. Basser, Fast, accurate 2D-MR relaxation exchange spectroscopy (REXSY): Beyond compressed sensing, J. Chem. Phys. 145 (2016), 154202.

[15] A. Lauerer, T. Binder, C. Chmelik, E. Miersemann, J. Haase, D. M. Ruthven and J. Kärger, Uphill diffusion and overshooting in the adsorption of binary mixtures in nanoporous solids, Nat. Comm. 6 (2015), 7697.

[16] Y. Yang, N. Burke, S. Ali, S. Huang, S. Lim and Y. Zhu, Experimental studies of hydrocarbon separation on zeolites, activated carbons and MOFs for applications in natural gas processing, RSC Adv. 7 (2017), 12629-12638.

[17] M. I. Ainte, NMR techniques for measuring transport phenomena in microporous materials, PhD Thesis, University of Cambridge, 2017.

[18] R. M. Cotts, M. J. R. Hoch, T. Sun and J. T. Marker, Pulsed field gradient stimulated echo methods for improved NMR diffusion measurements in heterogeneous systems, J. Magn. Reson. 83 (1989), 252-266.

[19] P. T. Callaghan. Translational Dynamics & Magnetic Resonance, Oxford University Press, 2011.

Chapter 8

Accelerating the estimation of 3D spatially resolved T_2 distributions

In Chapters 4, 5 and 6, 7 the benefits of optimising the reconstruction technique or the sampling pattern, independently, for the most accurate estimation of parameters in particular ill-conditioned inverse problems in NMR applications, were investigated. In practice, the optimisation of the sampling pattern and the reconstruction technique are not entirely separate processes. In this chapter, it is shown that the most accurate parameter estimation is obtained through an appropriate combination of a good sampling pattern and reconstruction technique. This is illustrated with respect to the extraction of spatially resolved T_2 distributions, referred to as T_2 maps.

8.1. Introduction

The principle of obtaining a T_2 map of dimension d (d = 1, 2 or 3 for a 1D, 2D or 3D T_2 map) is similar across the plethora of different experimental techniques. A series of p T_2 weighted images of dimension d is acquired and the T_2 distribution is extracted pixel-wise using an Inverse Laplace Transform algorithm, as discussed in Section 3.1 and Chapters 4 and 5, or some other single or multi-exponential fitting method. The experimental techniques differ in the way the T_2 weighted images are acquired, and can be categorized into experimental techniques which use either a variable echo time or keep the echo time fixed and vary the length of the train of echoes. The most widely used technique with a variable echo time is the Single Echo Spin Echo (SESE) pulse sequence [1, 2]. Although the SESE sequence is very robust, the acquisition time is prohibitively long. A single line of k-space is sampled at each echo and a single echo is formed for each excitation. Therefore, the acquisition time scales approximately linearly with $n^{d-1}p$, with n being the length of each image dimension. The leading experimental technique with a fixed echo time is the Multi Echo Spin Echo (MESE) pulse sequence [3-6]. A single line of k-space

is sampled at each echo, but multiple echoes (an echo train) are formed for each excitation. Therefore, the acquisition time scales approximately linearly with n^{d-1}, which is a significant improvement over the acquisition time of the SESE method. The MESE pulse sequence is often referred to by different names depending on the application. For example, when used for the purpose of obtaining T_2 contrast images, it is referred to as a Rapid Acquisition with Relaxation Enhancement (RARE) pulse sequence, as discussed in Section 2.16. The main disadvantage of the MESE technique is the use of a train of slice-selective refocusing pulses, which, if not perfect, can lead to complex stimulated echo patterns [7, 8]. Another technique with a fixed echo time is the Spin Echo - Single Point Imaging (SE-SPI) pulse sequence [9-13]. The SE-SPI pulse sequence forms a train of p spin echoes for each excitation and at each echo a single point in k-space is sampled. As a result, the acquisition time scales approximately linearly with n^d which is significantly longer than the acquisition time of an MESE pulse sequence. However, its main advantage is the capability to use much shorter echo times; this has enhanced the capability to study fluids characterized by sub-millisecond T_2 relaxation time constants in porous media applications.

This chapter is concerned with obtaining 3D T_2 maps, which in application to porous media, would open up opportunities for characterising local pore structure, pore saturation and wettability. The major problem associated with obtaining 3D T_2 maps is the long acquisition time; the acquisition time for a fully sampled dataset sampled with an SE-SPI technique is in the order of days [12]. The system under investigation may not be stable during a long acquisition time, or one might be investigating phenomena which are changing with time. As a result, reducing the acquisition time is important in reducing the instability effects (i. e., temporal blurring) or capturing the time-dependent phenomena. Without under sampling, the MESE technique offers the shortest acquisition time compared to the SESE and SE-SPI techniques, as discussed above. However, in most cases, this is still prohibitively long; for the system under consideration in this chapter, acquiring a fully sampled data set with an MESE technique takes ~ 16 h. The focus of this chapter is on exploring under sampling patterns for MESE experiments and investigating different reconstruction techniques, based on the Compressed Sensing ideas discussed in Section 3.2, for obtaining 3D T_2 maps from the under sampled data.

In this chapter, an experimental implementation of an MESE pulse sequence to acquire a 3D T_2 map from under sampled MRI data is demonstrated. A conventional sampling

pattern where, for each line sampled in *k*-space, all echoes down the echo train are acquired (coherent sampling pattern) is compared to a proposed sampling pattern where the *k*-space line sampled at each echo in the echo train is independent of the *k*-space line sampled in the previous echo (incoherent sampling pattern). Further, the performance of the conventional reconstruction technique of Total Variation (TV) regularization, outlined in Section 3.2.2 is compared with the performance of the more recent reconstruction techniques of Nuclear Norm (NN) regularization [14] and Nuclear Total Generalized Variation (NTGV) regularization [15, 16]. It is shown that using an incoherent sampling pattern and NTGV regularization as the reconstruction technique, quantitative 3D T_2 maps are obtained even at sampling percentages of 3.1% of *k*-space, corresponding to a 32-fold decrease in acquisition time.

The chapter is structured as follows. Section 8.2 introduces the different reconstruction methods; pseudocodes and practical guidance in implementing the different reconstruction methods are also given. The experimental settings and acquisitions are described in Section 8.3. The results and discussion are presented in Section 8.4.

8.2. Mathematical treatment

Obtaining 3D spatially resolved T_2 maps from under sampled *k*-space data consists of two steps. The first step is the reconstruction of T_2-weighted images from the under sampled *k*-space data. In the second step, the T_2-weighted images serve as an input to an Inverse Laplace Transform algorithm, or a single or multi-exponential fitting procedure in order to obtain the T_2 distribution of each individual pixel. This chapter focuses on the first step. It is noted that reconstruction techniques that try to collapse these two steps into one step have started to be developed [17, 18], particularly in application to Magnetic Resonance Fingerprinting. However, these techniques typically suffer from partial volume effects; they assume that each voxel is characterized by one or more T_2 components. This may not be the case in rock imaging applications, where a continuous distribution of T_2 values is expected in each pixel. However, this area of research will be important in the future.

As discussed in Section 3.2, the acquired under sampled *k*-space data $\boldsymbol{Y}$, vectorised as $\boldsymbol{Y}_\mathrm{v}$, is related to the vectorised T_2-weighted images, $\boldsymbol{U}_\mathrm{v}$, the Fourier operator $\boldsymbol{\mathcal{F}}$, the under sampling matrix, $\boldsymbol{S}$, and the vectorised noise matrix, $\boldsymbol{E}_\mathrm{v}$, by:

$$\boldsymbol{Y}_{\mathrm{v}} = \boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_{\mathrm{v}} + \boldsymbol{E}_{\mathrm{v}}\ . \tag{8.1}$$

As discussed in Section 3.2.2, the simplest reconstruction technique for the T_2-weighted images is to fill the missing data points in the *k*-space with zeros and to perform the inverse Fourier transform, a process known as Zero-Filling (ZF):

$$\boldsymbol{U}_{\mathrm{v}} = \boldsymbol{\mathcal{F}}^{-1}\boldsymbol{S}^{\mathrm{T}}\boldsymbol{Y}_{\mathrm{v}}\ . \tag{8.2}$$

Zero-Filled reconstructions are typically sub-optimal because they are dominated by noise and under sampling artefacts. Better reconstructions can be obtained by incorporating prior information about $\boldsymbol{U}_{\mathrm{v}}$ in the reconstruction process in the form of regularization. Three types of regularization are considered: Total Variation (TV), Nuclear Norm (NN) and Nuclear Total Generalized Variation (NTGV) regularization.

8.2.1. TV regularization

The application of TV regularization is based on the prior knowledge that each individual T_2-weighted image, $\boldsymbol{U}_{\mathrm{v}}$, is sparse in the domain of pixel-wise differences in all three-dimensions of the image; the transformation into such a domain will be represented by the matrix $\boldsymbol{R}$ whose detailed construction is described in the work of Benning *et al*. [19]. TV reconstructions are then obtained from the following minimization problem:

$$\boldsymbol{U}_{\mathrm{v}} = \arg\min_{\boldsymbol{U}_{\mathrm{v}}} \left(\frac{1}{2}\|\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \alpha\|[\boldsymbol{R}\,\boldsymbol{U}_{\mathrm{v}}]_{\boldsymbol{3}}\|_{2,1}\right). \tag{8.3}$$

α is a regularization parameter controlling the amount of regularization imposed on the reconstruction. The operation $[\boldsymbol{R}\,\boldsymbol{U}_{\mathrm{v}}]_3$, reshapes the vector $\boldsymbol{R}\,\boldsymbol{U}_{\mathrm{v}}$ into a matrix with three columns; each column contains the pixel-wise differences of the T_2-weighted images in a particular direction.

The minimization problem in Eq. (8.3) is tackled numerically by first rewriting it as:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{v}) = \arg\min_{\boldsymbol{U}_{\mathrm{v}}} \max_{\boldsymbol{v}} \left(\frac{1}{2}\|\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \boldsymbol{v}^{\mathrm{T}}\boldsymbol{R}\,\boldsymbol{U}_{\mathrm{v}} - \chi'_{\alpha}(\boldsymbol{v})\right). \tag{8.4}$$

$\boldsymbol{v}$ is known as a dual variable and $\chi'_{\alpha}(\boldsymbol{v})$ is the indicator function defined as:

$$\chi'_{\alpha}(\boldsymbol{v}) = \begin{cases} 0 & \|[\boldsymbol{v}]_{\boldsymbol{3}}\|_{2,\infty} \le \alpha \\ +\infty & \|[\boldsymbol{v}]_{\boldsymbol{3}}\|_{2,\infty} > \alpha \end{cases}\ . \tag{8.5}$$

Eq. (8.5) is in the form of the primal-dual problem discussed in Section 3.3:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{v}) = \min_{\boldsymbol{U}_{\mathrm{v}}} \max_{\boldsymbol{v}} f(\boldsymbol{U}_{\mathrm{v}}) - g(\boldsymbol{v}) + \boldsymbol{v}^{\mathrm{T}} \boldsymbol{R}\, \boldsymbol{U}_{\mathrm{v}}\,, \tag{8.6}$$

where $f(\boldsymbol{U}_{\mathrm{v}}) = ½||\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}||_2^2$ and $g(\boldsymbol{v}) = \chi'_{\alpha}(\boldsymbol{v})$. The iteration scheme of the Primal-Dual Hybrid Gradient Method (PDGHM), outlined in Section 3.3, is then used to numerically solve the primal dual problem in Eq. (8.6). The proximal operator for $f(\boldsymbol{U}_{\mathrm{v}})$ was calculated Section 3.3.3 while the proximal operator for $g(\boldsymbol{v})$ was discussed in Section 5.2.2. A pseudocode for the numerical solution of the TV regularization problem is given in Table 8.1.

Table 8.1. Pseudocode to numerically solve Eq. (8.3), based on the PDHGM method. In the pseudocode, sum($\boldsymbol{A}$, i) sums the column or row vectors of matrix $\boldsymbol{A}$ along the i-th dimension and, repmat($\boldsymbol{A}$, k, i) stacks k copies of $\boldsymbol{A}$ in the i-th dimension.

Step 1. Choose algorithm parameters τ, σ and the regularization parameter, α.

Step 2. Set the convergence tolerance, *TOL*.

Step 3. Initialize $\boldsymbol{v}^{(0)} = \underline{\boldsymbol{v}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{U}_{\mathrm{v}}^{(0)} = \underline{\boldsymbol{U}}_{\mathrm{v}}^{(0)} \neq \boldsymbol{0}$.

Step 4. Initialize count number, $k = 0$ and convergence tracker, $\varepsilon^{(0)} = 0$.

Step 5. **while** $\varepsilon^{(k)} >$ TOL **do**

a. $\underline{\boldsymbol{U}}_{\mathrm{v}}^{(k+1)} \leftarrow \boldsymbol{U}_{\mathrm{v}}^{(k)} - \tau\, \boldsymbol{R}^{\mathrm{T}} \boldsymbol{v}^{(k)}$

b. $\boldsymbol{U}_{\mathrm{v}}^{(k+1)} \leftarrow \boldsymbol{\mathcal{F}}^{-1}\,((\boldsymbol{\mathcal{F}}\,\underline{\boldsymbol{U}}_{\mathrm{v}}^{(k+1)} + \tau\, \boldsymbol{S}^{\mathrm{T}}\boldsymbol{Y}_{\mathrm{v}})/(1 + \tau\, \mathrm{diag}(\boldsymbol{S}^{\mathrm{T}}\boldsymbol{S})))$

The division operation is performed element-wise.

c. $\underline{\boldsymbol{v}}^{(k+1)} \leftarrow \boldsymbol{v}^{(k)} + \sigma\,(2\boldsymbol{U}_{\mathrm{v}}^{(k+1)} - \boldsymbol{U}_{\mathrm{v}}^{(k)})$

d. $\boldsymbol{x} \leftarrow \sqrt{\mathrm{sum}([\underline{\boldsymbol{v}}^{(k+1)}]_3^2, 2)}$

All operations in this line are element-wise.

e. $\boldsymbol{v}^{(k+1)} \leftarrow \underline{\boldsymbol{v}}^{(k+1)}/\max(1, \mathrm{repmat}(\boldsymbol{x}, 3, 1)/\alpha)$

All operations in this line are element-wise.

f. $\varepsilon^{(k+1)} \leftarrow ||\boldsymbol{U}_{\mathrm{v}}^{(k+1)} - \boldsymbol{U}_{\mathrm{v}}^{(k)}||_2 \,/\, ||\boldsymbol{U}_{\mathrm{v}}^{(k)}||_2$

g. $k \leftarrow k + 1$

end while

The condition for the algorithm to converge is $\sqrt{\tau\sigma} \leq 1/||\boldsymbol{R}||_{2,2}$. The choice of $\tau = \sigma = 1/||\boldsymbol{R}||_{2,2}$ was made in this work, where $||\boldsymbol{R}||_{2,2} \approx 2.8$. The number of iterations required to arrive at a reasonable convergence is ~ 300. The time in which this convergence

is achieved for a 64 × 64 × 64 × 64 data set (3 spatial dimensions and 1 echo train dimension) with a 2.0 GHz Intel® Core™ i5-4590T CPU and 16.4 GB RAM, is ~ 6 min.

8.2.2. NN regularization

Other prior information that can be used in improving reconstructions is the fact that the pixel-wise magnitude decays between different T_2-weighted images are highly correlated [20-24]. This is mathematically translated into the matrix $[\boldsymbol{U}_\mathrm{v}]_p$ having a low rank, where p is the echo train length; the operation $[\boldsymbol{U}_\mathrm{v}]_p$, reshapes the vector $\boldsymbol{U}_\mathrm{v}$ into a matrix with p columns. It is a well-known result in linear algebra that the number of non-zero singular values of $[\boldsymbol{U}_\mathrm{v}]_p$ is equal to the rank of $[\boldsymbol{U}_\mathrm{v}]_p$. Therefore, it is in theory possible to improve the reconstructions by constraining $[\boldsymbol{U}_\mathrm{v}]_p$ to have a minimal number of non-zero singular values. In practice, this is computationally intractable, and the problem is relaxed by constraining $[\boldsymbol{U}_\mathrm{v}]_p$ to have a minimal sum of singular values. The sum of singular values is referred to as the nuclear norm, $\|[\boldsymbol{U}_\mathrm{v}]_p\|_*$, or the Schatten-1-norm, $\|[\boldsymbol{U}_\mathrm{v}]_p\|_{S1}$. For a matrix $\boldsymbol{A}$ with singular values $\sigma_i(\boldsymbol{A})$, the Schatten-p-norm is defined as:

$$\|\boldsymbol{A}\|_{S_p} = \left(\sum_i \sigma_i(\boldsymbol{A})^p\right)^{1/p} . \tag{8.7}$$

Nuclear Norm (NN) regularization reconstructions are therefore obtained from the following minimization problem:

$$\boldsymbol{U}_\mathrm{v} = \arg\min_{\boldsymbol{U}_\mathrm{v}} \left(\frac{1}{2}\|\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_\mathrm{v} - \boldsymbol{Y}_\mathrm{v}\|_2^2 + \alpha\left\|[\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}\right\|_{S_1}\right) . \tag{8.8}$$

To numerically solve Eq. (8.8), it is first noted that [25]:

$$\alpha\left\|[\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}\right\|_{S_1} = \max_{\boldsymbol{V}} \mathrm{tr}\left([\boldsymbol{V}^\mathrm{T}\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}\right) = \max_{\boldsymbol{V}} \langle \boldsymbol{V}, [\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}\rangle \quad \text{s.t. } \max_i \sigma_i(\boldsymbol{V}) \le \alpha, \tag{8.9}$$

where $\langle\cdot,\cdot\rangle$ refers to the inner product of matrices and $\boldsymbol{V}$ is known as a dual matrix. As a result, Eq. (8.8) is rewritten as:

$$\boldsymbol{U}_\mathrm{v} = \arg\min_{\boldsymbol{U}_\mathrm{v}} \max_{\boldsymbol{V}} \left(\frac{1}{2}\|\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_\mathrm{v} - \boldsymbol{Y}_\mathrm{v}\|_2^2 + \langle \boldsymbol{V}, [\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}\rangle - \chi''_\alpha(\boldsymbol{V})\right) , \tag{8.10}$$

where $\chi''_\alpha(\boldsymbol{V})$ is an indicator function defined as:

$$\chi''_\alpha(\boldsymbol{V}) = \begin{cases} 0 & \|\boldsymbol{V}\|_{S\infty} \le \alpha \\ +\infty & \|\boldsymbol{V}\|_{S\infty} > \alpha \end{cases} . \tag{8.11}$$

Eq. (8.10) is in the form of the primal-dual problem discussed in Section 3.3:

$$\left([\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}, \boldsymbol{V}\right) = \min_{[\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}} \max_{\boldsymbol{V}} f\left([\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}\right) - g(\boldsymbol{V}) + \langle \boldsymbol{V}, [\boldsymbol{U}_\mathrm{v}]_{\boldsymbol{p}}\rangle , \tag{8.12}$$

where $f([\boldsymbol{U}_v]_p) = ½||\boldsymbol{S}\,\boldsymbol{\mathcal{F}}\,\boldsymbol{U}_v - \boldsymbol{Y}_v||_2^2$ and $g(\boldsymbol{V}) = \chi''_\alpha(\boldsymbol{V})$. The iteration scheme of the PDGHM algorithm, outlined in Section 3.3, is then used to numerically solve the primal-dual problem in Eq. (8.12). The proximal operator for $f([\boldsymbol{U}_v]_p)$ was calculated Section 3.3.3. The proximal operator for $g(\boldsymbol{V})$ was calculated in the work Cai *et al*. [25]; it is equal to $\boldsymbol{V}$ reconstructed from its Singular Value Decomposition, with singular values > α set to be equal to α. A pseudocode for the numerical solution of the NN regularization problem is given in Table 8.2.

The condition for the algorithm to converge is $\tau\,\sigma \leq 1$. The choice of $\tau = \sigma = 1$ is made in this work. The number of iterations required to arrive at a reasonable convergence is ~ 200. The time in which this convergence is achieved for a 64 × 64 × 64 × 64 dataset with a 2.0 GHz Intel® Core™ i5-4590T CPU and 16.4 GB RAM, is ~ 4 min.

Table 8.2. Pseudocode to numerically solve Eq. (8.8), based on the PDHGM method. In the pseudocode, svd($\boldsymbol{A}$) calculates the Singular Value Decomposition of $\boldsymbol{A}$ as $\boldsymbol{A}_1\,\boldsymbol{A}_2\,\boldsymbol{A}_3^T$.

Step 1. Choose algorithm parameters τ, σ and the regularization parameter, α.

Step 2. Set the convergence tolerance, *TOL*.

Step 3. Initialize $\boldsymbol{V}^{(0)} = \underline{\boldsymbol{V}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{U}_v^{(0)} = \underline{\boldsymbol{U}}_v^{(0)} \neq \boldsymbol{0}$.

Step 4. Initialize count number, $k = 0$ and convergence tracker, $\varepsilon^{(0)} = 0$.

Step 5. **while** $\varepsilon^{(k)} >$ TOL **do**

a. $[\underline{\boldsymbol{U}}_v]_p^{(k+1)} \leftarrow [\boldsymbol{U}_v]_p^{(k)} - \tau\,\boldsymbol{V}^{(k)}$

b. $\boldsymbol{U}_v^{(k+1)} \leftarrow \boldsymbol{\mathcal{F}}^{-1}\,((\boldsymbol{\mathcal{F}}\,\underline{\boldsymbol{U}}_v^{(k+1)} + \tau\,\boldsymbol{S}^T\boldsymbol{Y}_v)/(1 + \tau\,\mathrm{diag}(\boldsymbol{S}^T\boldsymbol{S})))$

The division operation is performed element-wise.

c. $\underline{\boldsymbol{V}}^{(k+1)} \leftarrow \boldsymbol{V}^{(k)} + \sigma\,(2[\boldsymbol{U}_v]_p^{(k+1)} - [\boldsymbol{U}_v]_p^{(k)})$

h. $[\boldsymbol{A}_1, \boldsymbol{A}_2, \boldsymbol{A}_3] \leftarrow \mathrm{svd}(\underline{\boldsymbol{V}}^{(k+1)})$

d. $\boldsymbol{A}_2 \leftarrow \min(\alpha, \boldsymbol{A}_2)$

The operation in this line is element-wise

e. $\boldsymbol{V}^{(k+1)} \leftarrow \boldsymbol{A}_1\,\boldsymbol{A}_2\,\boldsymbol{A}_3^T$

f. $\varepsilon^{(k+1)} \leftarrow ||\boldsymbol{U}_v^{(k+1)} - \boldsymbol{U}_v^{(k)}||_2 \,/\, ||\boldsymbol{U}_v^{(k)}||_2$

g. $k \leftarrow k + 1$

end while

8.2.3. NTGV regularization

TV regularization uses prior information about the individual T_2-weighted images to improve the reconstructions, while NN regularization uses the relationship between the different T_2-weighted images to improve the reconstructions. The natural extension is to combine these regularization techniques such that the reconstructions inherit the best features from TV and NN regularization. One such technique is Nuclear Total Generalized Variation (NTGV) which consists in estimating the T_2-weighted images from the following minimization problem:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{W}) = \arg\min_{\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{W}} \left(\frac{1}{2} \|\boldsymbol{S}\, \boldsymbol{\mathcal{F}}\, \boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \alpha \left\| [\boldsymbol{U}_{\mathrm{v}}]_{\boldsymbol{p}} - [\boldsymbol{W}_{\mathrm{v}}]_{\boldsymbol{p}} \right\|_{S_1} + \beta \|[\boldsymbol{R}\, \boldsymbol{W}_{\mathrm{v}}]_{\boldsymbol{3}}\|_{2,1} \right), \tag{8.13}$$

where $\boldsymbol{W}$ is an auxiliary matrix. The second term enforces the reconstruction to inherit good features from NN regularization while the third term enforces the reconstruction to inherit good features from TV regularization. The regularization parameters α and β control the amount of regularization imposed. The idea behind NTGV regularization is similar to the idea behind MTGV regularization, which was discussed in Chapter 5.

To numerically solve Eq. (8.13), the ideas introduced in Chapter 5 and in Sections 8.2.1 and 8.2.2 are used to rewrite NTGV regularization as:

$$(\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{W}, \boldsymbol{V}, \boldsymbol{z}) = \arg\min_{\boldsymbol{U}_{\mathrm{v}}, \boldsymbol{W}} \left(\frac{1}{2} \|\boldsymbol{S}\, \boldsymbol{\mathcal{F}}\, \boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}\|_2^2 + \langle \boldsymbol{V}, [\boldsymbol{U}_{\mathrm{v}}]_{\boldsymbol{p}} - [\boldsymbol{W}_{\mathrm{v}}]_{\boldsymbol{p}} \rangle - {\chi''}_{\alpha}(\boldsymbol{V}) + \boldsymbol{z}^{\mathrm{T}} \boldsymbol{R}\, \boldsymbol{W}_{\mathrm{v}} - {\chi'}_{\beta}(\boldsymbol{z}) \right), \tag{8.14}$$

where ${\chi''}_{\alpha}(\boldsymbol{V})$ and ${\chi'}_{\beta}(\boldsymbol{z})$ were defined in Sections 8.2.1 and 8.2.2. The approach in Chapter 5 is then followed to write Eq. (8.14) in a primal-dual problem format:

$$\left([\boldsymbol{U'}_{\mathrm{v}}]_{\boldsymbol{p}}, \boldsymbol{V'} \right) = \min_{[\boldsymbol{U'}_{\mathrm{v}}]_{\boldsymbol{p}}} \max_{\boldsymbol{V'}} f\left([\boldsymbol{U'}_{\mathrm{v}}]_{\boldsymbol{p}} \right) - g(\boldsymbol{V'}) + \langle \boldsymbol{V'}, \boldsymbol{Q}\, \boldsymbol{U'}_{\mathrm{v}} \rangle , \tag{8.15}$$

where $f([\boldsymbol{U'}_{\mathrm{v}}]_p) = ½||\boldsymbol{S}\, \boldsymbol{\mathcal{F}}\, \boldsymbol{U}_{\mathrm{v}} - \boldsymbol{Y}_{\mathrm{v}}||_2^2$ and $g(\boldsymbol{V'}) = {\chi''}_{\alpha}(\boldsymbol{V}) + {\chi'}_{\beta}(\boldsymbol{z})$. $\boldsymbol{Q}$ is constructed as in Section 5.2.1. The iteration scheme of the PDHGM algorithm, outlined in Section 3.3, is then used to numerically solve the minimization problem. The proximal operator for $f([\boldsymbol{U'}_{\mathrm{v}}]_p)$ was derived in Section 3.3.3. The proximal operator of $g(\boldsymbol{V'})$ is obtained from the sum of the proximal operators of its constituent functions, each of which were discussed in Sections 5.2.1 and 8.2.2. A pseudocode for the numerical solution of the NTGV regularization problem is given in Table 8.3.

Table 8.3. Pseudocode to numerically solve Eq. (8.13), based on the PDHGM method. In the pseudocode, sum($\boldsymbol{A}$, i) sums the column or row vectors of matrix $\boldsymbol{A}$ along the i-th dimension, repmat($\boldsymbol{A}$, k, i) stacks k copies of $\boldsymbol{A}$ in the i-th dimension and, svd($\boldsymbol{A}$) calculates the Singular Value Decomposition of $\boldsymbol{A}$ as $\boldsymbol{A}_1 \boldsymbol{A}_2 \boldsymbol{A}_3^{\mathrm{T}}$.

Step 1. Choose algorithm parameters τ, σ and the regularization parameters, α, β.

Step 2. Set the convergence tolerance, TOL.

Step 3. Initialize $\boldsymbol{V}^{(0)} = \underline{\boldsymbol{V}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{z}^{(0)} = \underline{\boldsymbol{z}}^{(0)} = \boldsymbol{0}$, $\boldsymbol{U}_{\mathrm{v}}^{(0)} = \underline{\boldsymbol{U}}_{\mathrm{v}}^{(0)} \neq \boldsymbol{0}$, $\boldsymbol{W}_{\mathrm{v}}^{(0)} \neq \boldsymbol{0}$.

Step 4. Initialize count number, $k = 0$ and convergence tracker, $\varepsilon^{(0)} = 0$.

Step 5. **while** $\varepsilon^{(k)} >$ TOL **do**

a. $[\underline{\boldsymbol{U}}_{\mathrm{v}}]_p^{(k+1)} \leftarrow [\boldsymbol{U}_{\mathrm{v}}]_p^{(k)} - \tau\, \boldsymbol{V}^{(k)}$

b. $\boldsymbol{U}_{\mathrm{v}}^{(k+1)} \leftarrow \boldsymbol{\mathcal{F}}^{-1}\left((\boldsymbol{\mathcal{F}}\, \underline{\boldsymbol{U}}_{\mathrm{v}}^{(k+1)} + \tau\, \boldsymbol{S}^{\mathrm{T}}\boldsymbol{Y}_{\mathrm{v}})/(1 + \tau\, \mathrm{diag}(\boldsymbol{S}^{\mathrm{T}}\boldsymbol{S}))\right)$

The division operation is performed element-wise.

c. $[\boldsymbol{W}_{\mathrm{v}}]_p^{(k+1)} \leftarrow [\boldsymbol{W}_{\mathrm{v}}]_p^{(k)} + \tau\, (\boldsymbol{V}^{(k)} - \boldsymbol{Q}^{\mathrm{T}}\boldsymbol{z}^{(k)})$

d. $\underline{\boldsymbol{V}}^{(k+1)} \leftarrow \boldsymbol{V}^{(k)} + \sigma\, (2\boldsymbol{U}_{\mathrm{v}}^{(k+1)} - [\boldsymbol{U}_{\mathrm{v}}]_p^{(k)} - 2[\boldsymbol{W}_{\mathrm{v}}]_p^{(k+1)} + [\boldsymbol{W}_{\mathrm{v}}]_p^{(k)})$

e. $[\boldsymbol{A}_1, \boldsymbol{A}_2, \boldsymbol{A}_3] \leftarrow \mathrm{svd}(\underline{\boldsymbol{V}}^{(k+1)})$

f. $\boldsymbol{A}_2 \leftarrow \min(\alpha, \boldsymbol{A}_2)$

The operation in this line is element-wise

g. $\boldsymbol{V}^{(k+1)} \leftarrow \boldsymbol{A}_1\, \boldsymbol{A}_2\, \boldsymbol{A}_3^{\mathrm{T}}$

h. $\underline{\boldsymbol{z}}^{(k+1)} \leftarrow \boldsymbol{z}^{(k)} + \sigma\, \boldsymbol{Q}\, (2\boldsymbol{W}_{\mathrm{v}}^{(k+1)} - \boldsymbol{W}_{\mathrm{v}}^{(k)})$

i. $\boldsymbol{x} \leftarrow \sqrt{\mathrm{sum}([\underline{\boldsymbol{z}}^{(k+1)}]_3^2, 2)}$

All operations in this line are element-wise.

j. $\boldsymbol{v}^{(k+1)} \leftarrow \underline{\boldsymbol{z}}^{(k+1)}/\max(1, \mathrm{repmat}(\boldsymbol{x}, 3, 1)/\beta)$

All operations in this line are element-wise.

k. $\varepsilon^{(k+1)} \leftarrow ||\boldsymbol{U}_{\mathrm{v}}^{(k+1)} - \boldsymbol{U}_{\mathrm{v}}^{(k)}||_2 \,/\, ||\boldsymbol{U}_{\mathrm{v}}^{(k)}||_2$

l. $k \leftarrow k + 1$

end while

The condition for the algorithm to converge is $\tau\, \sigma \leq 1/12$. The choice of $\tau = \sigma \approx 0.28$ is made in this work. The number of iterations required to arrive at a reasonable convergence is ~ 1000. The time in which this convergence was achieved for a $64 \times 64 \times 64 \times 64$ dataset with a 2.0 GHz Intel® Core™ i5-4590T CPU and 16.4 GB RAM, is ~ 1 h.

8.2.4. Reconstruction quality measure

The reconstructed T_2-weighted images obtained from the different reconstruction techniques are compared using the peak-signal-to-noise ratio (PSNR) metric:

$$\mathrm{PSNR} = 20\log_{10}\left(\frac{\mathrm{length}(\boldsymbol{U}_{\mathrm{v}})\,\max(\boldsymbol{U}_{\mathrm{v,FS}})}{\|\boldsymbol{U}_{\mathrm{v}}-\boldsymbol{U}_{\mathrm{v,FS}}\|_2^2}\right) . \tag{8.16}$$

$\boldsymbol{U}_{\mathrm{v}}$ refers to the T_2-weighted images obtained from under sampled k-space data processed either using Zero-Filling, TV regularization, NN regularization or NTGV regularization, while $\boldsymbol{U}_{\mathrm{v,FS}}$ refers to the T_2-weighted images obtained from fully sampled k-space data. The reconstruction with the larger PSNR is the better reconstruction. The perfect reconstruction has PSNR = ∞. The accuracy of reconstruction from TV regularization and NN regularization depends on the value of the regularization parameter α, while the accuracy of reconstruction from NTGV regularization depends on the regularization parameters α and β. For all methods, a range of regularization parameters was considered, and the regularization parameter which gave the largest PSNR was chosen. The results reported in this chapter correspond to the reconstructions obtained using these parameter values.

For completeness, NTGV regularization could be turned into a one-parameter regularization technique by heuristically constraining α and β to be functions of each other, as was done with MTGV regularization in Chapter 5. For systems in which no ground truth is known, the remaining independent regularization parameter could then be chosen using any of the techniques described in the work of Mitchell *et al*. [26]. These extensions are not investigated in this chapter but could be the focus of future investigations.

The T_2 maps which result from the reconstructed T_2-weighted images are not compared quantitatively, because the step of obtaining the T_2 map commonly involves a regularization technique, which introduces some uncertainty because of the need to choose the regularization parameter. However, an illustration of the T_2 map obtained using L_1 regularization, described in Chapter 4, is given in Section 8.4.

8.3. Materials and methods

The experimental acquisitions were performed with the help of Dr. Daan de Kort of the Magnetic Resonance Research Centre, University of Cambridge.

The acquisition of a 3D spatially resolved T_2 map is performed on the sample illustrated in Fig. 8.1(a). The sample is composed of 7 tubes of 10 mm diameter, bound in a circular arrangement. Each tube is filled up to a height of 40 mm; in the bottom 20 mm by a gadolinium chloride ($GdCl_3.6H_2O$) solution and in the top 20 mm by polydimethylsiloxane, trimethylsiloxy terminated (Alfa Aesar), which will be referred to as PDMS oil. Gadolinium chloride solutions of three different concentrations are used: 0.82 mM (fluid A), 0.55 mM (fluid B) and 0.22 mM (fluid C), and the particular tube allocation is shown in Fig. 1(a). PDMS oils of three different molecular weights are used: 28 kDa (fluid D), 9 kDa (fluid E) and 2 kDa (fluid F). Again, the particular tube allocation for the PDMS oils is given in Fig. 8.1(a).

All experiments are conducted on a Bruker SWB superconducting magnet, operating at a resonant frequency of 300.88 MHz for ^{1}H observation, using a radio frequency (r.f.) coil of diameter 66 mm. The maximum magnetic field gradient strength available is 77 G cm^{-1}.

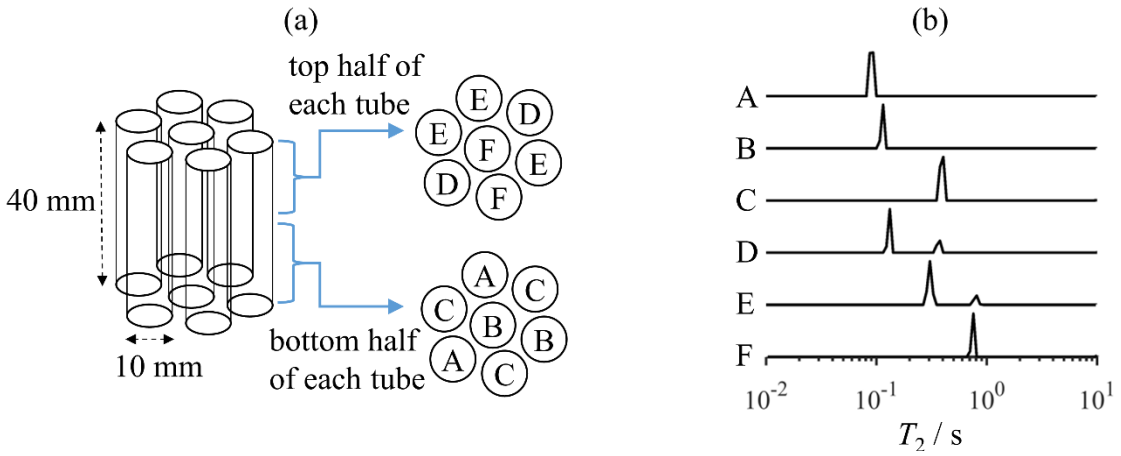

Fig. 8.1. (a) Schematic arrangement of the tubes used in the experiment (not to scale). Each tube is filled in the bottom half with a doped water solution (A, B or C) and on the top half with PDMS oil (D, E or F). (b) Bulk T_2 distributions of each fluid used in the experiment, A-F, obtained individually from single shot CPMG experiments.

The bulk T_2 distributions for each individual fluid used, obtained using a single shot CPMG experiment, as described in Section 2.10, with echo time 15 ms and processed using L_1 regularization, as described in Chapter 4, are shown in Fig. 8.1(b). L_1 regularization was selected based on the fact that the PDMS oils D and E did not show mono-exponential behaviour in the T_2 decay and the conventional Tikhonov regularization method was unable to distinguish the two peaks of fluid D in the T_2 distribution. The T_2 distribution for each individual fluid was unchanged before and after the experiment,

indicating that the gadolinium chloride remains largely dissolved only in the water solution and does not affect the T_2 distribution of the PDMS oils.

The MESE imaging technique used is similar to the RARE pulse sequence described in Section 2.16, adapted for 3D imaging by adding another phase magnetic field gradient to substitute for the slice selective magnetic field gradient. Hard pulses of duration 82.5 μs and 165 μs for 90° and 180° r.f. excitations, respectively, are used. The pulse sequence is designed for a field-of-view of 35 mm × 35 mm × 45 mm, and respective spatial resolution of 547 μm × 547 μm × 703 μm, with the first two dimensions being phase directions (k_{p1} and k_{p2}) and the third dimension being a read direction, k_r (which is parallel to the axis of the tubes). A dwell time of 5 μs, echo time of 15 ms and repetition time of 3.5 s are used. Each echo train is composed of 64 echoes (RARE factor of 64).

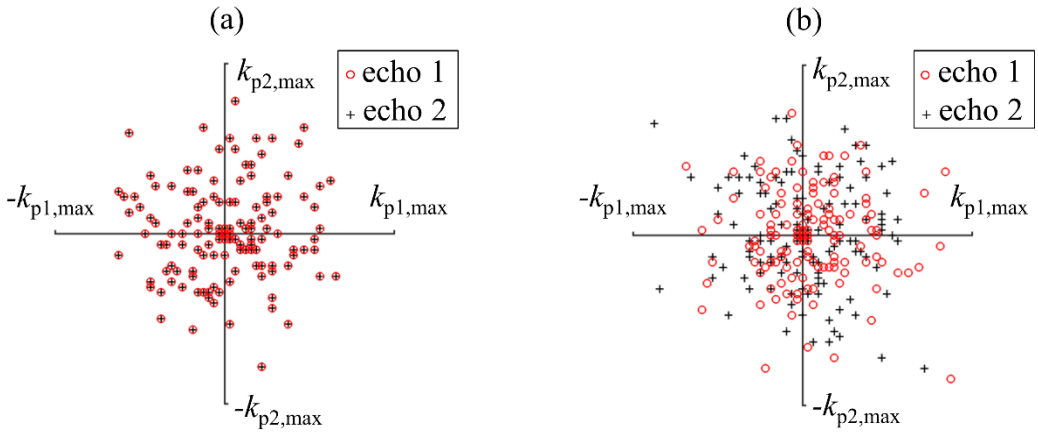

Fig. 8.2. The sampling pattern used for the first two echoes and for the two dimensions of *k*-space which can be randomly under sampled; the patterns are shown for 3.1% sampling of *k*-space. Patterns are shown for (a) a coherent sampling pattern and (b) an incoherent sampling pattern.

Under sampling is performed on the two phase directions of *k*-space and sampling percentages of 100% (fully sampled), 50%, 25%, 12.5%, 6.3% and 3.1% of *k*-space are investigated. The corresponding acquisition times are 16 h, 8 h, 4 h, 2 h, 1 h and 30 min. For the two under sampled dimensions, the 2D sampling pattern follows a bi-level approach [27]. The central 3 × 3 out of the 64 × 64 pixels of *k*-space are fully sampled while the subsequent points acquired are random, with the density of points following a Gaussian distribution from the centre *k*-space with a standard deviation of 15% of the largest *k*-space value. With these two under sampled dimensions (k_{p1} and k_{p2}) and a fully sampled third dimension (k_r), the sidelobe-to-peak ratio (SPR) [28] of the 3D sampling

pattern is 0, 0.029, 0.059, 0.091, 0.141 and 0.246 for the corresponding sampling percentages of 100%, 50%, 25%, 12.5%, 6.3% and 3.1% of *k*-space. Two sampling patterns, illustrated in Fig. 8.2, are investigated: a conventional coherent sampling pattern which, in practice, means that for each *k*-space line sampled in the read direction, all echoes down the echo train are acquired and; an incoherent sampling pattern which means that the *k*-space line sampled in the read direction at each echo in the echo train is independent of the *k*-space line sampled in the previous echo. Once the sampling pattern is selected, the choice of the specific *k*-space lines to be acquired during an echo train is random. This is done to distribute the gradient dissipation homogenously and, therefore, to reduce the effect of the diffusive attenuation of the signal.

8.4. Results and discussion

The results of the reconstructions of the T_2-weighted images from the under sampled MRI data using different techniques are now presented. For illustration, only the results corresponding to the experimental data acquired at the lowest sampling percentage, 3.1% of *k*-space, are shown.

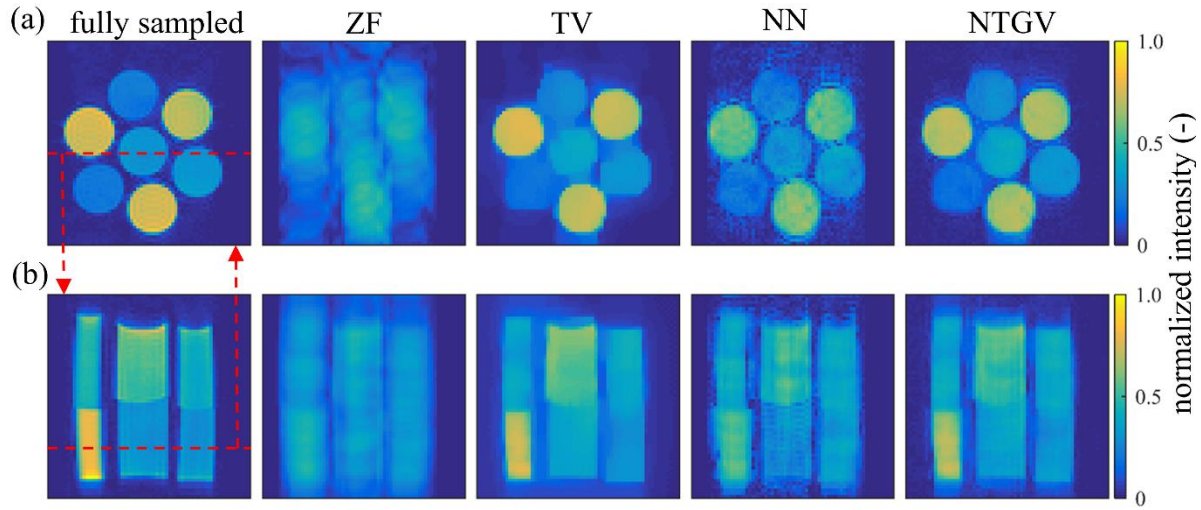

Fig. 8.3. Comparison of the T_2-weighted images obtained from a fully sampled dataset with the reconstructed T_2-weighted images obtained using ZF, TV regularization, NN regularization and NTGV regularization for a coherent sampling pattern. The data are shown for a sampling percentage of 3.1% of *k*-space. (a) An example of a transverse T_2-weighted image through the tubes at a position shown in (b) and echo number 10. (b) An example of a longitudinal T_2-weighted image through the tubes at a position shown in (a) and echo number 10.

Fig. 8.3 shows the reconstructions of the T_2-weighted images using Zero-Filling, TV regularization, NN regularization and NTGV regularization for a coherent sampling pattern. An example is given for two selected slices (one transverse and one longitudinal) through the 3D sample; data are shown for the 10^{th} echo. These images are also compared with the corresponding images obtained from the fully sampled MRI data. It is observed that the Zero-Filled reconstructions are poor and dominated by under sampling artefacts; it is not possible to distinguish between the different tubes present in the sample. The application of TV regularization gives a significant improvement, but the reconstructions have artificial stair-casing features, a well-known feature of over-regularization in TV regularization [19]. The boundary between some of the tubes is not very clear, particularly in the transverse slice image. The reconstructions obtained from NN regularization seem slightly worse than the reconstructions obtained from TV regularization; more noise is observed both within and outside the tubes, although the tubes are more clearly distinguished. NTGV regularization gives improved reconstructions over TV and NN regularization, although not noticeably. The NTGV regularization reconstructions inherit good features from TV and NN regularization: the noise level is reduced (a TV regularization property) and the tubes are clearly distinguished (a NN regularization property).

Fig. 8.4 shows the reconstructions of the T_2-weighted images using Zero-Filling, TV regularization, NN regularization and NTGV regularization for an incoherent sampling pattern. Again, Zero-Filled reconstructions are poor and dominated by under sampling artefacts. The reconstructions obtained from TV regularization are similar to the results from TV regularization in Fig. 8.3, where the sampling pattern was coherent. This is explained by the fact that TV regularization is imposed on each individual 3D T_2-weighted image independently of the other 3D T_2-weighted images. Therefore, whether the sampling pattern for the corresponding *k*-space of the different T_2-weighted images is the same or different does not make a noticeable difference. The reconstructions obtained from NN regularization show significant improvement both over TV and NN reconstructions obtained using a coherent sampling pattern. This significant improvement in the NN regularization performance is explained by the fact that NN regularization relies on the correlation between the different T_2-weighted images. Therefore, having different sampling patterns for the corresponding *k*-space of the different T_2-weighted images reduces the coherent artefacts that are introduced using a coherent sampling pattern.

NTGV regularization gives further improvement in performance over TV and NN regularization. NTGV reconstructions inherit good features from TV and NN regularization, similar to the case of a coherent sampling pattern.

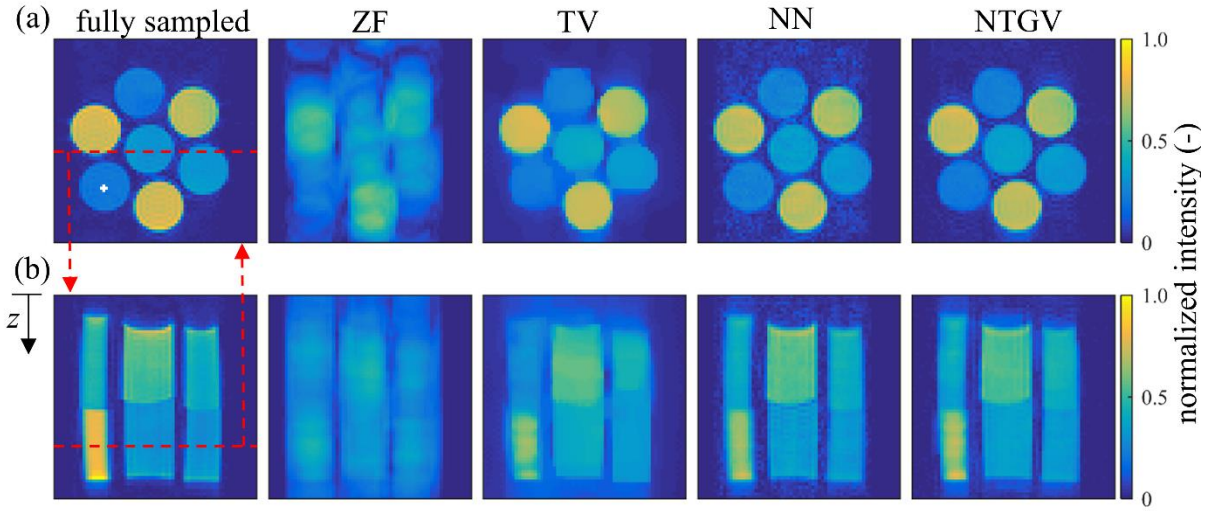

Fig. 8.4. Comparison of the T_2-weighted images obtained from a fully sampled dataset with the reconstructed T_2-weighted images obtained using ZF, TV regularization, NN regularization and NTGV regularization for an incoherent sampling pattern. The data are shown for a sampling percentage of 3.1% of *k*-space. (a) An example of a transverse T_2-weighted image through the tubes at a position shown in (b) and echo number 10. (b) An example of a longitudinal T_2-weighted image through the tubes at a position shown in (a) and echo number 10. The '+' marker in the fully sampled image of (a) denotes the projection of the line over which the T_2 mapping shown in Fig. 8.6 is taken. The distance '*z*' in (b) corresponds to the distance '*z*' used in Fig. 8.6.

A quantitative comparison using the PSNR metric between the reconstructions obtained using Zero-Filling, TV regularization, NN regularization and NTGV regularization is presented in Fig. 8.5. The comparison is made over the different sampling percentages studied: 50%, 25%, 12.5%, 6.3% and 3.1% of *k*-space. Fig. 8.5(a) shows the results for a coherent sampling pattern. At all sampling percentages, NTGV regularization outperforms TV regularization, which in turn outperforms NN regularization. All the regularization techniques perform better than Zero-Filling. These agree with the qualitative observations in Fig. 8.3. The performance gap between NTGV regularization and TV regularization is negligible, suggesting that when a coherent sampling pattern is used, there is only a small benefit in using NTGV regularization over the conventional TV regularization. Fig. 8.5(b) shows the results for an incoherent sampling pattern. The first observation is that NTGV regularization still remains the most accurate reconstruction method. The performance gap between NTGV regularization and the other regularization techniques has increased. In

addition, the performance of NTGV regularization for an incoherent sampling pattern is significantly better than for a coherent sampling pattern. These two observations suggest that the full potential of NTGV regularization is only utilized when an incoherent sampling pattern is used. This is consistent with the earlier discussion. It is also observed that the performance of NN regularization is significantly improved from the case of a coherent sampling pattern and its performance is better than TV regularization, unlike for the case of a coherent sampling pattern. Further, the performance of TV regularization and Zero-Filling is observed to be similar to when using a coherent sampling pattern, as expected.

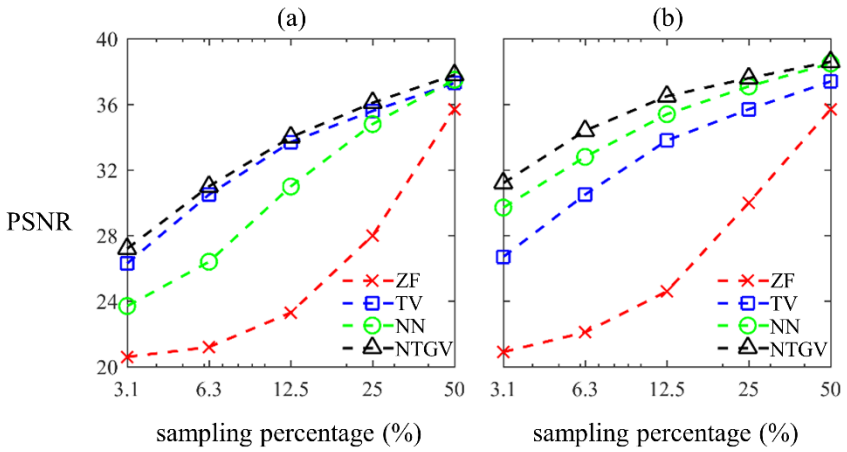

Fig. 8.5. Comparison of the different techniques: ZF, TV regularization, NN regularization and NTGV regularization in reconstructing the T_2-weighted images from under sampled MRI data at different sampling percentages. Data are shown for (a) a coherent and (b) an incoherent sampling pattern. The PSNR metric is defined in Eq. (8.16). The lines are included to guide the eye.

The results presented suggest that using an incoherent sampling pattern and NTGV regularization as the reconstruction technique gives the best reconstructions at all sampling percentages studied, according to the PSNR metric. In practice, there are other factors that one might want to consider when choosing the best reconstruction technique. An important factor is that NTGV regularization is a two-parameter regularization technique, as opposed to TV regularization and NN regularization which are one-parameter regularization techniques. Further, the convergence speed for the numerical solution of the NTGV regularization problem is slower than the convergence speed for the numerical solution of the TV regularization and NN regularization methods; typically 3-5 times more iterations are needed and the time per iteration is approximately double for NTGV regularization, as compared to TV and NN regularization. Both these factors make the processing time for

NTGV regularization method longer than the processing time for TV regularization and NN regularization. This observation is particularly important when the sampling percentage is large. For example, at sampling percentages of 25% and 50% of *k*-space, no visible difference between the reconstructions obtained from the different techniques was observed, although the PSNR metric reveals some differences. Therefore, at the high sampling percentages of 25% and 50% of *k*-space, the small gain in image quality from NTGV regularization over TV and NN regularization may be outweighed by the larger processing time of NTGV regularization.

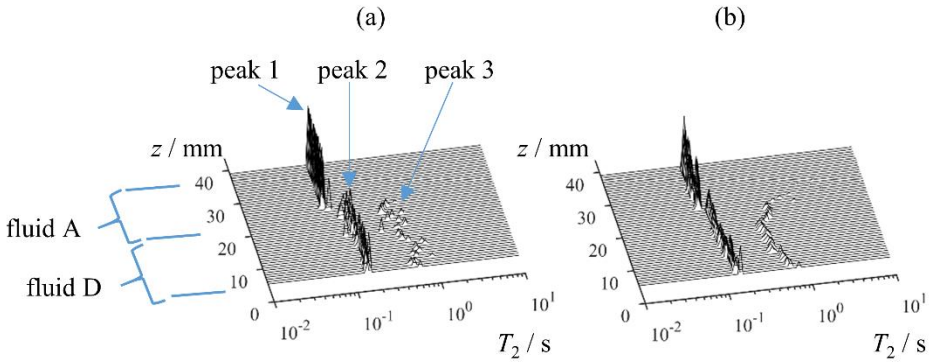

Fig. 8.6. 1D T_2 maps over the line denoted in Fig. 8.4(a). The 1D T_2 maps were obtained from (a) the fully sampled dataset and, (b) the data acquired at 3.1% sampling of *k*-space data using an incoherent sampling pattern and processed using NTGV regularization. The line over which the mapping is performed runs through fluids D and A. The distance *z* is defined in Fig. 8.4(b).

For completeness, the 3D T_2 maps obtained from the T_2-weighted images are now considered. Each pixel in the 3D image is assigned a T_2 distribution by using L_1 regularization, outlined in Chapter 4, to convert the decay of the pixel intensity into a T_2 distribution. As this step is not the focus of this chapter, only an example is presented; the optimisation of the L_1 regularization technique is not considered. Fig. 8.6 shows the spatially resolved T_2 distribution over a line through the 3D image (1D T_2 map), whose location is annotated in Fig. 8.4(a). The line over which the T_2 mapping is performed runs through the tube filled with fluids D and A, whose bulk T_2 distributions are given in Fig. 8.1(b). Fig. 8.6(a) shows the results obtained from the T_2-weighted images of the fully sampled data. Fig. 8.6(b) shows the results obtained from the T_2-weighted images reconstructed from the MRI data acquired at 3.1% sampling of *k*-space and an incoherent sampling pattern, processed using NTGV regularization. The position of the peaks in the T_2 distributions of fluids A (peak 1) and D (peaks 2 and 3), and the relative intensities of

peaks 2 and 3 are summarised in Table 8.4. It is observed that the 1D T_2 maps obtained from both the fully sampled and under sampled data are in good agreement with the bulk liquid measurements in Fig. 8.1(b). The positions of peaks 1 and 2, as obtained from Fig. 8.6(a) and 8.6(b) are < 15% different from the results in Fig. 8.1(b). The position of peak 3 is less well predicted from the 3D T_2 maps of both fully sampled and under sampled data; this is mainly due to the low intensity of this peak as compared to peak 2. For fluid D, the fraction of the total intensity represented by peak 2, obtained from Figs. 8.1(b), 8.6(a) and 8.6(b) are in good agreement; respectively 0.67, 0.82 ± 0.07 and 0.73 ± 0.05. The fact that the two peaks in the T_2 distribution of fluid D can be distinguished and quantified to a good degree even from highly under sampled data is an achievement of the proposed sampling scheme and processing method. Indeed, it was shown in Chapter 7 that bi-exponential decays where the population of one of the components is small are challenging to resolve even from 1D NMR experiments.

Table 8.4: Estimation of the position of the peaks in the T_2 distribution of fluids A (peak 1) and D (peaks 2 and 3), and the fraction of total intensity of fluid D represented by peak 2, determined using three different techniques. The bulk measurement is taken from Fig. 8.1(b). The measurement from the fully sampled dataset is taken from Fig. 8.6(a). The measurement from the under sampled dataset is taken from Fig. 8.6(b). The peaks are annotated in Fig. 8.6(a). The uncertainty shown corresponds to the spatial variation of the properties, as observed from Fig. 8.6.

	T_2 / ms			intensity fraction
	peak 1	peak 2	peak 3	peak 2 / (peak 2+peak 3)
bulk measurement	87	140	380	0.67
fully sampled dataset	88 ± 3	160 ± 10	600 ± 160	0.82 ± 0.07
under sampled dataset	90 ± 8	130 ± 10	540 ± 90	0.73 ± 0.05

8.5. Conclusions

This chapter described an experimental implementation of an MESE pulse sequence for the acquisition of under sampled MRI data, with the aim of obtaining a 3D T_2 map. Two classes of sampling patterns were compared: a conventional coherent sampling pattern where for each line sampled in *k*-space, all echoes down the echo train are acquired, and an incoherent sampling pattern where the *k*-space line sampled for each echo in the echo train is independent of the *k*-space line sampled in the previous echo. The conventional

reconstruction technique of TV regularization was compared to the more recent techniques of NN and NTGV regularization. It was observed that an incoherent sampling pattern is superior to a coherent sampling pattern and that NTGV regularization outperformed TV regularization and NN regularization at all sampling percentages studied. The full potential of NTGV regularization is particularly utilized when an incoherent sampling pattern is used. Using an incoherent sampling scheme and NTGV regularization as the reconstruction technique, quantitative results were obtained even at a sampling percentage of 3.1% of *k*-space, corresponding to a 32-fold decrease in the acquisition time.

References

[1] D. R. Bailes, I. R. Young, D. J. Thomas, K. Straughan, G.M. Bydder and R. E. Steiner, NMR imaging of the brain using spin-echo sequences, Clin. Radiol. 44 (1982), 395-414.

[2] M. A. Johnson, J. M. Pennock, G. M. Bydder, R. E. Steiner, D. J. Thomas, R. Hayward, D. R. T. Bryant, J. A. Payne, M. I. Levene, A. Whitelaw, L. M. S. Dubowitz and V. Dubowitz, Clinical NMR imaging of the brain in children: Normal and neurologic disease, Am. J. Roentgenol. 141 (1983), 1005-1018.

[3] C. S. Poon and R. M. Henkelman, Practical T_2 quantitation for clinical-applications, J. Magn. Reson. Imaging 2 (1992), 541-553.

[4] S. E. Harms, P. T. Siemens, P. Hildenbrand and G. Plum, Multiple spin echo magnetic resonance imaging of the brain, RadioGraphics 6 (1986), 117-134.

[5] D. A. Feinberg, C. M. Mills, J. P. Posin, D. A. Ortendahl, N. M. Hylton, L. E. Crooks, J. C. Watts, L. Kaufman, M. Arakawa, J. C. Hoenninger and M. Brant-Zawadzki, Multiple spin-echo magnetic resonance imaging, Radiology 155 (1985), 437-442.

[6] B. J. Dardzinski, T. J. Mosher, S. Li, M. A. van Slyke and M. B. Smith, Spatial variation of T_2 in human articulate cartilage, Radiology 205 (1997), 546-550.

[7] G. P. Liney, A. J. Knowles, D. J. Manton, L. W. Turnbull, S. J. Blackband and A. Horsman, Comparison of conventional single echo and multi-echo sequences with a fast spin-echo sequence for quantitative T_2 mapping: Application to the prostate, J. Magn. Reson. Imaging 6 (1996), 603-607.

[8] M. S. Sussman, L. Vidarsson, J. M. Pauly and H. M. Cheng, A technique for rapid single-echo spin-echo T_2 mapping, Magn. Reson. Med. 64 (2010), 536-545.

[9] O. Petrov, G. Ersland and B. J. Balcom, T_2 distribution mapping profile with phase-encode MRI, J. Magn. Reson. 209 (2011), 39-46.

[10] O. V. Petrov and B. J. Balcom, Two-dimensional T_2 distribution mapping in porous solids with phase encode MRI, J. Magn. Reson. 212 (2011), 102-108.

[11]R. E. Nechifor, K. Romanenko, F. Marica and B. J. Balcom, Spatially resolved measurements of mean spin-spin relaxation time constants, J. Magn. Reson. 239 (2014), 16-22.

[12]D. Xiao and B. J. Balcom, Two-dimensional T_2 distribution mapping in rock core plugs with optimal *k*-space sampling, J. Magn. Reson. 220 (2012), 70-78.

[13]D. Xiao and B. J. Balcom, Restricted *k*-space sampling in pure phase encode MRI of rock core plugs, J. Magn. Reson. 231 (2013), 126-132.

[14]J.-F. Cai, E. J. Candés and Z. Shen, A singular value thresholding algorithm for matrix completion, SIAM J. Optim. 20 (2010), 1956-1982.

[15]F. Knoll, M. Holler, T. Koesters, M. Cloos, R. Otazo, K. Bredies and D. S. Sodickson, Simultaneous multi-modality/multi-contrast image reconstruction with nuclear-norm TGV, Proc. Intl. Soc. Mag. Reson. Med. 24 (2016), 873.

[16]F. Knoll, M. Holler, T. Koesters, R. Otazo, K. Bredies and D. S. Sodickson, Joint MR-PET reconstruction using a multi-channel image regularizer, IEEE Trans. Med. Imaging 36 (2017), 1-16.

[17]M. Davies, G. Puy, P. Vandergheynst and Y. Wiaux, A compressed sensing framework for magnetic resonance fingerprinting, SIAM J. Imaging Sci. 7 (2014), 2623-2656.

[18]R. Duarte, A. Repetti, P. A. Gómez, M. Davies and Y. Wiaux, Greedy approximate projection for magnetic resonance fingerprinting with partial volumes (2018), arxiv: 1807.06912.

[19]M. Benning, L. F. Gladden, D. Holland, C.-B. Schönlieb and T. Valkonen, Phase reconstruction from velocity-encoded MRI measurements – A survey of sparsity - promoting variational approaches, J. Magn. Reson. 238 (2014), 26-43.

[20]F. H. Petzschner, I. P. Ponce, M. Blaimer, P. M. Jakob and F. A. Breuer, Fast MR parameter mapping using *k-t* principal component analysis, Magn. Reson. Med. 66 (2011), 706-716.

[21]L. Feng, R. Otazo, H. Jung, J. H. Jensen, J. C. Ye, D. K. Sodickson and D. Kim, Accelerated cardiac T_2 mapping using breath-hold multiecho fast spin-echo pulse sequence with *k-t* FOCUSS, Magn. Reson. Med. 65 (2011), 1661-1669.

[22]C. Huang, C. G. Graff, E. W. Clarkson, A. Bilgin and M. I. Altbach, T_2 mapping from highly under sampled data by reconstruction of principal component coefficient maps using compressed sensing, 67 (2012), 1355-1366.

[23]T. Zhang, J. M. Pauly and I. R. Levesque, Accelerating parameter mapping with a locally low rank constraint, Magn. Reson. Med. 73 (2015), 655-661.

[24]B. Zhao, W. Lu, T. K. Hitchens, F. Lam, C. Ho and Z.-P. Liang, Accelerated MR parameter mapping with low-rank and sparsity constraints, Magn. Reson. Med. 74 (2015), 189-498.

[25]A. Chambolle and T. Pock, An introduction to continuous optimization for imaging, Acta Numerica, 25 (2016), 161-319.

[26]J. Mitchell, T. C. Chandrasekera, and L. F. Gladden, Numerical estimation of relaxation and diffusion distributions in two dimensions, Prog. Nucl. Magn. Reson. Spectrosc. 62 (2012), 34-50.

[27]Z. Wang and G. R. Arce, Variable density compressed image sampling, IEEE Trans. Image Process. 19 (2010), 264-270.

[28]M. Lustig, D. Donoho and J. M. Pauly, Sparse MRI: The application of Compressed Sensing for rapid MR imaging, Magn. Reson. Med. 58 (2007), 1182-1195.

Chapter 9

Bubble burst hydrodynamics captured with MR velocimetry

It was shown in Chapter 8 how the appropriate combination of a sampling pattern and processing technique can be used to reduce the experimental acquisition time; this opens opportunities in studying systems which are unstable over a long acquisition time or in investigating fast-changing phenomena. In this chapter, under sampling of *k*-space data, a recently proposed processing method [1] and a carefully designed experimental system are used to highly accelerate the acquisition of MR velocimetry data, which allows the first experimental observation of the spatially resolved hydrodynamics of a bubble burst event.

9.1. Introduction

The burst of a bubble is an important event in many areas of science and engineering. A major transport mechanism of microorganisms [2, 3], salt [4-6] and charged aerosols [7] from the ocean to the atmosphere is through bubbles bursting at the water-air interface. In sparged bioreactors, the largest rate of cell death has been experimentally observed to occur at the free surface of the reactor [8-10]. This has been attributed to the entrapment of cells in the liquid film formed between bubbles and the free air surface, and their subsequent high acceleration during the film rupture and bubble burst [11-14]. Bubbles bursting at an oil-water interface lead to the formation sub-micrometer emulsions of oil in water [15], which is being used to produce nanomaterials with tunable rheological properties [16, 17]. The release and reach of flavours and aroma from carbonated drinks, such as champagne, are closely related to the burst of bubbles [18, 19].

When a bubble reaches a liquid-air interface, after an initial potential bouncing [20-23], the bubble rests on the surface, separated by a liquid film from the free air. The liquid film subsequently drains, a process which is explained by thinning film theories [24-26]. When the film reaches a critical thickness, a hole is punched in the film. If the hole is of a size

above approximately twice the thickness of the sheet, it expands [27]. Most of the retracting film is subsequently concentrated in a toroidal rim of liquid which bounds the unbroken film [28]. The rim motion follows to a good approximation the path defined by the original shape of the bubble [29], moving at a large retracting velocity of the order of 10 m s^{-1} [30, 31]. The process is characterised by film drop shedding [32-35], although this may not always be present [36]. The cavity collapses under the influence of converging capillary waves. The wave ripples formed in the cavity surface have been observed experimentally [37, 38] and captured by numerical methods [39-41]. At the convergence point of the capillary waves, a high capillary pressure point is formed which leads to the formation of an upward and downward jet. The upward jet moves in the empty space left by the bubble cavity and may lead to jet drops being ejected [42-47]. The angle of the downward jet to a vertical axis depends on the position in which the bubble gets punctured, which for millimetre sized bubbles is largely random [48]. Once the upward jet achieves the maximum height, it falls back into the liquid pool and, after some oscillations, liquid motion dies out.

The main experimental observation technique used to validate the predictions of numerical methods about the shape of the cavity, evolution of the upward jet and the phenomenon of jet drop shedding during bubble burst has been fast photography [5, 34, 37, 38, 40, 43-45, 49], although X-ray imaging [50] has also been used. Fast photography has also been used to validate the predictions of numerical methods about the ejection speed of the upward jet, by tracking the change of the shape of the upward jet between subsequent frames. However, apart from Dopplermeter measurements on a small volume element from MacIntyre [38], observations of the spatially resolved velocity field in the vicinity of the bursting bubble have been limited to flow visualization using a dye [37, 38] or seeded particles [40]. Measurement of the spatially resolved velocity field is critical in applications such as bioreactor design because it allows the calculation of the Energy Dissipation Rate (EDR) distribution in the liquid after the bubble burst [51]. EDR values above a certain threshold value have been linked to cell death [40, 52, 53]. The EDR distribution, together with information about the local concentration of cells in the vicinity of a bursting bubble can be used to quantify the rate of cell death. Further, measurement of the spatially resolved velocity field is important in validating the predictions of numerical methods, particularly in cases when different numerical methods disagree.

In this chapter, MR velocimetry is used to provide the first experimental spatially resolved velocity field of a bubble burst event. The velocity field map is then used to investigate a discrepancy between the predictions of two different numerical methods: the boundary integral numerical method used by Boulton-Stone and Blake [54] and the numerical solution to the full Navier-Stokes equations by Duchemin *et al*. [39]. The simulations of Boulton-Stoke and Blake [54] predicted a downward jet created by the detachment of a single vortex from the bubble base, while the simulations of Duchemin *et al*. [39] predicted a downward jet created by the detachment of two counter-rotating vortices.

The chapter is structured as follows. Section 9.2 describes the experimental system and the signal acquisition scheme. Section 9.3 describes a recently proposed processing technique for MR velocimetry data; the processing technique is validated against pipe laminar flow data. The experimental results of the bubble burst study and the discussion are presented in Section 9.4.

9.2. Materials and methods

A PTFE cylindrical pipe of inner diameter D = 15.4 mm is mounted vertically and filled up to a certain level with water, in which 17 mM $DyCl_3.6H_2O$ is added ($T_1 \approx T_2$ = 55 ms). This concentration of dopant agent ensures fast relaxation of the NMR signal, magnetic susceptibility matching of the solution and air and negligible effect on the surface tension of water [55]. The contact angle of water with the PTFE surface is > 90°, resulting in an upward curving meniscus being formed between water and air, as illustrated in Fig. 9.1(a). Air is slowly injected into the pipe using a Harvard Apparatus Model ‘22’ syringe pump. Air passes through a nozzle of 1 mm inner diameter, pinching off to consistently form bubbles of radius 2.4 mm. The bubble rises through water and rests on the free surface for some time, before bursting. The upward curving meniscus forces the bubble to stay in the centre of the pipe, being the point of highest liquid level [56]. The consistent size of the bubbles and their position at the centre of the pipe allows reproducible bubble burst experiments, which is critical in accelerating the acquisition of the experimental data, as discussed in the later parts of this section. The subsequent bubble burst hydrodynamics is studied using MR velocimetry.

All experiments are conducted on an AV-400 Bruker magnet, operating at a resonant frequency of 400.25 MHz for ^{1}H observation with a r.f. coil of 25 mm diameter. The

maximum magnetic field gradient amplitude available in each spatial direction is 146 G cm^{-1}.

The 2D MR spiral imaging technique developed by Tayler *et al*. [55], shown in Fig. 9.1(b), is employed for acquiring velocity images. The spiral is designed according to the method of Hargreaves *et al*. [57] for an image comprising of 64 × 64 pixels and of resolution 265 μm × 265 μm. The 2D images are acquired over a slice thickness of 150 μm. 1024 complex points spaced at 2 μs are acquired along the spiral, corresponding to 25% sampling percentage and imaging time of 2.05 ms. The spiral density and coverage of *k*-space is optimised according the method outlined in Section 9.3.1; the optimal spiral density is close to being constant and the spiral stretches to ~ 40% of k_{max}, as illustrated in Fig, 9.1(b). The velocity-encoding gradients are designed with δ = 0.26 ms, Δ = 0.49 ms and g = 29 G cm^{-1}. The r.f. excitation pulse duration is 200 μs and is designed for a flip angle of 11°, to allow for a fast return to equilibrium of the magnetization vector. The timings of the r.f. pulse, velocity-encoding gradients and imaging gradients allow for a repeat of the pulse sequence every 4.0 ms.

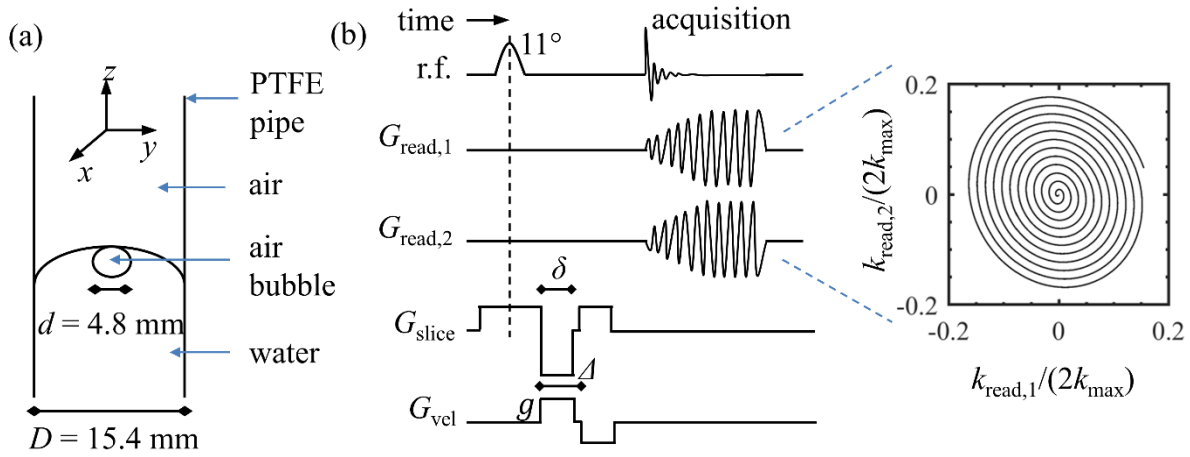

Fig. 9.1. (a) Schematic of experimental setup. (b) Pulse sequence used for MR velocimetry acquisitions and the corresponding *k*-space traversal.

Three component velocity maps on a 2D image are obtained for a transverse slice (perpendicular to the axis of the pipe) and a longitudinal slice (parallel to the axis of the pipe). Both slices cut through approximately the centre of the bubble. For a given slice direction (transverse or longitudinal) and a given direction of the velocity to be measured (x, y or z), the one component velocity map (v_x, v_y or v_z) is acquired by applying repeatedly every 4 ms the pulse sequence in Fig. 9.1(b) with the velocity-encoding gradient in the respective direction (x, y or z) and with alternating polarity between consecutive pulse

sequences (from $\pm g$ to $\mp g$). By taking the difference between the phase of the MRI images reconstructed from the acquired *k*-space data of consecutive pulse sequences, and with reference to a zero flow experiment, the one component velocity map at 4 ms temporal resolution is reconstructed, as discussed in Section 2.18. The maps of the other components of velocity are obtained in a similar way at temporal resolutions of 4 ms but with the pulse sequences applied to different bursting bubbles. The size, position and the hydrodynamics of different bursting bubbles at the free surface is consistent, as will be shown in Section 9.4. As a result, the one component velocity maps in different directions, v_x, v_y or v_z, are superimposed, by aligning the point in time at which the bubble burst starts, to effectively give three component velocity maps at a temporal resolution of 4 ms.

9.3. Reconstruction technique

The acquired *k*-space data using the spiral imaging technique is not sampled on a regular Cartesian grid. As a result, a Non-Uniform Fast Fourier Transform (NUFFT) operator [58], $\boldsymbol{\mathcal{F}}$, constructed with the aid of the spiral path measurement technique of Duyn *et al*. [59], outlined in Section 2.17, and Voronoi compensation diagrams [60] is used to relate the *k*-space data, $\boldsymbol{Y}_\mathrm{v}$, to the complex image, $\boldsymbol{U}_\mathrm{v}$ and the noise, $\boldsymbol{E}_\mathrm{v}$, according to:

$$\boldsymbol{Y}_\mathrm{v} = \boldsymbol{\mathcal{F}}\,\boldsymbol{U}_\mathrm{v} + \boldsymbol{E}_\mathrm{v}\ . \tag{9.1}$$

For MR velocimetry data, the complex image, $\boldsymbol{U}_\mathrm{v}$, is composed of the magnitude, $\boldsymbol{P}_\mathrm{v}$ and the phase, $\boldsymbol{\Phi}_\mathrm{v}$ such that $\boldsymbol{U}_\mathrm{v} = \boldsymbol{P}_\mathrm{v}\cdot\exp(\mathrm{i}\boldsymbol{\Phi}_\mathrm{v})$, where the dot notation refers to element-wise multiplication. The magnitude of the image is encoded with the ^{1}H density map, while the phase of the image is encoded with the velocity of the fluid. In this chapter, the focus is on the reconstruction of $\boldsymbol{\Phi}_\mathrm{v}$.

A naïve reconstruction technique for $\boldsymbol{\Phi}_\mathrm{v}$ is to first obtain $\boldsymbol{U}_\mathrm{v}$ from Zero-Filling:

$$\boldsymbol{U}_\mathrm{v} = \boldsymbol{\mathcal{F}}^{-1}\,\boldsymbol{Y}_\mathrm{v}\ , \tag{9.2}$$

and then extract $\boldsymbol{\Phi}_\mathrm{v}$ from the phase of $\boldsymbol{U}_\mathrm{v}$. This is typically sub-optimal for under sampled *k*-space data because the reconstructed $\boldsymbol{\Phi}_\mathrm{v}$ suffers from under sampling artefacts which, for spiral imaging, consist of spiral ripples throughout the image.

More commonly, $\boldsymbol{\Phi}_\mathrm{v}$ is obtained by first reconstructing the real and the imaginary part of $\boldsymbol{U}_\mathrm{v}$ separately through common regularization techniques, such as Total Variation (TV)

regularization, discussed in Section 8.2.1, and subsequently combining them to give $\boldsymbol{\Phi}_{\mathrm{v}}$ [61-63]. However, it is not clear why the reconstructed phase should be accurate when the real and imaginary parts are reconstructed separately.

Reconstruction methods which regularise the phase of the image, based on some prior knowledge of the phase, have recently been developed [1, 64-68]. A complication arises from the fact that the velocity images are generally acquired from the phase difference of different MR images, as discussed in Section 9.2. While there might be some prior knowledge about the velocity from fluid mechanics, there is generally no prior knowledge about the phase of individual MR images, since they are corrupted by, for example, field inhomogeneity artefacts. As a result, the regularisation needs to be imposed on the phase difference of different MR images, rather than separately on the phase of each MR image. The reconstruction method of Benning *et al*. [1], outlined below and used in this chapter, was found to be the most suitable in addressing this complication.

As discussed in Section 9.2, to acquire a one component velocity image, four MR images are needed, corresponding to: flow and $\pm g$ polarity for the velocity-encoding gradient ($\boldsymbol{P}_{\mathrm{v},1}\cdot\exp(\mathrm{i}\,\boldsymbol{\Phi}_{\mathrm{v},1})$); flow and $\mp g$ polarity the velocity-encoding gradient ($\boldsymbol{P}_{\mathrm{v},2}\cdot\exp(\mathrm{i}\,\boldsymbol{\Phi}_{\mathrm{v},2})$); no flow and $\pm g$ polarity for the velocity-encoding gradient and ($\boldsymbol{P}_{\mathrm{v},3}\cdot\exp(\mathrm{i}\,\boldsymbol{\Phi}_{\mathrm{v},3})$); and no flow and $\mp g$ polarity the velocity-encoding gradient ($\boldsymbol{P}_{\mathrm{v},4}\cdot\exp(\mathrm{i}\,\boldsymbol{\Phi}_{\mathrm{v},4})$).

The individual magnitudes, $\boldsymbol{P}_{\mathrm{v},j}$, are obtained through TV regularization, using the PDHGM method, as outlined in Section 8.2.1:

$$\boldsymbol{P}_{\mathrm{v},j} = \arg\min_{\boldsymbol{P}_{\mathrm{v},j}} \left(\frac{1}{2} \left\| \boldsymbol{\mathcal{F}}\, \boldsymbol{P}_{\mathrm{v},j} - \boldsymbol{Y}_{\mathrm{v}.j} \right\|_2^2 + \alpha \left\| \left[\boldsymbol{R}\ \ \boldsymbol{P}_{\mathrm{v},j} \right]_{\mathbf{2}} \right\|_{2,1} \right), \tag{9.3}$$

where $\boldsymbol{R}$ represents the TV transform operator.

Once the magnitude images are computed, the individual phase images, $\boldsymbol{\Phi}_{\mathrm{v},j}$, and the phase image corresponding to the velocity, $\boldsymbol{\Phi}_{\mathrm{v}} = ½\ [(\boldsymbol{\Phi}_{\mathrm{v},1} - \boldsymbol{\Phi}_{\mathrm{v},2}) - (\boldsymbol{\Phi}_{\mathrm{v},3} - \boldsymbol{\Phi}_{\mathrm{v},4})]$, are reconstructed using the prior knowledge that the velocity field has smooth spatial features. This is numerically achieved using the linearised Bregman iteration reconstruction technique developed by Benning *et al*. [1]. The iteration steps are as follows:

$$\begin{cases} \boldsymbol{\Phi}_{\mathrm{v}}{}^{(k+1)} = \arg\min_{\boldsymbol{\Phi}_{\mathrm{v}}} \left\{ \tau \langle \boldsymbol{\Phi}_{\mathrm{v}} - \boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}, \nabla H\left(\boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}\right) \rangle + D_J^{\boldsymbol{p}^{(k)}}\left(\boldsymbol{\Phi}_{\mathrm{v}}, \boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}\right) \right\} \\ \boldsymbol{p}^{(k+1)} = \boldsymbol{p}^{(k)} - \tau \nabla H\left(\boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}\right) \end{cases}, \tag{9.4a}$$

where:

$$H\left(\boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}\right) = \frac{1}{2}\sum_j \left\| \boldsymbol{\mathcal{F}}\left(\boldsymbol{P}_{\mathrm{v},j} \cdot \exp\left(\mathrm{i}\,\boldsymbol{\Phi}_{\mathrm{v},j}{}^{(k)}\right) - \boldsymbol{Y}_{\mathrm{v}.j}\right)\right\|_2^2, \tag{9.4b}$$

$$J\left(\boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}\right) = \frac{1}{2}\left\|\boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}\right\|_2^2 + \alpha \left\| \left[\boldsymbol{Q}\ \ \boldsymbol{\Phi}_{\mathrm{v}}{}^{(k)}\right]_{\mathbf{2}}\right\|_{2,2}. \tag{9.4c}$$

The Bregman distance for a function J is defined as:

$$D_J^{\boldsymbol{p}}(\boldsymbol{x}_1, \boldsymbol{x}_2) = J(\boldsymbol{x}_1) - J(\boldsymbol{x}_2) - \langle \boldsymbol{p}, \boldsymbol{x}_1 - \boldsymbol{x}_2 \rangle\,, \tag{9.4d}$$

with $\boldsymbol{p} \in \partial J(\boldsymbol{x}_2)$, and $\partial J(\boldsymbol{x}_2)$ being the so-called sub-differential, which is a generalisation of the classical differential for convex functions. The second order derivative matrix, $\boldsymbol{Q}$, in Eq. (9.4c) enforces smoothness in the phase image corresponding to the velocity, $\boldsymbol{\Phi}_{\mathrm{v}}$. The algorithmic parameter, τ, controls the convergence of the numerical solution. A convergence criteria for τ is found in the work of Benning *et al*. [1]; in this work, $\tau = 1.9$ gave a good compromise between stability and convergence speed. The regularization parameter, α, is chosen to be equal to the regularization parameter for the same sampling percentage of *k*-space (25%) in Section 9.3.1, where the proposed processing technique is validated against experimental data.

The code used for processing the data was provided by Dr Martin Benning of the Department of Applied Mathematics and Theoretical Physics, University of Cambridge. Two changes are made to the code. The original code was designed for the case when the phase image corresponding to the velocity is obtained from the difference of only two phase images; the code is adapted to deal with the case when the phase image corresponding to the velocity is obtained from the difference of four phase images. The second change is the design of the regularization matrix, $\boldsymbol{Q}$; $\boldsymbol{Q}$ is designed similarly to the first order differences in the work of Benning *et al*. [69], but with the first order differences substituted by second order differences.

9.3.1. Validation of the reconstruction technique

It is difficult to validate the reconstruction technique outlined in Section 9.3 against the MR velocimetry data of the bursting bubble because the ground truth of the velocity map is not known. Instead, the reconstruction technique is validated against experimental data of fully developed Hagen-Poiseuille flow in a cylindrical pipe, for which the ground truth for the axial velocity, v_z, is well known:

$$v_z = \begin{cases} 2v_{\mathrm{z,mean}}\left(1 - \frac{4(x^2+y^2)}{D^2}\right) & \text{for } x^2 + y^2 \le D^2/4 \\ 0 & \text{for } x^2 + y^2 > D^2/4 \end{cases}. \tag{9.5}$$

$v_{z,mean}$ is the mean axial velocity, x and y are the Cartesian coordinates from the centre of the pipe, and D is the diameter of the pipe.

The experimental system used to validate the results of the proposed reconstruction technique, outlined below, is separate from the system described in Section 9.2. Experiments are performed on a flow loop comprising of a pipe of inner diameter D = 16 mm. A 27 mM $GdCl_3.6H_2O$ solution ($T_1 \approx T_2$ = 100 ms) is used as the flowing fluid. The flow is driven by a Watson Marlow 505s peristaltic pump and is stabilised by a dampener to give a Reynolds number of Re = 250. The same magnet system as described in Section 9.2 is used. The spiral imaging pulse sequence illustrated in Fig. 9.1(b) is used to acquire 2D axial velocity maps on a transverse slice of 1 mm thickness. The spiral is designed according to the method of Hargreaves *et al*. [57] for an image comprising of 64 × 64 pixels and of resolution 280 μm × 280 μm. The velocity-encoding gradients are designed with δ = 0.67 ms, Δ = 1.7 ms and g = 29 G cm^{-1}. The r.f. excitation pulse duration is 256 μs and is designed for a flip angle of 90°. The capability of the processing technique proposed in Section 9.3 to reconstruct the axial velocity map is investigated for 16 different sampling percentages from 100% to 6.3%, corresponding to imagibg times from 8.1 ms to 0.5 ms. For each sampling percentage, variable density spirals [70] of the form:

$$\rho(|k|) = \begin{cases} c_1\left(1 - \frac{|k|}{c_4}\right)^{c_2} + c_3 & \text{for } |k| \le c_4 \\ 0 & \text{for } |k| > c_4 \end{cases}, \tag{9.6}$$

are optimised with respect to the constants c_i, such that they most closely reproduce the zero phase variation that characterises the phase associated with water under zero flow conditions. In Eq. (9.6), ρ refers to the density of sampled points and $|k|$ refers to the magnitude of the inverse space variable.

For the case of fully developed Hagen-Poiseuille flow, it is possible to use a very specific regularization matrix, $\boldsymbol{Q}$, in the reconstruction technique outlined in Section 9.3, rather than a general regularization matrix that enforces smoothness in the velocity map. Such a regularization matrix is the third derivative matrix, which is used in this work. This choice of the regularization matrix is appropriate for the case of Hagen-Poiseuille flow because:

$$\partial^3 v_z / \partial x^3 = \partial^3 v_z / \partial y^3 = 0 \, . \tag{9.7}$$

As a result, the velocity map, and therefore the phase of the image corresponding to the velocity map, is sparse in the third derivative domain. It is noted that such a choice of

regularization matrix is not appropriate for the case of the bubble burst, because the velocity map is more complicated; the best that can be done in that case is to have a regularization matrix that enforces smoothness.

The quality of the velocity images reconstructed from the under sampled *k*-space data using the proposed processing method is quantified using the PSNR metric, defined in Section 8.2.4, with the ground truth being the velocity profile in Eq. (9.5). The quality of the velocity images depends on the regularization parameter, α; a range of α is scanned and the reported results refer to the value of α which gave the largest PSNR value.

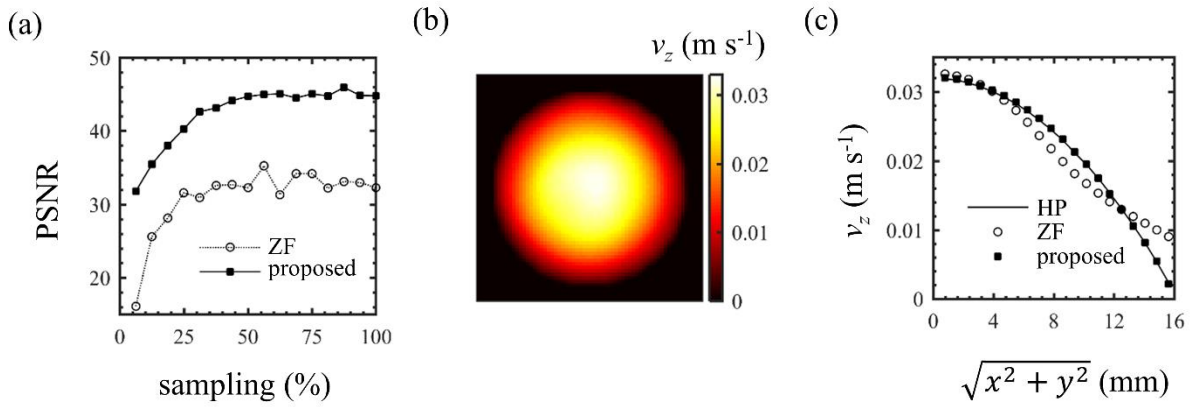

Fig. 9.2. (a) The PSNR metric of the reconstructed velocity maps of the fully developed Hagen-Poiseuille flow using Zero-Filling (ZF) and the proposed reconstruction technique at different sampling percentages. The lines are included to guide the eye. (b) Example of the reconstructed velocity map using the proposed processing technique at 6.3% sampling. The image is composed of 64 × 64 pixels and is of resolution 280 μm × 280 μm. (c) Comparison of the radially average velocity profile from the reconstructed velocity maps using Zero-Filling (ZF) and the proposed processing technique, at 6.3% sampling, with the theoretical Hagen-Poiseuille (HP) flow, as described by Eq. (9.5).

The results of the reconstructed velocity maps using the proposed processing methods, compared to Zero-Filling reconstructions, are summarised in Fig. 9.2(a). It is observed that at all sampling percentages investigated, reconstructions using the proposed processing technique are better than Zero-Filling reconstructions. It is noted that PSNR $\neq \infty$ even at a sampling percentage of 100% because the edges of *k*-space are never sampled during spiral imaging, and because of artefacts introduced by the NUFFT operator. An example of the reconstructed velocity map using the proposed reconstruction technique at a sampling percentage of 6.3% is shown in Figs. 9.2(b) and 9.2(c). The reconstructed velocity map agrees very well with the theoretical result of the Hagen-Poiseuille flow, thereby validating the results of the proposed processing method.

It is noted that such low sampling percentages (6.3%) give good reconstruction results only because, as discussed, it is possible to define a very specific regularization matrix based on the fluid dynamics of the system. For the case of the bubble burst hydrodynamics, where the regularization matrix is a general regularization matrix that enforces smoothness, such low sampling percentages are not possible. This is the reason for choosing a higher sampling percentage of 25%, as discussed in Section 9.2.

9.4. Results and discussion

Initially, the velocity maps of the bubble burst event are presented in Section 9.4.1. The velocity maps are subsequently used in Section 9.4.2 to shed light into a discrepancy between the predictions of different numerical methods about the vorticity map associated with the bubble burst event.

9.4.1. Bubble burst velocity maps

The evolution of the map of the velocity in the x-direction, v_x, for a transverse plane that cuts through the middle of the bubble is shown in Fig. 9.3. The evolution is given at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms, where t = 4 ms corresponds to the time when the bubble burst event starts. A strong inward flow towards the bursting bubble is observed, which is consistent with the boundary-layer surface flow theory of MacIntyre [37, 38]. The inward flow weakens with time and gets further away from the centre of the bubble burst, eventually reaching the pipe surface. The subsequent periodic reversal between inward and outward flow observed at t = 48 and 60 ms is caused by the periodic fall back into the liquid pool of the upward jet; the upward jet is not observed in the v_x map, but it is unambiguously distinguished in Figs. 9.5-9.7, which include the v_z map.

At a particular spatial location indicated in Fig. 9.3, v_x is recorded every 4 ms for a time interval 0-280 ms and the process is repeated for three other bubbles. The results are shown in Fig. 9.4. It is observed that the velocity evolutions for the different bubble burst events are very similar. The implication of this observation is that the bubble burst hydrodynamics is highly reproducible. As a result, if the map of the velocities in the x-direction, v_x, y-direction, v_y, and z-direction, v_z, are measured individually for three different bubbles, then, to a good approximation, these three components of the velocity

can be combined to construct a three-component velocity map. This is the method employed for constructing the three component velocity maps reported in this chapter.

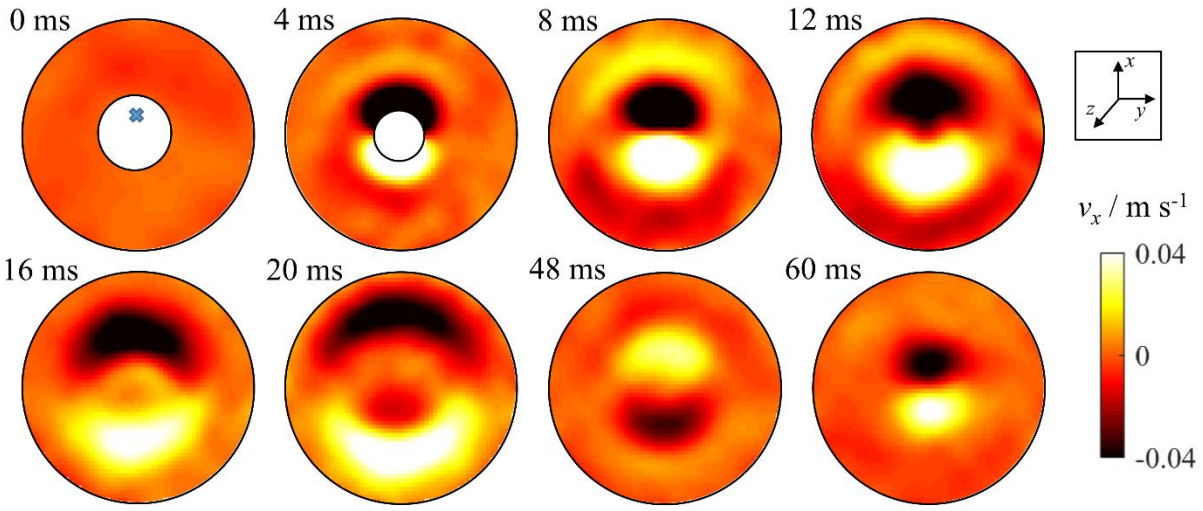

Fig. 9.3. The evolution of the map of velocity in the *x*-direction, v_x, for a transverse plane cutting through the middle of the bubble at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms; t = 4 ms corresponds to the start of the bubble burst event. The empty circles at t = 4 ms and 8 ms denote the bursting bubble. The spatial resolution in both dimensions is 265 μm and the pipe diameter is 15.4 mm. The pixel-wise uncertainty in the reported velocity is 2 mm s^{-1}, estimated from the pixel-wise standard deviation of v_x at t = 0 ms. The '×' mark on the velocity image at t = 0 ms denotes the spatial location at which the measurements reported in Fig. 9.4 are made.

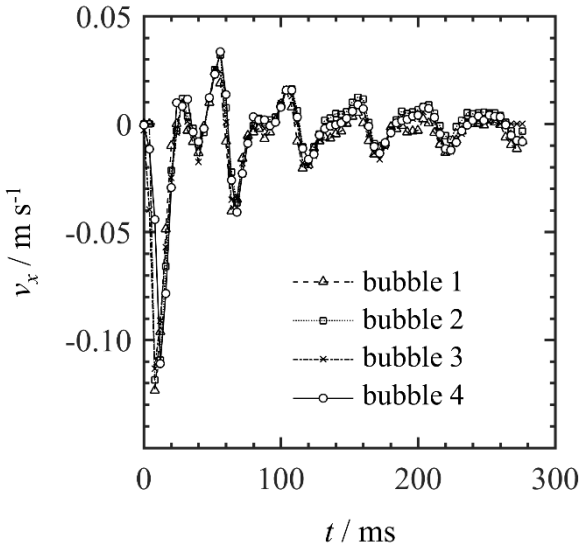

Fig. 9.4. The evolution of the velocity in the *x*-direction, v_x, at the spatial location marked by '×' in Fig. 9.3 for a time interval 0-280 ms; t = 4 ms corresponds to the start of the bubble burst event. The uncertainty in the reported velocity is 2 mm s^{-1}, estimated using the method described in Fig. 9.3.

The evolution of the three component velocity map at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms for a transverse plane that cuts through the middle of the bubble is shown in Fig. 9.5. Apart from the inward flow that was observed in Fig. 9.3, Fig. 9.5 shows a central region with an initial large velocity in the z-direction, which corresponds to an upward jet being ejected; the maximum ejection speed detected is 0.15 m s^{-1}. Once the jet achieves the maximum height at t = 12 ms (which is more clearly seen in Figs. 9.6 and 9.7), the jet falls back; this is most evident at t = 48 ms. Subsequently, the jet moves upwards again (t = 60 ms) and the process repeats periodically; this is also evident from the data in Fig. 9.4. The time between the start of the bubble burst and the reach of the maximum upward jet height, ≈ 8 ms, agrees with the time predicted by numerical works performed for bubbles of similar size [53]. The measured maximum ejection speed is a factor of 4-5 slower than the predictions of numerical methods [41] and photographic measurements [36]. This is attributed to the 4 ms time averaging and the 265 μm spatial averaging of the MR measurements. The numerical and photographic measurements do not report the time evolution of the ejection speed and it is, therefore, difficult to verify the hypothesis of the time and spatial averaging of the MR measurements.

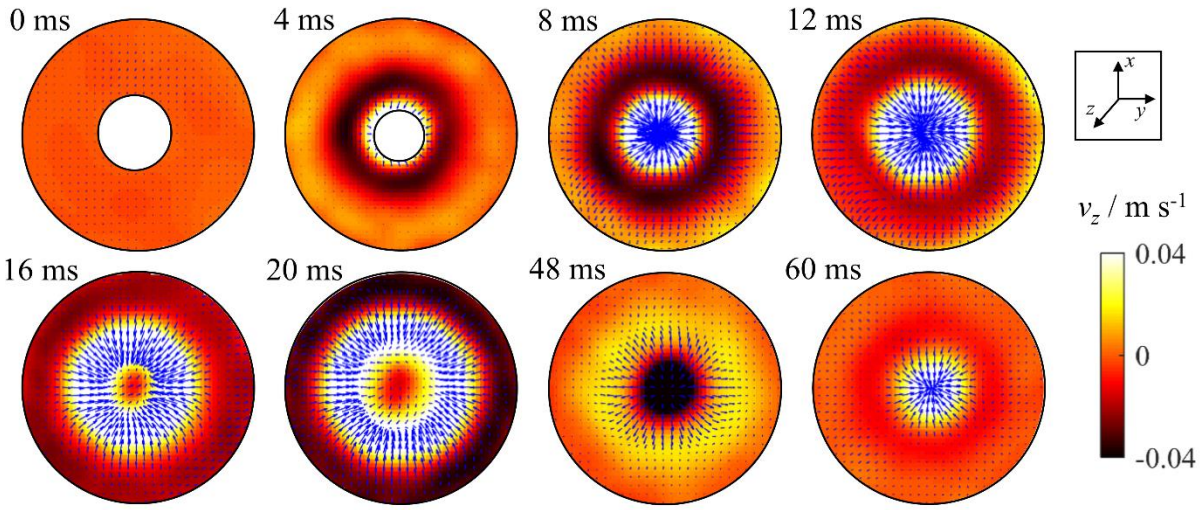

Fig. 9.5. The map of the out-of plane velocity, v_z, (colour coded) and the in-plane velocities, v_x and v_y (arrow coded) for a transverse plane cutting through the middle of the bubble at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms; t = 4 ms corresponds to the start of the bubble burst event. The empty circles at t = 4 ms and 8 ms denote the bursting bubble. The positive z-direction is the out-of-plane direction. The largest arrow corresponds to a velocity of 0.15 m s^{-1}. The largest reported ejection speed is v_z = 0.15 m s^{-1}. The spatial resolution in both dimensions is 265 μm and the pipe diameter is 15.4 mm. The uncertainty in the estimate of the v_x, v_y and v_z are 2, 2 and 1 mm s^{-1}, respectively, estimated from the pixel-wise standard deviation of v_x, v_y and v_z at t = 0 ms.

The evolution of the map of the velocity in the *z*-direction, v_z, for a longitudinal plane that cuts through the middle of the bubble is shown in Fig. 9.6. The evolution is given at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms. Apart from the features of bubble burst already observed in Fig. 9.5, the longitudinal velocity image clearly shows a downward jet which is created in conjunction with the upward jet (t = 12 ms). This agrees with all numerical works on bubble burst. The downward jet continues to exist as the upward jet starts to fall (t = 16 ms) and eventually joins the falling upward jet (t = 20 ms). The longitudinal velocity images also give information about the change in the cavity shape during the bubble collapse, although highly spatially averaged as compared to what fast photography can achieve, which can be in the order of a few microns.

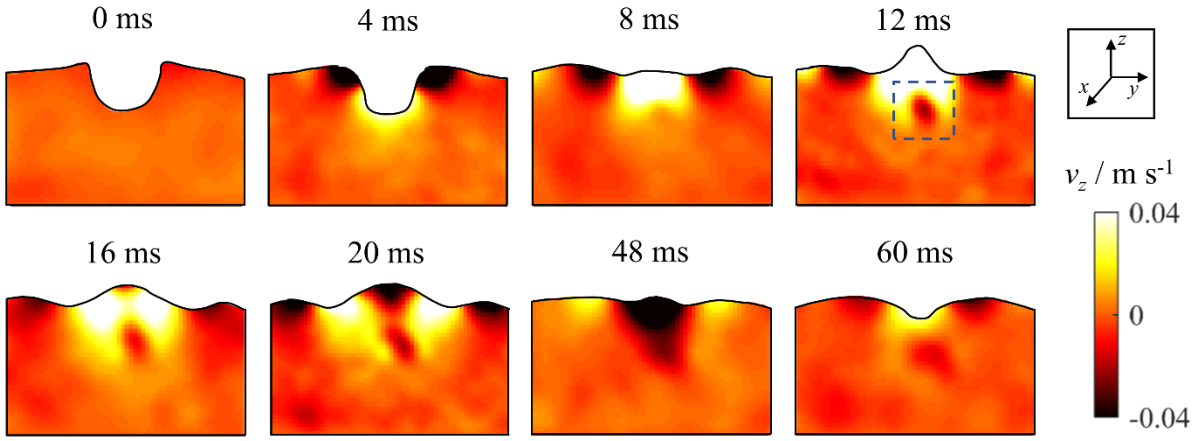

Fig. 9.6. The evolution of the map of velocity in the *z*-direction, v_z, for a longitudinal plane cutting through the middle of the bubble at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms; t = 4 ms corresponds to the start of the bubble burst event. The spatial resolution in both dimensions is 265 μm and the pipe diameter is 15.4 mm. The pixel-wise uncertainty in the stated velocity is 3 mm s^{-1}, estimated from the pixel-wise standard deviation of v_z at t = 0 ms. The highlighted feature at t = 12 ms corresponds to the downward jet.

As the upward jet achieves the maximum height (t = 12 ms), the formation of a jet drop was observed only occasionally. This is to be expected because the Bond number, $Bo = \rho_l g a^2 / \gamma$, and the Laplace number, $La = \rho_l \gamma a / \mu_l^2$, of the bubble are approximately 1.0 and 1.7 × 10^5 which correspond to an intermediate regime of jet drop formation / no jet drop formation [44]. In the definition of the Bond and Laplace number, ρ_l, μ_l, and γ are the density, viscosity and surface tension of the liquid, respectively, a is the radius of the bubble and g stands for the acceleration due to gravity. Although bubble burst events that are associated with jet drop formation are an area of ongoing scientific interest [44-47], the experimental work presented in this chapter corresponds only to bubble burst events

that are not associated with jet drop formation. The reason for this choice is that the scope of the chapter is to study the dynamics of the liquid during the bubble burst; this dynamics can be complicated and difficult to interpret if jet drops are also present.

Applying the same method as for the transverse velocity images, the three-component velocity map on a longitudinal plane is reconstructed and its evolution at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms is shown in Fig. 9.7. The main features of the three-component velocity maps were already observed in Figs. 9.3, 9.5 and 9.6. An important observation is that the out-of-plane velocity, v_x, is small compared to the other components of velocity. This is to be expected because of the symmetry of the flow. Therefore, the observed non-zero values of v_x arise from the imperfect experimental symmetry.

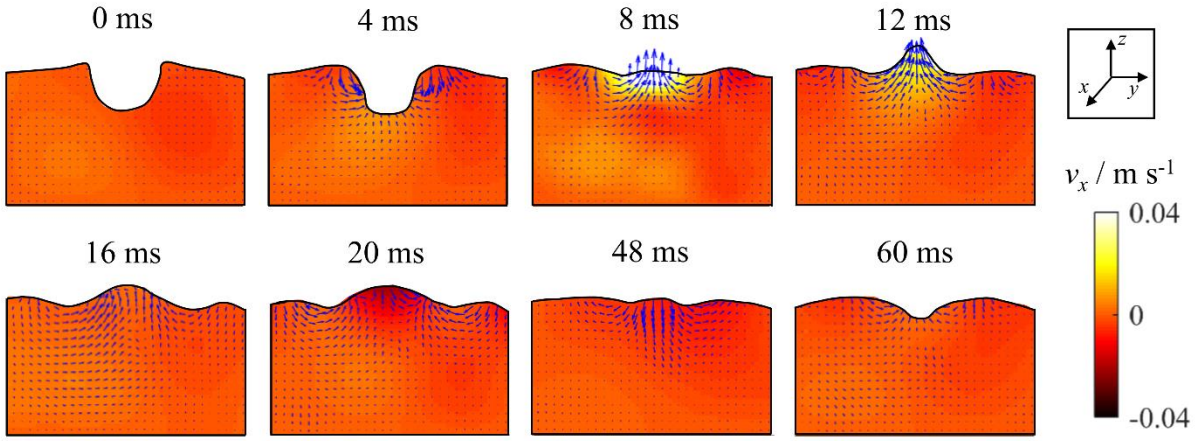

Fig. 9.7. The map of the out-of-plane velocity, v_x, (colour coded) and the in-plane velocities, v_y and v_z (arrow coded) for a longitudinal plane cutting through the middle of the bubble at times t = 0, 4, 8, 12, 16, 20, 48 and 60 ms; t = 4 ms corresponds to the start of the bubble burst event. The positive x-direction is the out-of- plane direction. The largest arrow corresponds to a velocity of 0.13 m s^{-1}. The largest reported ejection speed is v_z = 0.13 m s^{-1}. The spatial resolution in both dimensions is 265 μm and the pipe diameter is 15.4 mm. The uncertainty in the estimate of the v_x, v_y and v_z are 2, 2 and 3 mm s^{-1}, respectively, estimated from the pixel-wise standard deviation of v_x, v_y and v_z at t = 0 ms.

9.4.2. Bubble burst vorticity maps

The in-plane components of the velocity in the longitudinal plane, v_y and v_z, are now used to investigate a discrepancy between the predictions of the numerical works of Boulton-Stone and Blake [54] and Duchemin *et al.* [39], about the vortical structure generated by the bubble burst event for bubbles of similar size. Fig. 9.8 shows the evolution of the vorticity in the out-of-plane direction, ω_x, at times t = 4, 12, 20 and 60 ms, as calculated from the MR velocimetry data. The main observation is that at the point in

time in which the downward jet is observed to arise (t = 12 ms in Fig. 9.6), two counter-rotating vortices are observed at the location of the downward jet. This agrees with the work of Duchemin *et al*. [39] but disagrees with the work of Boulton-Stone and Blake [54]; the numerical work of Boulton-Stone and Blake [54] predicts a downward jet associated with the detachment of a single vortex from the base of the bubble. The main difference between the implementation of the two numerical works comes from the assumptions imposed: Boulton-Stoke and Blake [54] used an integral boundary method with the assumption of negligible viscosity effects; Duchemin *et al*. [39] relaxed the viscosity assumption and solved the full Navier-Stokes equations using finite differences.

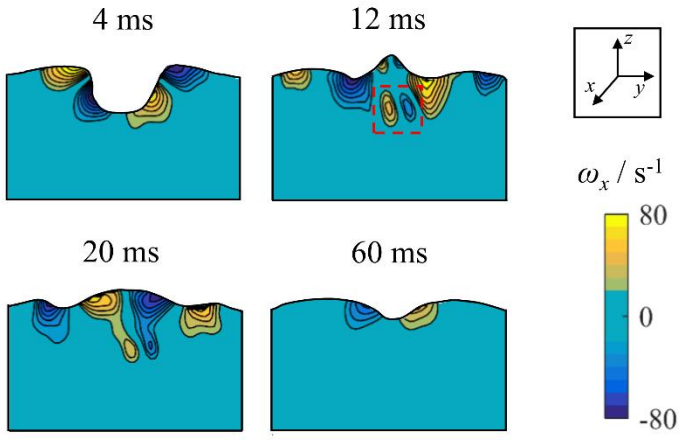

Fig. 9.8. Contour map of the vorticity in the out-of-plane direction, ω_x, for a longitudinal plane cutting through the middle of the bubble at times t = 4, 12, 20 and 60 ms; t = 4 ms corresponds to the start of the bubble burst event. The highlighted feature at t = 12 ms shows two counter-rotating vortices associated with the downward jet which is observed for the corresponding time frame in Fig. 9.6.

The results suggest that viscous effects make a significant contribution to the stresses in the downward jet. This is important for Computational Fluid Dynamics (CFD) simulations of the bubble burst event, which are critical in the design of bioreactors.

9.5. Conclusions

The first experimental measurement of spatially resolved velocity field generated in the liquid when a bubble bursts at a liquid/air interface was reported. Three-component velocity maps on 2D images obtained at 4 ms temporal resolution have revealed features which were up to now only predicted by numerical methods. The velocity maps were then used to establish that the downward jet ejected during the bubble burst event is created by

the detachment of two counter-rotating vortices; this agrees with solutions of the full Navier-Stokes equations using finite differences but disagrees with the numerical solutions of velocity maps using integral boundary methods which assume no viscosity effects.

References

[1] M. Benning, M. M. Betcke, M. J. Erhardt and C.-B. Schöenlib, Gradient descent in a generalised Bregman distance framework, (2016) arXiv:1612.02506v2.

[2] D. C. Blanchard and L. Syzdek, Mechanism for the water-to-air transfer and concentration of bacteria, Science 170 (1970), 626-628.

[3] E. R. Baylor, V. Peters and M. B. Baylor, Water-to-air transfer of virus, Science 197 (1977), 763-764.

[4] S. G. Boyce, Source of atmospheric salts, Science 113 (1951), 620-621.

[5] A. H. Woodcock, C. F. Kientzler, A. B. Arons and D. C. Blanchard, Giant condensation nuclei from bursting bubbles, Nature 172 (1953), 1144-1145.

[6] B. J. Mason, Bursting of air bubbles at the surface of sea water, Nature 174 (1954), 470-471.

[7] D. C. Blanchard, The electrification of the atmosphere by particles from bubbles in the sea, Prog. Oceanogr. 1 (1963), 73-202.

[8] A. Handa-Corrigan, A. N. Emery and R. E. Spier, Effect of gas-liquid interfaces on the growth of suspended mammalian cells: mechanisms of cell damage by bubbles, Enzyme Microb. Technol. 11 (1989), 230-235.

[9] K. T. Kunas and E. T. Papoutsakis, Damage mechanisms of suspended animal cells in agitated bioreactors with and without bubble entrainment, Biotechnol. Bioeng. 36 (1990), 476-483.

[10] Y. Christi, Animal-cell damage in sparged bioreactors, Trends Biotechnol. 18 (2000), 420-432.

[11] J. J. Chalmers and F. Bavarian, Microscopic visualization of insect cell-bubble interactions. II: The bubble film and bubble rupture, Biotechnol. Prog. 7 (1991), 151-158.

[12] R. S. Cherry and C. T. Hulle, Cell death in the thin films of bursting bubbles, Biotechnol. Prog. 8 (1992), 11-18.

[13]J. Wu, Mechanisms of animal cell damage associated with gas bubbles and cell protection by medium additives, J. Biotechnol. 43 (1995), 81-94.

[14]P. L. L. Walls, J. C. Bird and L. Bouroiba, Moving with bubbles: A review of the interaction between the bubbles and the microorganisms that surround them, Integr. Comp. Biol. 54 (2014), 1014-1025.

[15]J. Feng, M. Roché, D. Vigolo, L. N. Arnaudov, S. D. Stoyanov, T. D. Gurkov, G. G. Tsutsumanova and H. A. Stone, Nanoemulsions obtained via bubble-bursting at a compound interface, Nat. Phys. 10 (2014), 606-612.

[16]A. Gupta, H. B. Eral, T. A. Hatton and P. C. Doyle, Nanoemulsions: formation, properties and application, Soft Matter 12 (2016), 2826-2841.

[17]J. Feng, J. K. Nunes, S. Shin, J. Yan, Y. L. Kong, R. K. Prud'homme, L. N. Arnaudov, S. D. Stoyanov and H. A. Stone, A scalable platform for functional nanomaterials via bubble-bursting, Adv. Mater. 28 (2016), 4047-4052.

[18]G. Liger-Belair, G. Polidori and P. Jeandet, Recent advances in the science of champagne bubbles, Chem. Soc. Rev. 37 (2008), 2490-2511.

[19]G. Liger-Belair, C. Cilindre, R. D. Gougeon, M. Lucio and I. Gebefügi, Unravelling different chemical fingerprints between a champagne wine and its aerosols, PNAS 106 (2009), 16545-16549.

[20]T. Sanada, M. Watanabe and T. Fukano, Effects of viscosity on coalescence of a bubble upon impact with a free surface, Chem. Eng. Sci. 60 (2005), 5372-5384.

[21]J. Zawala, S. Dorbolo, D. Terwagne, N. Vandewalle and K. Malysa, Bouncing bubble on a liquid/gas interface resting or vibrating, Soft Matter 7 (2011), 6719-6726.

[22]J. Zawala, S. Dorbolo, N. Vandewalle and K. Malysa, Bubble bouncing at a clean water surface, Phys. Chem. Chem. Phys. 15 (2013), 17324-17332.

[23]S. Zhang, S. P. Wang and A. M. Zhang, Experimental study on the interaction between bubble and free surface using a high-voltage spark generator, Phys. Fluids 28 (2016), 032109.

[24]G. Marucci, A theory of coalescence, Chem. Eng. Sci. 24 (1969), 975-985.

[25]A. K. Chesters and G. Hofman, Bubble coalescence in pure liquids, Appl. Sci. Res. 38 (1982), 353-361.

[26]H. Kočárkova, F. Rouyer and F. Pigeonneau, Film drainage of viscous liquid on top of bare bubble: Influence of the Bond number, Phys. Fluids 25 (2013), 022105.

[27]G. I. Taylor and D. H. Michael, On making holes in a sheet of fluid, J. Fluid Mech. 58 (1973), 625-639.

[28]Lord Rayleigh, Some applications of photography, Nature 44 (1891), 249-254.

[29]A. B. Pandit and J. F. Davidson, Hydrodynamics of the rupture of thin liquid films, J. Fluid Mech. 212 (1990), 11-24.

[30]W. E. Ranz, Some experiments on the dynamics of liquid films, J. Appl. Phys. 30 (1959), 1950-1955.

[31]F. E. C. Culick, Comments on a ruptured soap film, J. Appl. Phys. 31 (1960), 1128-1129.

[32]F. Knelman, N. Dombrowski and D. M. Newitt, Mechanism of the bursting of bubbles, Nature 173 (1954), 261.

[33]D. C. Blanchard and L. D. Syzdek, Film drop production as a function of bubble size, J. Geophys. Res. 93 (1988), 3649-3654.

[34]D. Spiel, On the birth of film drops from bubbles bursting on seawater surfaces, J. Geophys. Res. 103 (1998), 24907-24918.

[35]H. Lhuissier and E. Villermaux, Bursting bubbles, Phys. Fluids, 21 (2009), 091111.

[36]E. Ghabache, G. Liger-Belair, A. Antkowiak and T. Séon, Evaporating of droplets in a Champagne aerosol, Sci. Rep. 6 (2016), 25148.

[37]F. MacIntyre, Bubbles: A boundary-layer “microtome” for micron-thick samples of a liquid surface, J. Phys. Chem. 72 (1968), 589-592.

[38]F. MacIntyre, Flow patterns in breaking bubbles, J. Geophys. Res. 77 (1972), 5521-5528.

[39]L. Duchemin, S. Popinet, C. Josserand and S. Zaleski, Jet formation in bubbles bursting at a free surface, Phys. Fluids 14 (2002), 3000-3008.

[40]P. L. L. Walls, O. McRae, V. Natarajan, C. Johnson, C. Antoniou and J. C. Bird, Quantifying the potential for bursting bubbles to damage suspended cells, Sci. Rep. 7 (2017), 15102.

[41]L. Deike, E. Ghabache, G. Liger-Belair, A. K. Das, S. Zaleski, S. Popinet and T. Séon, Dynamics of jets produced by bursting bubbles, Phys. Rev. Fluids 3 (2018), 013603.

[42]D. C. Blanchard, The size and height to which jet drops are ejected from bursting bubbles in seawater, J. Geophys. Res. 94 (1989), 10999-11002.

[43]E. Ghabache, A. Antkowiak, C. Josserand and T. Séon, On the physics of fizziness: How bubble bursting controls droplets ejection, Phys. Fluids 26 (2014), 121701.

[44]P. L. L. Walls, L. Henaux and J. C. Bird, Jet drops from bursting bubbles: How gravity and viscosity couple to inhibit droplet production, Phys. Rev. E 92 (2015), 021002(R).

[45]E. Ghabache and T. Séon, Size of the top jet produced by bubble bursting, Phys. Rev. Fluids 1 (2016), 051901.

[46]A. M. Gañán-Calvo, Revision of bubble bursting: Universal scaling laws of top jet drop size and speed, Phys. Rev. Lett. 119 (2017), 204502.

[47]C. F. Brasz, C. T. Barlett, P. L. L. Walls, E. G. Flynn, Y. E. Yu and J. C. Bird, Minimum size for the top jet drop from a bursting bubble, Phys. Rev. Fluids 3 (2018), 074001.

[48]F. MacIntyre, Additional problems in bubble and jet drop research, Limnol. Oceanogr. 23 (1978), 571-573.

[49]C. F. Kientzler, A. B. Arons, D. C. Blanchard and A. H. Woodcock, Photographic investigation of the projection of droplets by bubbles bursting at a water surface, Tellus 6 (1954), 1-7.

[50]J. S. Lee, B. M. Weon, S. J. Park, J. H. Je, K. Fezzaa and W.-K. Lee, Size limits the formation of liquid jets during bubble bursting, Nat. Commun. 2 (2011), 367.

[51]B. R. Bird, W. E. Stewart and E. N. Lightfoot, Transport Phenomena, John Wiley & Sons, 1960.

[52]M. Mollet, N. Ma, Y. Zhao, R. Brodkey, R. Taticek and J. J. Chalmers, Bioprocess equipment: Characterization of energy dissipation rate and its potential to damage cells, Biotechnol. Prog. 20 (2004), 1436-1448.

[53]M. A. Garcia-Briones, R. S. Brodkey and J. J. Chalmers, Computer simulations of the rupture of a gas bubble at a gas-liquid interface and its implications in animal cell damage, Chem. Eng. Sci. 49 (1994), 2301-2320.

[54]J. M. Boulton-Stone and J. R. Blake, Gas bubbles bursting at a free surface, J. Fluid Mech. 254 (1993), 437-466.

[55]A. B. Tayler, D. J. Holland, A. J. Sederman and L. F. Gladden, Exploring the origins of turbulence in multiphase flow using compressed sensing MRI, Phys. Rev. Lett. 108 (2012), 264505.

[56]D. Vella and L. Mahadevan, The 'Cheerios effect', Am. J. Phys. 73 (2005), 817-825.

[57]B. A. Hargreaves, D. G. Nishimura and S. M. Conolly, Time-optimal multi-dimensional gradient waveform design for rapid imaging, Magn. Reson. Med. 51 (2004), 81-92.

[58]J. A. Fessler and B. P. Sutton, Nonuniform fast Fourier transforms using min-max interpolation, IEEE Trans. Sign. Process. 51 (2003), 560-574.

[59]J. H. Duyn, Y. Yang, J. A. Frank and J. W. van der Veen, Simple correction method for k-space trajectory deviations in MRI, J. Magn. Reson. 132 (1998), 150-153.

[60]V. Rasche, R. Proksa, R. Sinkus, P. Börnert and H. Eggers, Resampling of data between arbitrary grids using convolution interpolation, IEEE Trans. Med. Imaging 18 (1999), 385-392.

[61]D. J. Holland, D. M. Malioutov, A. Blake, A. J. Sederman and L. F. Gladden, Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing, J. Magn. Reson. 203 (2010), 236-246.

[62]J. Trzasko and A. Manduca, Highly undersampled Magnetic Resonance image reconstruction via the homotopic l_0-minimization, IEEE Trans. Med. Imaging 28 (2009), 106-121.

[63]T. Roberts, N. Kingsbury and D. J. Holland, Sparse recovery of complex phase-encoded velocity images using iterative thresholding, IEEE Trans. Conf. Image Process. (2013), 350-354.

[64]J. A. Fessler and D. C. Noll, Iterative image reconstruction in MRI with separate magnitude and phase regularization, IEEE Int. Symp. Biomed. Imaging (2014), 209-212.

[65]M. V. W. Zibetti and A. R. de Pierro, Separate magnitude and phase regularization in MRI with incomplete data: Preliminary results, IEEE Int. Symp. Biomed. Imaging (2010), 736-739.

[66]F. Zhao, D. C. Noll, J.-F. Nielsen and J. A. Fessler, Separate magnitude and phase regularization via Compressed Sensing, IEEE Trans. Med. Imaging 28 (2012), 106-121.

[67]M. V. W. Zibetti and A. R. de Pierro, Improving compressive sensing in MRI with separate magnitude and phase priors, Multidim. Syst. Sign. Process. 28 (2017), 1109-1131.

[68] F. Ong, M. Uecker, U. Tariq, A. Hsiao, M. T. Alley, S. S. Vasanawala and M. Lustig, Robust 4D flow denoising using divergence-free wavelet transform, Magn. Reson. Med. 73 (2015), 828-842.

[69]M. Benning, L. F. Gladden, D. J. Holland, C.-B. Schöenlib and T. Valkonen, Phase reconstruction from velocity-encoded MRI measurements – A survey of sparsity-promoting variational approaches, J. Magn. Reson. 238 (2014), 26-43.

[70]C.-M. Tsai and D. G. Nishimura, Reduced aliasing artifacts using variable-density *k*-space, Magn. Reson. Med. 43 (2000), 452-458.

Chapter 10

Experimental evidence of velocity profile inversion in developing laminar flow at the entrance of pipes

The research presented in this chapter is somewhat unrelated to the main focus of the thesis; there is no development or application of novel signal sampling or processing techniques. Indeed, the magnetic resonance sampling and processing techniques used are well-established techniques. However, this chapter follows the theme of Chapter 9; MR velocimetry is used to gain new insights into a long-standing fluid mechanics problem which is relevant to many engineering systems.

10.1. Introduction

The subject of entrance laminar flow in pipes for Newtonian fluids has been extensively studied over the years. The main focus of experimental research in this area to-date has been to establish the entrance length required for fully developed Hagen-Poiseuille flow and the excess pressure drop incurred due to momentum change at the entrance [1-8]. The experimental data have shown that the entrance length required for fully developed Hagen-Poiseuille flow, z_{ent}, is well approximated by $z_{ent} = 0.06\, D\, Re_p$, and that the excess pressure drop incurred at the entrance, Δp_{exc}, is well approximated by $\Delta p_{exc} = 0.62\, \rho\, v_{mean}^2$, where D is the diameter of the pipe, Re_p is the Reynolds number based on the pipe diameter, ρ is the density of the fluid and v_{mean} is the mean velocity. These experimental results are in good agreement with the results of analytical [9-13] and numerical methods [14-21].

Given the aforementioned focus, experimental measurements have generally been made at relatively large distances from the entrance to the pipe and a discrepancy between the predictions of analytical and numerical methods about the development of the velocity

profile very close to the entrance to the pipe has not been investigated experimentally. Early approaches considering the development of the velocity profile very close to the entrance of the pipe were based on analytical solutions of approximate Navier-Stokes equations obtained by performing linearization of the inertial terms [9, 12, 22, 23] or using Prandtl's boundary layer assumptions and solving these approximate equations using integral methods [11, 13, 24], series expansions [10, 25-27] or numerical finite-difference methods [14, 15, 28]. These approximations were relaxed with the advent of digital computation, which enhanced the capability to solve the full Navier-Stokes equations by numerical methods. Using this approach, various workers reported numerical finite-difference solutions of the development of the velocity profile for laminar flows in a cylindrical pipe [29-33]. A salient feature has been captured by the numerical finite-difference solutions of the full Navier-Stokes equations which was not evident in the analytical solutions of approximate Navier-Stokes equations: Starting from a uniform velocity profile, the axial velocity attains a maximum at a position other than the centre of the pipe for a short distance downstream of the flow, with a local minimum at the centre of the pipe. The position of the maximum has been reported to move from the wall towards the centre of the pipe, staying at the centre thereafter. An illustration of the difference between the velocity profile development predicted from analytical solutions of approximate Navier-Stokes equations and numerical finite-difference solutions of the full Navier-Stokes equations is given in Fig. 10.1.

The phenomenon predicted from the numerical finite-difference solutions of the full Navier-Stokes equations is not constrained to the case of a uniform velocity profile at the entrance to the pipe, as was shown in the work dos Santos and Figueiredo [20], who studied numerically the entrance flow in a pipe through a contraction. However, the inlet boundary condition has a large influence on the exact development of the velocity profile [29, 30]. The developing axial velocity profile is accompanied by a non-zero radial velocity component, but this is expected to be negligible at $Re_p > 10$ for a starting uniform velocity profile [29, 30, 32]. The physical explanation that has been given for the maximum in the axial velocity profile not being at the centre of the pipe [34] is that the rate of flow development spreads inwards from the wall towards the centre of the pipe, due to larger shear stresses at the wall. Therefore, the acceleration needed to compensate for the deceleration at the wall is first felt close to the wall, causing a maximum in the axial velocity profile which moves towards the centre of the pipe. Since the phenomenon is

constrained to a short distance after the entrance to the pipe, questions have been raised whether this is a truncation artefact of the numerical finite-difference methods [29, 35] and whether it is possible to obtain sufficiently accurate experimental data to validate this behaviour [17].

This chapter presents the first experimental study of the axial velocity profile development for laminar flow of a Newtonian fluid very close to the entrance to a cylindrical pipe, with the aim of investigating whether the behaviour of the axial velocity profile is closer to the predictions of analytical solutions of approximate Navier-Stokes equations or to the predictions of numerical finite-difference solutions of the full Navier-Stokes equations. The velocity profile at the entrance to a pipe also determines the local pressure distribution and heat and mass transfer coefficients. Therefore, the experimentally measured velocity profiles can be used to validate or refute Computational Fluid Dynamics (CFD) results which are commonly used in the design of engineering systems where entrance flow is important, such as capillary rheometers [36], microfluidic channels [37], filtration equipment [38], injection moulding [39] and impinging jet systems [40]. The close relationship between the velocity profile and pressure distribution during developing laminar flow has also been proposed to aid the diagnosis and treatment of blood flow related diseases [41, 42].

The chapter is structured as follows. Section 10.2 describes the experimental setup and the materials used. The results, discussion and comparison with literature data are presented in Section 10.3.

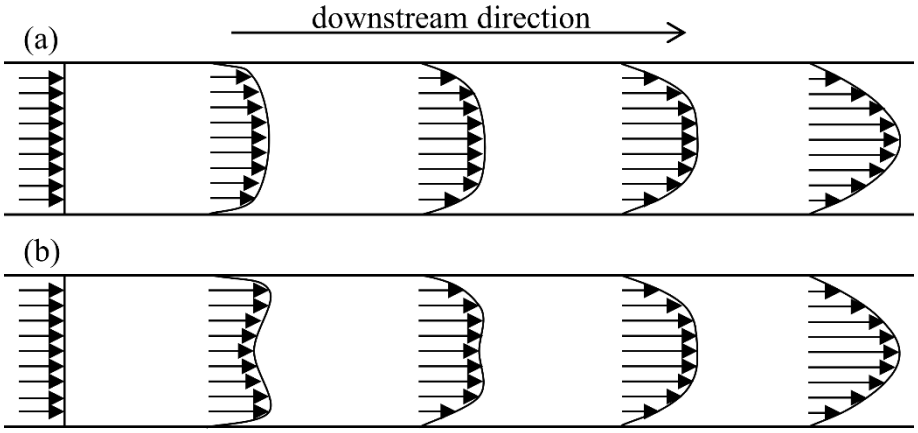

Fig 10.1. Schematic of the development of the axial velocity profile at the entrance to a cylindrical pipe for laminar flow, as predicted by (a) analytical solutions of approximate Navier-Stokes equations and by (b) numerical finite-difference solutions of the full Navier-Stokes equations.

10.2. Materials and methods

Experiments are performed on a flow loop comprising of a vertical cylindrical pipe of inner diameter D = 16 mm. The flow loop is designed such that a uniform velocity profile at the entrance to the pipe is produced by flowing the fluid through a regular square-channel monolith. Water with 0.36 mM $GdCl_3.6H_2O$ added $T_1 \approx T_2$ = 80 ms is used as the flowing fluid. Flow is driven upwards by a Watson-Marlow 505s peristaltic pump, with a dampener used to mitigate flow fluctuations. A 100 mm long cylindrical cordierite monolith with 0.4 mm^2 square channels running throughout its length is inserted in the pipe to obtain a uniform velocity profile.

All experiments are conducted on an AV-400 Bruker spectrometer, operating at a resonant frequency of 400.25 MHz for ^1H observation, with a r.f. coil of 25 mm diameter. The maximum magnetic field gradient amplitude available in each spatial direction is 146 G cm^{-1}.

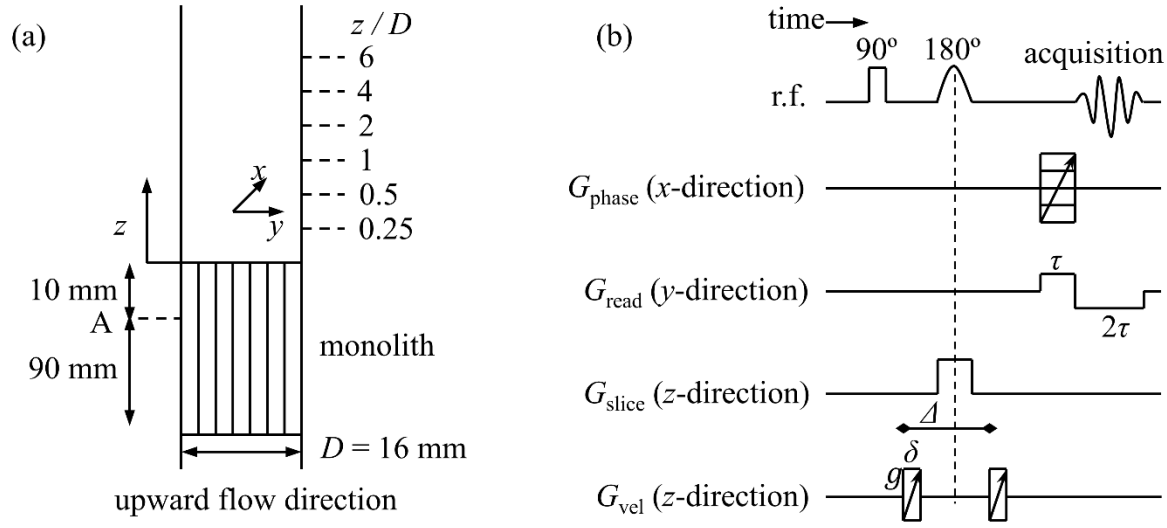

Fig. 10.2. (a) Schematic of the positions inside the monolith (position A) and at the exit of the monolith, or equivalently the entrance to the pipe, (z / D = 0.25, 0.5, 1, 2, 4, 6) where MR velocimetry acquisitions of the axial velocity in a transverse plane are obtained. (b) MR velocimetry pulse sequence used in this work, with the important acquisition parameters identified.

MR velocimetry acquisitions of the axial velocity in a transverse slice are performed at a distance of 10 mm before the exit of the monolith in order to establish whether the velocity profile at the exit of the monolith is uniform and at distances z from the exit of the monolith, z / D = 0.25, 0.5, 1, 2, 4, 6, to investigate the development of the velocity

profile beyond the monolith. The position $z / D = 0$ which identifies the exit of the monolith, also identifies the entrance to the pipe. These are illustrated in Fig. 10.2(a). Four Reynolds numbers, Re_p, are investigated in this study: 120 ± 10, 250 ± 10, 500 ± 20 and 1100 ± 50. The quoted uncertainty is the calculated standard deviation of the Reynolds numbers measured at the different z / D positions chosen and arises from the experimental variation in the flow rate with the axial position in the pipe. The corresponding mean velocities to the Reynolds numbers studied are 0.8, 1.6, 3.1 and 6.9 cm s^{-1}. The Reynolds number Re_p relates to the flow at the exit of the monolith and uses the diameter of the pipe as the characteristic length.

The MR velocimetry technique used is a combination of a conventional spin-echo based imaging technique with velocity-encoding magnetic field gradients, as illustrated in Fig. 10.2(b). For velocity imaging inside the monolith, the transverse slice thickness is 5 mm, images comprise of 512 × 512 pixels and are of resolution 0.039 mm × 0.039 mm, δ = 0.25 ms, Δ = 0.5 ms, τ = 0.51 ms; g is optimised to achieve a high signal-to-noise ratio, and the time between complex data points acquired in the read (y) direction is 1 μs. The uncertainty in the pixel-wise velocity measurement inside the monolith is 1.4 mm s^{-1}, as calculated from the standard deviation of the velocity image acquired under zero flow conditions. The spatial resolution of the velocity images inside the monolith is high in order to resolve velocities inside individual channels. For all velocity images beyond the exit of the monolith (i.e., within the pipe), the transverse slice thickness is 1 mm, images comprise of 64 × 64 pixels and are of resolution 0.313 mm × 0.313 mm, δ = 0.5 ms, Δ = 1 ms, τ = 1.28 ms; g is optimised to achieve a high signal-to-noise ratio, and the time between complex data points acquired in the read (y) direction is 20 μs. The uncertainty in the pixel-wise velocity measurement at the exit of the monolith is 0.3 mm s^{-1}, as calculated from the standard deviation of the velocity image acquired under zero flow conditions. The phase image used to calculate the velocity is obtained from the phase difference between two images with positive and negative velocity-encoding gradients ($+g$ and $-g$) and with correction from a zero velocity image, as discussed in Section 2.18.

10.3. Results and discussion

First, the results of velocity measurements inside the monolith are presented in Section 10.3.1, which show that the velocity profile immediately after the exit of the monolith is

well approximated by a uniform velocity profile. Then, the results of the velocity profile development at the entrance to the pipe are given in Section 10.3.2, which show clearly a maximum in the velocity profile whose position moves from the wall to the centre of the pipe with distance along the pipe. These data are then compared to available results from numerical finite-difference solutions of the full Navier-Stokes equations in the literature and the differences are critically discussed.

10.3.1. Velocity profile inside the monolith

In this section, results of the axial velocity measurements on a transverse slice positioned at a distance of 10 mm before the exit of the monolith are presented. Fig. 10.3(a) reports the magnitude image of the fluid in the monolith, which shows that the square channels are of similar size and they are distributed evenly throughout the cross-section of the pipe. Each channel comprises of ~ 16 × 16 pixels, which should give sufficient information about the velocity profile inside an individual channel. The velocity profile distribution for an expanded region corresponding to a single representative channel at Re_p = 500 is given in Fig. 10.3(b). The Reynolds number in each individual channel, Re_c, is estimated to be ~ 6, 10, 21 and 46 when the corresponding Reynolds number at the exit of the monolith, Re_p, is 120, 250, 500 and 1100, respectively. Re_c is calculated based on the hydraulic diameter of the channel.

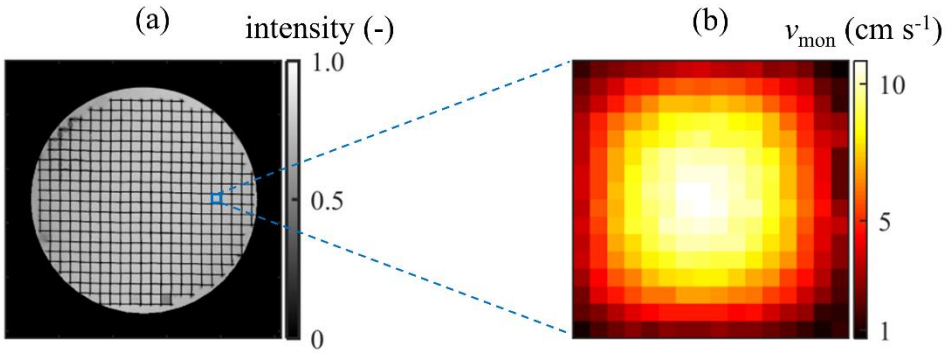

Fig. 10.3. (a) Cross-section magnitude image of the fluid in the monolith, composed of parallel square channels, at a distance of 10 mm before the exit of the monolith (position A in Fig. 10.2(a)). The fluid is MR-active while the monolith is MR-inactive and does not contribute to the MR signal. The image comprises of 512 × 512 pixels and is of resolution 0.039 mm × 0.039 mm. The thickness of the wall separating the channels is approximately 0.15 mm. (b) Axial velocity image of a single representative channel highlighted in (a) at Re_p = 500, Re_c =21. The image comprises of 16 × 16 pixels and is of resolution 0.039 mm × 0.039 mm. The maximum uncertainty in the pixel-wise reported velocity is 1.4 mm s^{-1}. The axial velocity in the monolith is denoted by v_{mon} to distinguish it from the axial velocity at the exit of the monolith, denoted by v.

At these low Reynolds numbers, the entrance length required for fully developed flow in an individual square channel is < 5 mm [43]. Therefore, at the imaging section positioned 10 mm before the exit of the monolith (position A in Fig. 10.2(a)), fully developed laminar flow is expected in each channel, even at the highest Reynolds number studied. At this distance before the exit of the monolith, the mean velocities in each channel are then radially averaged to give the velocity profile across the whole cross-section of the monolith.

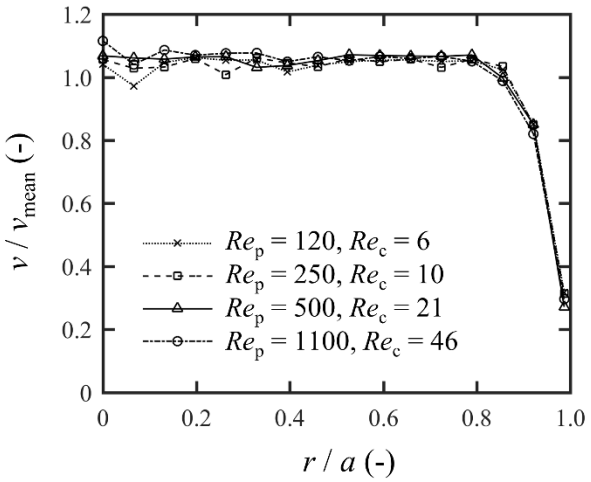

Fig. 10.4. Radially averaged axial velocity profile at a position of 10 mm before the exit of the monolith (position A in Fig. 10.2(a)) at the four Reynolds numbers studied. r refers to the radial distance, a refers to the radius of the pipe, and v_{mean} refers to the mean velocity. The lines are included to guide the eye. The graphs are constructed by radial averaging the mean velocities of the fluid within each channel. The uncertainty related to each data point is dominated by the circular asymmetry of the velocity profile, rather than the uncertainty related to pixel-wise velocity measurements. The maximum uncertainty of the reported v / v_{mean} values is 0.05. For reasons of clarity, error bars are not shown.

The radially averaged velocity profiles obtained at each of the Reynolds numbers investigated are shown in Fig. 10.4. It is observed that the velocity profile across the cross-section of the monolith is well approximated by a uniform velocity profile at all Reynolds numbers investigated. The slight deviation from the uniform velocity profile close to the wall of the pipe is caused by the non-ideality of the channels close to the wall of the pipe. The non-ideality of the channels extends for approximately one channel side length from the wall of the pipe. This region corresponds to r / a in the range 0.9-1.0, where r is the radial distance from the centre of the pipe and a is the radius of the pipe. It is in this region that the slight deviation from the uniform velocity profile is observed in Fig. 10.4. The effect that the deviation from the uniform velocity profile close to the wall

of the pipe has in the development of the velocity profile at the exit of the monolith is discussed in Section 10.3.2. Since the velocity profile at each individual channel is fully developed at 10 mm before the exit of the monolith, it is expected that the velocity profile across the cross-section of the monolith shown in Fig. 10.4 remains the same until the exit of the monolith. Therefore, the velocity profile immediately after the exit of the monolith and entrance to the pipe is well approximated by a uniform velocity profile.

10.3.2. Velocity profile development at the entrance to the pipe

The development of the axial velocity profile at the entrance to the pipe from the initial profile shown in Fig. 10.4 is now reported, with a particular focus on the location and magnitude of the maximum velocity in the profile.

The experimental results for the axial velocity profile at distances of z / D = 0.25, 0.5, 1, 2, 4 and 6 from the entrance to the pipe and Reynolds numbers, Re_p, of 120, 250, 500 and 1100 are presented in Fig. 10.5. The key observation from these results is that for a short distance from the entrance to the pipe (the extent of which depends on the Reynolds number), a maximum in the velocity profile which is not positioned at the centre of the pipe is clearly distinguished at all Reynolds numbers investigated. These experimental data therefore strongly support the predictions of the numerical finite-difference solutions of the full Navier-Stokes equations and do not agree with the predictions of analytical solutions of approximated Navier-Stokes equations.

Considering the distance from the entrance to the pipe, z / D, over which the phenomenon is visible, it is seen that the phenomenon extends to greater distances down the pipe as Re_p increases; the phenomenon is visible until $z / D \sim 0.5$ for $Re_p = 120$, $z / D \sim 2$ for $Re_p = 250$, $z / D \sim 4$ for $Re_p = 500$ and $z / D \sim 6$ for $Re_p = 1100$. This observation is in qualitative agreement with the observations of all the works reported which use numerical finite-difference methods to solve the full Navier-Stokes equations and can be explained in terms of the boundary-layer theory [25]; the boundary layer develops faster for low Reynolds number, which, as is well known, is why the entrance length for fully developed laminar flow is smaller for low Reynolds numbers.

Sufficient literature data for the quantitative comparison of the present experimental results are only available at $Re_p = 500$. Fig. 10.6 shows (a) the variation of the position of the maximum on the velocity profile, r_c / a, (b) the variation of the maximum velocity, v_{max} / v_{mean}, and (c) the variation of the velocity at the centre of the pipe, v_{centre} / v_{mean}, with

the distance from the entrance to the pipe, z / D, at $Re_p = 500$ for both the experimental data presented in this study and the results from the numerical finite-difference solutions of the full Navier-Stokes equations of Friedmann *et al.* [29], Wagner [30] and Dombrowski *et al.* [32].

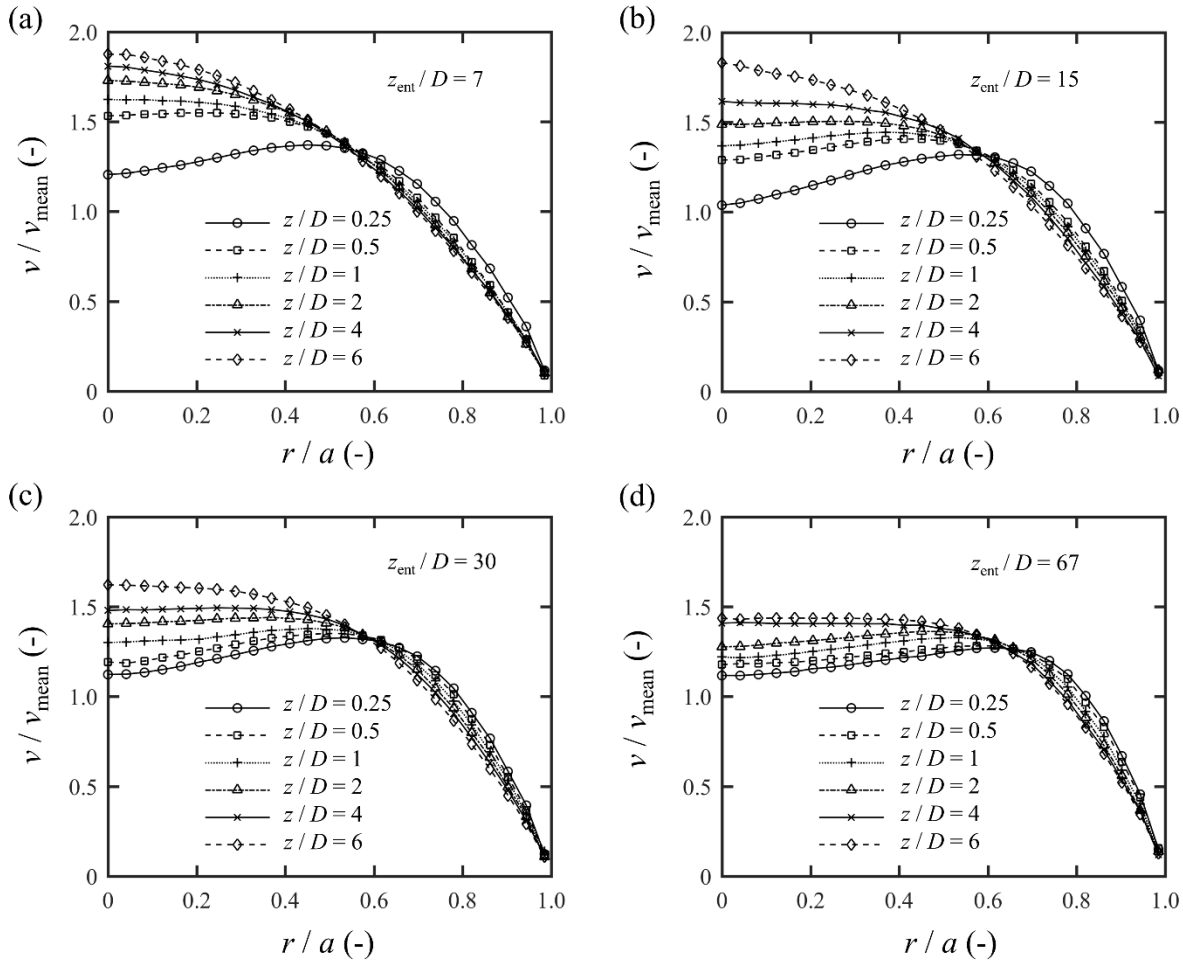

Fig. 10.5. Radially averaged axial velocity profiles at positions z / D = 0.25, 0.5, 1, 2, 4 and 6 from the entrance to the pipe and Reynolds numbers, Re_p, of (a) 120, (b) 250, (c) 500 and, (d) 1100. The lines are included to guide the eye. For each Reynolds number, the entrance length required for fully developed laminar flow, z_{ent} / D, is included in the figure. The uncertainty related to each data point is dominated by the circular asymmetry of the velocity profile, rather than the uncertainty related to pixel-wise velocity measurements. The maximum uncertainty of the reported v / v_{mean} values is 0.03. For reasons of clarity, error bars are not shown.

Considering the results from the numerical finite-difference solutions of the full Navier-Stokes equations, a major difference is observed between the results of Friedmann *et al.* [29] and Dombrowski *et al.* [32], as compared to the results of Wagner [30]. This is attributed to the use of different inlet boundary conditions; Friedmann *et al.* [29] and Dombrowski *et al.* [32] used a uniform velocity profile while

Wagner [30] used inlet boundary conditions corresponding to a piston moving at constant speed and solved the problem in a frame of reference in which the piston is stationary. The observation that the use of different inlet boundary conditions leads to significant quantitative differences in the measured r_c / a, v_{max} / v_{mean} and v_{centre} / v_{mean} is the main reason for not including in Fig. 10.6 the data from the numerical work of dos Santos and Figueiredo [20]. A minor difference is observed even between the results of Friedmann *et al*. [29] and Dombrowski *et al*. [32], showing that even with the use of the same inlet boundary conditions, different implementations of numerical methods can lead to different results, in terms of the predicted values of r_c / a, v_{max} / v_{mean} and v_{centre} / v_{mean}.

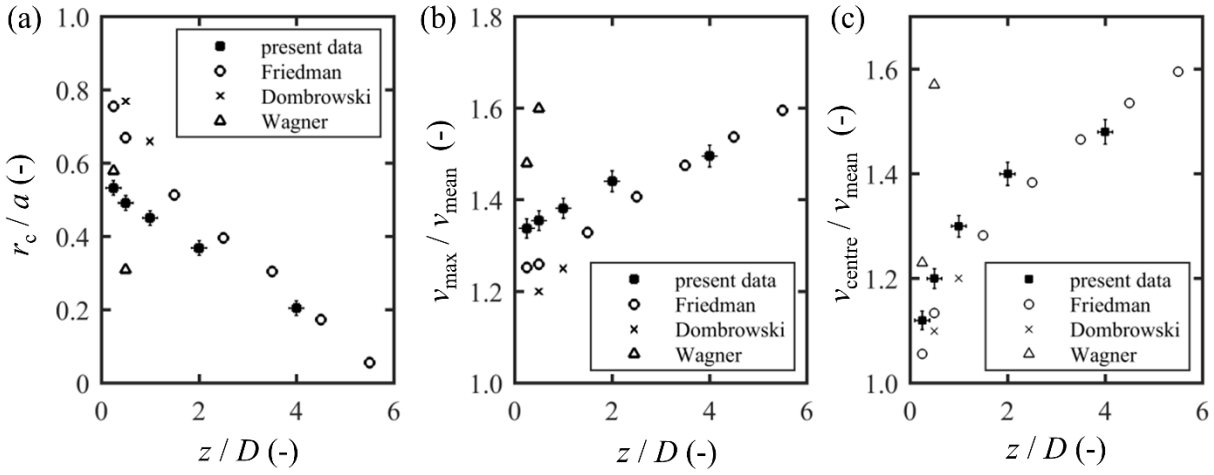

Fig. 10.6. Comparison of the present experimental results at $Re_p = 500$ with the results from finite-difference numerical methods solution of full Navier-Stokes equations from Friedmann *et al*. [29] and Dombrowski *et al*. [32] at $Re_p = 500$ and Wagner [30] at $Re_p = 400$. The graphs show the variation of (a) the critical radius r_c / a at which the maximum of the velocity occurs, (b) the maximum of the velocity v_{max} / v_{mean} and (c) the velocity at the centre of the pipe v_{centre} / v_{mean} with the distance from the entrance to the pipe z / D. The results are only shown for values of z / D at which the maximum in the velocity profile is not at the centre of the pipe. The uncertainty associated with the estimation of r_c / a is dominated by the spatial resolution of the MR image; the uncertainty associated with v_{max} / v_{mean} and v_{centre} / v_{mean} is dominated by the circular asymmetry of the velocity profile, rather than the uncertainty related to pixel-wise velocity measurements.

The inlet boundary condition in the present experimental work is closer to the inlet boundary conditions used by Friedmann *et al*. [29] and Dombrowski *et al*. [32], and it is observed from Fig. 10.6 that there is good agreement between the present experimental data and the results of those workers. However, the experimentally observed r_c / a is consistently lower, while v_{max} / v_{mean} and v_{centre} / v_{mean} are consistently higher than the values predicted by Friedmann *et al*. [29] and Dombrowski *et al*. [32]. These

discrepancies can be explained by the fact that the inlet boundary condition in the experiment is not a perfect uniform velocity profile (as it is assumed in the work of Friedmann *et al*. [29] and Dombrowski *et al*. [32]) but is as shown in Fig. 10.4. The consistently lower r_c / a observed experimentally is consistent with the flow already spreading inwards from the wall towards the centre of the pipe before the flow exits the monolith and enters the pipe; this is supported by the data shown in Fig. 10.4. The consistently higher v_{max} / v_{mean} and v_{centre} / v_{mean} follow from continuity.

10.4. Conclusions

The first experimental evidence of a maximum in the velocity profile positioned not at the centre of the pipe during laminar flow development of a Newtonian fluid at the entrance to a cylindrical pipe was presented. This phenomenon has been previously reported from numerical finite-difference solutions of the full Navier-Stokes equations, but has not been captured by previous analytical solutions of approximated Navier-Stokes equations or any other experimental study. The experimentally observed behaviour of the phenomenon as a function of the Reynolds number and distance downstream from the entrance to the pipe is in good agreement with numerical finite-difference solutions of the full Navier-Stokes equations. The minor differences can be explained in terms of the use of slightly different inlet boundary conditions.

References

[1] W. Rieman, The value of the Hagenbach factor in the determination of viscosity by the efflux method, J. Am. Chem. Soc. 50 (1928), 46-55.

[2] S. Goldstein, Modern Developments in Fluid Dynamics, Dover Publications, 1938.

[3] J. Nikuradse, Lows of flow in rough pipes, Tech. Rep. NASA Tech. Mem. 1292 (1950).

[4] S. T. McComas and E. R. G. Eckert, Laminar pressure drop associated with the continuum entrance region and for slip flow in a circular tube, J. Appl. Mech. 32 (1965), 765-770.

[5] B. Atkinson, Z. Kemblowski and J. M. Smith, Measurements of velocity profile in developing liquid flows, AIChE J. 13 (1967), 17-20.

[6] N. S. Berman and V. A. Santos, Laminar velocity profiles in developing flows using a Laser Doppler Technique, AIChE J. 15 (1969), 323-327.

[7] N. D. Sylvester and S. L. Rosen, Laminar flow in the entrance region of a cylindrical tube: Part I. Newtonian fluids, AIChE J. 16 (1970), 964-966.

[8] F. Durst, S. Ray, B. Ünsal and O. A. Bayoumi, The development lengths of laminar pipe and channel flows, J. Fluids Eng. 127 (2005), 1154-1160.

[9] H. L. Langhaar, Steady flow in the transition length of a straight tube, Trans. ASME: J. Appl. Mech. 9 (1942), 55-58.

[10] M. Collins and W. R. Schowalter, Laminar flow in the inlet section of a straight channel, Phys. Fluids 5 (1962), 1122-1124.

[11] W. D. Slattery and J. C. Slattery, Flow in the entrance region of a tube, J. Basic Eng. 85 (1963), 41-46.

[12] E. M. Sparrow, S. H. Lin and T. S. Lundgren, Flow development in the hydrodynamic entrance region of tubes and ducts, Phys. Fluids 7 (1964), 338-347.

[13] A. K. Mohanty and S. B. L. Asthana, Laminar flow in the entrance region of a smooth pipe, J. Fluid Mech. 90 (1978), 433-447.

[14]R. W. Hornbeck, Laminar flow in the entrance region of a pipe, Appl. Sci. Rec. 13 (1964), 224-232.

[15]E. B. Christiansen and H. E. Lemmon, Entrance region flow, AIChE J. 11 (1965), 995-999.

[16]J. S. Vrentas, J. L. Duda and K. G. Bargeron, Effect of axial diffusion of vorticity on flow development in circular conduits: Part I. Numerical solutions, AIChE J. 12 (1966), 837-844.

[17]B. Atkinson, M. P. Brocklebank, C. C. H. Card and J. M. Smith, Low Reynolds number developing flows, AIChE J. 15 (1969), 548-553.

[18]R. Y. Chen, Flow in the entrance region at low Reynolds numbers, J. Fluids Eng. 95 (1973), 153-185.

[19]R. C. Gupta, Laminar flow in the entrance of a tube, Appl. Sci. Res. 33 (1977), 1-10.

[20]R. G. dos Santos and J. R. Figueiredo, laminar elliptic flow in the entrance region of tubes, J. Braz. Soc. Mech. Sci. Eng. 29 (2007), 233-239.

[21]H. Kanda and K. Shimomukai, Numerical study of pressure distribution in entrance pipe flow, J. Complexity 25 (2009), 253-267.

[22]J. Boussinesq, Calcul de la moindre longeur que doit avoir un tube circulaire, èvasè á son entrèe, pour qu'un regime sensiblement uniforme s'y ètablisse, et de la dèpense de charge, qu'y entraine l'ètablissement de ce regime, Comptes Rendus 113 (1891), 49-51.

[23]C. L. Wiginton and R. L. Wendt, Flow in the entrance region of ducts, Phys. Fluids. 12 (1969), 465-466.

[24]L. Schiller, Die entwicklung der laminaren geschwindigkeitsverteilung, Z. Angew. Math. Mech. 2 (1922), 96-106.

[25]H. Schlichting, Boundary-Layer Theory, 7th Ed., McGraw-Hill, 1969.

[26]M. van Dyke, Entry flow in a channel, J. Fluid Mech. 44 (1970), 813-823.

[27]S. D. R. Wilson, Entry flow in a channel. Part 2, J. Fluid Mech. 46 (1971), 787-799.

[28]J. R. Bodoia and J. F. Osterle, Finite difference analysis of plan Poiseuille and Couette flow developments, Appl. Sci. Res. 10 (1961), 265-276.

[29]M. Friedmann, J. Gillis and N. Liron, Laminar flow in a pipe at low and moderate Reynolds numbers, Appl. Sci. Res. 19 (1968), 426-438.

[30]M. H. Wagner, Developing flow in circular conduits: transition from plug flow to tube flow, J. Fluid Mech. 72 (1975), 257-268.

[31]I. S. Goldberg and R. T. Folk, Solutions for steady and nonsteady entrance flow in a semi-infinite circular tube at very low Reynolds numbers, SIAM J. Appl. Math. 48 (1988), 770-791.

[32]N. Dombrowski, E. A. Foumeny, S. Ookawara and A. Riza, The influence of Reynolds number on the entry length and pressure drop for laminar pipe flow, Can. J. Chem. Eng. 71 (1993), 472-476.

[33]Z. Kountouriotis, M. Philippou and G. C. Georgiou, Development lengths in Newtonian Poiseuille flows with wall slip, Appl. Math. Comp. 291 (2016), 98-114.

[34]D. Fargie and B. W. Martin, Developing lamina flow in a pipe of circular cross-section, Proc. Roy. Soc. Lond. A 321 (1971), 461-476.

[35]J. Gillis, Viscous flow in the entry region of a straight channel, Phys. Fluids 12 (1969), II-280.

[36]M. R. Mackley and R. P. G. Rutgers, Capillary Rheometry, *in* Rheological Measurements (ed. A. A. Coyller and D. W. Clegg), Springer (1998), 167-189.

[37]T. Gervais and K. F. Jensen, Mass transport and surface reactions in microfluidics systems, Chem. Eng. Sci. 61 (2006), 1102-1121.

[38]V. M. Geraldes, V. A. Semião and M. N. Pinho, Nanofiltration mass transfer at the entrance region of a slit laminar flow, Ind. Eng. Chem. Res. 37 (1998), 4792-4800.

[39]Y. Kuo and M. R. Kamal, The fluid mechanics and heat transfer of injection mould filling of thermoplastic materials, AIChE J. 22 (1976), 661-669.

[40]M. J. Hancock and J. W. M. Bush, Fluid pipes, J. Fluid Mech. 466 (2002), 285-304.

[41]D. N. Ku, Blood flow arteries, Annu. Rev. Fluid Mech. 29 (1997), 399-434.

[42]S. S. Varghese, S. H. Frankel and P. F. Fischer, Direct numerical simulation of stentotic flows. Part 1. Steady flows, J. Fluid Mech. 582 (2007), 253-280.

[43]E. M. Sparrow, C. W. Hixon and G. Shavit, Experiments on laminar flow development in rectangular ducts, J. Basic Eng. 89 (1967), 116-123.

Chapter 11

Conclusions and future work

In this thesis, novel signal sampling and processing techniques were introduced in application to NMR and MRI problems. The proposed techniques are summarised in the following sections. It was shown that either through the separate or joint optimization of the processing techniques and the sampling patterns, the estimation of magnetic resonance parameters can be improved, or the experimental acquisition time can be reduced while keeping the accuracy of the estimated parameters unchanged. The presented methods were used to study experimental systems at higher temporal resolutions and in extracting features more accurately than was previously possible. The presented methods and the investigated experimental systems have suggested potential extensions and applications, which are also outlined in the following sections.

11.1. Optimising processing techniques

To obtain distributions of relaxation time constants and diffusion coefficients which are *a priori* known to be sparse, from 1D and 2D NMR experiments, it was proposed in Chapter 4 to use L_1 regularization as the processing technique. Using L_1 regularization, relaxation time constants and diffusion coefficients differing by as little as ~10% were resolved; the conventional processing technique of Tikhonov regularization can typically only resolve features differing by a factor of ~3. L_1 regularization reconstructions were shown to be stable at SNR < 20, while conventional reconstructions from the NNLS method are unstable at these noise levels. L_1 regularization was used to resolve inter and intra particle components of a hexane/dodecane mixture in porous silica beads. This capability is expected to be valuable in studying multi-component chemical mixtures in catalytic and rock systems.

When no prior information is available about the distributions of relaxation time constants and diffusion coefficients being extracted, it was proposed in Chapter 5 to process the 1D and 2D NMR data using Modified Total Generalized Variation (MTGV) regularization.

MTGV regularization retains both sparse and smooth features in the extracted distributions because it inherits good properties from both Tikhonov and L_1 regularizations. Over a range of simulated distributions and a range of SNR, MTGV regularization was shown to outperform Tikhonov and L_1 regularizations, apart from when the distributions were very discrete, in which case L_1 regularization performed slightly better.

It was shown in Chapter 8 that the recently proposed Nuclear Total Generalized Variation (NTGV) regularization is the best processing technique for obtaining 3D spatially resolved T_2 distributions from highly under sampled *k*-space data. NTGV regularization exploits simultaneously the sparsity of the individual T_2-weighted images in the Total Variation (TV) domain and the correlation of the pixel-wise T_2 decays. NTGV regularization was shown to outperform the conventional processing technique of TV regularization and the more recent Nuclear Norm (NN) regularization. NTGV regularization was used to experimentally obtain quantitative 3D spatially resolved T_2 maps with a 32-fold decrease in the acquisition time. The improved temporal resolution in the acquisition of spatially resolved parameter distributions opens opportunities in studying unstable systems over a long acquisition time and in investigating fast-changing phenomena.

To reconstruct MR velocity images from under sampled *k*-space data, it was proposed in Chapter 9 to perform regularization on the phase of the images, rather than on the magnitude of the images, which is the conventional processing technique. Using a recently proposed phase regularization method, it was possible to capture, for the first time, the hydrodynamics of a bubble burst event by acquiring three-component velocity maps on 2D images at a temporal resolution of 4 ms. The reconstructions showed a downward jet associated with the creation of two counter-rotating vortices. The observation was used to validate numerical method predictions. The work in Chapter 9 has opened opportunities to study fluid mechanics problems which remain a long-standing challenge, such as the hydrodynamics of bubble coalescence [1-4].

From the aforementioned results, it is concluded that the choice of the processing technique depends on the prior information available about the experimental system. In the following paragraphs, future research questions are suggested, and advice is given on the choice of the processing techniques.

Monitoring of reaction kinetics or other dynamic systems by a series of measurements in time of diffusion coefficients or relaxation time constant distributions [5-8] is an

alternative to spectroscopic monitoring when no spectroscopic resolution is available. One problem with processing these data sets is that the distributions at each time step are commonly reconstructed individually. The distributions reconstructed in this way may vary unrealistically between different time steps because of the automatic selection of the regularization parameter and the variation in the SNR. It is hypothesised that the reconstructed distributions may be improved by performing NN regularization in the time domain, similarly to method described in Section 8.2.2. The application of NN regularization is based on the knowledge that the distributions at different time steps are highly correlated. It would be interesting to compare NN regularization reconstructions with conventional reconstructions of experimental data.

In spatially resolved wettability study of rocks using NMR relaxation time constant distributions [9, 10], small changes in wettability can affect the extraction efficiency of hydrocarbons. The detection of these small changes using common conventional processing techniques for NMR relaxation data can be difficult because the small changes are typically superimposed on a very broad background distribution. A new rank-sparse regularization technique [11, 12] has been proposed to reconstruct a series of distributions which are composed of highly correlated backgrounds and sparse features appearing in some of the distributions. It is hypothesised that this regularization technique may be able to detect the small changes in wettability.

The regularization of the phase of the image for the reconstruction of MR velocity images from under sampled *k*-space data was done with a generic smoothing regularization matrix in Chapter 9. However, as discussed in Section 9.3.1, using a regularization matrix based on the knowledge of the fluid dynamics of the system may allow under sampling at unprecedented levels. Such a regularization matrix was easy to construct for the simple laminar flow case in Section 9.3.1. However, when the fluid dynamics is complex, the choice of the regularization matrix is challenging. Divergence-free regularization, which enforces continuity, has been previously proposed as a potential regularization matrix [13, 14]. It would be interesting to investigate whether a more elaborate regularization matrix can be constructed from the discretisation of the full Navier-Stokes equations.

In most MR velocimetry applications, the high under sampling in the *k*-space data is done to acquire velocity images at a high temporal resolution. Conventionally, the acquired velocity images are reconstructed individually. However, the high correlation between

subsequent velocity images [15, 16] can be used in the form of regularization to improve the reconstructions. Further, mass conservation between subsequent frames be cast as a regularization matrix and help in further improving the reconstructions.

MRI has been successfully used to characterize bubble size distributions [17, 18] and the liquid velocity field of bubbly flows [19, 20]. To study fast flowing systems, under sampling of the *k*-space data is needed. It is, therefore, critical to develop processing techniques for highly under sampled *k*-space data for the accurate reconstruction of the MRI magnitude images which would be subsequently used in calculating the bubble size distribution or studying the hydrodynamics of the system. An important reconstruction quality metric for the MRI magnitude images is the correct segmentation of the bubble-liquid interface. We are collaborating with the Cambridge Image Analysis group at the Department of Applied Mathematics and Theoretical Physics, University of Cambridge for the development of a joint segmentation-reconstruction regularization technique for this application [21].

11.2. Optimising sampling patterns

A novel method for optimising sampling patterns for 1D relaxation and diffusion NMR experiments, based on the Cramér-Rao Lower Bound (CRLB) theory, was outlined in Chapter 6. In Chapter 6, the findings of the CRLB theory were validated against NMR diffusion experimental data of an emulsion of toluene in water, with the aim of the most accurate estimation of the lognormal distribution parameters that characterise the emulsion droplet size distribution. In Chapter 7, the findings of the CRLB theory were validated against NMR diffusion experimental data of a gaseous mixture of methane/ethane adsorbed in a zeolite, with the aim of the most accurate estimation of the bi-exponential model parameters that characterise this system. Sampling patterns were optimised by minimizing an objective function that depends on the accuracy of estimation of the model parameters. In both examples, the predictions of the CRLB theory about the variation of the objective function with the sampling pattern were within 10% difference from the objective function calculated from the uncertainties of the model parameters obtained from the same experimental sampled points. It was shown that CRLB theory is applicable at SNR values of as low as ~10. The application of the presented method requires some prior knowledge of the parameters being estimated; a sensitivity analysis on the prior knowledge

was performed. The optimisation of the sampling patterns can be done with the objective of a more accurate estimation of the parameters of a model or with the objective of reducing the acquisition time while keeping the accuracy of the estimated parameters fixed. A case study of how the acquisition time can be reduced while keeping the accuracy of the estimated parameters fixed using the CRLB theory was given in Sections 6.4.4 and 7.4.2.

The CRLB-based method for optimising sampling patterns outlined in this thesis is currently being extended in application to 2D NMR correlation experiments. In collaboration with the Christodoulou Group, Department of Structural and Molecular Biology, University College London, the optimised sampling patterns for D-T_2 correlation experiments will be used to accelerate the experimental acquisition time, such that the evolution of protein aggregation can be observed.

The CRLB theory has limitations, which if overcome, could widen the applicability of the method in application to NMR and MRI problems. Future research questions are proposed in the following paragraphs.

The CRLB theory, as presented in Section 3.4, is only applicable when the parameters are estimated through an unbiased estimation technique. This means that the CRLB theory is not directly applicable to the problem of optimising sampling patterns for experiments whose data are processed using regularization. Since one of the most common techniques of reconstructing relaxation time constants and diffusion coefficient distributions from NMR data is through regularization, it would be valuable to extend the CRLB theory to the case of regularization. Fortunately, theoretical developments for biased CRLB estimations have been made [22-24]. It would be very interesting to see the biased CRLB theory used in optimising sampling patterns for improving the reconstruction of relaxation time constants and diffusion coefficient distributions from NMR data. An envisaged difficulty is the effect that the choice of the regularization parameter has on the optimisation of the sampling patterns.

The CRLB theory, as presented in Section 3.4, is only applicable when the signal is real-valued. As a result, the CRLB theory is not directly applicable to the problem of optimising sampling patterns for complex-valued signals, such as in under sampled MRI applications. Fortunately, theoretical developments for the application of the CRLB theory to complex-valued signal have been made [25, 26] and it would be interesting to see whether

the CRLB theory can bring a systematic approach in an area which is dominated by trial and error. Another approach to using CRLB theory for optimising sampling patterns in under sampled MRI applications, is to deal with the magnitude of *k*-space rather than the complex-valued *k*-space. We are currently developing this method for optimising sampling patterns for obtaining high resolution MRI images of rocks.

The optimised sampling pattern using the CRLB theory is fixed before the experimental acquisition is started and, therefore, assumes a reasonable prior knowledge of the system. In a fully automated experimental acquisition protocol, it would be ideal if a program could learn from the sampled points acquired and adjust the future sampled points accordingly. This is the idea behind the recent work of Song *et al*. [27] and it is expected to be an active research area in the future.

11.3. Joint optimisation of processing techniques and sampling patterns

The processing technique and the sampling pattern for a given NMR or MRI experiment can be optimised separately. However, the joint optimisation of the processing technique and the sampling pattern can lead to better reconstructions than the separate optimisations, as was observed in Chapter 8. It was shown in Chapter 8 that the full potential of NTGV regularization is realised when it is used in processing data sampled according to an incoherent sampling pattern rather than a coherent sampling pattern.

A systematic approach to the joint optimisation of the processing technique and the sampling pattern is, therefore, needed. This could be the focus of future investigations.

References

[1] S. Narayanan, L. H. J. Goossens and N. W. F. Kossen, Coalescence of two bubbles rising in line at low Reynolds numbers, Chem. Eng. Sci. 29 (1974), 2071-2082.

[2] T. Sanada, A. Sato, M. Shirota and M. Watanabe, Motion and coalescence of a pair of bubbles rising side by side, Chem. Eng. Sci. 64 (2009), 2659-2671.

[3] S. Orvalho, M. C. Ruzicka, G. Olivieri and A. Marzocchella, Bubble coalescence: Effect of bubble approach velocity and liquid viscosity, Chem. Eng. Sci. 134 (2015), 205-216.

[4] J. Feng, X. Li, Y. Bao, Z. Cai and Z. Gao, Coalescence and conjunction of two in-line bubbles at low Reynolds numbers, Chem. Eng. Sci. 141 (2016), 261-270.

[5] M. Nilsson, M. Khajeh, A. Botana, M. A. Bernstein and G. A. Morris, Diffusion NMR and trilinear analysis in the study of reaction kinetics, Chem. Comm. 1252 (2009), 1252-1254.

[6] M. Khajeh, A. Botana, M. A. Bernstein, M. Nilsson and G. A. Morris, Reaction kinetics studied using diffusion-ordered spectroscopy and multiway chemometrics, Anal. Chem. 82 (2010), 2102-2108.

[7] L. Buljubasich, B. Blümich and S. Stapf, Reaction monitoring of hydrogen peroxide decomposition by NMR relaxometry, Chem. Eng. Sci. 65 (2010), 1394-1399.

[8] B. F. Gomes, L. M. S. Nunes, C. M. S. Lobo, L. F. Cabeça and L. A. Colnago, *In situ* study of the magnetoelectrolysis phenomenon during copper electrodeposition using time domain NMR relaxometry, Anal. Chem. 86 (2014), 9391-9393.

[9] V. Bortolotti, P. Macini, E. N. Mesini, F. Srisuriyachai, P. Fantazzini and M. Gombia, Combined spatially resolved and non-resolved 1H-NMR relaxation analysis to assess and monitor wettability reversal in carbonate rocks, IPTC 13443 (2009).

[10] V. Bortolotti, P. Macini, E. N. Mesini, P. Fantazzini, M. Gombia and F. Srisuriyachai, Probing wettability reversal in carbonatic rocks by partially resolved and non-resolved 1H-NMR relaxation analysis, SPE Paper 133937 (2010).

[11] E. J. Candés, X. Li, Y. Ma and J. Wright, Robust principal component analysis? J. ACM, 58 (2011), 11.

[12] V. Chandrasearan, S. Sanghavi, P. A. Parrilo and A. S. Willsky, Rank-sparsity incoherence for matrix decomposition, SIAM J. Optim. 21 (2011), 572-596.

[13] F. Ong, M. Uecker, U. Tariq, A. Hsiao, M. T. Alley, S. S. Vasanawala and M. Lustig, Robust 4D flow denoising using divergence-free wavelet transform, Magn. Reson. Med. 73 (2015), 828-842.

[14] J. Mura, A. M. Pinto, J. Sotelo, I. Valverde, C. Tejos, M. E. Andia, P. Irarrázaval and S. Uribe, Enhancing the velocity data from 4D flow MR images by reducing its divergence, IEEE Trans. Med. Imaging 35 (2016), 2353-2364.

[15] J. Tsao, P. Boesiger and K. P. Pruessmann, *k-t* BLAST and *k-t* SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations, Magn. Reson. Med. 50 (2003), 1031-1042.

[16] H. Pedersen, S. Kozerke, S. Ringgaard, K. Nehrke and W. Y. Kim, *k-t* PCA: temporally constrained *k-t* BLAST reconstruction using principal component analysis, Magn. Reson. Med. 62 (2009), 706-716.

[17] D. J. Holland, A. Blake, A. B. Tayler, A. J. Sederman and L. F. Gladden, Bubble size measurement using Bayesian magnetic resonance, Chem. Eng. Sci. 84 (2012), 735-745.

[18] A. B. Tayler, D. J. Holland, A. J. Sederman and L. F. Gladden, Applications of ultra-fast MRI to high voidage bubbly flow: Measurement of bubble size distributions, interfacial area and hydrodynamics, Chem. Eng. Sci. 71 (2012), 468-483.

[19] D. J. Holland, D. M. Malioutov, A. Blake, A. J. Sederman and L. F. Gladden, Reducing data acquisition times in phase-encoded velocity imaging using compressed sensing, J. Magn. Reson. 203 (2010), 236-246.

[20] A. B. Tayler, D. J. Holland, A. J. Sederman and L. F. Gladden, Exploring the origins of turbulence in multiphase flow using compressed sensing MRI, Phys. Rev. Lett. 108 (2012), 264505.

[21] V. Corona, M. Benning, M. J. Erhardt, L. F. Gladden, R. Mair, A. Reci, A. J. Sederman, S. Reichelt and C.-B. Schöenlieb, Enhancing joint reconstruction and segmentation with non-convex Bregman iteration (2018), arXiv:1807.01660.

[22] Y. C. Eldar, Minimum variance in biased estimation: Bounds and asymptotically optimal estimators, IEEE Trans. Signal Process. 52 (2004), 1915-1930.

[23] Y. C. Eldar, Uniformly improving the Cramér-Rao Lower Bound and maximum-likelihood estimation, IEEE Trans. Signal Process. 54 (2006), 2943-2956.

[24] E. Song, Y. Zhu, J. Zhou and Z. You, Minimum variance in biased estimation with singular Fisher Information Matrix, IEEE Trans. Signal Process. 57 (2009), 376-381.

[25] S. F. Yau and Y. Bresler, A compact Cramér-Rao Lower Bound expression for parametric estimation of superimposed signals, IEEE Trans. Signal Process. 40 (1992), 1226-1230.

[26] A. van den Bos, A Cramér-Rao Lower Bound for complex parameters, IEEE Trans. Signal Process. 42 (1994), 2859.

[27] Y.-Q. Song, Y. Tang, M. D. Hürlimann and D. G. Cory, Real-time optimization of nuclear magnetic resonance experiments, J. Magn. Reson. 289 (2018), 72-78.

Matrix notation

Symbol	Denotes	Comment
$\boldsymbol{a}$	column vector, i.e. 1D matrix	bold, italic, lowercase letter
$\boldsymbol{A}$	matrix, i.e. two or higher dimensional matrix	bold, italic, uppercase letter
$\mathrm{tr}(\boldsymbol{A})$	trace of matrix $\boldsymbol{A}$	
$\mathrm{diag}(\boldsymbol{A})$	diagonal of matrix $\boldsymbol{A}$	
$\boldsymbol{A}^{\mathrm{T}}$	transpose	superscript 'T'
$\boldsymbol{A}^{-1}$	inverse of $\boldsymbol{A}$	
$\boldsymbol{A}^{-}$	generalized inverse of $\boldsymbol{A}$	
$\kappa(\boldsymbol{A})$	condition number of $\boldsymbol{A}$	
$\sigma_i(\boldsymbol{A})$	singular values of $\boldsymbol{A}$	
$\boldsymbol{A}_{\mathrm{v}}$	vectorised matrix, formed by stacking the columns of $\boldsymbol{A}$ into a large vector	subscript 'v'
$[\boldsymbol{a}]_n$	reshape of vector $\boldsymbol{a}$ into a matrix with n columns	square bracketed vector, subscript 'n'
a_i	i-th element of vector $\boldsymbol{a}$	
A_{ij}	element of matrix $\boldsymbol{A}$, positioned at (i, j)	
$\mathrm{sum}(\boldsymbol{A}, i)$	sums the column or row vectors of matrix $\boldsymbol{A}$ along the i-th dimension	
$\mathrm{svd}(\boldsymbol{A})$	Singular Value Decomposition of $\boldsymbol{A}$	
$\mathrm{repmat}(\boldsymbol{A}, k, i)$	stacks k copies of $\boldsymbol{A}$ in the i-th dimension	
$\|\boldsymbol{a}\|_p$	L_p-norm of vector $\boldsymbol{a}$, defined in Section 3.1.1	
$\|\boldsymbol{A}\|_{p,q}$	$L_{p,q}$-norm of matrix $\boldsymbol{A}$, defined in Section 3.3.4	
$\|\boldsymbol{A}\|_{Sp}$	Schatten-p-norm, defined in Section 8.2.2	
$\otimes$	Kronecker product	
$\langle \boldsymbol{A}, \boldsymbol{B} \rangle$	inner product between matrices $\boldsymbol{A}$ and $\boldsymbol{B}$	
$\boldsymbol{A} \cdot \boldsymbol{B}$	element-wise product of matrices $\boldsymbol{A}$ and $\boldsymbol{B}$	central dot