Multi-scale modelling of diffusion magnetic resonance imaging in heterogeneous media

Chun-Ting Lau

University College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity Term 2022

Abstract

Diffusion magnetic resonance imaging (diffusion MRI) is an imaging technique which is sensitive to the micrometre scale movement of water molecules within a medium. It is routinely used in clinical practice, as well as increasingly in biomedical research as a quantitative tool. However, mathematical relations describing the image signal in a heterogeneous medium are poorly established. There is therefore a need for a better mathematical foundation for diffusion MRI which accounts for the heterogeneous nature of biological tissue, both in terms of molecular transport and magnetic properties.

By using a novel distribution function formulation describing diffusion-weighted spin-echo (DW-SE), which is the most common form of diffusion MRI, it is shown that the local modulation wave-vector, known as the q-vector, and the time scale define the operating regime of a DW-SE pulse sequence. Different locations in the q-t parameter space therefore correspond to respective asymptotic models describing DW-SE.

The effects of micro-scale magnetic heterogeneities are then analysed using multiple scales. It is found that in the long-time regime of DW-SE, local induced variations in the q-vector are typically of comparable magnitude to the macro-scale applied q-vector, whereas they are negligible in the short-time regime. Therefore, in the subsequent analysis of the long-time regime, the spatial q-vector variation is considered simultaneously with transport heterogeneities. Using multiple scales homogenisation, a multi-compartmental effective medium model has been derived. The effective diffusion tensors can be calculated by solving a cell problem over a periodic cell of the microstructure, however it is shown that the leading order effect of the spatial q-vector variation integrates exactly to zero.

Meanwhile, the short-time regime of DW-SE is analysed using a boundary layer model. It is first considered using an unphysical assumption of instantaneous modulation of spins, with results agreeing with literature. Taking advantage of the reduced problem complexity, the effects of realistic gradient pulses are then numerically computed. The image signal is found to vary approximately linearly with the pulse duration of a rectangular pulse, thus leading to a proposed two-point extrapolation method for correction. Meanwhile, for more general irregular pulse shapes, it is found that as long as they are symmetric, they correspond to an equivalent rectangular pulse with easily computable parameters.

Finally, for the intermediate-time regime, numerical solutions to the full problem are sought, using simple model geometries and a microscopy-derived realistic microstructure. The results agree with the two derived models at the respective asymptotes, with a transitional region of about a decade in the q-value. Additionally, the transitional region occurs at smaller q-values for isolated intra-cellular spaces compared to the connected extra-cellular space. This finding can inform future experiment design and modelling, particularly in relation to separating and analysing the intra-cellular signal component.

Acknowledgements

I would like to express my deepest gratitude to:

- The Clarendon Fund, the Oxford-Radcliffe Graduate Scholarship and the Croucher Foundation for providing the funding,
- The Oxford-Nottingham Biomedical Imaging (ONBI) Centre for Doctoral Training, which this degree is part of, for their continued support,
- My supervisors Stephen Payne and Jon Chapman, for their guidance and feedback throughout the research,
- Tianyou Xu and Karla Miller for providing the transmission electron microscopy image and segmentation data, and
- Ching Lok Chong for useful discussions on the distribution function and co-rotating formulations.

Contents

Chapter 1

Introduction

1.1 Background

Diffusion-weighted magnetic resonance imaging (diffusion MRI) refers to a family of imaging methods which are sensitive to the microscopic motion of MRI-visible molecules, usually water, within a medium. In biomedical applications, the power of diffusion MRI lies in its ability to non-invasively probe micro-scale tissue structure and potentially allow the inference of micro-scale biophysical processes. In clinical practice, it has found important applications in the diagnosis and management of acute stroke patients. As diffusion MRI is able to detect early changes to microscopic water transport properties in the brain before macro-scale structural change, valuable time is gained to determine the optimal treatment to minimise permanent damage. Furthermore, diffusion MRI has become a standard tool for neuroscience research, where it is used to characterise the white-matter fibre structure and the connectivity between different areas of the brain.

In recent years, there is an increasing trend of using diffusion MRI as a quantitative rather than qualitative technique. The quantitative use of diffusion MRI originates all the way back from the initial diffusion spin-echo experiment by Stejskal and Tanner [108], where the diffusion coefficient of a homogeneous sample of dry glycerol. However, the complex nature of biological tissues presents many difficult challenges in the quantitative

interpretation of diffusion MRI measurements. The microstructure of tissue introduces heterogeneities both in terms of tracer transport and magnetic properties, such as cell membranes introducing impermeable barriers, or diamagnetic substances such as deoxyhaemoglobin and lipids causing magnetic inhomogeneities. Many modelling methods have been proposed to infer properties of tissue microstructure and hence their clinical implications from diffusion MRI measurements, for example the apparent diffusion coefficient for the assessment of acute cortical ischaemia [72, 112], the q-space concept for investigating axon diameter and direction in the white matter [4, 21, 47, 54], diffusion tensor imaging (DTI) to measure anisotropic diffusion [10], and the intravoxel incoherent motion (IVIM) concept to model blood flow in vascularised tissue [64, 66, 114]. However, since these modelling methods arise from separate clinical needs or interests, they require bespoke assumptions either explicitly or implictly, and the relationship between model parameters and the physical anatomy is often unclear. There is therefore a need for a better mathematical foundation for these techniques.

1.2 Asymptotic methods in multi-scale problems

Many engineering problems can be described and solved by differential equations. However, such equations do not necessarily possess analytical solutions, making them difficult to solve exactly. For practical purposes, mathematical methods are therefore required to obtain meaningful approximate solutions. In practice, these problems are often posed in multi-scale settings. For example, co-existing physical processes may have magnitudes of different scales, meaning that some are dominant over others. Heterogeneous structures within the domain of definition may also occur at multiple length or time scales. By studying the limit behaviour as the scale separation tends to infinity, asymptotic methods provide a powerful approach to derive approximate solutions to these problems. These methods disentangle scale-separated phenomena in a systematic manner, thus breaking down a complex multi-scale problem into several simpler ones.

One type of asymptotic methods which is particularly relevant to biological tissue is multi-scale homogenisation. This method applies to problems with a scale separation presented by a heterogeneous microstructure which is much smaller than the macroscopic observation scale; examples of such multi-scale heterogeneous media include composite materials and porous media. The aim of homogenisation is to derive an approximate continuum model at the macro-scale, based on the asymptotic behaviour of the heterogeneous medium at the large scale separation limit. Although the method was originally developed in engineering disciplines such as geotechnics, it has also found application in modelling biological tissues, for example vascular tumours [102], myocardium [106] and the cerebral cortex [36]. Typically, the scale separation between cellular or microvascular structures and the macroscopic scale is exploited to apply multi-scale homogenisation.

Another type of asymptotic methods which is relevant is boundary layer analysis. This method applies when there is a scale separation between the local processes near a boundary and the bulk behaviour in a domain. The essence of the method is to isolate the heterogeneous behaviour in the local scale by considering a simpler problem within the boundary layer, so its effect in the context of the global domain can then be summarised by an integration through the layer. This method is also commonly used in engineering disciplines, for example in fluid mechanics and thermodynamics.

The motivation to use multi-scale asymptotic methods is twofold. Firstly, at a practical level it greatly reduces the complexity of the problem, particularly when further numerical approximation of the solution is required. If numerical methods are directly applied on a multi-scale problem, the discretisation scheme needs to resolve all scales simultaneously. This comes with a huge computational cost – for example a typical MRI voxel size is about 1 mm while an axon has a diameter of order 1 μm, so to directly discretise a voxel-sized volume of white matter would require at least 1 billion elements already. Secondly and perhaps more importantly, at a theoretical level it systematically extracts the important parameters which govern the overall physical behaviour. This allows the comparison and analysis of solutions from different situations in a mathematically justified manner. Take

civil engineering as an example – the flow field of water seeping through one particular soil sample can indeed be solved numerically, if one has the computational resources, but arguably it is more meaningful to describe it using the macroscopic Darcy's law. This concept is equally applicable in clinical imaging, as information about the collective, overall behaviour of tissue is more likely to be diagnostically useful, compared to the observation of every single living cell in one's body. The work in this thesis therefore aims to use this methodology to investigate the link between diffusion MRI and multi-scale media. By deriving mathematically justified models to explain the MRI signal, they can inform us about how the measured signals relate to microstructure properties, and allow for better quantitative imaging methods to be developed.

1.3 Outline of thesis

In Chapter 2, a literature review on the basic concepts of diffusion MRI and its current developments is presented. The underlying physics and the technical aspects of diffusion MRI implementation are reviewed with focus on the mathematical aspects, in particular the validity of a partial differential equation description using the Bloch-Torrey equations. The need for modelling in the quantitative interpretation of diffusion MRI is discussed. A review on current diffusion MRI modelling methods then reveals a gap in the literature in the mathematical understanding of diffusion MRI in general heterogeneous media and appropriate models to describe it.

In Chapter 3, a novel distribution function formulation of diffusion-weighted spin-echo (DW-SE), which is the most common form of diffusion MRI, is proposed. Using this formulation, the relevant scales defining an operating regime of DW-SE are identified. These are the local spatial modulation wave-vector (the q-vector) and the time scale respectively, leading to the use of the q-t parameter space to characterise diffusion MRI experiments.

In Chapter 4, the local perturbation of the q-vector arising from magnetic hetero-

geneities is modelled using multi-scale homogenisation. This leads to Chapter 5, which covers the long-time regime by using multi-scale homogenisation to model transport and magnetic heterogeneities simultaneously. The short-time regime is then covered in Chapter 6 with the use of boundary layer analysis. The intermediate-time regime is also considered in Chapter 7 through the use of numerical simulations. Finally, Chapter 8 concludes the thesis with some final remarks and outlines the future research directions arising from this thesis.

Chapter 2

Principles of diffusion-weighted MRI

2.1 Introduction

This chapter covers the basic principles of diffusion-weighted MRI and its current developments, which include the theoretical aspects of the underlying physics, the technical aspects to implement a transport-sensitive MRI experiment, and finally the modelling aspects of quantitatively interpreting these images. The aim of this chapter is to critically review the existing literature on these topics, so as to identify gaps of current knowledge as discussed in the previous chapter.

2.2 Nuclear magnetic resonance in diffusive and advective media

The underlying physics that is responsible for diffusion-weighted MRI is the behaviour of nuclear magnetic resonance in media which is diffusive and/or advective. This section aims to give a brief overview of this topic, with particular focus on its mathematical descriptions at length scales relevant to biological tissue.

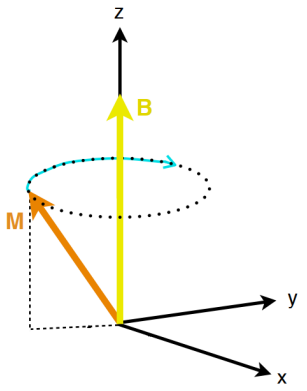

Figure 2.1: Larmor precession of magnetisation vector **M** about the direction of a strong magnetic field **B** for a positive gyromagnetic ratio, assuming no relaxation effects. Adapted from [90].

2.2.1 The phenomenon of nuclear magnetic resonance

Central to the field of MRI physics is the phenomenon of nuclear magnetic resonance (NMR), which occurs when certain atomic nuclei are placed in a strong magnetic field. Its discovery dates back to 1946 when two groups of researchers independently reported the NMR phenomenon experimentally [15, 95], with its importance recognised by the 1952 Nobel Prize for Physics shared between Felix Bloch and Edward Purcell. This topic is covered comprehensively in many texts, for example an exposition at a level appropriate to imaging applications has been given by Jezzard and Clare [52], while detailed descriptions of its physical origins can be found in Chizhik et al. [26] and Levitt [70].

Nuclear magnetic resonance is observed when atomic nuclei possessing a magnetic moment, such as those of 1-hydrogen (^{1}H), are placed in a strong magnetic field and perturbed from the magnetic equilibrium. At the atomic level, such behaviour is quantum mechanical and is related to the splitting of energy states of the particles under a magnetic field. At a bulk continuum level, however, the collective behaviour of an ensemble of particles can be described using a phenomenological approach, as suggested by Bloch [14]. At equilibrium, nuclear magnetic moments preferentially align to the strong magnetic field, resulting in a net magnetisation in the field direction. When they are perturbed from equilibrium, the misalignment causes magnetic moments to precess about the field

direction at the Larmor frequency, which is proportional to the magnetic flux density. The time-variation of the net magnetisation concentration $\mathbf{M}$ under a magnetic field $\mathbf{B}$, assuming no relaxation, is therefore described by:

$$\frac{\mathrm{d}\mathbf{M}}{\mathrm{d}t} = \gamma \mathbf{M} \times \mathbf{B}, \tag{2.1}$$

where γ is the gyromagnetic ratio, determined by the quantum properties of the atomic nuclei in consideration. If a static magnetic field is applied in the positive z-direction: $\mathbf{B} = (0, 0, B_0)^T$, the solution to (2.1) takes the form of a rotation about the z-axis:

$$\mathbf{M} = \begin{pmatrix} a\cos\omega_0 t + b\sin\omega_0 t \\ -a\sin\omega_0 t + b\cos\omega_0 t \\ c \end{pmatrix},$$

where a, b, c are some constants and $\omega_0 = \gamma B_0$; the rotation is clockwise if the gyromagnetic ratio is positive (Figure 2.1). For the nucleus of ^{1}H, which is responsible for almost all MRI applications, the value of γ is approximately 267.52×10^6 rad s^{-1} T^{-1}. For modern clinical MRI scanners, the static field strength is in the range of 1.5 to 7 Teslas, so the nominal Larmor frequency ($f_0 = \gamma B_0 / 2\pi$) is in the radiofrequency (RF) range of 63.9 to 298 MHz.

2.2.2 Relaxation and transport of nuclear magnetisation

Equation (2.1) describes the immediate non-equilibrium behaviour of magnetisation, however over time the spins have a tendency to restore a state of dynamic equilibrium through irreversible, random processes. This is known as the relaxation of magnetisation. When there is a strong magnetic field with flux density B_0, conventionally applied in the z-direction, a net magnetisation M_0 in the same direction is induced at equilibrium. Equation (2.1) can then be modified to model relaxation as exponentially-decaying processes

[14]:

$$\frac{\mathrm{d}M_x}{\mathrm{d}t} = \gamma(M_y B_z - M_z B_y) - \frac{M_x}{T_2}, \tag{2.2}$$

$$\frac{\mathrm{d}M_y}{\mathrm{d}t} = \gamma(M_z B_x - M_x B_z) - \frac{M_y}{T_2}, \tag{2.3}$$

$$\frac{\mathrm{d}M_z}{\mathrm{d}t} = \gamma(M_x B_y - M_y B_x) - \frac{M_z - M_0}{T_1}, \tag{2.4}$$

where T_1 and T_2 are the longitudinal and transverse relaxation time constants respectively. These equations are collectively known as the Bloch equations, and are valid when the magnetic field perturbation is negligible compared to the static B_0 field, so the equilibrium magnetisation is only ever slightly perturbed from $(0, 0, M_0)^T$. Under a static field $\mathbf{B} = (0, 0, B_0)^T$, the solution to (2.2)-(2.4) is in the form:

$$\mathbf{M} = \begin{pmatrix} e^{-t/T_2}(a\cos\omega_0 t + b\sin\omega_0 t) \\ e^{-t/T_2}(-a\sin\omega_0 t + b\cos\omega_0 t) \\ M_0 - ce^{-t/T_1} \end{pmatrix},$$

which precesses at the Larmor frequency but exponentially decays towards the equilibrium state.

Similar to precession, the physical origin of longitudinal and transverse relaxations at an atomic level is not classical. They are driven by random thermal processes that stochastically affect the discrete energy states of the nuclei, providing a pathway to restore equilibrium eventually. The relaxation time constants depend on the rate in which mixing occur between the energy states, therefore they vary with both the magnetic field

Tissue	$B_0 =$ 1.5T[1]	3.0T[2,3]	4.0T[4]
White matter T1	640	860	1040
Grey matter T1	880	1200	1410
White matter T2	80	80	50
Grey matter T2	80	110	50

[1] MacFall *et al.* (1987). *Magn. Reson. Imaging* **5**, 209-220.
[2] Wansapura *et al.* (1999). *J. Magn. Reson. Imaging* **9**, 531-538.
[3] Clare and Jezzard *et al.* (2001). *Magn. Reson. Med.* **45**, 630-634.
[4] Duewell *et al.* (1996). *Radiology* **199**, 780-786.

Table 2.1: Average values of T1 and T2 (units: ms) in the human brain, cited in [52]

strength B_0 and the environment in which the nuclei are situated in, as shown in Table 2.1. In biological tissues, the dominant contribution of transverse relaxation is the irreversible decoherence in Larmor frequencies among the ensemble of nuclei, caused by random fluctuations in local magnetic field induced by neighbouring tumbling molecules. However, this effect does not contribute to longitudinal relaxation, therefore generally T2 is shorter than T1 [52, 69].

A completely separate mechanism, which causes an apparent loss of transverse magnetisation, is often attributed as a type of 'relaxation' termed T2* relaxation [52]. This is caused by spatial inhomogeneities in the static field B_0, leading to a *bulk* rather than random decoherence of Larmor frequencies. Therefore, this effect is not a true relaxation, as the process involved is *reversible* and is not responsible for the restoration of equilibrium. Indeed, the T2* reversibility is exploited in the spin-echo experiment [44], in which a pulse of rotating magnetic field is used to flip the magnetisation by 180 degrees, thus reversing the bulk dephasing effect and refocusing the magnetisation. In fact, the T2* effect is only observed due to the lack of spatial specificity to resolve an inhomogeneous static B_0 field, whereas true transverse (T2) relaxation occurs among the ensemble of nuclei at every infinitesimal volume of space, so it is observed irrespective of spatial resolution.

In the Bloch equations (2.2)-(2.4), the magnetisation is assumed to only vary in time, ignoring any effects caused by spatial variation. However, spatial variation of magnetisation concentration introduces another pathway to dynamic equilibrium, which is through the irreversible, diffusive transport of these resonating nuclei. At a molecular level, spatial transport occurs as Brownian motion as postulated by kinetic theory. This effect was first observed and described by Hahn [44]; the modelling of this random walk behaviour of nuclei in magnetic resonance was subsequently proposed by Carr and Purcell [23]. However, at a continuum level, the effects of spatial magnetisation transport can be described by considering its associated flux quantities over an infinitesimal volume. This approach is first adopted by Torrey [113], who modified the Bloch equations by adding diffusive terms

which obey Fick's law:

$$\frac{\partial M_x}{\partial t} = \nabla \cdot (D\nabla M_x) + \gamma(M_y B_z - M_z B_y) - \frac{M_x}{T_2}, \tag{2.5}$$

$$\frac{\partial M_y}{\partial t} = \nabla \cdot (D\nabla M_y) + \gamma(M_z B_x - M_x B_z) - \frac{M_y}{T_2}, \tag{2.6}$$

$$\frac{\partial M_z}{\partial t} = \nabla \cdot (D\nabla M_z) + \gamma(M_x B_y - M_y B_x) - \frac{M_z - M_0}{T_1}, \tag{2.7}$$

where D is the diffusion coefficient of the nuclei in the medium, and $\mathbf{M}(x, y, z, t)$ is the magnetisation *density* which depends on both space and time. In the original formulation of Torrey the diffusion behaviour is assumed to be isotropic, therefore the diffusion coefficient is a positive scalar; for anisotropic media this can be generalised by replacing it with a diffusion tensor. The same as the Bloch equations, here the magnetic field perturbation is assumed to be negligible so the equilibrium state is effectively constant. Equations (2.5)-(2.7) are collectively known as the Bloch-Torrey equations, and are partial differential equations in both space and time variables, rather than ordinary differential equations in just the time variable in (2.2)-(2.4). A logical, simple extension to (2.5)-(2.7), as proposed by Stejskal [107], is the addition of advective transport terms caused by a bulk flow of velocity $\mathbf{u}$:

$$\frac{\partial M_x}{\partial t} = \nabla \cdot (D\nabla M_x - \mathbf{u}M_x) + \gamma(M_y B_z - M_z B_y) - \frac{M_x}{T_2}, \tag{2.8}$$

$$\frac{\partial M_y}{\partial t} = \nabla \cdot (D\nabla M_y - \mathbf{u}M_y) + \gamma(M_z B_x - M_x B_z) - \frac{M_y}{T_2}, \tag{2.9}$$

$$\frac{\partial M_z}{\partial t} = \nabla \cdot (D\nabla M_z - \mathbf{u}M_z) + \gamma(M_x B_y - M_y B_x) - \frac{M_z - M_0}{T_1}. \tag{2.10}$$

The detectable signal of NMR originates from the transverse components of magnetisation, which resonates nominally at the Larmor frequency during free precession. Therefore, it is conventional to denote the transverse component of $\mathbf{M}$ as a complex phasor: $M_{xy} = M_x + iM_y$. It is also often convenient to express the T2 relaxation term using the rate constant $R_2 = 1/T_2$. Since the transverse components of the magnetic field (B_x and B_y) are negligible during free resonance, (2.8) and (2.9) can be simplified to a more

compact form:

$$\frac{\partial M_{xy}}{\partial t} = \nabla \cdot (D\nabla M_{xy} - \mathbf{u}M_{xy}) - i\gamma B_z M_{xy} - R_2 M_{xy}. \tag{2.11}$$

2.2.3 Validity of deterministic continuum formulation

As explained above, at the atomic or molecular level (nanometre scale) both magnetic resonance and spatial transport of nuclear magnetic moments are stochastic in nature. However, as observed by Einstein [35], the time evolution and spatial variation of probability distributions arising from the stochastic processes satisfy some deterministic partial differential equations. Macroscopic quantities, which are some expected values of many samples from these distributions, are therefore governed by deterministic laws. Subsequently, with the introduction of stochastic calculus and its analysis, this equivalence between stochastic and deterministic points of view has been justified with mathematical rigour [8]. Since the number of stochastic steps taken is large at a continuum scale, the nanoscale stochastic effects are averaged out and contained in the coefficients of the PDE. In the case of resonating spins in an infinite homogeneous diffusive medium, this can be explained using the central limit theorem, as illustrated by Carr and Purcell [23]. The mathematical equivalence also holds when the domain contains boundaries, as it can be demonstrated that absorption, instantaneous reflection and partial reflection conditions in the stochastic sense correspond to Dirichlet, Neumann and Robin boundary conditions in the PDE sense respectively [98, 105, 117].

For the description of NMR in advective-diffusive media, the validity of a deterministic continuum formulation is dependent on the scaling condition of stochastic collisions being frequent. This can be assessed by the Knudsen number, which is the mean free path between collisions divided by the length scale of interest. In liquids, the mean free path is comparable to the molecular size (<1nm), while cellular structures of biological tissue is at the micrometre scale. The Knudsen number therefore is very small, implying that a large number of molecular collisions, responsible for relaxation and transport, hap-

pen over the length scale of interest, thus justifying the deterministic approach. In fact, the use of the Bloch equations already implicitly assumes this scale separation anyway, through the bundling of resonance and relaxation processes into the gyromagnetic ratio and relaxation time constants. Delving into the underlying stochastic processes is only insightful when inferring nanoscale properties, such as chemical composition or molecular structure, from observed continuum parameters. For the purpose of modelling the effect of cellular structures in MRI, the stochastic and deterministic approaches are mathematically equivalent. The PDE formalism of (2.8)-(2.10) together with appropriate boundary conditions is therefore sufficient, contrary to the suggestion in some NMR/MRI texts that the Bloch-Torrey equations do not hold in heterogeneous media and hence the necessary language of atomic stochastic processes in modelling [19, 20].

2.3 Implementing transport-sensitive MRI

This section outlines the implementation of diffusion-weighted MRI, which are MRI experiments that are sensitive to magnetisation transport. The aim of this section is to provide the necessary context for the mathematical description of diffusion-weighted MRI, rather than covering the technical details such as hardware consideration or pulse sequence design. A more comprehensive discussion of these topics is available in the literature [104, 122].

2.3.1 Hardware setup for gradient fields and RF pulses

The generic hardware setup of an MRI scanner is schematically shown in Figure 2.2. The outermost layer is the superconducting magnet, which generates a strong, static magnetic field B_0 in the longitudinal (z-) direction so that magnetic resonance can occur inside the bore. The next layer comprises three sets of gradient coils which, when powered up, generate magnetic fields with longitudinal components that vary linearly in the x-, y- and z-directions respectively. The innermost layer, just surrounding the imaged

volume, are sets of RF coils which are tuned to the nominal Larmor frequency. Depending on whether they are in transmitting or receiving mode, the RF coils either generate or detect oscillatory magnetic fields in the transverse direction. The gradient and RF coils are connected, via amplifiers and electronics, to a computer console which coordinates their inputs and outputs according to a programmed pulse sequence, and subsequently reconstruct an MR image from the acquired information.

The role of the gradient coils is to spatially modulate the resonance frequency in a volume by introducing spatial variation of longitudinal field strength B_z. By superimposing some combination of the three gradient fields onto the static field, the Larmor frequency as a function of space becomes: $\omega = \gamma(B_0 + G_x x + G_y y + G_z z)$, with the origin set at the middle of the bore. A spatial gradient of Larmor frequency can therefore be specified at any direction by setting an appropriate combination of inputs (G_x, G_y, G_z).

Meanwhile, the role of the RF coils in magnetic resonance is more complex. In receiving mode, they are used to detect the resonating magnetisation by passively sensing the oscillatory induced voltages. In transmitting mode, the coils are responsible for generating pulses of oscillatory transverse magnetic field, usually referred to as RF pulses, that can rotate the magnetisation vector about an axis on the transverse plane. RF pulses are essential in any magnetic resonance experiment, as the spins need to be perturbed away from their equilibrium longitudinal alignment. They can also be incorporated into more sophisticated pulse sequences to manipulate magnetic resonance in specific ways. To

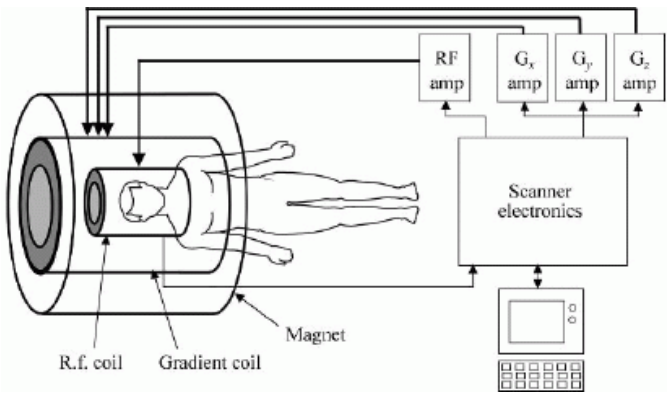

Figure 2.2: Block diagram of hardware arrangement in a generic MRI scanner, reproduced from [52].

illustrate the principle, consider a rotating RF magnetic field of magnitude B_1, which can be generated by a quadrature drive, superimposed on a strong static field of magnitude B_0:

$$B_x = B_1 \cos(\omega t), B_y = -B_1 \sin(\omega t), B_z = B_0.$$

The duration of this pulse is much shorter than relaxation and transport time scales, so its effect on magnetisation can be described by the Bloch equation without decay terms (2.1). Using a change of variable from the stationary frame (x, y, z) to the rotating frame (x', y', z) with respect to the driving frequency:

$$M_{x'} = M_x \cos \omega t - M_y \sin \omega t,$$
$$M_{y'} = M_x \sin \omega t + M_y \cos \omega t,$$

Equation (2.1) then becomes:

$$\frac{\mathrm{d}M_{x'}}{\mathrm{d}t} = -(\omega - \gamma B_0) M_{y'}, \tag{2.12}$$
$$\frac{\mathrm{d}M_{y'}}{\mathrm{d}t} = \gamma B_1 M_z + (\omega - \gamma B_0) M_{x'}, \tag{2.13}$$
$$\frac{\mathrm{d}M_z}{\mathrm{d}t} = -\gamma B_1 M_{y'}. \tag{2.14}$$

In the case of excitation, the equilibrium magnetisation is the initial condition: $\mathbf{M}(0) = (0, 0, M_0)^T$. Defining $\omega_{\text{eff}} = \sqrt{\gamma^2 B_1^2 + (\omega - \gamma B_0)^2}$, the solution to (2.12)-(2.14) can be written exactly, and then approximated for $|\omega - \gamma B_0| \ll |\gamma B_1|$, as:

$$M_{x'} = M_0 \frac{\gamma B_1 (\omega - \gamma B_0)}{\omega_{\text{eff}}^2} (1 - \cos \omega_{\text{eff}} t) \qquad \approx 0, \tag{2.15}$$
$$M_{y'} = M_0 \frac{\gamma B_1 \sin \omega_{\text{eff}} t}{\omega_{\text{eff}}} \qquad \approx M_0 \sin \gamma B_1 t, \tag{2.16}$$
$$M_{z'} = M_0 \left(\frac{\gamma^2 B_1^2 \cos \omega_{\text{eff}} t}{\omega_{\text{eff}}^2} + \frac{(\omega - \gamma B_0)^2}{\omega_{\text{eff}}^2} \right) \qquad \approx M_0 \cos \gamma B_1 t. \tag{2.17}$$

As both the amplitude and frequency are approximated in the above expressions, they are only valid for small t and diverges after time. The approximations are exact when $\omega = \gamma B_0$, which is if the driving frequency exactly matches the Larmor frequency. In this case, the magnetisation vector rotates about the x'-axis at an angular rate of $-\gamma B_1$

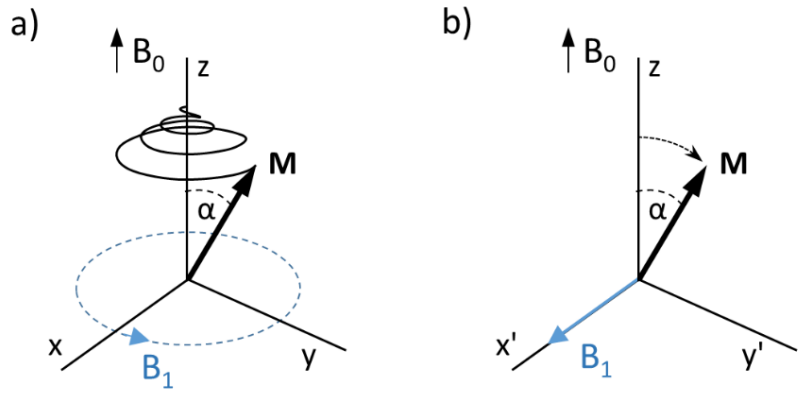

Figure 2.3: Rotation of magnetisation vector by an RF pulse shown in (a) stationary and (b) rotating frames, reproduced from [48].

(Figure 2.3), therefore if the pulse has a duration of $t = \pi/(2\gamma B_1)$, the magnetisation will be rotated by 90° ending up on the transverse plane. Similarly, if either the duration or the amplitude is doubled, the magnetisation will be rotated by 180°. These two pulses are particularly useful in MRI: the 90° pulse generates the maximum amount of transverse magnetisation and hence signal from equilibrium, whereas the 180° pulse completely reverses the magnetisation and can be used for refocusing transverse magnetisation or sensitising contrast in T1 relaxation.

In the simple example, the rotating B_1 field only has one specific frequency component, therefore the excitation is maximum when it matches the Larmor frequency and drops

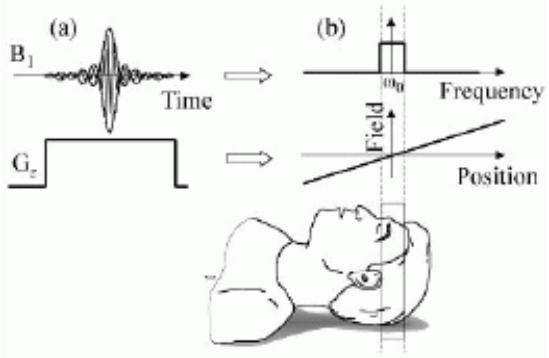

Figure 2.4: Frequency-selective excitation using a sinc-function modulated RF pulse. When applied in conjunction with a linear field gradient, spatial selectivity of excitation (slice selection) is achieved. Reproduced from [52].

off either side. This is because the approximations in (2.15)-(2.17) worsen as $|\omega - \gamma B_0|$ increases. In practice, the frequency content of RF pulses can be carefully designed to achieve a certain excitation profile over a range of Larmor frequencies, trading off a number of factors such as selectivity required, pulse duration or energy deposition. For example, a sinc-function pulse envelope can be used for an approximate 'top-hat' excitation profile (Figure 2.4), with the bandwidth of frequencies excited determined by the sinc envelope duration. This is often used in conjunction with a magnetic field gradient to achieve spatial excitation selectivity, also known as slice selection as the pulse only rotates magnetisation within a slice of the volume.

2.3.2 Diffusion-weighted spin-echo (DW-SE) pulse sequence

The simplest, and by far the most common, pulse sequence that is sensitive to spatial transport of magnetisation is the diffusion-weighted spin-echo (DW-SE) sequence with pulsed gradients [61]. It is a simple modification of the common spin-echo technique for T2-weighted imaging, which involves a 90° pulse followed by a refocusing 180° pulse (Figure 2.5), by adding pairs of large field gradients either side of the 180° pulse (Figure 2.6). Other less common transport-sensitive pulse sequences exist as well, such as stimulated echo [74, 75, 111] and steady-state free precession (SSFP) [59, 63, 76] methods, although they are beyond the scope of this thesis.

In the spin-echo sequence without diffusion weighting, the signal undergoes T2* decay after the initial excitation by a 90° pulse, due to bulk decoherence of spins caused by B_0

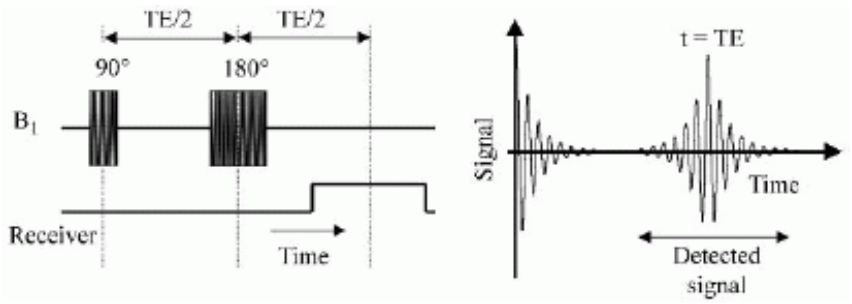

Figure 2.5: The spin-echo pulse sequence, reproduced from [52].

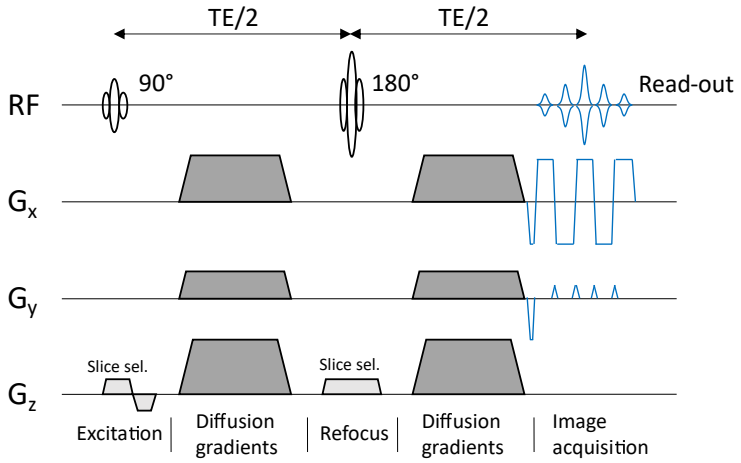

Figure 2.6: Diagram of a typical diffusion-weighted spin-echo (DW-SE) pulse sequence.

inhomogeneities. The spin dephasing is reversed by applying a 180° pulse at the midpoint of the sequence, so the magnetisation is refocused and produces an echo in the detected signal, the amplitude of which now reflects the true transverse (T2) relaxation over the echo time (TE). By introducing spatial encoding during signal detection, a T2-weighted image can be reconstructed. The details of the spatial encoding and image reconstruction methods depend on the image acquisition protocol (see [52]), although for the purpose of modelling and interpreting DW-SE it can be treated as a black box providing an image, which effectively is an array of numbers containing voxel-averaged signal values.

Meanwhile for the diffusion-weighted sequence, extra signal attenuation is introduced by the pair of field gradients, which are usually referred to as the 'diffusion gradients'. The first lobe of diffusion gradients before the 180° pulse causes a spatial modulation of magnetisation phase along the applied gradient direction. If there is no translational motion of spins, they would be perfectly demodulated by the second lobe of diffusion gradients, which has the same area under the curve as the first lobe. In a diffusive and/or advective medium, however, spins can mix with each other through spatial transport in the time between the two diffusion gradients (the mixing period), leading to further signal attenuation on top of the intrinsic T2 decay when demodulated. Therefore, when the image is acquired and reconstructed, the signal values are weighted by both T2 and transport-induced decay. In clinical science, the transport-induced decay is usually attributed solely

to diffusion with the advective behaviour ignored, hence the transport-induced component of signal attenuation is commonly referred to as 'diffusion-weighting'. However, the role of advection in DW-SE has also been recognised, with the more general and accurate terminology of 'intra-voxel incoherent motion' (IVIM) preferred by some authors [62, 65, 114].

There is a huge variety of DW-SE implementations in the literature. For example, the diffusion gradient pulses may be comparable to the diffusion time [65, 66, 114], or much shorter than the diffusion time to allow for a delta-function modelling approach [21, 107, 108]. More general waveforms of pulsed diffusion gradients, such as oscillatory waveforms (OGSE), have also been proposed [85, 115, 119]. To achieve the same level of diffusion sensitisation, short or oscillatory pulses require stronger hardware, therefore they are more common for research rather than clinical use. In most clinical applications, DW-SE is used as a stand-alone pulse sequence to produce images that are T2- and diffusion-weighted. Less commonly, it can also act as a module integrated into more complex pulse sequences, for example the magnetisation can be prepared, prior to DW-SE, by sequences such as arterial spin labelling (ASL) [100], resulting in an ASL image which is subsequently diffusion-weighted. Despite the variety of implementations (and their associated acronyms), the fundamental principles and mechanisms are the same. The action of any DW-SE pulse sequence can be formulated by the extended Bloch-Torrey equations (2.8)-(2.10), with the various implementations being subsets of the general DW-SE description.

2.4 Current methods of diffusion MRI modelling

In this section, the quantitative aspects of modelling and interpreting diffusion-weighted MRI data are covered. Firstly, the concept of quantitative imaging as an inverse problem is explained, demonstrating the need for modelling using suitable assumptions. The models found in the current literature are then summarised and categorised, critically reviewing

their associated assumptions and identifying the links and gaps between them.

2.4.1 The need for modelling: quantitative imaging as an inverse problem

With the insertion of the pair of diffusion gradients, the acquired image signal of a DW-SE pulse sequence is weighted by both T2 and transport-induced decay. The extra transport-induced attenuation qualitatively relates to the diffusive and advective rates, as well as the magnitude, direction and duration of the diffusion gradients. Quantitatively, the action of DW-SE on magnetisation can be formulated by the extended Bloch-Torrey equations (2.8)-(2.10) with appropriate boundary conditions, which describes the behaviour of NMR in a general diffusive-advective medium from a continuum perspective, as explained in Section 2.2. Equations such as (2.8)-(2.10) therefore serve as the link between the *physical quantities* and the observed *signal quantities* that an MRI scanner picks up.

When we have full knowledge about the physical properties of the imaged medium, such as the configuration of its boundaries, diffusion coefficient, advection velocities etc., we can predict the eventual image signal resulting from a pulse sequence by solving the Bloch-Torrey equations. This is known as the *forward* problem, which solves for a prediction of signal quantities given the full knowledge of physical properties. However, the setting of quantitative imaging is the *inverse* problem: we aim to infer some unknown physical properties of a medium through some measurement of signal quantities. Unlike the forward problem, the inverse problem usually is not well-posed. Without even considering the effects of noise, the finite number of obtained signal values is not enough to determine the exact configuration of the medium, for example the geometry of domain boundaries, which has essentially infinitely many degrees of freedom. Therefore, assumptions need to be made about the unknown medium, thus representing the physical problem by a model with a finite number of parameters. These model parameters, which have a known relation to the actual physical quantities when the assumptions are satisfied, can then be determined by regression against the signal data gathered experimentally.

An ideal choice of model for a particular diffusion MRI application should reflect the physical reality of the expected imaged medium, while providing the right level of quantitative information required. For example, the simplest, two-parameter modelling approach to DW-SE is the characterisation using only T2 and the diffusion coefficient, assuming the medium to be infinite, homogeneous and isotropic with no advection. Commonly, the effect of T2 relaxation is not estimated, with its effect being considered as common mode signal attenuation and factored out by comparing images taken at the same echo time, thus effectively leading to a one-parameter estimation model. This approach is used extensively for diffusion coefficient quantification in chemistry [22, 110], as the physical assumption is indeed justified for a still, homogeneous chemical sample placed in an NMR magnet. However, the assumption no longer holds in a heterogeneous medium. Although the same one-parameter model can be applied to data collected from a heterogeneous medium, and indeed produces a good fit under some limiting conditions, the data-derived 'apparent diffusion coefficient', a commonly used quantity as proposed by Stejskal and Tanner [108] and Le Bihan et al. [65], can no longer be interpreted strictly as the physical diffusion coefficient of the constituent molecules.

There exists some confusion in current MRI literature between *physical* and *signal* quantities. In a recent review article, Novikov et al. [86] claimed that diffusion tensor imaging (DTI, the tensor generalisation of the aforementioned apparent diffusion coefficient concept) is not a model, and thus reflects a 'time-dependent physical quantity' in the diffusion tensor with no assumption. However, rather than being a true physical quantity describing the behaviour of magnetisation within the medium, the 'time-dependent diffusion tensor' in the definition of Novikov et al. is really a signal quantity determined by polynomial fitting of the logarithm of observed signal magnitudes. Indeed, in this logic the deduction of such quantities involves no assumption, as they are merely some particular summary of the signal data. However, the interpretation of the signal quantity as physical quantities itself involves the use of some model and assumptions, either explicitly or implicitly.

2.4.2 Apparent homogeneous continuum modelling

This modelling approach of an apparent homogeneous continuum has been studied and implemented extensively, and includes the concepts of apparent diffusion coefficient (ADC) [65, 114] and diffusion tensor imaging (DTI) [10, 11, 67, 93, 94]. As mentioned in the previous section, the simplest approach to model a DW-SE sequence is to assume the medium to be homogeneous and isotropic, with no advection and to have far away boundaries. The transverse magnetisation in each voxel can therefore be characterised by its transverse relaxation time constant and diffusion coefficient. However, in clinical imaging there is inevitably some motion of the patient during the pulse sequence, leading to some bulk coherent motion of the medium, therefore a constant, spatially coherent advection velocity should also be included in the model. Recall the Bloch-Torrey equation for transverse magnetisation in the absence of RF fields (2.11), which is valid either side of the 180° refocusing pulse of the DW-SE sequence:

$$\frac{\partial M_{xy}}{\partial t} = \nabla \cdot (D\nabla M_{xy} - \mathbf{u}M_{xy}) - i\gamma B_z M_{xy} - R_2 M_{xy}. \tag{2.18}$$

Since the medium is assumed to be magnetically homogeneous, the magnetic field generated by the static magnet and gradient coils can be written as a function of space and time:

$$B_z(\mathbf{x}, t) = B_0 + \mathbf{g}(t) \cdot \mathbf{x}$$

where B_0 is the uniform static magnetic field and $\mathbf{g}(t)$ is the field gradient, assuming that the field variation generated by the gradient coils is linear in space. Equation (2.18) can then be written in the rotating frame with respect to the nominal Larmor frequency $\omega_0 = \gamma B_0$:

$$\frac{\partial M_\perp}{\partial t} = \nabla \cdot (D\nabla M_\perp - \mathbf{u}M_\perp) - i\gamma \mathbf{g}(t) \cdot \mathbf{x} M_\perp - R_2 M_\perp. \tag{2.19}$$

where $M_\perp = M_{xy} \exp(i\gamma B_0 t)$. To take into account the phase reversal effect of the 180° pulse, which can be assumed to be instantaneous, the following effective quantities can be

defined and used in (2.19) so that the equation holds continuously across the 180° pulse:

$$M_\perp^{\text{eff}} = \begin{cases} M_\perp & \text{if } 0 < t < t_{180}, \\ \overline{M_\perp} & \text{if } t > t_{180}, \end{cases} \qquad \mathbf{g}^{\text{eff}} = \begin{cases} \mathbf{g} & \text{if } 0 < t < t_{180}, \\ -\mathbf{g} & \text{if } t > t_{180}. \end{cases}$$

With the assumption that D, $\mathbf{u}$, γ and R_2 are all constants and the initial condition $M_\perp^{\text{eff}}(\mathbf{x}, 0)$ is spatially homogeneous, the solution to (2.19) when posed on an infinite domain is given by:

$$M_\perp^{\text{eff}}(\mathbf{x}, t) = M_0 \exp\left(-R_2 t - i\mathbf{x}\cdot\mathbf{q} - D\int_0^t |\mathbf{q}|^2 \,\mathrm{d}\tau + i\mathbf{u}\cdot\int_0^t \mathbf{q}\,\mathrm{d}\tau\right), \tag{2.20}$$

where $\mathbf{q}(\tau) = \int_0^\tau \gamma \mathbf{g}^{\text{eff}}(\tau')\,\mathrm{d}\tau'$. Since the pair of diffusion gradients have the same time integral, $\mathbf{q}$ equals zero when the signal is measured at echo time, meaning that the spatial modulation is perfectly cancelled. Therefore, the image signal given by (2.20) has two amplitude attenuation components, from T2 and diffusion-induced decay respectively, and a phase shift arising from the uniform advection velocity. The effect of the diffusion gradients on the signal amplitude can be written compactly as:

$$|M_{xy}| = M_0 \exp(-R_2 t)\exp(-bD), \tag{2.21}$$

where $b = \int_0^t |\mathbf{q}|^2 \,\mathrm{d}\tau$, which has units of time per length squared, is known as the b-value of the DW-SE pulse sequence, with $b = 0$ corresponding to the normal spin-echo without diffusion weighting.

Using this model, the diffusion coefficient can be inferred by acquiring images at two or more b-values and fitting the magnitude data with an exponential curve. However, the fitted coefficient only reflects the true, physical diffusion coefficient when the model assumptions are satisfied, which is the case for homogeneous chemical samples but not for cellular biological tissues. Therefore, in clinical imaging the fitted parameter is referred to as the apparent diffusion coefficient (ADC) [65], which is dependent on many effects that violate the model assumptions, such as incoherent advection, diffusion or magnetic heterogeneity, existence of diffusion barriers etc. The ADC can also vary depending on

the b-value used to solve the inverse problem [61].

Diffusion tensor imaging (DTI) is a simple extension to the ADC concept by relaxing the isotropic assumption of the medium, whilst keeping all other assumptions including homogeneity in place. The diffusion coefficient is therefore replaced by the symmetric, second-order diffusivity tensor $\mathbf{D}$. The derivation of the model is conceptually identical to the isotropic case, with (2.20), the solution to the Bloch-Torrey equation, modified as:

$$M_{\perp}^{\text{eff}}(\mathbf{x}, t) = M_0 \exp\left(-R_2 t - i\mathbf{x}\cdot\mathbf{q} - \mathbf{D} : \int_0^t \mathbf{q}\otimes\mathbf{q}\,\mathrm{d}\tau + i\mathbf{u}\cdot\int_0^t \mathbf{q}\,\mathrm{d}\tau\right), \tag{2.22}$$

where the tensor product of two vectors is defined as $(\mathbf{v}\otimes\mathbf{w})_{ij} = v_i w_j$ and the double-dot product of two second-order tensors is defined as $\mathbf{A} : \mathbf{B} = \sum_i \sum_j A_{ij}B_{ij}$. The b-value in (2.23) is therefore replaced by the b-tensor or b-matrix: $\mathbf{b} = \int_0^t \mathbf{q}\otimes\mathbf{q}\,\mathrm{d}\tau$, resulting to the signal magnitude expression for DTI:

$$|M_{xy}| = M_0 \exp(-R_2 t)\exp(-\mathbf{b} : \mathbf{D}). \tag{2.23}$$

In a three-dimensional domain, the diffusivity tensor has six degrees of freedom. Therefore, together with the degree of freedom of the T2 time constant, they can be uniquely determined by seven MR images with different encoding directions by the diffusion gradients [12]. If the sequence of diffusion gradients is designed such that the b-matrix is isotropic: $\mathbf{b} = b\mathbf{I}$, then (2.23) is only sensitive to the trace of the diffusivity tensor. This allows the apparent diffusion trace ($\operatorname{tr}\mathbf{D}$), an empirical measure which may be clinically useful, to be evaluated from as few as two MR images taken with different isotropic b-matrices.

Qualitatively, the apparent diffusivity coefficient or tensor derived from the continuum approach relates to the translational mobility of the spins within the medium. However, these apparent quantities can be interpreted as a true physical diffusivity only when the model assumptions are satisfied. The quantitative relation between the observed apparent diffusivity and the physical properties of a heterogeneous medium, such as its microstructure geometry or local diffusivity variations, requires further modelling with

relaxed assumptions to reflect the physical reality. For the restricted diffusion behaviour in interconnected porous domains such as the extracellular tissue space, the concept of tortuosity has been proposed in the literature [60, 79, 81, 92]. However, there is a lack of mathematical justification of the tortuosity factor, which is usually arbitrarily defined as the ratio between the actual path length and the net displacement of a particle. Little is also mentioned about the assumptions required for its valid use in modelling. Meanwhile, the recently proposed effective medium theory [84] is an attempt to mathematically relate the ADC with local diffusivity variations, although it only considers a very specific scenario when the local diffusivity is slightly perturbed from the mean diffusion coefficient, rather than the generalised case with considerable differences in diffusivity or the existence of impermeable boundaries within the microstructure. There is therefore a gap in understanding of the mathematical relations between general microstructure properties and observed apparent parameters and their required assumptions. These will be addressed in the work of this thesis, which aims to derive mathematically these apparent models using a set of assumptions which are as general as possible, while reflecting the physical reality of porous media such as biological tissue.

2.4.3 Multi-compartmental modelling

The multi-compartmental approach is widespread in literature for modelling substance transport in biological tissue, for example water, oxygen or drug molecules, by assigning compartment-specific properties to the different physical environments of intracellular, extracellular or vascular spaces. For diffusion MRI, a common approach is a two-compartment model which modifies the diffusion-induced attenuation in (2.21) to a bi-exponential function [78, 80, 101]:

$$|M_{xy}| = M_0 e^{-R_2 t}(f e^{-bD_1} + (1-f)e^{-bD_2}), \tag{2.24}$$

where D_1, D_2 are different apparent diffusion coefficients (ADCs) and f is the volume fraction. This can be extended to n compartments in the form:

$$|M_{xy}| = M_0 e^{-R_2 t} \sum_{i=1}^{n} f_i e^{-bD_i}, \tag{2.25}$$

where f_i is the volume fraction of the i^{th} compartment totalling to one.

The theoretical justification of multi-compartmental models has been controversial and is the subject of much debate in literature. The compartmentalisation is likely to be due to microstructure boundaries in the medium, as almost all ^{1}H nuclei in biological tissue exist as water molecules, so it is unlikely for multiple species of ^{1}H nuclei which have different diffusion coefficients to coexist. Therefore, multi-compartmental models face the same challenges as apparent continuum models, in terms of the quantitative relation between microstructure boundaries and observed apparent diffusion coefficients, with the additional complexity of possible coupling between these compartments. For the two-compartment model, the 'slower' and 'faster' diffusion compartments have been attributed to uncoupled intra- and extracellular spaces by some authors [30], although this has been contested by a number of experimental investigations. For example the volume fractions fitted from this approach are found to be approximately 0.25 and 0.75 for the slow and fast compartments respectively, disagreeing with accepted volume fractions of intra- and extracellular water which are pretty much the opposite way round [2]. Moreover, the ADCs of the two compartments are observed to be substantially different, although the physical diffusion coefficient are measured to be of similar magnitude using compartment-specific techniques [34, 103], suggesting a dominant role of microstructure boundaries rather than the biochemical environment in reducing the ADCs. The mathematical explanation of the observed ADCs and volume fractions, based on the physical anatomy of biological tissues, remains an open question.

2.4.4 Q-space imaging and the concept of propagators

The concept of q-space imaging (QSI) has been proposed as another way to characterise signal attenuation behaviour observed in porous media [21, 73, 99], which cannot be described by the mono-exponential decay of expressions (2.21) or (2.23). Central to the q-space formalism is the assumption that diffusion gradients are strong but instantaneous, in other words they are modelled as Dirac delta functions. This clearly is an unphysical assumption, as it requires an infinite strength of gradient magnetic fields; in practice, the limitations of gradient coils hardware means the diffusion gradients take the form of finite lobes instead. There has been limited literature on the effect and validity of the short pulse assumption in practice. Even though some authors claim to use this formalism to model generalised gradients [18, 97], they are still modelled as a train of delta functions. Such pulse trains can be thought of as many two-pulse sequences patched together, so the modelling principles for the typical two-pulse gradients apply.

In terms of the mathematical principle of q-space imaging, the q-vector is defined as the wave-vector of the spatial modulation of spins:

$$\mathbf{q}(t) = \int_0^t \gamma \mathbf{g}^{\text{eff}}(\tau)\, \mathrm{d}\tau. \tag{2.26}$$

Since the narrow pulse spatially modulate or de-modulate the phase of spins instantaneously, $\mathbf{q}(t)$ is a piecewise constant function of time – for two-pulse PGSE it steps up from zero to a constant during the first pulse and steps back down to zero during the second. In the literature, the notation for $\mathbf{q}(t)$ sometimes varies by a factor of 2π if the q-vector is defined as the spatial frequency instead of the wave-vector; throughout this thesis the wave-vector convention is used. The pulse sequences can therefore be characterised by the wave-vector ($\mathbf{q}$) and the duration between the two gradient pulses (Δ), which is also known as the mixing or diffusion time. This means that if the medium is magnetically homogeneous (B_0 and R_2 being constant), magnetisation precesses and relaxes during the mixing period at the same rate independent of its position. Therefore, the precession and

relaxation terms of the transverse Bloch-Torrey equation (2.18) become separable from the transport terms, effectively reducing it to a classical diffusion-advection problem.

In the literature, q-space imaging data are most often modelled and interpreted by the propagator formalism [31]. The transport propagator essentially is the probability distribution of a stochastic transport process given an initial position, or the Green's function (also known as the fundamental solution) of the diffusion-advection equation with some boundary conditions; as explained in Section 2.2.3 the two interpretations are equivalent. Define $P(\mathbf{x}, \mathbf{x}_0, t - t_0)$ as the propagator, or Green's function, of the following diffusion-advection equation with some boundary conditions:

$$\frac{\partial c}{\partial t} = \nabla \cdot (\mathbf{D}\nabla c - \mathbf{u}\, c), \tag{2.27}$$

where $\mathbf{x}_0$, t_0, $\mathbf{x}$, t are the initial and final position and time variables respectively and $t > t_0$. When $t = t_0$ the Green's function is the Dirac delta function i.e. $P(\mathbf{x}, \mathbf{x}_0, 0) = \delta(\mathbf{x} - \mathbf{x}_0)$. Since the diffusion-advection problem is time-invariant, its Green's function is a function of the time difference $(t - t_0)$ rather than both t and t_0. If the domain is infinite and homogeneous in space, then the propagator is a scaled Gaussian function which only depends on $(\mathbf{x} - \mathbf{x}_0)$ and $(t - t_0)$, while if the domain has boundaries or the equation is heterogeneous in space, then it depends on both $\mathbf{x}$ and $\mathbf{x}_0$. For the diffusion-advection behaviour of spins during the mixing period, the initial condition is given by:

$$M_{xy}(\mathbf{x}, t_0) = M_0(\mathbf{x}) \exp(-i\mathbf{q} \cdot \mathbf{x}) \tag{2.28}$$

where $M_0(\mathbf{x})$ is the complex transverse magnetisation density monetarily before the instantaneous modulation by the first gradient pulse. The magnetisation density after the mixing time and just before the demodulation gradient pulse ($t = t_0 + \Delta^-$) is then given by a convolution of the transport propagator with the initial condition, multiplied by the coherent precession and relaxation:

$$M_{xy}(\mathbf{x}, t_0 + \Delta^-) = \exp(-i\gamma B_0 \Delta - R_2 \Delta) \int M_0(\mathbf{x}_0) \exp(-i\mathbf{q} \cdot \mathbf{x}_0)\, P(\mathbf{x}, \mathbf{x}_0, \Delta)\, \mathrm{d}\mathbf{x}_0. \tag{2.29}$$

Momentarily after the instantaneous demodulation by the second gradient pulse ($t = t_0 + \Delta^+$), the magnetisation density becomes:

$$\begin{aligned} M_{xy}(\mathbf{x}, t_0 + \Delta^+) &= \exp(i\mathbf{q}\cdot\mathbf{x})\, M_{xy}(\mathbf{x}, t_0 + \Delta^-) \\ &= \exp(-i\gamma B_0 \Delta - R_2\Delta) \int M_0(\mathbf{x}_0) \exp(i\mathbf{q}\cdot(\mathbf{x}-\mathbf{x}_0))\, P(\mathbf{x}, \mathbf{x}_0, \Delta)\, \mathrm{d}\mathbf{x}_0. \end{aligned} \tag{2.30}$$

The concept of an averaged propagator has been proposed to interpret the signal attenuation, initially in the spectroscopy context [32, 53]. However, the mathematical basis of such averaging is somewhat unclear. The transport-induced signal attenuation for a particular $(\mathbf{q}, \Delta)$ is given by:

$$E(\mathbf{q}, \Delta) = \iint M_0(\mathbf{x}_0) \exp(i\mathbf{q}\cdot(\mathbf{x}-\mathbf{x}_0))\, P(\mathbf{x}, \mathbf{x}_0, \Delta)\, \mathrm{d}\mathbf{x}_0 \mathrm{d}\mathbf{x}, \tag{2.31}$$

which is the integral of the transport part in (2.30) over a volume. For spectroscopy, which has no spatial resolution, both $\mathbf{x}$ and $\mathbf{x}_0$ are integrated over the entire domain; however for imaging $\mathbf{x}$ should be integrated over a voxel instead. The proposed averaged propagator involves a change of variable from $(\mathbf{x}, \mathbf{x}_0)$ to $(\mathbf{R} = \mathbf{x} - \mathbf{x}_0, \mathbf{x})$:

$$\overline{P}(\mathbf{R}, \Delta) = \int M_0(\mathbf{x}_0) P(\mathbf{R}, \mathbf{x}_0, \Delta)\, \mathrm{d}\mathbf{x}_0, \tag{2.32}$$

so (2.31) becomes [21, 53]:

$$E(\mathbf{q}, \Delta) = \int \overline{P}(\mathbf{R}, \Delta) \exp(i\mathbf{q}\cdot\mathbf{R})\, \mathrm{d}\mathbf{R}. \tag{2.33}$$

A problem arises in this expression as the domain of integration for the displacement variable $\mathbf{R}$ is only well defined when both $\mathbf{x}$ and $\mathbf{x}_0$ are variables over a domain without boundaries, for example an infinite domain. However, the q-spaced approach is most often used to probe heterogeneous media with irregular boundaries, in which case the justification of this formalism is unclear, and likely requires some non-trivial scale separation arguments to allow a valid approximation by an infinite domain.

An alternative approach to make sense of (2.33) is to interpret $\overline{P}(\mathbf{R}, \Delta)$ as the Green's

function of an apparent space-invariant continuum model, rather than an averaged propagator of a heterogeneous domain, so it is well defined by the displacement variable $\mathbf{R}$ irrespective of its location in the domain. This effectively uses the assumptions for ADC and DTI models discussed in Section 2.4.2, with such equivalence recognised by Basser [9]. In DTI, the transport behaviour of the apparent continuum is modelled using the second-order diffusion-advection equation, so the transport propagator is restricted to be a shifted and scaled Gaussian function with its second moment tensor (covariance) equating to $2\mathbf{D}^{\text{app}}\Delta$. This restriction can be relaxed to allow more general forms of Green's functions in the model, for example a non-zero fourth moment (kurtosis) can be allowed, equivalent to modelling the apparent continuum with a fourth-order differential equation. This method is known as diffusional kurtosis imaging (DKI) [51], with some suggestions that the apparent kurtosis may be a clinically useful biomarker [25, 58]. In theory, the propagator can take the form of any function in some function space if an infinite number of moments are considered, with the truncation at the n^{th}-order moment corresponding to an n^{th}-order apparent continuum model. This can be approached using cumulant expansions, which some authors claim to be a universal mathematical description of diffusion-weighted signals [56, 86]. However, as explained above this approach does not hold for heterogeneous media because of the fundamental assumption of an apparent space-invariant continuum; as for ADC and DTI models the mathematical relation between these apparent parameters and the physical heterogeneities remains unresolved.

To summarise q-space imaging, it is a formalism that characterises pulsed gradient spin-echo (PGSE) in magnetically homogeneous media by the modulation wave-vector ($\mathbf{q}$) and mixing time (Δ), with the assumption that diffusion gradients are very narrow pulses modelled as delta functions. The modelling of QSI in current literature is through the notion of the transport propagator, which is the Green's function of the diffusion-advection equation with appropriate boundary conditions. However, in heterogeneous domains the mathematical treatment of the propagator can be cumbersome, or involves non-trivial assumptions that are yet to be clarified.

2.4.5 "Diffusion-diffraction" modelling

A closely related modelling concept to propagators in q-space imaging is the apparent phenomenon of "diffusion-diffraction". When there is no advection, the diffusion propagator in heterogeneous media has also been modelled explicitly as a function of both $\mathbf{x}$ and $\mathbf{x}_0$ without using the displacement variable [17, 91], through a standard eigenmode decomposition of the Green's function of the diffusion equation [88]:

$$P(\mathbf{x}, \mathbf{x}_0, t - t_0) = \sum_{m=0}^{\infty} e^{-\lambda_m (t-t_0)} \phi_m(\mathbf{x}) \phi_m(\mathbf{x}_0) \tag{2.34}$$

where (ϕ_m, λ_m) solve the Helmholtz eigenvalue problem under some boundary conditions:

$$\nabla^2 \phi_m + \lambda_m \phi_m = 0, \tag{2.35}$$

with $\lambda_0 \leq \lambda_1 \leq \lambda_2 \leq \cdots$. The non-zero eigenfunctions ϕ_m are respectively scaled such that:

$$\delta(\mathbf{x} - \mathbf{x}_0) = \sum_{m=0}^{\infty} \phi_m(\mathbf{x}) \phi_m(\mathbf{x}_0). \tag{2.36}$$

Since the Helmholtz equation (2.35) is also used to solve for steady-state wave propagation, in MRI literature this is referred to as 'diffusion-diffraction' despite having no physical connection to diffraction phenomena. However, since there are infinitely many exponentially-decaying eigenmodes in (2.34), this approach is only insightful at the long time limit ($t \to \infty$) which reduces the propagator to the dominant mode of (ϕ_0, λ_0). For pure Neumann problems in a bounded domain $\lambda_0 = 0$ and ϕ_m is constant, reflecting the non-decaying steady state governed by material conservation. This greatly simplifies the calculations – indeed this is the setting usually considered in MRI literature. However, for other boundary conditions λ_0 is not necessarily zero, making calculations cumbersome with the additional complexity of selecting a timescale long enough for the $m = 0$ mode to be dominant, but short enough so it does not decay away and becomes effectively undetectable.

2.4.6 Apparent time and frequency dependency of diffusion

In the more recent literature on diffusion MRI, the notion of diffusion in heterogeneous media being 'time-' or 'frequency-dependent' has gained significant popularity. Initially proposed by Mitra et al. [77] and Stepišnik [109] respectively, these two interrelated formalisms have since been extensively used in experimental and numerical studies of diffusion MRI [5, 57, 68, 82, 85, 119, 120]. It must be stressed, however, that the temporal dependency is *apparent*, as it refers to the dependency of the apparent diffusion coefficient or tensor on the observation timescale, rather than the physical diffusion behaviour in the medium not being time-invariant.

In the original description of the formalism, Mitra et al. [77] used an averaged propagator ansatz to model diffusion in porous media, and reasoned that while the second moment tensor of the diffusion propagator is equal to $2\mathbf{D}t$ in an infinite homogeneous medium, this linear relationship with time does not necessarily hold when the medium is heterogeneous. Therefore, the observed apparent diffusion tensor varies depending on the observation timescale. In particular, they suggested that at short timescales, the apparent diffusion coefficient is sensitive to the surface-to-volume ratio, while at long timescales the connectivity and tortuosity of the pore space is encapsulated in an asymptotic diffusion tensor.

Mitra et al. [77] defined the time-dependent apparent diffusion coefficient as:

$$D_{\text{app}}(t) = \frac{\langle |\mathbf{r}(t) - \mathbf{r}(0)|^2 \rangle}{2n_d t}, \tag{2.37}$$

where $n_d = 3$ is the number of spatial dimensions. They suggested that at the $t \to 0$ short-time limit, the quantity of (2.37) asymptotically equates to:

$$D_{\text{app}}(t) = D_0 \left(1 - \frac{4}{3n_d\sqrt{\pi}} \frac{|S|}{|V|} \sqrt{D_0 t} + O(D_0 t) \right), \tag{2.38}$$

showing its dependence on the $|S|/|V|$ surface-to-volume ratio of the porous medium. Equation (2.38) is frequently cited in MRI literature and is considered a standard and

universal expression for the short-time regime [19, 43, 87].

Meanwhile, the apparent 'diffusion frequency spectrum' formalism applies to oscillatory gradient spin-echo (OGSE) methods, which use the timescale of the oscillatory gradients to probe the apparent time-dependency of diffusion. The diffusion frequency spectrum is usually defined in a stochastic formulation by the Fourier transform of the velocity autocorrelation function in time [19, 109]:

$$\mathbf{D}(\omega) = \frac{1}{2}\int_{-\infty}^{\infty} \langle \mathbf{v}(t) \otimes \mathbf{v}(0) \rangle \, e^{i\omega t} \mathrm{d}t, \tag{2.39}$$

where $\mathbf{v}(t)$ is the velocity probability distribution. However, the two-sided Fourier transform in this definition is ill-defined due to the time-irreversible nature of diffusion processes, since the integrand of (2.39) is undefined for negative time with no information about any behaviour before the specified initial condition at $t = 0$. This problem can be avoided by multiplying the autocorrelation function with a step function [85], effectively reducing (2.39) to a one-sided Fourier transform for positive time, in which case the transport-induced attenuation is suggested to be:

$$E(T) = \exp\left(-\frac{T^2}{2\pi}\int_{-\infty}^{\infty} \mathbf{q}(-\omega) \cdot \mathbf{D}(\omega)\mathbf{q}(\omega)\, \mathrm{d}\omega\right), \tag{2.40}$$

where $\mathbf{q}(\omega) = \frac{1}{T}\int_0^T \mathbf{q}(t)\exp(i\omega t)\, \mathrm{d}t$ is the approximate frequency spectrum in time of $\mathbf{q}(t)$, the spatial modulation wave-vector defined in (2.26).

In general, since the diffusion equation poses a time-irreversible initial value problem, the more appropriate transform method for the time variable is the Laplace transform, with exponentially-decaying non-equilibrium behaviour contained in the real part of the complex Laplace variable (s). The Laplace transform relates to the one-sided Fourier transform by taking the $s \to i\omega$ limit, which by the final value theorem describes the frequency-dependent behaviour of a system at the $T \to \infty$ limit. Therefore, the formalism of the diffusion frequency spectrum $\mathbf{D}(\omega)$ requires implicitly a time scale long enough for non-equilibrium initial conditions to decay away, similar to the 'diffusion-diffraction' analysis described previously.

2.4.7 Modelling advection in microstructure

As discussed in Section 2.3.2, diffusion MRI is sensitive to any spatial transport of magnetisation, including both diffusion and advection. When the advection velocity is constant in space, it causes a coherent phase shift to the transverse magnetisation signal as shown in Section 2.4.2. However, if the advection velocity is incoherent within a voxel, there is a spread in phase shift experienced by spins at different position. Therefore, in an analogous way to T2* decay (see Section 2.2.2), additional attenuation is observed when the signal density is integrated over the voxel, leading to the terminology of 'intra-voxel incoherent motion' (IVIM) proposed by Le Bihan et al. [65, 66].

In the work of Le Bihan et al. [66], a two-compartment model has been proposed to quantify the effect of microvascular perfusion in DW-SE images, in which the blood flow in the capillary bed is modelled as an effective random walk. The underlying physical assumption of this approach is the vascular compartment being a collection of capillaries, which are chains of straight segments of random lengths and orientations (Fig 2.7). Each capillary chain contains a plug flow of constant speed and is assumed to not interact with other capillaries. Under these conditions, the translational motion of spins in the vascular compartment can be modelled as a pseudo-diffusion process, while the extravascular compartment is assumed to be an isotropic apparent continuum with no advection. The overall IVIM attenuation is then described by the bi-exponential model of (2.24):

$$|M_{xy}| = M_0 e^{-R_2 t}(f e^{-bD^*} + (1-f)e^{-bD}), \tag{2.41}$$

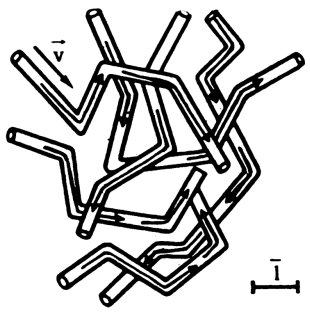

Figure 2.7: Microvascular geometry in the IVIM perfusion model, reproduced from [66].

where f and D^* are the volume fraction and pseudo-diffusion coefficient of the vascular compartment respectively. However, contrary to the physical assumption of the IVIM perfusion model, more recent studies on the capillary bed using microscopy has revealed an interconnected mesh-like structure [24], leading to the suggestion that microcirculation behaves as a pressure-driven flow through a porous medium [36]. Coincidentally, the phantom study conducted in the same work of Le Bihan et al. used a gravity-driven flow through a porous chromatographic column, which consisted of packed hydrogel microspheres, to validate the IVIM perfusion model. When the bi-exponential model of (2.41) is applied to their phantom data, the fitted perfusion volume fraction was reported to increase from 4% to 7% when the flow rate changed only slightly from 3.7 to 4.2 mL/min, inconsistent with the physical volume fraction of the porous space which is expected to be constant. Questions are therefore raised over the validity of the IVIM perfusion model for pressure-driven flow in interconnected porous media, similar to the microvasculature of biological tissues.

2.4.8 Effects of magnetic susceptibility heterogeneities

In all of the above models, the resultant longitudinal magnetic field (B_z) from the static (B_0) and gradient ($\mathbf{g} \cdot \mathbf{x}$) fields has been assumed to bear a perfectly linear profile within the medium. However, in biological tissues local magnetic susceptibility variations, or even discontinuities, may exist, leading to spatially-varying perturbations to the static B_0 field. In particular, these magnetic field perturbations occur at multiple length scales. At the bulk tissue level, this can occur at tissue-air interfaces, where there is a discontinuous change in magnetic susceptibility between the two media. Meanwhile, at the microstructure level, certain cellular structures may also contain substantial amounts of weakly magnetic substances. Notably, in the brain where diffusion MRI finds most clinical applications, the myelin sheath of white matter is very lipid-heavy, whereas deoxygenated red blood cells contain deoxyhaemoglobin; both of these substances are diamagnetic and cause magnetic field perturbations [89, 121].

It has been suggested that such local magnetic field fluctuations effectively act as local diffusion gradients [29], which can couple with magnetisation transport and contribute to spin-echo signal attenuation as well. The models in the existing literature on such local field perturbation effects have focused on the case when the globally-applied diffusion gradient is absent [50, 55, 83], although it has been suggested that transport-induced attenuation components from local and global gradient fields are coupled together [19]. The general mathematical description of such cross-coupled behaviour in diffusion MRI, in particular when microstructure transport barriers are also present, is under-explored.

2.5 Concluding remarks

To summarise this chapter, the basic principles of diffusion-weighted MRI and its current developments in modelling methods have been presented. At length scales relevant to biological tissues, the behaviour of nuclear magnetic resonance (NMR) in diffusive-advective media can be described by the extended Bloch-Torrey equations (2.8)-(2.10) and appropriate boundary conditions. The diffusion-weighted spin-echo (DW-SE) pulse sequence, which utilises a pair of strong gradient fields to spatially modulate and demodulate the phase of transverse magnetisation, is a common method to introduce sensitivity to transport of spins. Given the knowledge of the physical properties of the imaged medium, the resulting DW-SE signal can be calculated by solving the Bloch-Torrey equations as a forward problem. However, quantitative imaging is an inverse problem which is ill-posed in general. There is therefore a need for modelling to characterise the unknown imaged medium with a finite number of parameters, which can be obtained through regression of DW-SE measurements.

A number of DW-SE modelling methods for heterogeneous media are found in the literature and have been reviewed in Section 2.4. These models come in various levels of complexity, although there is a lack of rigorous mathematical or physical justifications. In particular, the majority of the models involve a common assumption of an apparent

homogeneous continuum, however the validity of these apparent continuum parameters and their mathematical relations to the physical heterogeneous medium remain unclear. Most models have also ignored the effect of magnetic susceptibility heterogeneities, which may cross-couple into the signal attenuation weighting from the DW-SE gradients. There is currently a gap in the literature in the understanding of diffusion MRI in general heterogeneous media and how they can be represented by mathematically justified models. This thesis therefore aims to address this gap by considering DW-SE in a general setting and systematically study the Bloch-Torrey equations in certain asymptotic regimes.

Chapter 3

Relevant scales and parameter space of DW-SE

3.1 Introduction

This chapter seeks to identify the relevant scales which characterise a diffusion-weighted spin-echo (DW-SE) MRI experiment. The motivation of doing so is to deduce a parameter space which show the different possible operating regimes of DW-SE, based on the relative strengths of physical processes. This is particularly important for the subsequent asymptotic analysis, as the parameter space gives a clear definition to the asymptotic limits being considered, as well as their physical interpretation. These will lay the foundation for the content in the rest of the thesis.

Prior to the definition of the parameter space, a novel interpretation of spin-echo MRI based on a distribution function formulation is first proposed in Section 3.2. Even though this formulation is not strictly necessary for the subsequent work in this thesis, it provides an easy, intuitive insight to the mechanism of DW-SE, so its relevant scales and parameters can be easily explained. This leads to the identification of the key parameters characterising DW-SE in Section 3.3, namely the instantaneous modulation wave-vector ($\mathbf{q}$) and the time duration (t). The comparison of these parameters with the physical

properties of an imaged medium, i.e. the characteristic diffusivity and microstructure length, thus defines the operating regime of a DW-SE pulse sequence. Then, a brief review of the literature is conducted to associate existing experimental and clinical studies to the locations within the q-t parameter space.

The q-t characterisation of diffusion MRI itself is not novel, with a similar concept having been proposed by [39] with more detailed discussions in [38]. However, the novelty of this chapter lies with the distribution function formulation, which allows for more intuitive interpretation of DW-SE signal behaviour, and allow for easier mathematical analysis by drawing a direct analogy to the classical class of diffusion-advection problems. This treatment also leads to the concept of the $\mathbf{q}$-vector being a spatially-varying variable reflecting the local micro-structural magnetic variation, instead of being a constant vector through space, which is a novel formulation.

3.2 Distribution function formulation for spin-echo MRI

3.2.1 Transport of tracers of different phases in a volume of fluid

At a given point $\mathbf{x}$ in space and at time t in a three-dimensional volume Ω, consider $P(\mathbf{x}, t, \theta)$ to be the concentration distribution of MR tracers that has transverse magnetisation at phase θ. This distribution shall satisfy the θ-periodic condition:

$$P(\mathbf{x}, t, \theta) = P(\mathbf{x}, t, \theta + 2\pi) \text{ for all } \theta.$$

The change of P over time can be summarised by four simultaneous processes, namely (i) the advection of tracers in space by the fluid bulk flow, (ii) the diffusion of tracers in space within the fluid, (iii) Larmor precession of phase in the presence of a strong longitudinal magnetic field, and (iv) the decoherence of phase due to random processes.

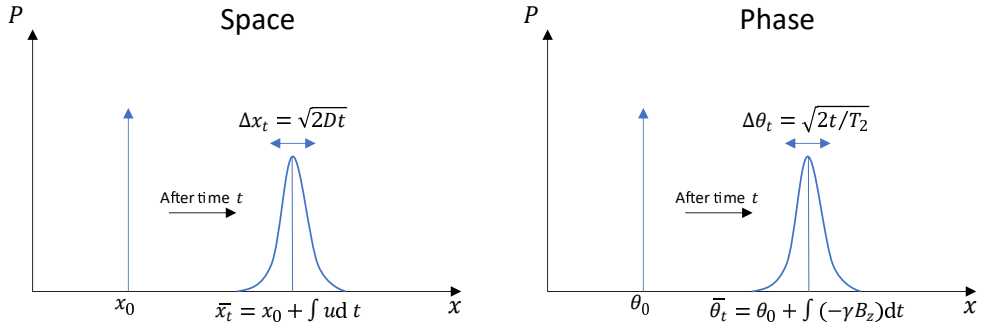

Figure 3.1: The change of a point concentration over time, i.e. the Green's function, in space or phase of an infinite homogeneous domain, showing their equivalence.

These processes are modelled individually with the following equations:

$$\frac{\partial P}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\mathbf{u}P) \qquad \text{(Advection in space)}, \tag{3.1}$$

$$\frac{\partial P}{\partial t} = \nabla_{\mathbf{x}} \cdot (\mathbf{D}\nabla_{\mathbf{x}} P) \qquad \text{(Diffusion in space)}, \tag{3.2}$$

$$\frac{\partial P}{\partial t} = \gamma B_z \frac{\partial P}{\partial \theta} \qquad \text{(Precession of phase)}, \tag{3.3}$$

$$\frac{\partial P}{\partial t} = R_2 \frac{\partial^2 P}{\partial^2 \theta} \qquad \text{(Decoherence of phase)}. \tag{3.4}$$

where $\mathbf{u}$, $\mathbf{D}$, γ, B_z and R_2 are the fluid velocity, diffusivity tensor of the tracer in the fluid, gyromagnetic ratio, longitudinal magnetic field strength and decoherence rate respectively. It is assumed that all these coefficients have no dependence on the phase θ, but that they may be heterogeneous in the space or time variables $(\mathbf{x}, t)$.

Equation (3.1) describes the flux of material being carried by the bulk fluid flow, whereas (3.2) is Fick's law of diffusion. In an infinite, homogeneous domain, these two processes combine to describe a shifted Gaussian profile of material displacement by advection and diffusion. Equations (3.3) and (3.4) are analogous to (3.1) and (3.2) respectively, describing the collective shift of phase and a Gaussian profile of decoherence over time. The equivalence between the spatial and phase equations in terms of their respective Green's functions is shown in Fig 3.1.

In the case when there is no creation or destruction of tracers, such as from radiofrequency pulses, the combined effect of the processes (i)-(iv) can be described by the fol-

lowing equation:

$$\frac{\partial P}{\partial t} = -\nabla_{\mathbf{x}} \cdot (\mathbf{u}P) + \nabla_{\mathbf{x}} \cdot (\mathbf{D}\nabla_{\mathbf{x}} P) + \gamma B_z \frac{\partial P}{\partial \theta} + R_2 \frac{\partial^2 P}{\partial^2 \theta} \text{ in } \Omega. \tag{3.5}$$

3.2.2 Consistency with equations existing in literature

The formulation of equation (3.5) is consistent with both the classical diffusion-advection equation for mass transport and the Bloch-Torrey equation. Consider the total tracer concentration at $(\mathbf{x}, t)$, which is given by:

$$C(\mathbf{x}, t) = \int_0^{2\pi} P(\mathbf{x}, t, \theta) \mathrm{d}\theta. \tag{3.6}$$

When (3.5) is integrated over a 2π period of θ, the θ-derivatives vanishes due to the periodicity condition. Since $\mathbf{D}$ and $\mathbf{u}$ are θ-independent, the order of spatial operators and the θ-integration can be interchanged. This results in the classical diffusion-advection equation:

$$\frac{\partial C}{\partial t} = \nabla_{\mathbf{x}} \cdot (\mathbf{D}\nabla_{\mathbf{x}} C - \mathbf{u}C). \tag{3.7}$$

Meanwhile, the total transverse magnetisation, which is measured by MR imaging, is the integral of all the magnetisation phasors:

$$M(\mathbf{x}, t) = \int_0^{2\pi} P(\mathbf{x}, t, \theta) e^{i\theta} \mathrm{d}\theta. \tag{3.8}$$

which is equivalent to taking the first Fourier moment of $P(\mathbf{x}, t, \theta)$ over the phase variable θ. Using integration by parts along with the θ-periodicity condition of P, the following identities hold:

$$\int_0^{2\pi} e^{i\theta} \frac{\partial P}{\partial \theta} \mathrm{d}\theta = -i \int_0^{2\pi} e^{i\theta} P \mathrm{d}\theta,$$
$$\int_0^{2\pi} e^{i\theta} \frac{\partial^2 P}{\partial \theta^2} \mathrm{d}\theta = -\int_0^{2\pi} e^{i\theta} P \mathrm{d}\theta.$$

Therefore when (3.5) is multiplied by $e^{i\theta}$ and then integrated over a 2π period of θ, using the above identities it reduces to the extended Bloch-Torrey equation for transverse

magnetisation (Chapter 2, Equation 2.11):

$$\frac{\partial M}{\partial t} = \nabla_{\mathbf{x}} \cdot (\mathbf{D}\nabla_{\mathbf{x}} M - \mathbf{u}M) - i\gamma B_z M - R_2 M. \tag{3.9}$$

3.2.3 Treatment of boundary conditions

To ensure uniqueness of the solution to the governing differential equation, boundary conditions are needed on the surface(s) that bound the domain of definition, in this case a four-dimensional volume in the $(\mathbf{x}, \theta)$ space. In the phase dimension, θ-periodicity of P implies a periodic boundary condition. Meanwhile for the spatial boundaries, it is assumed that tracers at all phases are subject to the same boundary conditions. This means the θ-independence assumption that has been made for tracer behaviour (i.e. the coefficients $\mathbf{u}$, $\mathbf{D}$, γ, B_z and R_2) shall also apply on the boundary conditions.

As an example, one simple boundary condition is the no-flux condition, where no tracers of any phase can enter or leave the volume across the boundary. This is expressed as:

$$\mathbf{n} \cdot (-\mathbf{D}\nabla_{\mathbf{x}} P + \mathbf{u}P) = 0 \text{ for all } \theta, \tag{3.10}$$

where $\mathbf{n}$ is the normal vector of the boundary surface. In the case when there is also no bulk flow across the boundary, we have $\mathbf{n} \cdot \mathbf{u} = 0$ on the boundary, reducing (3.10) to:

$$\mathbf{n} \cdot (-\mathbf{D}\nabla_{\mathbf{x}} P) = 0 \text{ for all } \theta. \tag{3.11}$$

Integrating (3.11) over a 2π period of θ using either (3.6) or (3.8), the boundary conditions for the total tracer concentration and transverse magnetisation can also be recovered:

$$\mathbf{n} \cdot (-\mathbf{D}\nabla_{\mathbf{x}} C) = 0, \tag{3.12}$$

$$\mathbf{n} \cdot (-\mathbf{D}\nabla_{\mathbf{x}} M) = 0. \tag{3.13}$$

The same treatment can be applied for other more complicated spatial boundary conditions regarding the transport of tracers across the boundary, by assuming its independence of θ.

3.2.4 Augmented phase dimension formulation

The formulation of equation (3.5) introduces an extra dimension θ to the degrees of freedom of the tracers, in addition to the three spatial dimensions of $\mathbf{x}$. In particular, the processes that occur in the spatial dimensions and phase dimension are equivalent mathematically. The first derivatives in either space or phase represent the collective movement of tracers in the respective dimensions due to an externally applied velocity or magnetic field, while the second derivatives represent the decoherence of the distribution over time in either space or phase. Therefore, (3.5) can be rewritten in a four-dimensional form by augmenting the phase variable θ as an extra dimension alongside the three spatial variables of $\mathbf{x}$, grouping the differential operators together as:

$$\frac{\partial P}{\partial t} + \nabla_{\mathbf{x},\theta} \cdot \left(- \begin{bmatrix} \mathbf{D} & 0 \\ 0 & R_2 \end{bmatrix} \nabla_{\mathbf{x},\theta} P + \begin{bmatrix} \mathbf{u} \\ -\gamma B_z \end{bmatrix} P \right) = 0 \text{ in } \Omega_{\mathbf{x},\theta}, \tag{3.14}$$

where $\Omega_{\mathbf{x},\theta}$ denotes the augmented four-dimensional volume in which the differential equation holds.

The resulting partial differential equation (3.14) has the same form as the classical diffusion-advection equation (3.7) , except that the divergence and gradient operators are now applied over all four dimensions of $\mathbf{x}$, θ rather than the usual three. The diffusion tensor $\mathbf{D}$ and velocity vector $\mathbf{u}$ are also augmented to account for the movement of tracers in the extra θ dimension.

The same augmentation formulation can also be used for boundary conditions. For example, the boundary condition of (3.10) can be written as:

$$\mathbf{n} \cdot \left(- \begin{bmatrix} \mathbf{D} & 0 \\ 0 & R_2 \end{bmatrix} \nabla_{\mathbf{x},\theta} P + \begin{bmatrix} \mathbf{u} \\ -\gamma B_z \end{bmatrix} P \right) = 0 \text{ on } \partial\Omega_{\mathbf{x},\theta}, \tag{3.15}$$

where surface $\partial\Omega_{\mathbf{x},\theta}$ is the boundary of four-dimensional volume $\Omega_{\mathbf{x},\theta}$ with normal vector $\mathbf{n}$. Since the boundary is the same for all θ, the boundary surface is always parallel to the θ-axis. Hence, the θ-component of $\mathbf{n}$ is always zero. This implies that only the

x-components of (3.15) need to be considered:

$$\mathbf{n} \cdot \left(- \begin{bmatrix} \mathbf{D} & 0 \end{bmatrix} \nabla_{\mathbf{x},\theta} P + \mathbf{u} P \right) = 0 \text{ on } \partial\Omega_{\mathbf{x}}.$$

The same form of augmentation can be applied to more complicated boundary conditions as well, as long as the θ-independence condition is satisfied.

3.2.5 Augmented dimension formulation in the spin-echo MRI context

The augmentation of an extra phase dimension onto a spatial domain, as described in (3.14) and (3.15), allows us to consider the interaction of tracers with the fluid flow and the applied magnetic field altogether in one framework. This can be interpreted as 'stacking up' identical copies of the spatial domain in an extra phase dimension θ, as shown in Figure 3.2, where for ease of visualisation the spatial domain is in 2D. Although the domain is heterogeneous in space, the augmented volume $\Omega_{\mathbf{x},\theta}$ and its associated boundaries $\partial\Omega_{\mathbf{x},\theta}$ have a structure parallel to the θ-axis due to their phase independence.

A diffusion-weighed spin-echo MRI experiment typically consists of the following sequence of events: at $t = 0$, a pulse of RF transverse magnetic field rotates longitudinal spins by 90° onto the transverse plane. The transverse magnetisation then undergoes Larmor precession due to the longitudinal B_z magnetic field, which can vary in both time and space. At $t = TE/2$, the midpoint of the echo time, another RF pulse rotates the spins by 180° about an axis on the transverse plane, thus reversing the precession phase. Finally at $t = TE$, the end of the echo time, the transverse magnetisation is imaged. The implication of these events in terms of the tracer phase distribution $P(\mathbf{x}, t, \theta)$ are explained below.

3.2.5.1 90° pulse as an initial condition

At the start of the pulse sequence, a 90° pulse is applied to rotate the spins from the longitudinal direction into into the transverse plane. In terms of transverse magnetisation,

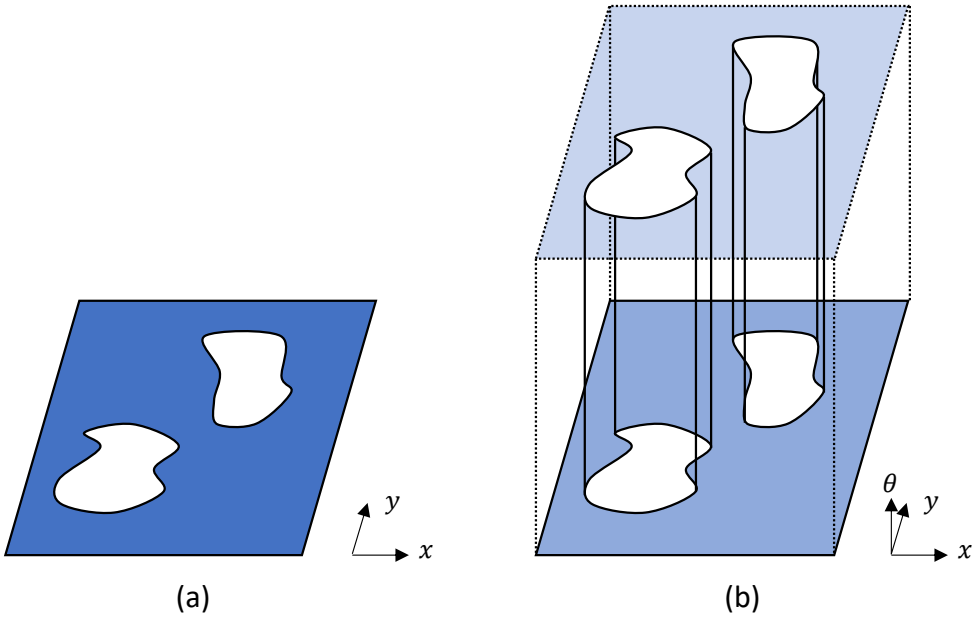

Figure 3.2: (a) A heterogeneous spatial domain consisting of volume Ω and boundaries Γ, and (b) the augmentation of the heterogeneous spatial domain in the phase dimension, showing the parallel structures. For ease of visualisation, the augmentation of a 2D domain is shown; the same concept applies for the 3D case.

this means an initial distribution of tracers which are coherent in phase are created. Therefore, the 90° pulse acts as an *initial condition* for the augmented diffusion-advection problem in the $(\mathbf{x}, \theta)$ space. Denote $M_0(\mathbf{x})$ as the magnitude of longitudinal magnetisation in space prior to the 90° pulse, and $\alpha(\mathbf{x})$ as the spatial selectivity of the 90° pulse (for an ideal non-selective pulse, $\alpha = 1$ for all $\mathbf{x}$). Without loss of generality, the initial phase of the spins can be set as zero. The initial condition for P is then given by:

$$P(\mathbf{x}, t = 0, \theta) = \sum_{n=-\infty}^{\infty} M_0(\mathbf{x})\alpha(\mathbf{x})\delta(\theta - 2n\pi), \tag{3.16}$$

where δ denotes the unit impulse distribution.

3.2.5.2 Spatially varying $\mathbf{B_z}$ field during echo time

In the absence of the rotating transverse magnetic field, the augmented diffusion-advection model of (3.14), (3.15) applies. The effect of a spatially varying B_z field can be interpreted as a "velocity" gradient advecting on P in the θ-direction, moving tracers to varying magnetisation phase depending on the location in space.

3.2.5.3 180° pulse as a mirror image operation

An ideal, instantaneous 180° RF magnetic field pulse completely reverses the phase of magnetisation. Therefore, for a 180° pulse that occurs at $t = t_{180}$, the phase distribution is flipped to its mirror image about the θ direction:

$$P(\mathbf{x}, t = t_{180}^{+}, \theta) = P(\mathbf{x}, t = t_{180}^{-}, -\theta) \tag{3.17}$$

where t_{180}^{-} and t_{180}^{+} represent the time momentarily before and after the pulse respectively.

For notational convenience, it is often easier to consider magnetisation to be continuous in time. Using the equivalent perspective of reversing the sign of B_z rather than the phase, consider the following effective quantities:

$$P^{\text{eff}}(\mathbf{x}, t, \theta) = \begin{cases} P(\mathbf{x}, t, \theta) & \text{if } 0 < t < t_{180}, \\ P(\mathbf{x}, t, -\theta) & \text{if } t > t_{180}. \end{cases} \quad B_z^{\text{eff}}(\mathbf{x}, t) = \begin{cases} B_z(\mathbf{x}, t) & \text{if } 0 < t < t_{180}, \\ -B_z(\mathbf{x}, t) & \text{if } t > t_{180}. \end{cases}$$

With P^{eff} and B_z^{eff} replacing P and B_z respectively, the augmented diffusion-advection model of (3.14), (3.15) now holds continuously over time before and after the 180° pulse. Throughout the thesis, the use of effective quantities P^{eff} and B_z^{eff} is assumed without explicit notation.

3.2.5.4 Signal measurement as a volume integral

Finally, at the end of a spin-echo MR experiment, the spatial distribution of total transverse magnetisation M is measured. However, since the imaging is done in a finite resolution, the eventual intensity of each voxel is inevitably a weighted volume average of M, with the weighting function dependent on the exact imaging method that is implemented. The signal intensity of each voxel is thus represented by the following expression:

$$I = \iiint_V w(\mathbf{x}) M(\mathbf{x}, t) \mathrm{d}V = \iiint_V \int_{-\pi}^{\pi} w(\mathbf{x}) P(\mathbf{x}, t, \theta) e^{i\theta} \mathrm{d}\theta \mathrm{d}V \tag{3.18}$$

where $w(\mathbf{x})$ is the weighting function for the corresponding voxel.

3.2.5.5 Signal attenuation induced from tracer transport

From expression (3.18), it can be seen that signal attenuation can arise from two fundamentally different modes: *pointwise* or *spatial* phase decoherence. Pointwise phase decoherence occurs when $P(\mathbf{x}, t, \theta)$ spreads out in the θ direction at each point in space; the R_2 term of the Bloch-Torrey equation is one, but not the only, origin. Since this mode of decoherence causes a pointwise reduction in magnetisation M, the resulting attenuation is independent of the voxel size and weighting function. Spatial phase decoherence, however, occurs when $P(\mathbf{x}, t, \theta)$ is not constant over a voxel, leading to loss from volume averaging when the signal is measured. This mode of decoherence can arise from unbalanced B_z field inhomogeneities, spatial variation of bulk flow, etc., and its associated signal attenuation increases with voxel size.

"Magnetic advection" by the B_z field on P *alone* is a time-reversible process. This means by applying an equal but opposite "velocity" gradient, tracers can be brought back into phase with no net loss of coherence, in the absence of spatial tracer transport. As explained above, in spin-echo MRI this is achieved by the 180° pulse in the midpoint of the echo time.

However, diffusion is a *time-irreversible* transport process. When there is a B_z field gradient, it introduces a concentration gradient of P in space, thus causing diffusive losses. Pointwise phase decoherence is enhanced through the diffusive exchange of tracers in space. Therefore, even though the spatial coherence of mean phase can be recovered eventually by the 180° pulse, it is not lossless and the measured signal has a diffusion-enhanced attenuation component. This loss mechanism is due to the decoherence in the phase dimension, and therefore is observed irrespective of spatial resolution.

Meanwhile, advective transport poses a different mode of attenuation. Although advection is time-reversible, in practice we have no control over the flow field to reverse it, unlike the effective B_z magnetic field that can exploit the phase-reversing 180° pulse. When both the flow and magnetic fields are spatially inhomogeneous, signal attenuation is caused by spatial phase decoherence. Similar to T2* decay, this signal loss is observed

only because of the lack of spatial specificity to resolve the inhomogeneous flow field.

In practice, signal attenuation is a combination of both modes of decoherence. The dominant mode is governed by the relative strengths of the diffusive and advective terms in (3.14) at the voxel length scale. However, in either mechanism the signal attenuation originates from the coupling between a B_z field gradient and spatial tracer transport. This therefore motivates the use of a coordinate system that follows the action of the spatially varying B_z field on the mean phase distribution, which is presented in the following section.

3.2.6 Transformation to co-rotating frame of reference

Define the co-rotating phase variable: $\phi = \theta + \int_0^t \gamma B_z \mathrm{d}\tau$. The coordinate system (x, ϕ) then represents the co-rotating frame of reference, as the locus of $\phi = 0$ is exactly the phase accumulation, caused by the action of B_z via the gyromagnetic ratio, at any given time t. Using the chain rule and the assumption that B_z is θ-independent, equation (3.14) becomes:

$$\frac{\partial P}{\partial t} + \nabla_{\mathbf{x},\phi} \cdot \left(- \begin{bmatrix} \mathbf{D} & \mathbf{Dq} \\ \mathbf{q}^T\mathbf{D} & \mathbf{q}^T\mathbf{Dq} + R_2 \end{bmatrix} \nabla_{\mathbf{x},\phi} P + \begin{bmatrix} \mathbf{u} \\ \mathbf{q}^T\mathbf{u} \end{bmatrix} P \right) = 0 \text{ in } \Omega_{\mathbf{x},\phi}, \tag{3.19}$$

with the no flux boundary condition along the $\mathbf{x}$-direction becoming:

$$\mathbf{n} \cdot \left(- \begin{bmatrix} \mathbf{D} & \mathbf{Dq} \end{bmatrix} \nabla_{\mathbf{x},\phi} P + \mathbf{u} P \right) = 0 \text{ on } \partial\Omega_{\mathbf{x},\phi}. \tag{3.20}$$

The vector $\mathbf{q}$ is defined as:

$$\mathbf{q}(\mathbf{x}, t) = \nabla_{\mathbf{x}} \phi = \int_0^t \nabla_{\mathbf{x}} (\gamma B_z) \mathrm{d}\tau. \tag{3.21}$$

This transformation vector can be understood by considering the time-dependent curvilinear co-ordinate system of $(\mathbf{x}, \phi)$. As shown in Figure 3.3. the co-ordinate system $(\mathbf{x}, \phi)$ keeps track of the phase accumulation at any given time t, so the local infinitesimal element $(\mathrm{d}x, \mathrm{d}\phi)$ is sheared. The transformation of each infinitesimal element in the $(\mathbf{x}, \phi)$ frame

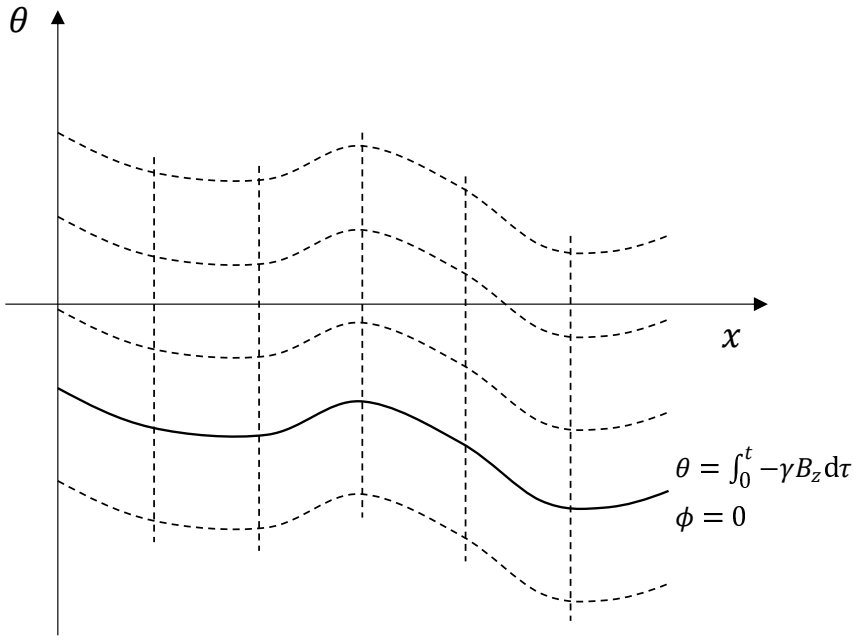

Figure 3.3: Transformation between the stationary co-ordinate sytem (x, θ) and the time-dependent curvilinear co-ordinate system (x, ϕ), which follows the phase modulation action by the magnetic field B_z. The q-vector is therefore the local gradient of the phase accumulation at time t.

back to the stationary frame $(\mathbf{x}, \theta)$ is therefore the gradient of the phase accumulation, which is the amount of shearing that has occurred.

The vector $\mathbf{q}$ defined in (3.21) is closely related to the q-vector used in the q-space imaging (QSI) modelling concept, which is reviewed in Chapter 2, Section 2.4.4, which is the reason for the choice of the same notation. The only difference, however, is that the q-vector is assumed to be constant in the q-space imaging formalism. However, in the definition of $\mathbf{q}$ here, any generic local variation of the magnetic field B_z is also allowed, hence it is possible that $\mathbf{q}$ is a function of both space and time. In terms of physical interpretation, the generalised q-vector defined here is the local, instantaneous spatial modulation wave-vector, which can be spatially varying for a generic non-uniform B_z.

In the co-rotating frame of reference, the magnetisation M_{rot} can be defined by integrating P over a period of ϕ:

$$M_{\text{rot}}(\mathbf{x}, t) = \int_0^{2\pi} P(\mathbf{x}, t, \phi) e^{i\phi} \mathrm{d}\phi = M(\mathbf{x}, t) e^{\int_0^t \gamma B_z \mathrm{d}\tau}. \tag{3.22}$$

Integrating (3.19), (3.20) over ϕ, or substituting (3.22) into (3.9), (3.13), yields:

$$\frac{\partial M_{\text{rot}}}{\partial t} + (\nabla_{\mathbf{x}} - i\mathbf{q}) \cdot (-\mathbf{D}(\nabla_{\mathbf{x}} - i\mathbf{q})M_{\text{rot}} + \mathbf{u}M_{\text{rot}}) + R_2 M_{\text{rot}} = 0 \text{ in } \Omega, \tag{3.23}$$

with no-flux boundary condition:

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla_{\mathbf{x}} - i\mathbf{q})M_{\text{rot}} + \mathbf{u}M_{\text{rot}}) = 0 \text{ on } \partial\Omega. \tag{3.24}$$

If the phase decoherence rate R_2 is constant, its contribution to magnetisation decay in (3.23), (3.24) can be separated out by introducing the co-decaying attenuation variable $S = e^{-R_2 t} M_{\text{rot}}$, resulting in the **co-rotating, co-decaying Bloch-Torrey equation**:

$$\frac{\partial S}{\partial t} + (\nabla_{\mathbf{x}} - i\mathbf{q}) \cdot (-\mathbf{D}(\nabla_{\mathbf{x}} - i\mathbf{q})S + \mathbf{u}S) = 0 \text{ in } \Omega, \tag{3.25}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla_{\mathbf{x}} - i\mathbf{q})S + \mathbf{u}S) = 0 \text{ on } \partial\Omega. \tag{3.26}$$

In theory, the co-rotating and stationary frames are equivalent at both $t = 0$ and $t = TE$, as the action of γB_z either side of the 180° pulse should perfectly cancel out, so $M_{\text{rot}} = M$ from (3.22). However, in practical diffusion MRI settings, some bulk movement of the domain Ω through the magnetic field gradient is inevitable. The phase of the measured signal is very sensitive to this motion, but since motion phase accumulates uniformly over a voxel, the signal magnitude is unaffected. The use of the co-rotating Bloch-Torrey equation to model the attenuation magnitude is therefore justified.

3.2.7 Relevant scales in the co-rotating, co-decaying formulation

In the co-rotating, co-decaying Bloch-Torrey equation, the different terms represent the various processes that are involved within the imaged medium. Therefore, the respective magnitudes of the terms specify the relative scale of these processes.

Consider a medium which has a microstructure of characteristic length scale d_c and diffusivity D_c. Given these two physical quantities, the magnitude of the $\nabla_{\mathbf{x}} \cdot \mathbf{D}\nabla_{\mathbf{x}} S$ term is set, thus defining a set of scales which can be used a benchmark to compare the other terms in the equation. These comparisons can then be summarised in a set

of dimensionless groups which define the operating regime given certain characteristic parameters.

Firstly, the time scales can be compared by considering the relative magnitude of the time-derivative term. The ratio between that and the benchmark $\nabla_{\mathbf{x}} \cdot \mathbf{D} \nabla_{\mathbf{x}} S$ term yields the first dimensionless group:

$$\pi_1 = \frac{t_c D_c}{d_c^2}. \tag{3.27}$$

In other words, the quantity $d_c^2 D_c^{-1}$ defines a time scale based on the diffusivity and geometry of the medium itself, which can be interpreted physically as the characteristic time for diffusion to occur over the microstructure length. Its comparison with the characteristic time t_c of a DW-SE pulse sequence therefore defines whether the pulse sequence is in a long- or short-time regime.

Secondly, the scale related to the phase modulation can be compared to the microstructure length scale by comparing the spatial derivatives and the $\nabla_{\mathbf{x}} \cdot \mathbf{D} \nabla_{\mathbf{x}} S$ benchmark. This yields the second dimensionless group:

$$\pi_2 = q_c d_c. \tag{3.28}$$

This means q_c^{-1} defines a length scale associated with the phase modulation of spins. Its comparison with the geometric length scale therefore defines whether the DW-SE pulse sequence is operating in a high q-value or low q-value regime.

Finally, the comparison between the advection and diffusion terms can be summarised by the third dimensionless group, which is the Péclet number:

$$\pi_3 = \mathrm{Pe} = \frac{u_c d_c}{D_c}. \tag{3.29}$$

In practice, unless perfusion from the vasculature is considered, in most biomedical applications of diffusion MRI the advection is negligible, meaning that the Péclet number is very small. Indeed, when advection is not present at all, the Péclet number simply reduces to zero and the relevant term can be ignored.

One benefit of using the co-rotating formalism to understand DW-SE is that the

scales relevant to the local signal behaviour become very clear. If we were considering the original Bloch-Torrey equation, since the longitudinal magnetic field B_z is dominated by the static field B_0, the underlying coupling between diffusion and the magnetic field gradient is somewhat masked. However, once we move into the co-rotating frame, it becomes clear that it is the local q-vector which is the relevant quantity to consider, as demonstrated by the dimensionless group above.

To conclude, the operating regime of a DW-SE pulse sequence can be summarised by the three dimensionless groups listed above, however in practice the Péclet number can effectively be ignored due to the negligible advection in most diffusion MRI applications. The other two dimensionless groups reflect the strength of the q-vector, and the time scale of the DW-SE pulse sequence. Therefore, the space of DW-SE operating regimes is defined by the q-t parameter space, which will be elaborated in the following section.

3.3 The q-t parameter space

3.3.1 Overview

As explained in the previous section, the two parameters which define the operating regime of DW-SE are the characteristic values of $|\mathbf{q}|$ and t. A parameter space can therefore be drawn out to explain the various possible regimes of DW-SE, given a medium with certain microstructure properties. This is shown in Figure 3.4 as a log-log plot, where the horizontal and vertical axes denote the scaling of time and q-value respectively. The central region denotes the regime where all terms in the co-rotating Bloch-Torrey equation are balanced, as $t_c \sim d_c^2 D_c^{-1}$ and $q_c \sim d_c^{-1}$. As we move away from the central region, some terms in the equation will start to dominate over others. This leads to some scale separation in the problem, which can be indexed by a small parameter ε. Asymptotic behaviour in these regimes can then be modelled by analysing the $\varepsilon \to 0$ limit.

3.3.2 Implication of practical considerations

So far, when presenting the q-t parameter space, there has not been any consideration as to whether all locations in the q-t parameter space can be reached in practice. This section therefore looks at the implication of some practical considerations on the reachable regimes in the parameter space.

Firstly, a major limitation to the attainable q-value of a DW-SE sequence is the strength of the gradient field that can be generated by the gradient coils. Consider a hardware system that has a maximum possible gradient strength of $|\mathbf{g}|_{\max}$. Its maximum contribution to the q-vector over a time period of t is therefore $|\mathbf{q}|_{\max} = \gamma|\mathbf{g}|_{\max}t$. Hence, the bound on the q-t parameter space due to the maximum gradient strength can be represented by a straight line on the q-t parameter space, as shown in Figure 3.5.

Secondly, another practical consideration is the signal-to-noise ratio (SNR) and the contrast-to-noise ratio (CNR) in the resulting image data. Generally, if we consider the solution to the co-rotating Bloch-Torrey equation, the final image signal intensity

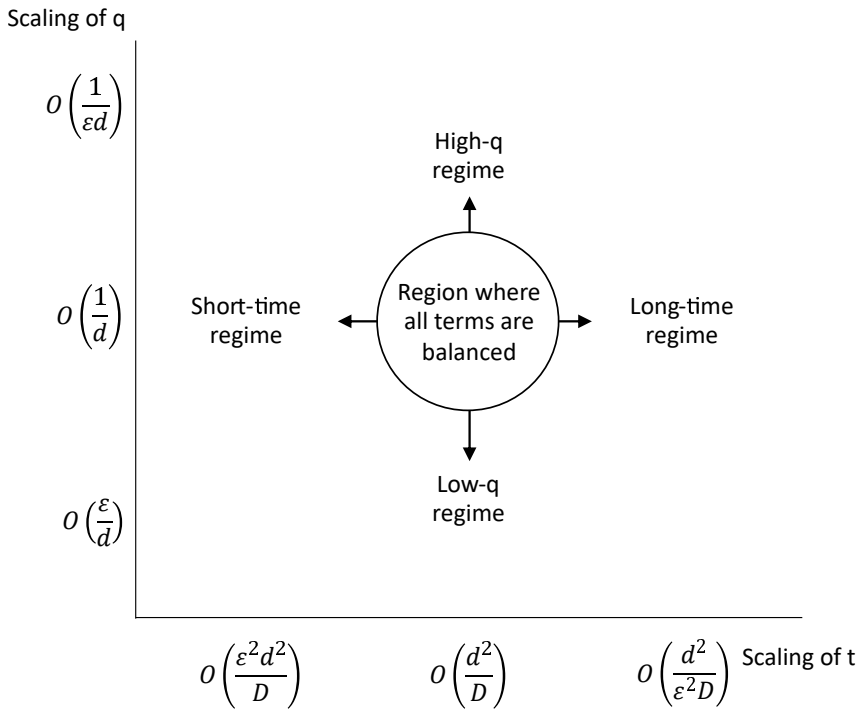

Figure 3.4: The q-t parameter space.

is expected to scale as $\exp(-|\mathbf{q}|^2 Dt)$. The implication of this is that there will be a compromise between the signal intensity and diffusion-related contrast in the DW-SE image. If $q_c^2 D_c t_c$ is small, there is better signal as it has not decayed away, but the contrast is small as all the diffusion-related information effectively ends up being a small perturbation to the signal. Meanwhile, if $q_c^2 D_c t_c$ is large, the diffusion-related contrast will be amplified, at a cost of lower signal. There may be exceptions to this in the case of isolated pores leading to restricted diffusion, which may preserve the signal intensity at longer times, however in any case the signal at long times is limited by T_2 relaxation too. As shown in Figure 3.6, there is therefore a region, in the form of a diagonal strip stretching from top-left to bottom-right through the central region, which offers both $O(1)$ signal and $O(1)$ contrast.

3.3.3 Review of experimental and clinical studies

In this section, a brief review of some diffusion MRI studies is conducted. The purpose of this review is to locate existing diffusion MRI protocols on the q-t parameter space, so that relevant regimes can be visualised and identified.

The five studies reviewed are listed in Table 3.1. Their approximate locations on the q-t parameter space, as well as their associated theoretical bounds based on the gradient strength used, are then plotted in Figure 3.7. On the q-t parameter space, the regions associated with characteristic microstructure lengths of $d = 2$ to 10 μm and diffusion coefficient of $D = 3.0$ μm^2 ms^{-1} are also shown, with the length scales chosen based on white matter axon diameters reported in [5].

Table 3.1: Details of selected diffusion MRI studies. Gradient strength units: mT m^{-1}.

MRI study	Year	Imaged medium	Study type	$\lvert\mathbf{g}\rvert_{\max}$
Le Bihan et al. [65]	1986	Human brain, *in vivo*	Experimental	4.86
Warach et al. [116]	1995	Human brain, *in vivo*	Clinical	30
Ferizi et al. [37]	2015	Human brain, *in vivo*	Experimental	285
Howard et al. [46]	2022	Macaque brain, post-mortem	Experimental	320
King et al. [54]	1994	Mouse brain, *in vivo*	Experimental	5859

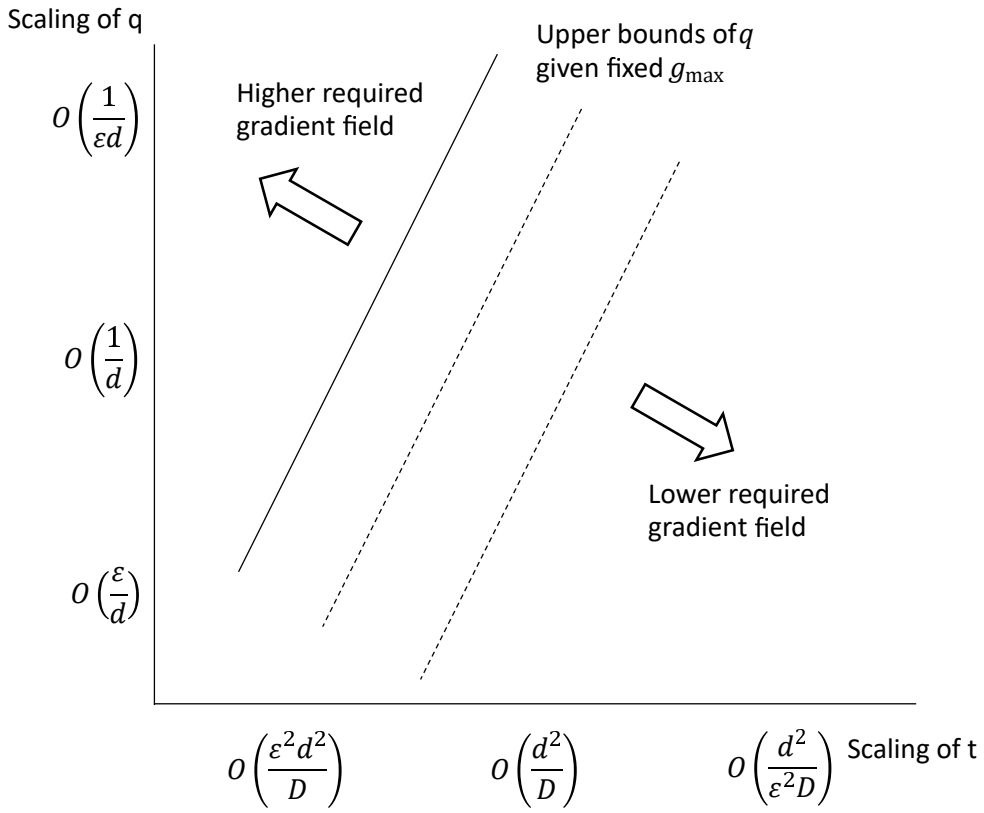

Figure 3.5: Bounds on the q-t parameter space due to maximum gradient strength.

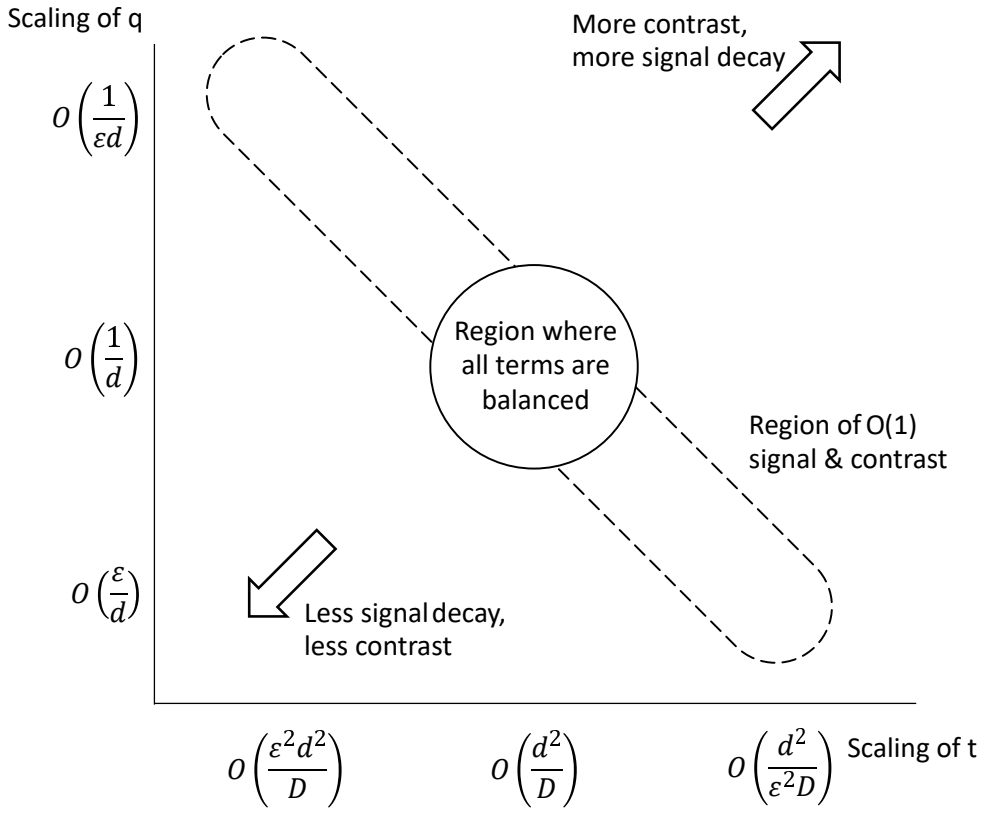

Figure 3.6: Relative signal intensity and contrast on the q-t parameter space.

The first three reviewed studies all involve diffusion MRI of the human brain *in vivo*. The first study by Le Bihan et al. [65] is one of the earliest example of diffusion MRI of biological tissue. The hardware at that time were relatively limited, therefore the study could only be limited to the long-time, low q-value regime. The second study by Warach et al. [116] is a clinical study of the use of diffusion MRI as a diagnosis tool for acute ischaemic stroke. Compared to the first study, the gradient field strength had improved by an order of magnitude, thus pushing the location in the q-t parameter space closer to the central region. This study is particularly relevant, as it is a clinical study rather than an experimental one, so it reflects the use of diffusion MRI in day-to-day clinical care. The third study by Ferizi et al. [37] is a more recent one, which takes advantage of recent hardware developments to gather vast amount of diffusion MRI data at multiple q-values and diffusion times, as shown by the scatter of points on the plot.

Meanwhile, the fourth and fifth studies are animal model studies. The fourth study by Howard et al. [46] operates with a maximum gradient strength similar to the third study, but instead on a macaque brain post-mortem. This study is unique in the sense that the same sample is used in other imaging and microscopy modalities, so a dataset relating the diffusion MRI data and the microstructure has been obtained. Finally, the study by King et al. [54] utilised ultra-high gradient field strengths to image the mouse brain *in vivo*. This is possible as the bore size for the mouse MR system is much smaller than that of a whole-body human scanner, so the gradient field strength can get to nearly twenty times that of the third study. The study only utilised a pair of very short (1 ms) gradient pulses which were spaced up to 201 ms apart, hence the location in the q-t parameter space is on the top-right. However, if the spacing between the strong gradient pulses are reduced, the gradient hardware in this study allows the short-time, high q-value regime to be reached as well, as depicted by the solid line in Figure 3.7.

To summarise these studies, with the advancement of modern hardware in human MR systems, the operating regimes of diffusion MRI are now pushing from the long-time, low q-value region into the central mixed region on the q-t parameter space. Meanwhile, in

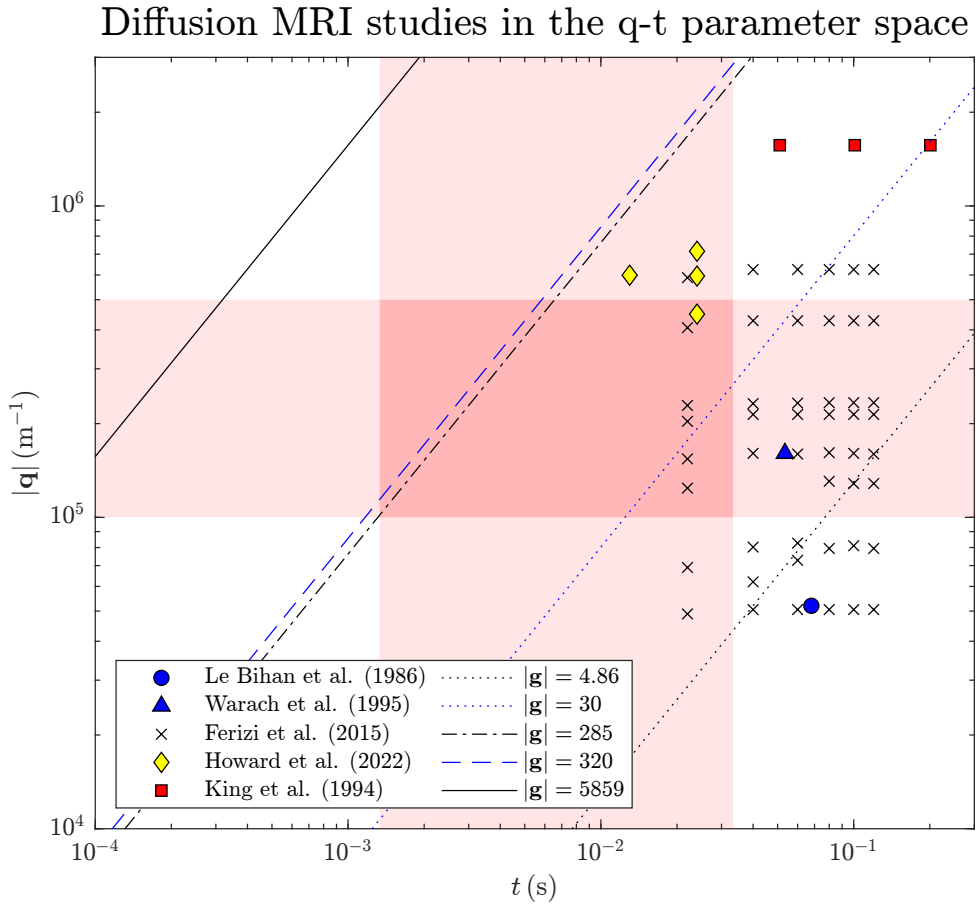

Figure 3.7: Approximate locations of selected experimental and clinical diffusion MRI studies in the q-t parameter space. Shaded regions denote the times and q-values associated with characteristic microstructure lengths of $d = 2$ to 10 μm and diffusion coefficient of $D = 3.0$ $\mu\mathrm{m}^2$ ms^{-1}. Respective lines show the theoretical bounds in the q-t space based on the maximum gradient strength (units: mT m^{-1}) used in the studies.

animal research, the smaller sizes of MR systems mean that much higher gradient field strengths are possible, therefore even the short-time, high q-value regime can potentially be reached.

3.4 Summary

In this chapter, a novel formalism of DW-SE based on the phase distribution function of spins has been proposed. By considering the change of co-ordinates in the phase space into the co-rotating frame of reference, the image signal can be described by the co-rotating, co-decaying Bloch-Torrey equation and its boundary condition (Equations

(3.25) and (3.26)), in which the action of the magnetic field is summarised in the *local* spatial modulation wave-vector. This wave-vector is a generalisation of the q-vector in existing literature, as it also accounts for possible local magnetic field variations.

Dimensional analysis of the co-rotating, co-decaying Bloch-Torrey equation then reveals three dimensionless groups which define the operating regime of DW-SE, respectively relating to the characteristic q-vector, time scale and advection strength. In practice, advection is likely to be negligible and can be ignored. Therefore, the different operating regimes can be described by regions in the q-t parameter space. Practical considerations, such as hardware limitations, and image signal and contrast, have also been considered in the context of the q-t parameter space. Finally, a brief review of the literature has shown that *in vivo* diffusion MRI of the human brain is found mainly in the long-time, small q-value regime, with a recent trend of pushing into the central mixed region. Meanwhile, the stronger gradient field strength available on animal MR systems means virtually all regimes on the q-t parameter space can be reached in *in vivo* animal research. This therefore motivates us to consider asymptotic modelling approaches for both the small-time and long-time regimes, as well as numerical simulations for the central mixed regime.

Chapter 4

Magnetic field perturbation due to susceptibility variations

4.1 Introduction

In this chapter, the perturbation to the magnetic field within a magnetically heterogeneous sample is analysed. As discussed in the literature review (Chapter 2, Section 2.4.8), biological tissue contains magnetic susceptibility variations at multiple length scales. These cause local variations in the magnetic field, which may interfere with the applied magnetic field gradient, which is essential to diffusion MRI pulse sequences. Therefore, the aim of this chapter is to systematically analyse such magnetic field perturbations using a multi-scale approach, so that the relative significance of these perturbations, depending on the operating regime, can be evaluated in the context of diffusion MRI. The work in this chapter has close relation to that of Jenkinson et al. [49], however here we extend their perturbation approach to account for magnetic susceptibility variations at both macro-scale, such as tissue-air interfaces, and micro-scale, such as around myelin sheath and other cellular structures.

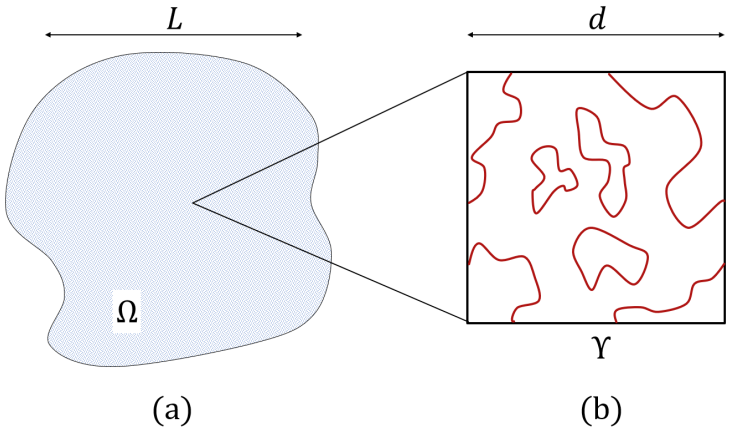

Figure 4.1: Schematic of porous medium: (a) macroscopic domain Ω, (b) periodic reference cell Υ, or the representative elementary volume (REV), over which the magnetic permeability varies.

4.2 Model formulation

4.2.1 Governing equation and boundary condition

In this model, we consider a heterogeneous, weakly magnetic sample being placed in a scanner, as shown in Figure 4.1. The sample, which is denoted by $\Omega \in \mathbb{R}^3$, is modelled to contain periodic micro-scale variations in magnetic permeability, whereas the air surrounding the sample has a uniform, isotropic permeability. The periodicity length scale d is assumed to be much smaller than the characteristic length scale of the sample L, with the ratio denoted by a small parameter $\zeta \ll 1$, thus leading to the following expression for the magnetic permeability:

$$\mu(\mathbf{x}) = \begin{cases} \mu_{\text{sample}}(\mathbf{x}/\zeta) & \text{in } \Omega, \\ \mu_{\text{air}} & \text{in } \mathbb{R}^3 \backslash \Omega, \end{cases} \tag{4.1}$$

where $\mu_{\text{sample}}(\mathbf{x}/\zeta)$ is a periodic function posed over the reference cell Υ of the microstructure (Figure 4.1b).

Assuming non-conductivity in the medium, the magnetic field is governed by the

relevant Maxwell's equations:

$$\nabla \times \mathbf{H} = 0, \tag{4.2}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{4.3}$$

with the magnetic field strength $\mathbf{H}$ related to the flux density $\mathbf{B}$ via the magnetic permeability:

$$\mathbf{B} = \mu \mathbf{H}. \tag{4.4}$$

By defining a scalar potential via $\mathbf{H} = \nabla\Phi$, the two differential equations can be combined into a single elliptic PDE:

$$\nabla \cdot (\mu \nabla \Phi) = 0. \tag{4.5}$$

The boundary condition for the governing equation (4.5) originates from the modelling of the scanner hardware which induces the magnetic field within the sample medium. Such hardware (the superconducting magnet and gradient coils) is modelled as a system acting at the far field which is decoupled from the characteristics of the medium. Therefore, a far field boundary condition for the scalar potential $\Phi \to \Phi_{\text{ff}}$ is prescribed for the equation. The hardware is calibrated such that when the sample is absent, the magnetic field generated is dominant in the z-direction, with the z-component of the magnetic flux density linearly varying in space, i.e. $B_z(\mathbf{x}) = B_0 + \mathbf{g} \cdot \mathbf{x}$ (see Section 2.3.1). This suggests that the far-field applied potential should be set as:

$$\Phi_{\text{ff}}(\mathbf{x}) = \frac{1}{\mu_{\text{air}}} \left(B_0 x_3 + g_1(x_1 x_3) + g_2(x_2 x_3) + g_3 \left(\frac{x_3^2}{2} - \frac{x_1^2 + x_2^2}{4} \right) \right), \tag{4.6}$$

where the indices $i = 1, 2, 3$ denote the three Cartesian vector components. When $\mu = \mu_{\text{air}}$ uniformly, this then satisfies Equation (4.5) itself and gives rise to the desired B_z field.

We note that since the magnetic permeability μ is expected to be discontinuous across boundaries, both between the sample and the air and within the sample across inhomogeneities, the strong derivatives in the governing equation (4.5) may not exist. Nevertheless, the problem remains well posed if the derivatives are considered in a weaker,

distributional sense, with appropriate decay conditions at the far field for $\Phi - \Phi_{\text{ff}} \to 0$.

4.2.2 Non-dimensionalisation and relevant scales

In order to represent the relative scaling between variables in the model as non-dimensional parameters, so that their asymptotic limits can be studied, the model is normalised with respect to relevant characteristic physical quantities:

$$\mathbf{x} = L\mathbf{x}^*,\ \mu = \mu_{\text{air}}\mu^*,\ \Phi = \frac{B_0 L}{\mu_{\text{air}}}\Phi^*,\ \mathbf{g} = \frac{B_0}{L}\mathbf{g}^*. \tag{4.7}$$

In this section, the asterisk notation is explicitly used to notate the non-dimensional form of variables. Substituting these into the dimensional model leads to the non-dimensional form of Equation (4.5) and its far-field boundary condition:

$$\nabla_{\mathbf{x}^*} \cdot (\mu^* \nabla_{\mathbf{x}^*} \Phi^*) = 0,\ \text{with } \Phi^* - \Phi^*_{\text{ff}} \to 0 \text{ at } |\mathbf{x}^*| \to \infty. \tag{4.8}$$

The far-field prescribed magnetic potential, in its non-dimensional form, is given by:

$$\Phi^*_{\text{ff}}(\mathbf{x}^*) = x_3^* + g_1^*(x_1^* x_3^*) + g_2^*(x_2^* x_3^*) + g_3^* \left(\frac{x_3^{*2}}{2} - \frac{x_1^{*2} + x_2^{*2}}{4} \right). \tag{4.9}$$

Since the magnetic field component generated by the gradient coils is much weaker than the static field B_0, the magnitude of the normalised gradient field is small, i.e. $g^* = |\mathbf{g}^*| \ll 1$.

For the magnetic permeability, the condition that the medium is weakly magnetic is normally denoted by a small magnetic susceptibility χ, which originates from normalising μ against the magnetic permeability of free space μ_0 through the relation $\mu = \mu_0(1 + \chi)$. However, as the far field condition for our physical model is calibrated against air rather than vacuum, here it is simpler to normalise μ against μ_{air} instead. Since air itself is also weakly magnetic ($\chi_{\text{air}} = 3.6 \times 10^{-7} \ll 1$), the difference $\mu_{\text{sample}} - \mu_{\text{air}}$ is small relative to both μ_0 and μ_{air} as well. Therefore, the non-dimensional magnetic permeability can be

written as:

$$\mu^*(\mathbf{x}^*) = \begin{cases} 1 + \eta f(\mathbf{x}^*/\zeta) & \text{in } \Omega^*, \\ 1 & \text{in } \mathbb{R}^3 \backslash \Omega^*. \end{cases} \tag{4.10}$$

Here η is the characteristic relative variation in $\mu^*(\mathbf{x}^*)$, which is small, and $f(\mathbf{x}^*/\zeta)$ is a periodic function over a non-dimensional reference cell Υ^* which is scaled to $O(1)$ in both magnitude and periodicity.

Therefore, the non-dimensionalised model is constituted by Equation (4.8), supplemented by the far-field boundary condition (4.9) and the non-dimensionalised definition of magnetic permeability (4.10). For ease of presentation, the asterisk notation for non-dimensional variables will be dropped, as we will only use non-dimensional variables from here on in the derivations, so explicit notation is no longer needed.

4.3 Asymptotic analysis

As we can see, in the non-dimensionalised problem there are three small parameters, namely ζ, η and g, which respectively represent the scale separation between the macro- and microstructure, the weak magnetic permeability variations, and the weak gradient field strength relative to background static field. Based on this model, we investigate the asymptotic behaviour of the magnetic field as these parameters tend to zero.

Since the three small parameters are independent of each other, their limits can be evaluated in any order. Here we first consider the $\zeta \to 0$ limit, which is a singular perturbation in the length scale variable, using multi-scale homogenisation. This is done by assuming a two-scale series expansion Ansatz for Φ within Ω:

$$\Phi = \Phi^{(0)}(\mathbf{x}, \mathbf{y}) + \zeta \Phi^{(1)}(\mathbf{x}, \mathbf{y}) + \zeta^2 \Phi^{(2)}(\mathbf{x}, \mathbf{y}) + \cdots . \tag{4.11}$$

In this expansion, $\mathbf{x}$ is the macro-scale variable over Ω, and $\mathbf{y} = \mathbf{x}/\zeta$ is the micro-scale variable over the scaled periodic reference cell Υ. These two spatial variables are treated as independent variables, with $\mathbf{x}$ describing the location in the macroscopic medium,

while $\mathbf{y}$ describes the precision location with respect to the periodic reference cell. To ensure uniqueness of this expansion, an additional zero-mean constraint in the micro-scale variable is required for all non-leading terms:

$$\frac{1}{|\Upsilon|}\int_{\Upsilon}\Phi^{(n)}(\mathbf{x},\mathbf{y})\,\mathrm{d}y = 0 \;\text{ for all } n \geq 1. \tag{4.12}$$

The a priori assumption of the ansatz is standard in formal homogenisation methods for elliptic problems [7, 33], and can be justified analytically using multi-scale notions of convergence [3, 28].

To derive relations which solve for the $\Phi^{(n)}$ terms, the expansion (4.11) is substituted into the governing equation (4.8), replacing the spatial derivatives with $\nabla_{\mathbf{x}} + \zeta\nabla_{\mathbf{y}}$ using the chain rule at the same time. Balancing each power of ζ then yields a cascade of elliptic problems in $\mathbf{y}$ posed over the periodic reference cell Υ, beginning with:

$$\nabla_{\mathbf{y}}\cdot(\mu\nabla_{\mathbf{y}}\Phi^{(0)}) = 0, \tag{4.13}$$

$$\nabla_{\mathbf{y}}\cdot(\mu\nabla_{\mathbf{y}}\Phi^{(1)}) = -\nabla_{\mathbf{y}}\cdot(\mu\nabla_{\mathbf{x}}\Phi^{(0)}) - \nabla_{\mathbf{x}}\cdot(\mu\nabla_{\mathbf{y}}\Phi^{(0)}), \tag{4.14}$$

$$\nabla_{\mathbf{y}}\cdot(\mu\nabla_{\mathbf{y}}\Phi^{(2)}) = -\nabla_{\mathbf{y}}\cdot(\mu\nabla_{\mathbf{x}}\Phi^{(1)}) - \nabla_{\mathbf{x}}\cdot(\mu\nabla_{\mathbf{y}}\Phi^{(1)}) - \nabla_{\mathbf{x}}\cdot(\mu\nabla_{\mathbf{x}}\Phi^{(0)}). \tag{4.15}$$

Periodicity in the leading order problem (4.13) means $\Phi^{(0)}$ is constant in the micro-scale variable $\mathbf{y}$ and is a function of only $\mathbf{x}$. Using this information and the linearity of (4.14), we can deduce that $\Phi^{(1)}$ relates to $\Phi^{(0)}$ in the form of:

$$\Phi^{(1)}(\mathbf{x},\mathbf{y}) = \mathbf{a}(\mathbf{y})\cdot\nabla_{\mathbf{x}}\Phi^{(0)}. \tag{4.16}$$

The zero-mean, periodic vector function $\mathbf{a}(\mathbf{y})$ uniquely solves the cell problem, defined over a periodic cell:

$$\nabla_{\mathbf{y}}\cdot(\mu(\mathbf{y})(\nabla_{\mathbf{y}}\mathbf{a} + \mathbf{I})) = 0, \tag{4.17}$$

with $\mathbf{I}$ denoting the identity tensor. Substituting (4.16) into the next order yields:

$$\nabla_{\mathbf{y}}\cdot(\mu\nabla_{\mathbf{y}}\Phi^{(2)}) = -\nabla_{\mathbf{y}}\cdot(\mu\nabla_{\mathbf{x}}\Phi^{(1)}) - \nabla_{\mathbf{x}}\cdot(\mu(\nabla_{\mathbf{y}}\mathbf{a} + \mathbf{I})\nabla_{\mathbf{x}}\Phi^{(0)}). \tag{4.18}$$

To ensure well-posedness of (4.18), the compatibility condition requires the integral of

the right-hand side over Υ to vanish. Using the divergence theorem on the $\Phi^{(1)}$ term, the condition simplifies to:

$$\nabla_{\mathbf{x}} \cdot \left(\int_{\Upsilon} \mu(\nabla_{\mathbf{y}} \mathbf{a} + \mathbf{I}) \, \mathrm{d}y \right) \nabla_{\mathbf{x}} \Phi^{(0)} = 0. \tag{4.19}$$

Equation (4.19) thus constitutes a macro-scale governing equation for the leading order potential $\Phi^{(0)}$ within Ω. In particular, it takes the exact same form as the original equation (4.8), but with a homogenised effective permeability instead. Therefore, by defining the effective permeability tensor as:

$$\mu_{\text{eff}} = \begin{cases} |\Upsilon|^{-1} \int_{\Upsilon} \mu(\nabla_{\mathbf{y}} \mathbf{a} + \mathbf{I}) \, \mathrm{d}y & \text{in } \Omega, \\ \mathbf{I} & \text{in } \mathbb{R}^3 \backslash \Omega, \end{cases} \tag{4.20}$$

the homogenised problem including the surrounding air and the far-field boundary condition can be written compactly as:

$$\nabla_{\mathbf{x}} \cdot (\mu_{\text{eff}} \nabla_{\mathbf{x}} \Phi^{(0)}) = 0, \quad \text{with } \Phi^{(0)} \to \Phi_{\text{ff}} \text{ at } |\mathbf{x}| \to \infty. \tag{4.21}$$

We now consider the asymptotic behaviour with respect to η for each ζ-expansion term. Since the material is weakly magnetic throughout, we expect μ_{eff} to be weakly perturbed from $\mathbf{I}$ as well. Indeed, this can be confirmed by looking at the cell problem (4.17) in terms of η. From the formulation of μ in (4.10) we have $\mu = \mathbf{I} + O(\eta)$, therefore (4.17) can be expressed as:

$$\nabla_{\mathbf{y}} \cdot (\nabla_{\mathbf{y}} \mathbf{a} + \mathbf{I} + O(\eta)(\nabla_{\mathbf{y}} \mathbf{a} + \mathbf{I})) = 0. \tag{4.22}$$

The leading-order problem in η, which is $\nabla^2_{\mathbf{y}} \mathbf{a} = 0$, has a trivial vanishing solution. Hence, it can be concluded that $\mathbf{a}(\mathbf{y}) = O(\eta)$, and as corollaries, $\Phi^{(1)} = O(\eta)$ and $\mu_{\text{eff}} = \mathbf{I} + O(\eta)$. Considering $\Delta\Phi = \Phi^{(0)} - \Phi_{\text{ff}}$, the homogenised problem (4.21) can be expanded in orders of η as:

$$\nabla_{\mathbf{x}} \cdot (\mathbf{I} + O(\eta)) \nabla_{\mathbf{x}} (\Delta\Phi) = -\nabla_{\mathbf{x}} \cdot (\mathbf{I} + O(\eta)) \nabla_{\mathbf{x}} \Phi_{\text{ff}}, \quad \text{with } \Delta\Phi \to 0 \text{ at } |\mathbf{x}| \to \infty. \tag{4.23}$$

Again, since the applied potential satisfies $\nabla_{\mathbf{x}}^2 \Phi_{\text{ff}} = 0$, the leading order of (4.23) in η becomes $\nabla_{\mathbf{x}}^2 \Delta\Phi = 0$ which has a trivial vanishing solution, hence $\Delta\Phi = O(\eta)$. Moving on to Equation (4.15), we have:

$$\nabla_{\mathbf{y}} \cdot (\mathbf{I} + O(\eta)) \nabla_{\mathbf{y}} \Phi^{(2)} = O(\eta) - \nabla_{\mathbf{x}} \cdot (\mathbf{I} + O(\eta)) \nabla_{\mathbf{x}} (\Phi_{\text{ff}} + O(\eta)). \tag{4.24}$$

Using the same argument, we can conclude that $\Phi^{(2)} = O(\eta)$. The subsequent cascade of $\Phi^{(n)}$ problems all follow the same form of (4.15) for $n \geq 2$, so all higher order terms scale as $O(\eta)$ as well.

As discussed in the previous chapter, the relevant quantity in DW-SE MRI is the instantaneous phase modulation wave-vector, or q-vector, within the sample Ω. This is given by the following expression:

$$\mathbf{q} = \gamma \int_0^t \pm \nabla B_z \mathrm{d}\tau = \gamma \int_0^t \pm \nabla (\mathbf{e}_3 \cdot \mu \nabla \Phi) \mathrm{d}\tau, \tag{4.25}$$

with $\mathbf{e}_3$ denoting the unit vector in the z-direction. The sign of the integrand depends on which side of the 180° reflection pulse of the spin-echo sequence. The multiple scales expansion in ζ for the magnetic field gradient term $\nabla(\mathbf{e}_3 \cdot \mu \nabla \Phi)$, simplified using the fact that $\Phi^{(0)}$ is constant in $\mathbf{y}$, is given by:

$$\begin{aligned} \nabla(\mathbf{e}_3 \cdot \mu \nabla \Phi) = \zeta^{-1} &\nabla_{\mathbf{y}} (\mathbf{e}_3 \cdot \mu (\nabla_{\mathbf{y}} \Phi^{(1)} + \nabla_{\mathbf{x}} \Phi^{(0)})) \\ &+ \nabla_{\mathbf{y}} (\mathbf{e}_3 \cdot \mu (\nabla_{\mathbf{y}} \Phi^{(2)} + \nabla_{\mathbf{x}} \Phi^{(1)})) \\ &+ \nabla_{\mathbf{x}} (\mathbf{e}_3 \cdot \mu (\nabla_{\mathbf{y}} \Phi^{(1)} + \nabla_{\mathbf{x}} \Phi^{(0)})) + \cdots . \end{aligned} \tag{4.26}$$

Using (4.16) and the information on the η-scaling of $\Phi^{(n)}$, the above expression in orders of η can be written as:

$$\nabla(\mathbf{e}_3 \cdot \mu \nabla \Phi) = \zeta^{-1} \nabla_{\mathbf{y}} (\mathbf{e}_3 \cdot \mu (\nabla_{\mathbf{y}} \mathbf{a} + \mathbf{I})(\nabla_{\mathbf{x}} \Phi_{\text{ff}} + O(\eta))) + \nabla_{\mathbf{x}} (\mathbf{e}_3 \cdot \nabla_{\mathbf{x}} \Phi_{\text{ff}}) + O(\eta). \tag{4.27}$$

The definition of Φ_{ff} from (4.6) gives the following:

$$\nabla_{\mathbf{x}} \Phi_{\text{ff}} = \mathbf{e}_3 + O(g), \quad \nabla_{\mathbf{x}} (\mathbf{e}_3 \cdot \nabla_{\mathbf{x}} \Phi_{\text{ff}}) = \mathbf{g}. \tag{4.28}$$

Therefore, by substituting the above into (4.27), the expression reduces to:

$$\nabla(\mathbf{e}_3 \cdot \mu\nabla\Phi) = \zeta^{-1}\nabla_{\mathbf{y}}(\mu(\nabla_{\mathbf{y}}\mathbf{a} + \mathbf{I}))_{33} + \mathbf{g} + \text{ H.O.T. in } \Omega. \tag{4.29}$$

If we consider the first term in this expression, since $\mu = \mathbf{I} + O(\eta)$ and $\mathbf{a} = O(\eta)$, it can be concluded that the tensor $\nabla_{\mathbf{y}}(\mu(\nabla_{\mathbf{y}}\mathbf{a}+\mathbf{I}))$, which is dependent solely on the microstructure via the cell problem solution, scales as $O(\eta)$. This leads to:

$$\nabla(\mathbf{e}_3 \cdot \mu\nabla\Phi) = \zeta^{-1}\eta\,\mathbf{b} + \mathbf{g} + \text{ H.O.T. in } \Omega, \tag{4.30}$$

with the microstructure information contained in the vector $\mathbf{b}(\mathbf{y}) = \eta^{-1}\nabla_{\mathbf{y}}(\mu(\nabla_{\mathbf{y}}\mathbf{a}+\mathbf{I}))_{33}$ which has been normalised to $O(1)$.

It is now clear that the two leading order terms of the local magnetic field gradient scale as $O(\zeta^{-1}\eta)$ and $O(g)$ respectively. To consider this back in physical terms, the dimensional form of expression (4.30) is given by:

$$\mathbf{g}_{\text{loc}}(\mathbf{x}, \mathbf{y}, t) = d^{-1}B_0\,\eta\mathbf{b}(\mathbf{y}) + \mathbf{g} + \text{ H.O.T.} \tag{4.31}$$

Consequently, if the local q-vector is considered, it is given by the following expression in dimensional quantities:

$$\mathbf{q}(\mathbf{x}, \mathbf{y}, t) = d^{-1}\int_0^t \pm\gamma B_0\,\eta\mathbf{b}(\mathbf{y})\,\mathrm{d}\tau + \mathbf{q}_{\text{app}} + \text{ H.O.T.}, \tag{4.32}$$

where $\mathbf{q}_{\text{app}} = \int_0^t \pm\gamma\mathbf{g}\,\mathrm{d}\tau$ is the "applied" q-vector assuming a homogeneous medium.

In this expression for the local q-vector, there is a perturbation term due to the microstructure-level susceptibility variation, which induces some local magnetic field variations. If we consider the magnetic field *alone*, this perturbation is $O(\eta)$ and hence is small; however this is amplified if the magnetic field *gradient* is considered, as such variations occur at a much smaller length scale and become more significant once the spatial derivative is taken. This amplification means that the $O(\zeta^{-1}\eta)$ perturbation term can become comparable, or even take over, the original $O(g)$ q-vector term. Therefore, both of these terms will need to be considered as leading order in the expression.

4.4 Discussion

In the previous sections, multi-scale analysis has been performed for the heterogeneous magnetic susceptibility problem in the context of diffusion MRI. In particular, the perturbation to the local q-vector due to the induced magnetic field variations has been analysed, resulting to the expression in (4.32). This reveals two leading order terms for the q-vector, with an $O(\zeta^{-1}\eta)$ term originating from the microstructure-induced magnetic field variation and an $O(g)$ term from the magnetic field gradient applied from the far-field.

In practice, these two competing leading order terms may become comparable in magnitude, depending on the parameters of the DW-SE pulse sequence. Using the Warach et al. [116] study as a benchmark, the values of the nominal static field and q-vector are 1.5 T and 1.6×10^5 m^{-1} respectively, and the DW-SE sequence has a duration of ~ 50 ms. A typical cell diameter within the white matter of a human brain is ~ 5 μm [5], whereas the local susceptibility variation is estimated to be ~ 100 ppb [118], so $\eta \sim 10^{-7}$. Using these characteristic values yields $d^{-1}\gamma B_0\, \eta t \sim 2.7 \times 10^5$ m^{-1}, which is comparable to the applied q-vector.

The above example is at a relatively long-time, low q-value regime. However, if we consider a shorter time and higher q-value, assuming B_0 stays the same, the first term in (4.32) reduces in magnitude due to the shorter time, whereas the second term for the applied q-vector increases in magnitude. Therefore, it is anticipated that the applied q-vector will dominate in this regime, with the local induced magnetic field perturbation becoming negligible. These will be the assumptions employed in the asymptotic models of Chapters 5 and 6 respectively.

Chapter 5

Homogenisation model of DW-SE for the long-time regime

5.1 Introduction

In this chapter, the long-time regime of diffusion-weighted spin-echo (DW-SE) in heterogeneous media is considered. The long-time regime is characterised by the duration of the pulse sequence being significantly longer than the characteristic time needed for diffusion to occur over the microstructure length scale. As shown in Figure 5.1, this regime is located in the q-t parameter space to the right of the central region where all scales are balanced. Associated with this scaling of time is also a natural scaling of the q-vector. Since the nominal signal scales as $\exp(-|\mathbf{q}|^2 Dt)$ after the diffusion-induced decay in a DW-SE sequence, in order for the scaling of the exponent to remain as $O(1)$, the q-value needs to be small when the characteristic time is long. This is represented by the diagonal strip in the q-t parameter space. Therefore, the regime that is of interest in this chapter is represented by the bottom-right corner in Figure 5.1, which is the intersection of the diagonal strip with the long-time region.

In addition to diffusion and magnetic heterogeneities, an additional weak advection term is considered in the homogenisation analysis. In practice, this advection term is

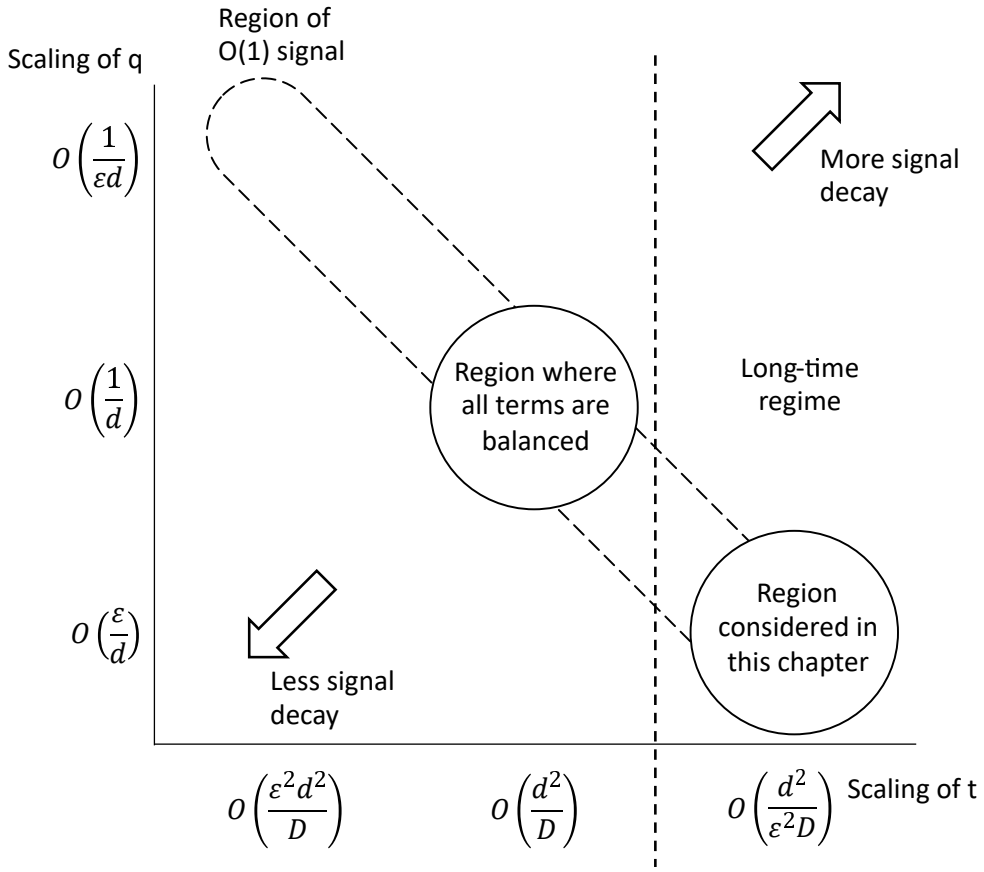

Figure 5.1: Definition of the long-time regime on the q-t parameter space.

likely to be negligible in DW-SE experiments, nevertheless the optional advection term is still included and can be omitted subsequently. In the context of biological tissues, this analysis can be used to model the intra- and extra-cellular spaces, as the impermeable cell membranes lead to diffusion and advection heterogeneities, as well as magnetic heterogeneities in the case of lipid-rich structures such as the myelin sheath. The multiscale homogenisation approach for the diffusion-advection equation in Auriault et al. [7] is extended for this problem, which aims to asymptotically derive a continuum model for porous media at a scale separation limit. In particular, the novel modelling methodology presented here accounts for microscale heterogeneities in both transport and magnetic behaviour simultaneously.

Finally, the homogenisation results are evaluated numerically. This is first done on a simple white matter model to demonstrate its principles. The numerical method is then applied on realistic geometries, obtained from transmission electron microscopy (TEM)

data from mouse white matter from a previously published study [121]. This gives us some insight to what the model might mean in realistic heterogeneous media.

5.2 Model formulation and relevant scales

5.2.1 Structure of porous medium

In this model we consider a periodic porous medium with characteristic length L (Fig. 5.2a), which is denoted as $\Omega \subset \mathbb{R}^3$ and is composed of multiple fluid-filled regions, with the i^{th} region denoted as Ω_i. The interfaces between compartments $\Gamma_{ij} = \partial\Omega_i \cap \partial\Omega_j$ are modelled as impermeable boundaries. The periodic reference cell of the microstructure, denoted as Υ with compartments Υ_i, has a characteristic length d much shorter than L (Fig. 5.2b). The spatially periodic assumption required in this approach is justified for random media with scale separation [6], as long as the periodic reference cell, referred to as the representative elementary volume (REV), contains sufficient instances of the random heterogeneities to allow a good approximation to the homogenised behaviour.

Since behaviours in the compartments are decoupled from each other, the derivations for each compartment are identical. It therefore suffices to consider one particular compartment Ω_α with internal boundaries Γ_α, and then models of the same form but with compartment-specific parameters can be combined into a complete description of the porous medium.

5.2.2 Formulation at micro-scale

As described in the previous chapter, the behaviour of the complex transverse magnetisation density $M_{xy}(\mathbf{x}, t)$ within the compartment can be described by the extended Bloch-Torrey equation given some initial condition:

$$\frac{\partial M_{xy}}{\partial t} = \nabla \cdot (\mathbf{D}\nabla M_{xy} - \mathbf{u}M_{xy}) - i\gamma B_z M_{xy} - R_2 M_{xy} \quad \text{in} \quad \Omega_\alpha, \tag{5.1}$$

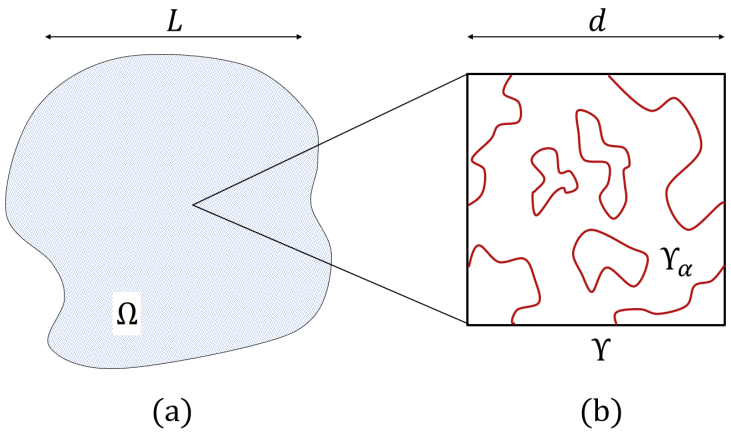

Figure 5.2: Schematic of porous medium: (a) macroscopic domain Ω, (b) periodic reference cell Υ, or the representative elementary volume (REV), of the microstructure with a particular compartment labelled as Υ_α.

with a no-flux boundary condition introduced by the impermeable boundaries:

$$\mathbf{n} \cdot (\mathbf{D}\nabla M_{xy} - \mathbf{u}M_{xy}) = 0 \quad \text{on} \quad \Gamma_\alpha, \tag{5.2}$$

where $\mathbf{n}$ is the outward normal of the boundary surface. In the above equations, γ and R_2 are the gyromagnetic ratio and transverse relaxation rate respectively, which are both assumed to be real constants. The diffusivity tensor $\mathbf{D}$, which is real and positive definite, varies with the microstructure and therefore is Υ-periodic. For the derivatives in (5.1)–(5.2) to exist in a strong sense, the diffusivity tensor is normally required to be continuous. However, by interpreting the derivatives in a distributional sense with an appropriate specification of function spaces, the problem can be relaxed without losing well-posedness to allow $\mathbf{D}$ to be discontinuous, equivalent to modelling multiple sub-compartments with perfect transmission interfaces within Υ_α. Meanwhile, assuming incompressible flow and no-flux internal boundaries, the Υ-periodic advection velocity $\mathbf{u}$ satisfies:

$$\nabla \cdot \mathbf{u} = 0 \quad \text{in} \quad \Omega_\alpha, \quad \mathbf{n} \cdot \mathbf{u} = 0 \quad \text{on} \quad \Gamma_\alpha. \tag{5.3}$$

The longitudinal magnetic field B_z is modelled by the following expression:

$$B_z(\mathbf{x}, t) = \langle B_0 \rangle_\alpha + \Delta B_0(\mathbf{x}) + \mathbf{g}(t) \cdot \mathbf{x}, \tag{5.4}$$

where $\langle B_0 \rangle_\alpha$ and ΔB_0 are the mean and varying parts in space of the static magnetic field (B_0) over Υ_α, and $\mathbf{g}(t)$ is the strength of the globally applied field gradients. Since the variation in B_0 originates from magnetic susceptibility heterogeneities in the microstructure, ΔB_0 is modelled to be Υ-periodic as well. To account for the phase-reversing 180° RF pulse (see Section 2.3), the effective quantities:

$$(M_{xy}^{\text{eff}}, B_0^{\text{eff}}, \mathbf{g}^{\text{eff}}) = \begin{cases} (M_{xy}, B_0, \mathbf{g}) & \text{if } 0 < t < t_{180}, \\ (\overline{M_{xy}}, -B_0, -\mathbf{g}) & \text{if } t > t_{180}, \end{cases}$$

are used in place of M_{xy}, B_0 and $\mathbf{g}$, so equation (5.1) holds continuously over the entire echo time. For clarity of presentation, the use of effective quantities are assumed throughout this chapter without explicit notation.

Since the attenuation of magnetisation originates from the spatial heterogeneity of NMR behaviour, for example introduced by an applied diffusion gradient, it is useful to introduce the co-rotating and co-decaying magnetisation variable:

$$S(\mathbf{x}, t) = \exp\left(\int_0^t (i\gamma B_z + R_2)\, \mathrm{d}\tau\right) M_{xy}. \tag{5.5}$$

With the phase refocusing effect of the 180° pulse, at the echo time of the spin-echo sequence $S(\mathbf{x}, TE) = \exp(R_2 t) M_{xy}$. The variable S therefore only contains the effect of transport-induced attenuation, with the coherent R_2 relaxation separated out. Substituting (5.5) into (5.1)–(5.2) yields the **co-rotating, co-decaying Bloch-Torrey equation** and its boundary conditions:

$$\frac{\partial S}{\partial t} + (\nabla - i\mathbf{q}) \cdot (-\mathbf{D}(\nabla - i\mathbf{q})S + \mathbf{u}\, S) = 0 \quad \text{in} \quad \Omega_\alpha, \tag{5.6}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla - i\mathbf{q})S + \mathbf{u}\, S) = 0 \quad \text{on} \quad \Gamma_\alpha. \tag{5.7}$$

Equations (5.6)–(5.7) are the governing equations which form the basis of the following asymptotic analysis. Here we are interested in the behaviour in the interior of the porous medium away from the external boundary on $\partial\Omega$, therefore we only consider the conditions on internal boundaries Γ_α. In the equations, $\mathbf{q}(\mathbf{x}, t) = \nabla \int_0^t \gamma B_z \, \mathrm{d}\tau$ is the *local* modulation

wave-vector, which can also be expressed as mean and Υ-periodic varying parts using expression (5.4):

$$\mathbf{q}(\mathbf{x},t) = \langle \mathbf{q} \rangle + \Delta\mathbf{q} = \int_0^t \gamma \mathbf{g}(\tau)\,\mathrm{d}\tau + \int_0^t \nabla(\gamma \Delta B_0)\,\mathrm{d}\tau. \tag{5.8}$$

If the medium is magnetically homogeneous i.e. $\Delta B_0 = 0$, then $\mathbf{q}$ reduces to its spatially-constant mean part, which is identical to the q-vector defined in existing literature (see Section 2.4.4). However, if the medium is magnetically heterogeneous the spatially-varying part cannot be ignored.

5.2.3 Non-dimensionalisation and relevant scales

To estimate the relative dominance of various terms in the governing equations of (5.6)–(5.7), the model is non-dimensionalised with respect to some characteristic quantities:

$$\mathbf{x} = l_c \mathbf{y}^*,\ t = t_c t^*,\ \mathbf{D} = D_c \mathbf{D}^*,\ \mathbf{u} = u_c \mathbf{u}^*,\ \mathbf{q} = q_c \mathbf{q}^*, S = M_c S^*, \tag{5.9}$$

where the subscript c and asterisk are used to denote characteristic quantities and non-dimensionalised variables respectively. The non-dimensional forms of (5.6), (5.7) are then given by:

$$\frac{\pi_2^2}{\pi_1}\frac{\partial S^*}{\partial t^*} + (\nabla_{\mathbf{y}^*} - \pi_2 i \mathbf{q}^*) \cdot (-\mathbf{D}^*(\nabla_{\mathbf{y}^*} - \pi_2 i \mathbf{q}^*)S + \mathrm{Pe}\,\mathbf{u}^* S^*) = 0 \ \text{ in } \ \Omega_\alpha^*, \tag{5.10}$$

$$\mathbf{n} \cdot (-\mathbf{D}^*(\nabla_{\mathbf{y}^*} - \pi_2 i \mathbf{q}^*)S^* + \mathrm{Pe}\,\mathbf{u}^* S^*) \ \text{ on } \ \Gamma_\alpha^*, \tag{5.11}$$

which contain three dimensionless groups, namely:

$$\pi_1 = q_c^2 D_c t_c,\ \pi_2 = q_c l_c,\ \mathrm{Pe} = \frac{u_c l_c}{D_c}, \tag{5.12}$$

where Pe is the Péclet number, the ratio between characteristic advective and diffusive rates.

The first dimensionless group π_1 relates the q-value and time scale of the DW-SE pulse sequence, and is independent of the choice of length scale. If the medium is homogeneous,

π_1 scales with the logarithm of the diffusion-induced signal attenuation (see Section 2.4.2), so it is expected that $\pi_1 \sim 1$. Meanwhile, π_2 and Pe depend on the length scale of interest. In particular, π_2 is the ratio between the characteristic length scale of the domain and an additional length scale introduced by the reciprocal of the characteristic q-value of the pulse sequence ($1/q_c$). In this model, we consider the particular regime of small q-value and weak advection from the microstructure perspective ($l_c = d$), meaning that $\pi_{2,d} = q_c d \ll 1$ and $\mathrm{Pe}_d = \frac{u_c d}{D_c} \ll 1$. The small magnitudes of $\pi_{2,d}$ and Pe_d therefore present some scale separation in the problem.

For ease of presentation, the asterisk notation for non-dimensional variables will be dropped in the rest of the chapter, as we will only use non-dimensional variables from here on in the derivations, so explicit notation is no longer needed.

5.3 Multi-scale homogenisation

As described in the introduction chapter, multi-scale homogenisation is an asymptotic method which aims to approximate the physical problem by its scale separation limit. This involves parameterising the scale separation by a small parameter ε, and then studying the asymptotic limit of $\varepsilon \to 0$. Here, the limit behaviour is derived by assuming a series expansion Ansatz for the solutions to the ε-problem and truncating it at the leading order, which is a common method for homogenisation.

5.3.1 Parameterising the scale separation

In multi-scale homogenisation, the scale separation concerned is the one between the microstructure length scale d and a larger macroscopic scale, which in this problem can be either L of the physical domain Ω, or $1/q_c$ introduced by the DW-SE pulse sequence. In practical applications, L is likely to be larger than $1/q_c$, i.e. $d \ll 1/q_c < L$. In order to derive an upscaled model which is valid at both length scales of $1/q_c$ and L, the smaller separation is chosen as the scaling limit of interest. Therefore, the small parameter ε is

given by:

$$\varepsilon = q_c d \ll 1. \tag{5.13}$$

In terms of ε, the regime chosen for this model can be represented by the following scaling of dimensionless groups:

$$\pi_1 = O(1),\ \pi_{2,d} = O(\varepsilon),\ \mathrm{Pe}_d = O(\varepsilon). \tag{5.14}$$

Using the micro-scale perspective, the scaling can be reflected by setting characteristic quantities such that:

$$l_c = d,\ t_c = \frac{1}{q_c^2 D_c},\ u_c = q_c D_c, \tag{5.15}$$

which then scale the non-dimensional variables in (5.9) to $O(1)$, as well as parameterising equations (5.10), (5.11) as:

$$\varepsilon^2 \frac{\partial S}{\partial t} + (\nabla_{\mathbf{y}} - \varepsilon i\mathbf{q}) \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} - \varepsilon i\mathbf{q})S + \varepsilon\, \mathbf{u}S) = 0 \ \text{ in } \ \Omega_\alpha, \tag{5.16}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} - \varepsilon i\mathbf{q})S + \varepsilon\, \mathbf{u}S) = 0 \ \text{ on } \ \Gamma_\alpha. \tag{5.17}$$

Through homogenisation, we aim to find the asymptotic behaviour of the solution to (5.16)–(5.17) as $\varepsilon \to 0$. Physically, this corresponds to studying a family of similar porous media sharing the same normalised reference cell Υ but with different values of d. The asymptotic limit being considered is when the microstructure scale becomes finer and finer ($d \to 0$), while keeping all other characteristic quantities constant.

5.3.2 Multiple scales series expansion

We look for the solution to (5.16), (5.17) in the form of a multiple scales series expansion:

$$S = S^{(0)}(\mathbf{x}, \mathbf{y}, t) + \varepsilon S^{(1)}(\mathbf{x}, \mathbf{y}, t) + \varepsilon^2 S^{(2)}(\mathbf{x}, \mathbf{y}, t) + \cdots \tag{5.18}$$

where $\mathbf{x}$ and $\mathbf{y}$ are the non-dimensional macro-scale and micro-scale variables related by $\mathbf{x} = \varepsilon\mathbf{y}$. The two spatial variables are treated as independent variables, with $\mathbf{x} \in \Omega$ describing the location in the macroscopic medium and $\mathbf{y} \in \Upsilon_\alpha$, a variable on a periodic

domain, describing the precise location within the periodic reference cell. The properties $\mathbf{D}$ and $\mathbf{u}$ are governed by the microstructure and are therefore functions of a sole variable $\mathbf{y}$. In particular, the flow properties of (5.3) are expressed in the micro-scale variable as:

$$\nabla_{\mathbf{y}} \cdot \mathbf{u} = 0 \ \text{ in } \ \Upsilon_\alpha, \ \ \mathbf{n} \cdot \mathbf{u} = 0 \ \text{ on } \ \partial\Upsilon_\alpha. \tag{5.19}$$

A change of variable is performed on the spatial derivatives in (5.16)–(5.17) using the chain rule:

$$\nabla_{\mathbf{y}} \to \varepsilon \nabla_{\mathbf{x}} + \nabla_{\mathbf{y}}.$$

Together with the multiple scales expansion of (5.18), these are then substituted into the non-dimensional equations (5.16)–(5.17); due to constraint of space the fully expanded equation is not shown here. Balancing coefficients at each order of ε in the expanded equations yields a series of PDE problems; in the following analysis they are treated as PDEs in the micro-scale variable $\mathbf{y}$, with functions of $\mathbf{x}$ and their derivatives treated as parameters. This allows us to derive some results about the micro-scale behaviour given some arbitrary macro-scale conditions, which can then be averaged locally to become a macro-scale homogenised model.

At order ε^0 of the series expansion of (5.16)–(5.17), the corresponding problem is given by:

$$\nabla_{\mathbf{y}} \cdot (-\mathbf{D} \nabla_{\mathbf{y}} S^{(0)}) = 0 \ \text{ in } \ \Upsilon_\alpha, \tag{5.20}$$

$$\mathbf{n} \cdot (-\mathbf{D} \nabla_{\mathbf{y}} S^{(0)}) = 0 \ \text{ on } \ \partial\Upsilon_\alpha. \tag{5.21}$$

It is a standard result that (5.20)–(5.21) only admit solutions constant in $\mathbf{y}$, i.e. $S^{(0)} = S^{(0)}(\mathbf{x}, t)$; its relation to the macro-scale variable will be determined later. Using this result and the flow properties given by (5.19), the problem for order ε^1 can be written as:

$$\nabla_{\mathbf{y}} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} S^{(1)} + \nabla_{\mathbf{x}} S^{(0)} - i\mathbf{q} S^{(0)})) = 0 \ \text{ in } \Upsilon_\alpha, \tag{5.22}$$

$$\mathbf{n} \cdot -\mathbf{D}(\nabla_{\mathbf{y}} S^{(1)} + \nabla_{\mathbf{x}} S^{(0)} - i\mathbf{q} S^{(0)}) = 0 \ \text{ on } \partial\Upsilon_\alpha. \tag{5.23}$$

This is a modified Laplace's problem solving for $S^{(1)}$, which contains a driving potential

term $\nabla_{\mathbf{x}} S^{(0)} - i\mathbf{q}S^{(0)}$. Again, (5.22)–(5.23) have a unique solution up to an additive constant, so it can be written as the sum of a mean part and fluctuating contributions from $\nabla_{\mathbf{x}} S^{(0)}$ and $S^{(0)}$:

$$S^{(1)}(\mathbf{x}, \mathbf{y}, t) = \langle S^{(1)} \rangle_{\alpha} + \mathbf{w}(\mathbf{y}) \cdot \nabla_{\mathbf{x}} S^{(0)} + k(\mathbf{y}) S^{(0)}. \tag{5.24}$$

The zero-mean, periodic vector and scalar functions $\mathbf{w}(\mathbf{y})$ and $k(\mathbf{y})$ are the unique solutions to the respective *cell problems* over the porous reference cell, given by:

$$\nabla_{\mathbf{y}} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} \mathbf{w} + \mathbf{I})) = 0 \quad \text{in} \quad \Upsilon_{\alpha}, \tag{5.25}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} \mathbf{w} + \mathbf{I})) = 0 \quad \text{on} \quad \partial\Upsilon_{\alpha}, \tag{5.26}$$

$$\nabla_{\mathbf{y}} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} k - i\mathbf{q})) = 0 \quad \text{in} \quad \Upsilon_{\alpha}, \tag{5.27}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} k - i\mathbf{q})) = 0 \quad \text{on} \quad \partial\Upsilon_{\alpha}. \tag{5.28}$$

Since the physical quantity $\mathbf{q}$ originates as the gradient of the accumulated phase precession:

$$\mathbf{q}(\mathbf{x}, t) = \langle \mathbf{q} \rangle + \Delta\mathbf{q} = \int_0^t \gamma \mathbf{g}(\tau) \, \mathrm{d}\tau + \int_0^t \nabla(\gamma \Delta B_0) \, \mathrm{d}\tau, \tag{5.29}$$

a particular integral $\xi(\mathbf{y})$, which is zero-mean and periodic, can be found for the micro-scale fluctuating part of $\mathbf{q}$ in (5.27)–(5.28), i.e. $\nabla_{\mathbf{y}} \xi = \Delta\mathbf{q}$. Substituting this into (5.27)–(5.28) then allows the cell problem for k to be related to that for $\mathbf{w}$ with the following relation:

$$k = -i\mathbf{w} \cdot \langle \mathbf{q} \rangle + i\xi. \tag{5.30}$$

Therefore, the $S^{(1)}$ solution can be written in terms of $\mathbf{w}$ as:

$$S^{(1)}(\mathbf{x}, \mathbf{y}, t) = \langle S^{(1)} \rangle_{\alpha} + \mathbf{w}(\mathbf{y}) \cdot (\nabla_{\mathbf{x}} S^{(0)} - i\langle \mathbf{q} \rangle S^{(0)}) + i\xi(\mathbf{y}) S^{(0)}. \tag{5.31}$$

The significance of the cell problem solving for $\mathbf{w}$, given by Equations (5.25)–(5.26), will be discussed in a later section.

Proceeding to the next order of the expansion, the ε^2 problem is given by:

$$\begin{aligned}\frac{\partial S^{(0)}}{\partial t} &+ \nabla_{\mathbf{y}} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} S^{(2)} + \nabla_{\mathbf{x}} S^{(1)} - i\mathbf{q} S^{(1)}) + \mathbf{u} S^{(1)}) \\ &+ (\nabla_{\mathbf{x}} - i\mathbf{q}) \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} S^{(1)} + \nabla_{\mathbf{x}} S^{(0)} - i\mathbf{q} S^{(0)}) + \mathbf{u} S^{(0)}) = 0 \ \text{ in } \Upsilon_\alpha,\end{aligned} \tag{5.32}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} S^{(2)} + \nabla_{\mathbf{x}} S^{(1)} - i\mathbf{q} S^{(1)}) + \mathbf{u} S^{(1)}) = 0 \ \text{ on } \partial\Upsilon_\alpha. \tag{5.33}$$

Being a pure Neumann problem solving for $S^{(2)}$ as a function of $\mathbf{y}$, (5.32)–(5.33) introduces a solvability condition which requires $S^{(0)}$ to relate to the macro-scale and time variables $\mathbf{x}$ and t in a certain way. This is revealed by taking the volume average of (5.32) over the periodic domain Υ_α and applying the divergence theorem with (5.33), which yields the following condition:

$$\frac{\partial S^{(0)}}{\partial t} + \frac{1}{|\Upsilon_\alpha|} \int_{\Upsilon_\alpha} (\nabla_{\mathbf{x}} - i\mathbf{q}) \cdot (-\mathbf{D}(\nabla_{\mathbf{y}} S^{(1)} + \nabla_{\mathbf{x}} S^{(0)} - i\mathbf{q} S^{(0)}) + \mathbf{u} S^{(0)}) \, \mathrm{d}\mathbf{y} = 0. \tag{5.34}$$

The $\nabla_{\mathbf{y}} S^{(1)}$ term can be related to $S^{(0)}$ by taking the gradient of the expression in (5.31):

$$\nabla_{\mathbf{y}} S^{(1)} = \nabla_{\mathbf{y}} \mathbf{w} \cdot (\nabla_{\mathbf{x}} S^{(0)} - i\langle \mathbf{q} \rangle S^{(0)}) + i\Delta\mathbf{q} S^{(0)}. \tag{5.35}$$

Substituting this into (5.34) leads to a relation only involving the leading order coefficient $S^{(0)}(\mathbf{x}, t)$ and its derivatives in terms of the time and macro-scale variables:

$$\frac{\partial S^{(0)}}{\partial t} + \frac{1}{|\Upsilon_\alpha|} \int_{\Upsilon_\alpha} (\nabla_{\mathbf{x}} - i\mathbf{q}) \cdot \left(-\mathbf{D}(\nabla_{\mathbf{y}} \mathbf{w} + \mathbf{I})(\nabla_{\mathbf{x}} - i\langle \mathbf{q} \rangle) S^{(0)}) + \mathbf{u} S^{(0)}\right) \mathrm{d}\mathbf{y} = 0. \tag{5.36}$$

This relation governs the value of $S^{(0)}$ at the macroscopic scale. The effects of the mean and fluctuating parts of the q-vector in (5.36) can be broken down as:

$$\begin{aligned}\frac{\partial S^{(0)}}{\partial t} &+ (\nabla_{\mathbf{x}} - i\langle \mathbf{q} \rangle) \cdot (-\langle \mathbf{F} \rangle_\alpha (\nabla_{\mathbf{x}} - i\langle \mathbf{q} \rangle) S^{(0)} + \langle \mathbf{u} \rangle_\alpha S^{(0)}) \\ &+ i\langle \Delta\mathbf{q} \cdot \mathbf{F} \rangle_\alpha (\nabla_{\mathbf{x}} - i\langle \mathbf{q} \rangle) S^{(0)} - i\langle \Delta\mathbf{q} \cdot \mathbf{u} \rangle_\alpha S^{(0)} = 0,\end{aligned} \tag{5.37}$$

where $\langle \star \rangle_\alpha = \frac{1}{|\Upsilon_\alpha|} \int_{\Upsilon_\alpha} \star \, \mathrm{d}\mathbf{y}$ denotes the averaging over a reference cell, and $\mathbf{F}(\mathbf{y}) = \mathbf{D}(\nabla_{\mathbf{y}} \mathbf{w} + \mathbf{I})$ is a tensor field in the periodic domain of Υ_α obtained from the cell problem. In particular, the tensor $\langle \mathbf{F} \rangle_\alpha$ is identical to the effective diffusion tensor $\mathbf{D}_\alpha^{\text{eff}}$ derived from homogenising the diffusion equation in a porous domain; such equivalence will be

discussed later. Meanwhile, using the fact that $\Delta\mathbf{q}$ is the gradient of an exact integral, the $\langle\Delta\mathbf{q}\cdot\mathbf{F}\rangle_\alpha$ term can be reduced to zero using integration by parts:

$$\begin{aligned}\int_{\Upsilon_\alpha}\Delta\mathbf{q}\cdot\mathbf{F}\mathrm{d}\mathbf{y}=\int_{\Upsilon_\alpha}\nabla\xi\cdot\mathbf{F}\mathrm{d}\mathbf{y}&=\int_{\Upsilon_\alpha}\nabla\cdot(\xi\mathbf{F})\,\mathrm{d}\mathbf{y}-\int_{\Upsilon_\alpha}\xi\nabla\cdot\mathbf{F}\mathrm{d}\mathbf{y}\\&=\int_{\partial\Upsilon_\alpha}\xi(\mathbf{n}\cdot\mathbf{F})\,\mathrm{d}\mathbf{s}-\int_{\Upsilon_\alpha}\xi(\nabla\cdot\mathbf{F})\,\mathrm{d}\mathbf{y}=0.\end{aligned}\tag{5.38}$$

Here we have used the divergence-free property of $\mathbf{F}$ and its associated no-flux boundary condition, as constituted by Equations (5.25)–(5.26). The same properties apply to the advection velocity $\mathbf{u}$ as well, hence the $\langle\Delta\mathbf{q}\cdot\mathbf{u}\rangle_\alpha$ term reduces to zero too, simplifying (5.37) to:

$$\frac{\partial S^{(0)}}{\partial t}+(\nabla_{\mathbf{x}}-i\langle\mathbf{q}\rangle)\cdot(-\langle\mathbf{F}\rangle_\alpha(\nabla_{\mathbf{x}}-i\langle\mathbf{q}\rangle)S^{(0)}+\langle\mathbf{u}\rangle_\alpha S^{(0)})=0.\tag{5.39}$$

Equation (5.39) therefore serves as the leading-order continuum approximation to the original problem, as we expect the exact solution to the original problem S to converge to the leading-order variable $S^{(0)}$ at the asymptotic limit of $\varepsilon\to 0$. Since Equation (5.39) no longer has any dependence on the micro-scale spatial variable $\mathbf{y}$, it serves as a macro-scale description of a homogenised continuum, with all the relevant effect from the microstructure contained in the effective parameters of $\langle\mathbf{F}\rangle_\alpha$, $\langle\mathbf{u}\rangle_\alpha$ and $\langle\mathbf{q}\rangle$.

As we can see in Equation (5.39), the homogenised model for the macro-scale variable $S^{(0)}$ takes the same form as the original equation (Eq. 5.6) . This means at this particular asymptotic limit considered, the Bloch-Torrey equation is indeed a valid approximation for a single porous compartment, even with heterogeneities in both transport and magnetic behaviour. However, as shown in the derivation above, the relations between the effective parameters in the homogenised model and micro-scale heterogeneities are not trivial, in particular the computation of the effective diffusion tensor requires solving a cell problem over the microstructure. Moreover, even though the homogenised effective parameters are all some form of averaged quantities, they are averaged over different domains. The effective transport parameters, namely $\langle\mathbf{F}\rangle_\alpha$ and $\langle\mathbf{u}\rangle_\alpha$, involve averaging across only the porous compartment considered, whereas the effective q-vector $\langle\mathbf{q}\rangle$, somehow counter-

intuitively, comes from the average across all compartments in the periodic cell, the reason being that this is the correct quantity such that the remaining micro-scale fluctuations can be exactly integrated (see Eq. 5.29).

5.3.3 Initial condition for the homogenised problem

Being an initial value problem, both the physical and homogenised problems require some initial conditions at $t = 0$. For the homogenised problem given by (5.39), the variable $S^{(0)}(\mathbf{x}, t)$ does not depend on the micro-scale variable $\mathbf{y}$, i.e. it is always constant locally over a reference cell. However, the initial condition for the physical problem given by (5.16)–(5.17) does not necessarily satisfy this condition. For example, the excitation efficiency of transverse magnetisation can be selective to the Larmor frequency (see Section 2.3.1), leading to variable excitation at the micro-scale if the medium is magnetically heterogeneous. Therefore, an additional relation is required to relate the physical initial condition to one that is appropriate for the homogenised problem.

The locally-varying initial condition for (5.16)–(5.17), which equates to the initial transverse magnetisation M_0, can be formulated as a function of both macro- and micro-scale variables:

$$S|_{t=0} = S_0(\mathbf{x}, \mathbf{y}) = M_0(\mathbf{x}, \mathbf{y}). \tag{5.40}$$

This physical initial condition can be related to the homogenised quantity $S^{(0)}(\mathbf{x}, t)$ using boundary layer analysis. Since the highest-order time derivative in (5.16)–(5.17) is small as it scales with ε^2, a thin boundary layer in the time variable is expected near $t = 0$, in which S varies at a very fast rate. Its behaviour can be approximated by defining an inner time variable: $t_{\text{in}} = \varepsilon^{-2} t$. Within the inner region, terms which contain $\mathbf{q}$, $\mathbf{u}$ or an $\mathbf{x}$-derivative in the expansion of (5.16)–(5.17) become negligible at the $\varepsilon \to 0$ limit, while the scaled time-derivative does not vanish. The inner problem is therefore given by:

$$\frac{\partial S_{\text{in}}}{\partial t_{\text{in}}} + \nabla_{\mathbf{y}} \cdot (-\mathbf{D} \nabla_{\mathbf{y}} S_{\text{in}}) = 0 \quad \text{in} \quad \Upsilon_\alpha, \tag{5.41}$$

$$\mathbf{n} \cdot (-\mathbf{D} \nabla_{\mathbf{y}} S_{\text{in}}) = 0 \quad \text{on} \quad \partial\Upsilon_\alpha, \tag{5.42}$$

with the initial condition $S_{\text{in}}|_{t_{\text{in}}=0} = S_0(\mathbf{x}, \mathbf{y})$. Equations (5.41)–(5.42) constitute a diffusion problem on a periodic domain with purely Neumann boundaries, leading to the standard result:

$$\lim_{t_{\text{in}}\to\infty} S_{\text{in}} = \langle S_0 \rangle_\alpha. \tag{5.43}$$

Meanwhile, in the outer region the homogenised model of (5.37) applies. Using the concept of asymptotic matching, the inner and outer solutions should match when the respective limits of $t_{\text{in}} \to \infty$ and $t_{\text{out}} \to 0$ are taken. The limit solution for the inner region given by (5.43) is consistent with the outer region model, as it is indeed constant over the micro-scale variable $\mathbf{y}$. Therefore, the appropriate initial condition for the homogenised problem relates to the physical initial condition as:

$$S^{(0)}|_{t=0} = \langle S_0 \rangle_\alpha = \langle M_0 \rangle_\alpha. \tag{5.44}$$

5.3.4 The cell problem and pore connectivity

In the homogenised model of (5.37), a key quantity that contains the information about the micro-scale behaviour is the non-dimensional tensor field $\mathbf{F}(\mathbf{y}) = \mathbf{D}(\nabla\mathbf{w} + \mathbf{I})$, which is obtained by solving the cell problem (5.25)–(5.26) for the zero-mean periodic function $\mathbf{w}(\mathbf{y})$. For the ease of the reader, the cell problem is restated here, with subscripts for the derivatives now unnecessary since $\mathbf{y}$ is the only spatial variable:

$$\nabla \cdot (-\mathbf{D}(\nabla\mathbf{w} + \mathbf{I})) = 0 \quad \text{in} \quad \Upsilon_\alpha, \tag{5.45}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla\mathbf{w} + \mathbf{I})) = 0 \quad \text{on} \quad \partial\Upsilon_\alpha. \tag{5.46}$$

It is noted that this cell problem, arising from the homogenisation of (5.6)–(5.7), is identical to that from the homogenisation of the diffusion equation [7]. Physically, $\mathbf{F}(\mathbf{y})$ can be interpreted as the tensor mapping a macro-scale potential, applied across the microstructure reference cell, to the local induced flux within the microstructure. When $\mathbf{F}(\mathbf{y})$ is averaged over the microstructure reference cell, the resulting tensor then maps the macro-scale potential to the macro-scale averaged flux, in other words $\langle \mathbf{F} \rangle_\alpha$ is the

effective diffusivity tensor $\mathbf{D}_\alpha^{\text{eff}}$. The origin of the macro-scale potential differs between the two homogenisation problems – for the diffusion equation it originates solely from a (real-valued) concentration gradient, whereas in DW-SE the concentration gradient combines with the q-vector to generate a complex-valued potential. Nevertheless, the role of the cell problem, and more importantly the resulting tensor field $\mathbf{F}(\mathbf{y})$, is equivalent in both problems.

Mathematically, the cell problem (5.45)–(5.46) solving for $\mathbf{w}(\mathbf{y})$ should be posed on the space of zero-mean functions with square-integrable first derivatives, over which the existence and uniqueness of $\mathbf{w}(\mathbf{y})$ follow from the Lax-Milgram theorem [33]. However, the form of $\mathbf{w}(\mathbf{y})$ can differ vastly depending on the domain topology of Υ_α, which in turn relates to the physical pore connectivity of the compartment being considered. Evaluating the component of $\mathbf{w}$ along any direction $\mathbf{e}_\lambda \in \mathbb{R}^3$, the resultant $w_\lambda = \mathbf{w} \cdot \mathbf{e}_\lambda$ is the unique solution to the following equations:

$$\nabla \cdot (-\mathbf{D}(\nabla w_\lambda + \mathbf{e}_\lambda)) = 0 \ \text{ in } \ \Upsilon_\alpha, \tag{5.47}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla w_\lambda + \mathbf{e}_\lambda)) = 0 \ \text{ on } \ \partial\Upsilon_\alpha. \tag{5.48}$$

The trivial solution is given by $w_\lambda = -\mathbf{e}_\lambda \cdot \mathbf{y}$, which leads to $\mathbf{F}\mathbf{e}_\lambda = 0$. However, this solution is only well defined if the domain Υ_α does not contain any closed path $\mathcal{P}$ on which the path integral $\int_\mathcal{P} \mathbf{e}_\lambda \cdot \mathrm{d}\mathbf{s}$ does not vanish, i.e. there are no paths crossing a periodic boundary in that direction. If such a path $\mathcal{P}$ exists, the square-integrable first derivative condition requires w_λ to be continuous, but the function $-\mathbf{e}_\lambda \cdot \mathbf{y}$ cannot be single-valued and continuous simultaneously on the path $\mathcal{P}$, which goes round periodically. In this case, the solution is non-trivial and $\mathbf{F}\mathbf{e}_\lambda$ is no longer constantly zero.

The physical explanation to the solution triviality is related to the pore connectivity within the medium. When a macro-scale potential gradient is applied along $\mathbf{e}_\lambda$ but the pores are not connected in that direction, it cannot generate a flux within the microstructure due to the blockage, therefore $\mathbf{F}\mathbf{e}_\lambda = 0$. However, if there exists a path which goes round the reference cell Υ_α periodically along the direction of $\mathbf{e}_\lambda$, which is equivalent to

having a percolation path in the macro-scale porous medium, a non-zero flux which is $\mathbf{Fe}_\lambda$ is generated within the microstructure. In particular, if a compartment consists of isolated domains which are disconnected in all directions, the resulting tensor field $\mathbf{F}$ is constantly zero in all directions, meaning that the compartment has no effective diffusivity at all.

5.3.5 Homogenised multi-compartmental model

To summarise the homogenisation result, the leading order continuum equation for a particular compartment is given by (5.37) with initial condition (5.44). Returning to dimensional quantities, the magnetisation for a particular compartment is described by the homogenised model:

$$\frac{\partial S_\alpha}{\partial t} + (\nabla - i\langle\mathbf{q}\rangle) \cdot (-\langle\mathbf{F}\rangle_\alpha(\nabla - i\langle\mathbf{q}\rangle)S_\alpha + \langle\mathbf{u}\rangle_\alpha S_\alpha) = 0 \quad \text{in} \quad \Omega, \tag{5.49}$$

with the initial condition given by $S_\alpha = \langle S_0\rangle_\alpha$ at $t = 0$, and some boundary conditions specified on the external boundary $\partial\Omega$. Comparing this to the original equation, the first term in (5.49) has the same form of (5.6) with the averaged quantities $\langle\mathbf{q}\rangle$, $\langle\mathbf{F}\rangle_\alpha$ and $\langle\mathbf{u}\rangle_\alpha$ replacing $\mathbf{q}$, $\mathbf{D}$ and $\mathbf{u}$ respectively.

If the initial condition is constant in the macroscopic scale, which is a reasonable assumption when the RF pulses used in the DW-SE sequence have a low spatial selectivity, the homogenised model can be further simplified by eliminating the spatial derivatives:

$$\frac{\mathrm{d}S_\alpha}{\mathrm{d}t} + \langle\mathbf{q}\rangle \cdot \langle\mathbf{F}\rangle_\alpha\langle\mathbf{q}\rangle S_\alpha - i\langle\mathbf{q}\rangle \cdot \langle\mathbf{u}\rangle_\alpha S_\alpha = 0. \tag{5.50}$$

Solving this equation yields a compartment-specific expression describing the co-rotating, co-decaying magnetisation density:

$$\frac{S_\alpha(t)}{\langle S_0\rangle_\alpha} = \exp\int_0^t (-\langle\mathbf{q}\rangle \cdot \langle\mathbf{F}\rangle_\alpha\langle\mathbf{q}\rangle + i\langle\mathbf{q}\rangle \cdot \langle\mathbf{u}\rangle_\alpha)\,\mathrm{d}\tau. \tag{5.51}$$

From this expression, the full multi-compartmental homogenised model can be assembled. Due to the refocusing effect of the spin-echo sequence, at echo time the relation $M_{xy} =$

$\exp(-R_2 t)S(t)$ can be used to return to the stationary frame variable. The overall signal density is then obtained by averaging $S_i(t)$ over all compartments with volume fraction weightings:

$$\langle M_{xy} \rangle = \sum_{i=1}^{n} \frac{|\Upsilon_i|}{|\Upsilon|} \exp(-R_{2,i} t) S_i(t). \tag{5.52}$$

Dividing this by the initial average magnetisation and using the relation $M_0 = S_0$ yields a multi-exponential expression for the signal attenuation behaviour:

$$\frac{\langle M_{xy} \rangle}{\langle M_0 \rangle} = \sum_{i=1}^{n} \frac{|\Upsilon_i|}{|\Upsilon|} \frac{\langle M_0 \rangle_i}{\langle M_0 \rangle} \exp\left(-R_{2,i} t + \int_0^t \left(-\langle \mathbf{q} \rangle \cdot \langle \mathbf{F} \rangle_i \langle \mathbf{q} \rangle + i \langle \mathbf{q} \rangle \cdot \langle \mathbf{u} \rangle_i \right) \mathrm{d}\tau \right). \tag{5.53}$$

The signal attenuation expression (5.53) resembles multi-compartmental models in the literature (see Section 2.4.3). However, by considering a more general problem using the scale separation assumption, the homogenisation derivation gives a more detailed explanation of the exponents and signal fractions in the multi-exponential model. The exponent indeed takes the same form as the existing models, with the diffusion and advection parts being quadratic and linear relative to the global q-vector $\langle \mathbf{q} \rangle$ respectively. Meanwhile, the signal fractions in (5.51) derived through homogenisation include not only the volume fractions, but also account for the compartment-specific transverse magnetisation excitation efficiency and T2 relaxation rates. When the latter two effects are constant over all compartments, then the signal fractions reduce to volume fractions, as in existing multi-compartmental models. These ignored effects may potentially explain the inconsistency between fitted signal fractions and accepted volume fraction values suggested in the current literature.

5.4 Numerical solution of the cell problem

In this section, a simple model of white matter tissue is used to apply the homogenisation results derived above. As explained previously, the action of diffusion MRI within the microstructure is captured by the solution to the cell problem over a periodic cell, in particular the tensor field $\mathbf{F}(\mathbf{y}) = \mathbf{D}(\nabla \mathbf{w} + \mathbf{I})$, from which the parameters of the homogenised

continuum model can be computed. Here we present a methodology based on the finite element method to numerically solve the cell problem.

The finite element method is well suited to elliptic problems, such as the cell problem, even when posed with discontinuous data. Its mathematical theory is well studied and can be found in numerous texts, for example [16, 27]. The method is common in many engineering disciplines, with its implementation available as various software packages, both closed- and open-source. In line with scientific reproducibility principles, the open-source packages `gmsh` [41] and FEniCS [71] are utilised as the mesh generator and solver respectively, with the implementation of the solver based on modifications to the example codes provided with the FEniCS package. Initially, a simple model is chosen to demonstrate the methodology. This is then applied to realistic geometries based on microscopy data.

5.4.1 The finite element method

The finite element method for a well-posed, linear problem can be loosely summarised by the following steps: Firstly, the PDE and its boundary conditions are formulated as a variational equality in the form:

$$\text{Find } u \in V \text{ such that } a(u, v) = L(v) \text{ for all } v \in V. \tag{5.54}$$

Here we seek our solution in an appropriate function space V which reflects the boundary conditions. Secondly, a finite-dimensional subspace V_h is chosen from the infinite-dimensional function space V, based on the discretisation of the domain into a mesh of finite elements. This leads to the Galerkin approximation of the original problem, which is described by:

$$\text{Find } u_h \in V \text{ such that } a(u_h, v_h) = L(v_h) \text{ for all } v_h \in V_h. \tag{5.55}$$

The approximation problem inherits well-posedness from the original problem. Finally, this finite-dimensional problem is assembled into a sparse system of linear equations and

solved numerically.

5.4.2 Variational formulation of the cell problem

Consider the component of $\mathbf{w}$ along a unit vector $\mathbf{e}_\lambda \in \mathbb{R}^3$, then the cell problem for the resultant $w_\lambda = \mathbf{w} \cdot \mathbf{e}_\lambda$ is written as:

$$\nabla \cdot (-\mathbf{D}(\nabla w_\lambda + \mathbf{e}_\lambda)) = 0 \text{ in } \Upsilon_\alpha, \tag{5.56}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla w_\lambda + \mathbf{e}_\lambda)) = 0 \text{ in } \partial\Upsilon_\alpha. \tag{5.57}$$

Here we consider the components in Cartesian coordinates: $\mathbf{e}_\lambda = \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$, which is sufficient to determine $\mathbf{w}$. However, if the domain has specific symmetries, for instance hexagonal tessellation, other choices of $\mathbf{e}_\lambda$ are possible to simplify the problem. The equivalent variational form of (5.56)–(5.57) is to solve for w_λ over the function space $\mathcal{W}_{\text{per}}(\Upsilon_\alpha) \subset H^1_{\text{per}}(\Upsilon_\alpha)$, i.e. the zero-mean subspace of functions with square-integrable first derivatives on the periodic domain Υ_α, such that:

$$\int_{\Upsilon_\alpha} (\mathbf{D}\nabla w_\lambda) \cdot \nabla v \, \mathrm{d}\mathbf{y} = -\int_{\Upsilon_\alpha} (\mathbf{D}\mathbf{e}_\lambda) \cdot \nabla v \, \mathrm{d}\mathbf{y}, \ \forall v \in \mathcal{W}_{\text{per}}(\Upsilon_\alpha). \tag{5.58}$$

The solution to the variational equality minimises the quadratic energy associated with the system, which is:

$$J(w_\lambda) = \frac{1}{2}\int_{\Upsilon_\alpha} (\mathbf{D}\nabla w_\lambda) \cdot \nabla w_\lambda \, \mathrm{d}\mathbf{y} + \int_{\Upsilon_\alpha} (\mathbf{D}\mathbf{e}_\lambda) \cdot \nabla w_\lambda \, \mathrm{d}\mathbf{y}. \tag{5.59}$$

From the energy perspective, it is easy to see why the zero-mean condition is essential: the energy is not strictly convex if the variational problem is posed on $H^1_{\text{per}}(\Upsilon_\alpha)$ instead. However, the implementation in FEniCS does not allow the zero-mean condition to be explicitly enforced in the choice of function space. Instead, this constraint is implemented in the form of a Lagrange multiplier c, which modifies the energy of the variational problem to become:

$$J(w_\lambda, c) = \frac{1}{2}\int_{\Upsilon_\alpha} (\mathbf{D}\nabla w_\lambda) \cdot \nabla w_\lambda \, \mathrm{d}\mathbf{y} + \int_{\Upsilon_\alpha} (\mathbf{D}\mathbf{e}_\lambda) \cdot \nabla w_\lambda \, \mathrm{d}\mathbf{y} + \int_{\Upsilon_\alpha} c w_\lambda \, \mathrm{d}\mathbf{y}. \tag{5.60}$$

The corresponding variational problem solving for $(w_\lambda, c) \in H^1_{\text{per}}(\Upsilon_\alpha) \times \mathbb{R}$ is given by:

$$\int_{\Upsilon_\alpha} ((\mathbf{D}\nabla w_\lambda) \cdot \nabla v + cv + w_\lambda d)\, \mathrm{d}\mathbf{y} = -\int_{\Upsilon_\alpha} (\mathbf{D}\mathbf{e}_\lambda) \cdot \nabla v \, \mathrm{d}\mathbf{y}, \ \forall (v, d) \in H^1_{\text{per}}(\Upsilon_\alpha) \times \mathbb{R}. \quad (5.61)$$

The solution to (5.61) can be shown to satisfy $c = 0$ and $\int_{\Upsilon_\alpha} w_\lambda \, \mathrm{d}\mathbf{y} = 0$, thus recovering the solution to (5.58). This variational problem is now ready to be discretised on a meshed domain and solved using FEniCS.

5.4.3 Results using a simple test case

To demonstrate the implementation, we consider a simple model of the cellular structure of white matter, which contains parallel, cylindrical axons in the z-direction of diameter $0.5d$ arranged in a square grid of spacing d, as shown in Fig. 5.3a. The cell membranes of the axons are assumed to have negligible thickness and to be impermeable. In both intra- and extra-axonal spaces, the diffusivity D is modelled to be uniform, isotropic and equal, hence $\mathbf{D} = \mathbf{I}$ after non-dimensionalisation. The advection velocity $\mathbf{u}$ is assumed to be constantly zero, and the excitation efficiency is assumed to be constant across both compartments.

Since the medium is homogeneous without any boundaries in the z-direction, w_3 and all z-derivatives are constantly zero, hence $\mathbf{F}\mathbf{e}_3 = (0, 0, 1)^T$. Therefore, it suffices to solve equation (5.61) over a 2D section to determine w_1 and w_2. For the intra-axonal compartment, each individual axon is isolated from each other, therefore the trivial solutions discussed in Section 5.3.4 apply to w_1 and w_2. The tensor field $\mathbf{F}_{\text{int}}$ is therefore constant and equal to the non-dimensionalised effective diffusivity tensor, which is given by:

$$\mathbf{D}^{\text{eff}}_{\text{int}} = \mathbf{F}_{\text{int}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (5.62)$$

Meanwhile for the extra-axonal compartment, numerical solution to the cell problem is required. A 2D mesh (Fig. 5.3b) is generated by triangulation for the non-dimensionalised

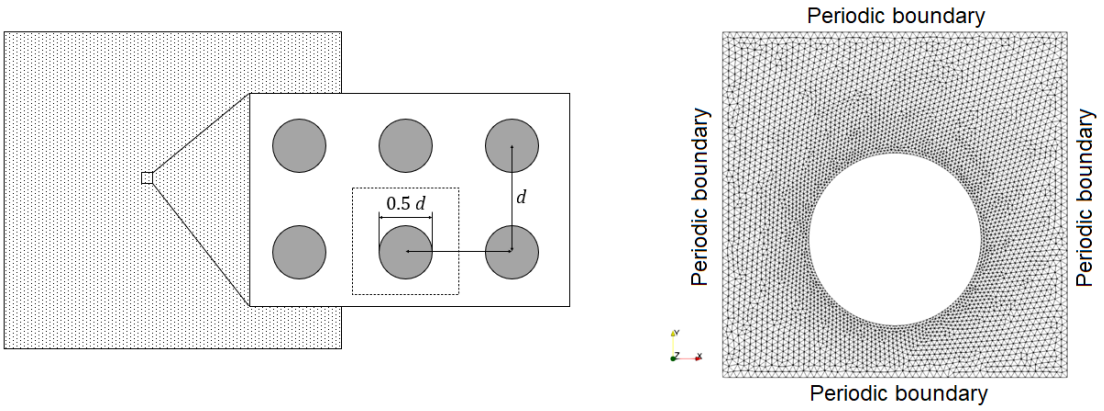

Figure 5.3: (a) Schematic of the porous medium showing periodic cell in microstructure. (b) Two-dimensional mesh generated for the extra-axonal compartment in a non-dimensionalised periodic cell.

periodic cell. The centre of the cylinder is deliberately offset by 0.1 in the y-direction to check if the solver is implemented correctly. The maximum element size in the domain is set to 0.02, with refinement on the curved boundary set to 0.01, which is 4% of the radius of curvature. Based on the mesh, Lagrange elements of degree 1 are used to define the finite-dimensional space V_h. This consists of all periodic, continuous functions which are piecewise linear on each element. The variational equality (5.61) is then implemented in FEniCS to solve for w_1 and w_2, with the respective solutions being shown as colour plots in Fig. 5.4.

From the computed solution for w_1, w_2 and the fact that $w_3 = 0$, the tensor field $\mathbf{F}_{\text{ext}}$ can be computed. The vector fields given by $\mathbf{F}_{\text{ext}}\mathbf{e}_i$ represent the local distribution of flux when a macro-scale potential gradient of unit magnitude is applied in direction $\mathbf{e}_i$. For the x- and y-directions, $\mathbf{F}_{\text{ext}}\mathbf{e}_1$ and $\mathbf{F}_{\text{ext}}\mathbf{e}_2$ are shown as vector plots in Fig. 5.4, whereas the domain homogeneity in the z-direction implies $\mathbf{F}_{\text{ext}}\mathbf{e}_3 = (0, 0, 1)^T$. The volume average of $\mathbf{F}_{\text{ext}}$ over the reference cell yields the effective diffusivity tensor for the extra-axonal

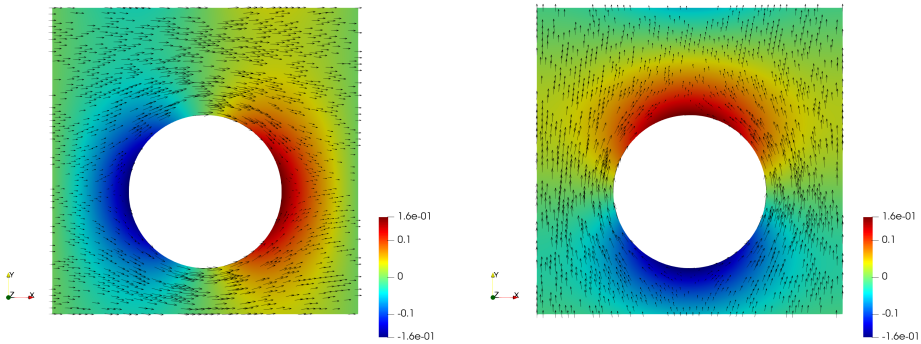

Figure 5.4: x- (left) and y-components (right) of the solution to the cell problem in the extra-axonal compartment (w_1 and w_2) shown as colour plots, with the induced flux distribution ($\mathbf{F}_{\text{ext}}\mathbf{e}_1$ and $\mathbf{F}_{\text{ext}}\mathbf{e}_2$) overlaid as vector plots.

compartment normalised against D:

$$\mathbf{D}_{\text{ext}}^{\text{eff}} = \langle \mathbf{F}_{\text{ext}} \rangle = \begin{bmatrix} 0.8359 & 0 & 0 \\ 0 & 0.8359 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{5.63}$$

From expression (5.53), the overall signal density is therefore given by:

$$\frac{\langle M_{xy} \rangle}{\langle M_0 \rangle} = f_{\text{int}}\, e^{-R_{2,\text{int}}} \langle \mathbf{q} \rangle \cdot \mathbf{D}_{\text{int}}^{\text{eff}} D \langle \mathbf{q} \rangle + (1 - f_{\text{int}})\, e^{-R_{2,\text{ext}}} \langle \mathbf{q} \rangle \cdot \mathbf{D}_{\text{ext}}^{\text{eff}} D \langle \mathbf{q} \rangle, \tag{5.64}$$

where $f_{\text{int}} = 0.1963$ is the volume fraction of the intra-axonal compartment, and the effective diffusion tensors $\mathbf{D}_{\text{int}}^{\text{eff}}, \mathbf{D}_{\text{ext}}^{\text{eff}}$ are given in (5.62), (5.63) respectively.

5.5 Effective diffusion tensors of real white matter geometries

In this section, the numerical method described in Section 5.4 is applied to realistic white matter geometries derived from microscopy data from mouse white matter. This allows us to investigate the effective diffusion tensors of the various compartments in an actual biological sample, which is representative of the heterogeneous media encountered in a clinical imaging setting.

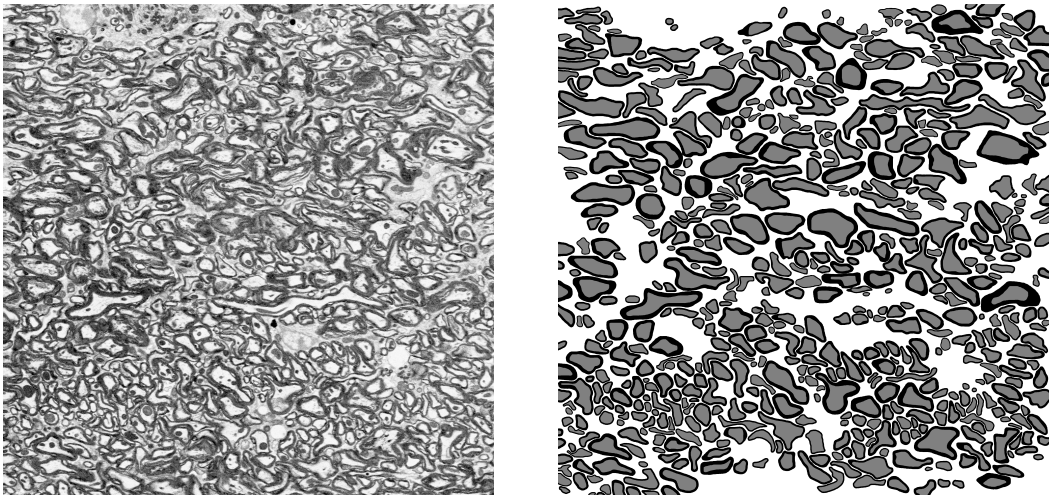

(a) TEM image (b) Segmentation map

Figure 5.5: Transmission electron microscopy image of mouse white matter, and its hand segmentation by Xu et al. [121] for the extra-axonal, myelin sheath, and intra-axonal compartments. Image matrix: 4000 ×4000, resolution: 7.1 nm isotropic.

5.5.1 Data preparation and mesh generation

The data have been collected and previously published by Xu et al. [121] in a study on a different topic of susceptibility-weighted MR imaging. The image data, obtained using transmission electron microscopy (TEM), were collected in a 4000 ×4000 matrix at an isotropic resolution of 7.1 nm, and were subsequently hand segmented into extra-axonal, myelin sheath and axonal compartments by the original authors, as shown in Figure 5.5. In this thesis, both the image and its segmentation map are obtained with the original authors' permission, and are further processed for use in numerical calculations of diffusion MRI. Since the microscopy data is only available as one 2D cross-section slice, similar to the previous section, we assume the microstructure to be uniform and infinite in the longitudinal direction. This also means it again suffices to solve the cell problem in two dimensions, as the longitudinal components of the diffusion tensor are trivial.

From the full 4000 ×4000 segmented map, 100 smaller 1000 ×1000 samples were cropped out for mesh generation. The reason for cropping out multiple smaller samples, rather than using the whole image, is two-fold. Firstly, since only one microscopy image is available, if the full image is used the result will only be one single diffusion tensor with

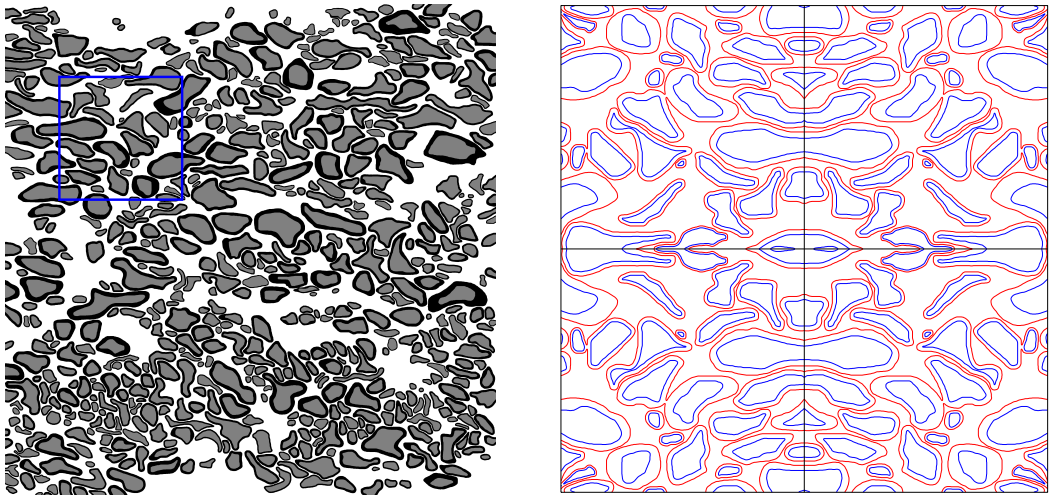

(a) Cropped area (b) Domain augmentation

Figure 5.6: An example of a 1000 ×1000 sample being cropped from the original image, and then augmented by appending its mirror image to form a larger periodic domain. The bottom-right of (b) corresponds to the cropped area in (a).

no comparison possible. Therefore, there is limited meaningful analysis which could be done. Secondly, due to limited computing resources, the mesh quality has to be reduced if the full image is used. For the 1000 ×1000 samples, the minimum mesh size can go to as low as 0.5px near the boundaries, so the features of the boundary contours can be resolved, especially when they are in close vicinity to each other. However, due to the limited amount of memory, this is not possible if the full image is used. Meanwhile, if multiple smaller samples are chosen from the full image, a bigger data set is effectively created, to allow for at least some form of analysis and statistics to gain some insight about the data.

For each cropped sample, the boundaries between the compartments are obtained by smoothening and thresholding the segmentation map image. In order to enforce the periodic boundary conditions, the samples are augmented by appending its own mirror images, so that a larger periodic domain four times its size is formed, as shown in Figure 5.6. The motivation of this process is to generate periodic domains which are still representative of the realistic geometry. Unfortunately, in 21 out of the 100 cropped areas, the intra-axonal space of some cells cut across a corner of the domain, meaning that after including the mirror images they form a ring instead. Not only does this lead to mesh

generation failure, it is also not a realistic topology for cells. These samples are therefore discarded, leaving a total of $n = 79$ samples to be considered. For more details of the automated processing and mesh generation pipeline using `gmsh`, please see Appendix A.

5.5.2 Numerical results and discussion

The cell problem is then solved over the 79 meshed domains using the FEniCS implementation described in Section 5.4, in order to solve for the effective diffusion tensors. In each of the meshed domains, only the extra-axonal compartments is considered, as the myelin sheath and intra-axonal compartments are not connected in both x- and y-directions, hence their associated effective diffusion tensor is trivial.

For the extra-axonal space, the statistics of the volume fraction and the two principal components of the effective diffusion tensor are shown in Figure 5.7. As we can see, there is some considerable variation in terms of the values of the effective parameters over the set of samples, however this is to be expected given the nature of the data. By inspecting

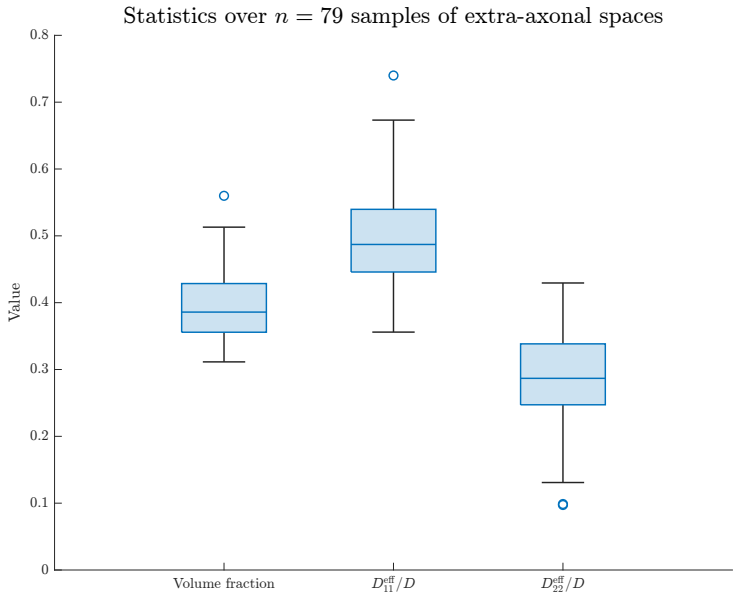

Figure 5.7: Statistics over $n = 79$ samples for the volume fractions and principal components of effective diffusion tensor of the extra-axonal compartment. Apart from the outliers (blue dot), defined by data more than 1.5 times the interquartile range away from a quartile, the bars denote the maximum and minimum values, whereas the box shows the median and lower/upper quartiles.

the full microscopy image, one can see some spatial variation in terms of packing density of the axons across the 4000 $\times$4000 sample, which leads to a corresponding variation in the effective diffusion tensors in the smaller samples, depending on their locations. However, if the sample size could be increased, the variation in these parameters were expected to decrease, as the samples should include information from a larger range of locations and average them out. Unfortunately, due to the limited data available, such analysis has not been possible. Nevertheless, one can still conclude from this limited dataset that a 1000 $\times$1000 sample size, which equates to 7.1 μm $\times$7.1 μm, is unlikely to be large enough to be a good representative elementary volume (REV) for the white matter microstructure. This is as would be expected, as these samples often only cover as few as three cells in one direction. It is also noted that this size is still much smaller than the characteristic length scale of an imaging voxel, and thus in practice much more effective homogeneity can be expected over the voxel scale.

5.6 Summary and conclusion

To summarise this chapter, multi-scale homogenisation has been applied to model diffusion MRI in heterogeneous media in the long-time, small q-value regime. Although advection is not expected to play a part in practical scenarios, a weak advection term has been considered in the model as well. This novel modelling approach allows for a simultaneous consideration of both transport and magnetic heterogeneities in the medium. The result of the homogenisation analysis reveals that the effect of magnetic heterogeneities integrate exactly to zero, whereas the heterogeneous diffusion barriers average out into compartment-specific effective diffusion tensors. The overall diffusion MRI signal is therefore described by a multi-exponential decay expression, which is similar to existing multi-compartmental models.

In addition, the homogenisation analysis reveals the mathematical origin of the effective diffusion tensors, an area which is under-explored in existing literature. The link

between microstructure properties and the macro-scale multi-exponential signal behaviour is the cell problem, which is a PDE posed over a representative elementary volume. This cell problem can be solved numerically by a finite element method, as presented in Section 5.4. This method is applied on realistic microstructure geometries obtained from transmission electron microscopy data of mouse white matter. However, the limited amount of data means that there are limited meaningful comparisons and quantitative conclusion that can be drawn from the numerical results. Nevertheless, a working example of a method to deduce the effective diffusion tensors from a segmented image has been presented. Given greater availability of data in the future, this work could easily be extended to evaluating the effective diffusion tensors from multiple samples, at different locations or of tissues with different pathology, allowing for a basis of more extensive biophysical modelling.

Chapter 6

Boundary layer model of DW-SE for the short-time regime

6.1 Introduction

In this chapter, the short-time regime of diffusion-weighted spin-echo (DW-SE) pulse sequences in porous media is considered. The definition of this short-time regime is that the time scale between the spatial modulation and demodulation of spins by the gradient pulses is shorter than the characteristic time for diffusion to occur over the microstructure length scale. In terms of the q-t parameter space, this corresponds to the region to the left of the central region where all scales from various processes are balanced, as shown in Figure 6.1.

In the current literature, the main model, which is frequently quoted as an universal expression for short diffusion times, is based on a diffusion propagator approach and taking a $t \to 0$ limit (see Chapter 2, Section 2.4.6). However, there are two implicit assumptions from this approach which do not reflect a practical DW-SE experiment. Firstly, by solely taking a $t \to 0$ limit but keeping all other scales constant, particularly that of the q-vector, the resulting contrast is poor since the signal magnitude scales as $\exp(-q^2Dt) = 1 - O(t)$. As explained in Chapter 3, Section 3.3.2, in order for both the signal and contrast to

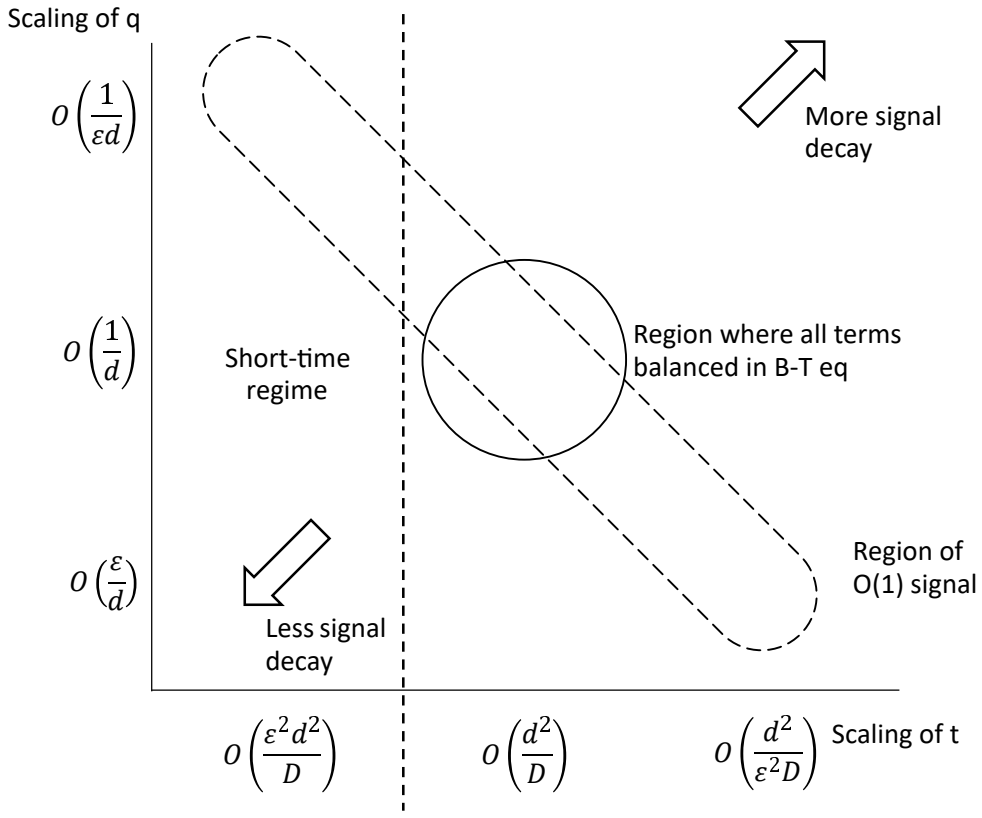

Figure 6.1: Definition of the small-time regime on the q-t parameter space.

scale as $O(1)$, the product q^2Dt needs to scale as $O(1)$ too, which is represented by by the top-left to bottom-right strip in the q-t parameter space (Figure 6.1). Therefore, at the same time as taking the $t \to 0$ limit, one should concurrently adjust the scale of the q-vector. Secondly, the diffusion propagator approach assumes a q-vector which is constant, in other words spins are spatially modulated and demodulated instantaneously by an infinitely strong but short gradient pulse. This clearly is an unphysical assumption, as the rate of modulation is limited by the maximum gradient that the scanner hardware can offer.

In the following sections, these two limitations of current models will be addressed. Instead of approaching the problem using a diffusion propagator approach, we start from the Bloch-Torrey equation itself and consider its asymptotic behaviour as we take the short time limit. A boundary layer type of behaviour is expected in this regime, as the effect of the heterogeneities, which are the diffusion barriers on the boundary surfaces, should not have the time to travel beyond the local boundary layer. We first investigate

the effect of the first limitation whilst keeping the assumption of constant q-vector, and then time-varying q-vector profiles will be considered next.

6.2 Model formulation

6.2.1 Governing equations and scale separation parametrisation

We start from the full, dimensional form of the co-rotating, co-decaying Bloch-Torrey equation, with no-flux boundary conditions and constant initial condition:

$$\frac{\partial S}{\partial t} + (\nabla - i\mathbf{q}) \cdot (-\mathbf{D}(\nabla - i\mathbf{q})S) = 0 \quad \text{in } \Omega, \tag{6.1}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla - i\mathbf{q})S) = 0 \quad \text{on } \partial\Omega, \tag{6.2}$$

$$S = M_0 \quad \text{at } t = 0. \tag{6.3}$$

In the above equations, $S(\mathbf{x}, t)$ is the magnetisation quantity under the co-rotating, co-decaying transformation; at $t = 0$ the initial magnetisation M_0 is assumed to be constant. The q-vector, $\mathbf{q}(t)$, is the spatial modulation wave-vector which is assumed to be spatially uniform. Although the diffusivity of the medium $\mathbf{D}$ can be in general anisotropic in the Bloch-Torrey equation, here in this chapter it is modelled to be constant and isotropic, i.e. $\mathbf{D} = D\mathbf{I}$, for ease of analysis.

In this chapter, we are interested in the behaviour of this equation at a time scale which is much shorter than that associated with the microstructure. For a characteristic length scale corresponding to the microstructure pore sizes $l_c = d$, and characteristic time scale t_c defined by the MRI pulse sequence itself, the scale separation condition for the short time regime can be expressed as the following relation:

$$t_c = \varepsilon^2 \frac{d^2}{D}. \tag{6.4}$$

The small parameter $\varepsilon \ll 1$ indexes the scale separation between the time scales; the short-time regime therefore corresponds to the $\varepsilon \to 0$ asymptotic limit. Meanwhile, as for

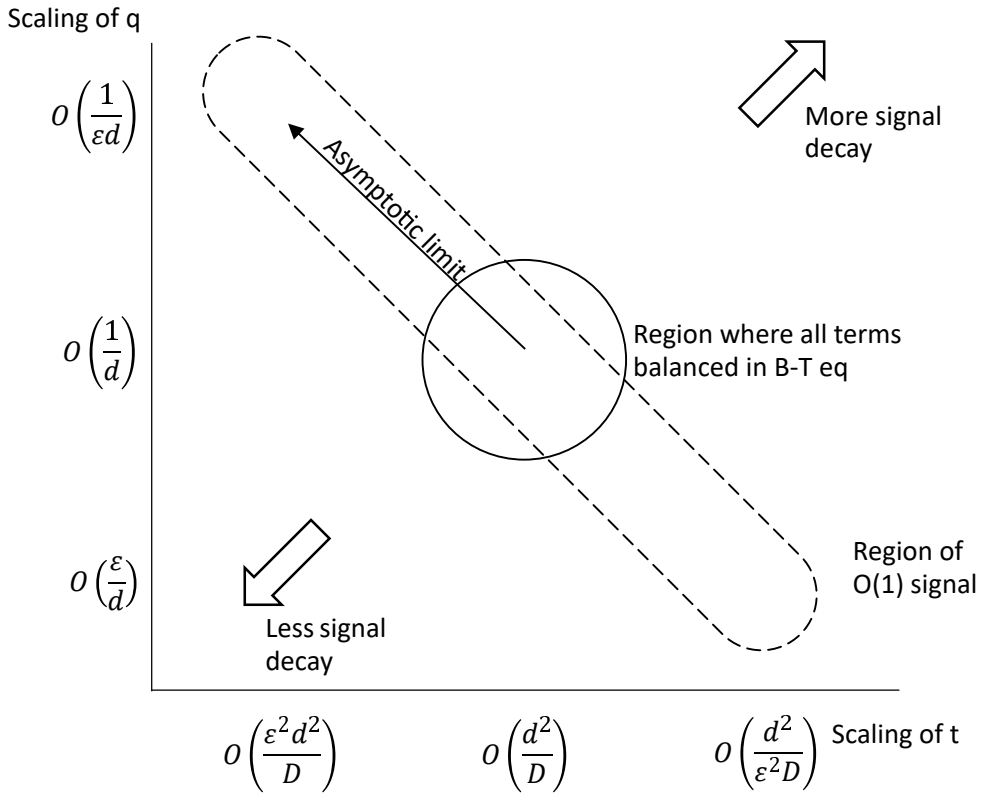

Figure 6.2: The asymptotic limit depicted on the q-t parameter space. The central region corresponds to the regime when all terms in the original equation are balanced; the asymptotic limit considered lies along the path which results in an O(1) signal.

the q-vector, its natural scaling is:

$$q_c = \varepsilon^{-1} d^{-1} = (D t_c)^{-1/2}. \tag{6.5}$$

This choice leads to the time-derivative and $\mathbf{q}$ terms in (6.1) to be similar in magnitude. It also means the signal decay rate, which is proportional to $|\mathbf{q}|^2 Dt$, is at $O(1)$ magnitude. Finally, the magnetisation is normalised against the initial condition, i.e. $S_c = M_0$. After rescaling, the non-dimensionalised equation and its boundary and initial conditions, with the relative ratio of magnitudes parametrised by ε, are given by:

$$\frac{\partial S}{\partial t} - (\varepsilon\nabla - i\mathbf{q}) \cdot (\varepsilon\nabla - i\mathbf{q}) S = 0 \text{ in } \Omega, \tag{6.6}$$

$$\mathbf{n} \cdot (\varepsilon\nabla - i\mathbf{q}) S = 0 \text{ on } \partial\Omega, \tag{6.7}$$

$$S = 1 \text{ at } t = 0. \tag{6.8}$$

This asymptotic regime can be plotted on the digram of the q-t parameter space, as shown in Figure 6.2. In the diagram, the central region corresponds to the regime when all terms in the Bloch-Torrey equation are balanced in magnitude; the $\varepsilon \to 0$ limit in (6.6)–(6.8) thus corresponds to the arrow going diagonally to the top-left direction.

6.2.2 Asymptotic boundary layer

In the non-dimensional equation (6.6), the highest-order spatial derivative is multiplied by the small parameter ε^2, contributing to a singular perturbation at the $\varepsilon \to 0$ limit. This suggests a boundary layer behaviour in space, where there is a thin region near the boundary in which the solution $S(\mathbf{x}, t)$ varies rapidly. Indeed, for the bulk of the spatial domain, the asymptotic solution can be obtained by ignoring the boundary condition and setting $\varepsilon = 0$ in (6.6), leading to the 'outer' problem:

$$\frac{\partial S_{\text{out}}}{\partial t} + \mathbf{q} \cdot \mathbf{q}\, S_{\text{out}} = 0, \tag{6.9}$$

$$S_{\text{out}} = 1 \text{ at } t = 0. \tag{6.10}$$

which has solution:

$$S_{\text{out}}(\mathbf{x}, t) = \exp\left(- \int_0^t |\mathbf{q}|^2 \, \mathrm{d}\tau \right). \tag{6.11}$$

It can be noted that in the outer problem, all the spatial derivatives vanish under the $\varepsilon = 0$ approximation, which corresponds to the medium behaving as if it is homogeneous. This means in the outer region, the behaviour of the signal is not influenced by the effects of inhomogeneities caused by impermeable boundaries.

However, this outer solution clearly does not satisfy the boundary condition given by (6.7). Instead, in the close vicinity to the boundary $\partial\Omega$, the presence of the boundary condition modifies the solution. In this 'inner' region, which is also known as the boundary layer, the solution varies rapidly in space so the spatial derivative terms in (6.6) only become comparable to other terms in the equation after being multiplied by ε.

To analyse the behaviour of the solution within the inner boundary layer, we can

switch to a local spatial co-ordinate system and then rescale the spatial derivatives to become $O(1)$ in magnitude. Consider a point on the boundary surface $\partial\Omega$. Using that point as the origin, we can define the local spatial co-ordinates as $(x_{\text{loc}}, y_{\text{loc}}, z_{\text{loc}})$ with x_{loc} pointing in the direction of the inner normal. A scaled inner spatial coordinate system can then be defined:

$$(x_{\text{in}}, y_{\text{in}}, z_{\text{in}}) = \left(\frac{x_{\text{loc}}}{\varepsilon}, \frac{y_{\text{loc}}}{\varepsilon}, \frac{z_{\text{loc}}}{\varepsilon}\right).$$

This choice of rescaling is to make the rescaled spatial derivatives to be of a similar magnitude to the other terms in the equation. Assuming the boundary surface $\partial\Omega$ is smooth, then under the above rescaling, the local surface element asymptotically approaches a flat surface. Therefore, to model the boundary layer behaviour, the original problem (6.6)–(6.8) can be rescaled, and then posed on the $x_{\text{in}} > 0$ semi-infinite half-space, with an additional far-field condition that as $x_{\text{in}} \to \infty$ the inner solution should asymptotically match the outer bulk solution at that point. This leads to the following inner problem:

$$\frac{\partial S_{\text{in}}}{\partial t} - (\nabla_{\text{in}} - i\mathbf{q}) \cdot (\nabla_{\text{in}} - i\mathbf{q}) S = 0, \tag{6.12}$$

$$\mathbf{e}_n \cdot (\nabla_{\text{in}} - i\mathbf{q}) S_{\text{in}} = 0 \text{ at } x_{\text{in}} = 0, \tag{6.13}$$

$$S_{\text{in}} \to S_{\text{out}} \text{ at } x_{\text{in}} \to \infty, \tag{6.14}$$

$$S_{\text{in}} = 1 \text{ at } t = 0, \tag{6.15}$$

where $\mathbf{e}_n$ is the unit vector of the inward normal at the boundary surface. Since $\mathbf{q}(\mathbf{x}, t)$ only varies in the global scale but not locally in the boundary layer, in the inner problem $\mathbf{q}$ is only a function of time and does not depend on the rescaled inner spatial co-ordinates. Meanwhile, instead of acting as an independent variable, the global spatial variable parametrises the inner problem through $\mathbf{q}$ and S_{out}.

Furthermore, it can be seen that none of the equations and boundary conditions in (6.12)–(6.15) has any dependence on the y_{in} and z_{in} variables, in other words within the boundary layer the inner solution only varies in the direction normal to the boundary surface. Therefore, the spatial derivative operators can be simplified as $\nabla_{\text{in}} = \mathbf{e}_n(\partial/\partial x_{\text{in}})$.

The inner problem therefore reduces to a problem for $S_{\text{in}}(x_{\text{in}}, t)$, a function of just two variables:

$$\frac{\partial S_{\text{in}}}{\partial t} - \left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right)\left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right) S_{\text{in}} + q_t^2 S_{\text{in}} = 0, \tag{6.16}$$

$$\left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right) S_{\text{in}} = 0 \text{ at } x_{\text{in}} = 0, \tag{6.17}$$

$$S_{\text{in}} \to \exp\left(-\int_0^t (q_n^2 + q_t^2)\,\mathrm{d}\tau\right) \text{ at } x_{\text{in}} \to \infty, \tag{6.18}$$

$$S_{\text{in}} = 1 \text{ at } t = 0. \tag{6.19}$$

Here q_n and q_t respectively represent the normal and tangential components of $\mathbf{q}$ with respect to the boundary surface. If we further use a substitution $U(x_{\text{in}}, t) = S_{\text{in}} \exp\left(\int_0^t q_t^2 \mathrm{d}\tau\right)$, we reduce it to the following problem:

$$\frac{\partial U}{\partial t} - \left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right)\left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right) U = 0, \tag{6.20}$$

$$\left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right) U = 0 \text{ at } x_{\text{in}} = 0, \tag{6.21}$$

$$U \to \exp\left(-\int_0^t q_n^2\,\mathrm{d}\tau\right) \text{ at } x_{\text{in}} \to \infty, \tag{6.22}$$

$$U = 1 \text{ at } t = 0. \tag{6.23}$$

This effectively means that the boundary layer behaviour can be characterised by a family of problems in the two variables of x_{in} and t, parametrised by $q_n(t)$ which is the normal component of the q-vector locally at the boundary surface.

After solving for the normalised inner problem, a model can be formulated for the volume-integrated signal over a voxel, in conjunction with the solution in the outer region. In the absence of boundaries, the signal will be given by S_{out} throughout the voxel; for each surface element on the boundary $\partial\Omega$, its associated boundary layer causes a deficit in signal decay. Such signal deficit per unit surface area can be computed through integration of the inner solution:

$$\int_{\text{BL}} (S_{\text{in}} - S_{\text{out}})\mathrm{d}x_{\text{loc}} = \varepsilon \exp\left(-\int_0^t q_t^2 \mathrm{d}\tau\right) \int_0^\infty (U - U_\infty)\,\mathrm{d}x_{\text{in}}. \tag{6.24}$$

The resulting imaging signal over a voxel, assuming that it is a uniform integration over the voxel volume, is therefore:

$$S_{\text{voxel}} = \int_{\Omega} S_{\text{out}}\,\mathrm{d}V + \int_{\partial\Omega}\left(\varepsilon \exp\left(-\int_0^t q_t^2 \mathrm{d}\tau\right)\int_0^{\infty}(U-U_{\infty})\,\mathrm{d}x_{\text{in}}\right)\mathrm{d}A. \tag{6.25}$$

6.2.3 Relating boundary layer variables to dimensional quantities

So far in the above sections, the boundary layer model has been formulated in terms of non-dimensional variables, which eventually characterises the signal behaviour near a boundary surface by the 'inner problem', described by the partial differential equations of (6.20)–(6.23). It is useful to establish the relations between the dimensional quantities and the non-dimensional variables, both in the inner and outer co-ordinates, as this allows us to go back into describing the image signal in dimensional quantities after obtaining the solution to the inner problem.

In particular, regarding the inner variables (x_{in}, t), they relate to the dimensional space and time variables $(x_{\text{dim}}, t_{\text{dim}})$ by:

$$x_{\text{in}} = \varepsilon^{-1}x = (q_c d)x = q_c x_{\text{dim}}, \tag{6.26}$$

$$t = t_c^{-1}t_{\text{dim}} = \varepsilon^{-2}d^{-2}Dt_{\text{dim}} = q_c^2 Dt_{\text{dim}}. \tag{6.27}$$

The relevant dimensionless groups for the inner problem are therefore $(q_c x_{\text{dim}}, q_c^2 Dt_{\text{dim}})$, which are dependent on the characteristic q-value and the diffusion coefficient, but not the geometrical length-scale of the microstructure. This is a direct result of the length scale separation assumption between the characteristic diffusion length and the much longer microstructure length-scale.

Regarding the total signal over a voxel in terms of dimensional quantities, Equation (6.25) can be re-dimensionalised by multiplying corresponding characteristic quantities to

become:

$$S_{\text{voxel}} = M_0 \int_\Omega \exp\left(-D\int_0^{t_{\text{dim}}} |\mathbf{q}_{\text{dim}}|^2 \,\mathrm{d}\tau\right) \mathrm{d}V + M_0 \int_{\partial\Omega} q_c^{-1} \exp\left(-D\int_0^{t_{\text{dim}}} q_t^2 \mathrm{d}\tau\right) \left(\int_0^\infty (U - U_\infty)\,\mathrm{d}(x_{\text{in}})\right) \mathrm{d}A. \quad (6.28)$$

6.3 Solutions to inner problem for constant q-vector

In this section, the specific case when the q-vector is constant in time is considered. This corresponds to a short-pulse gradient assumption on the DW-SE sequence, as the q-vector jumps instantaneously between 0 and its constant value at both the start and the end of the pulse sequence. Clearly, this assumption is not physical, as the finite magnitude of magnetic fields does not allow for the instantaneous spatial modulation of spins. Nevertheless, this is the assumption employed in the vast majority of existing modelling results of surface-to-volume effects in diffusion imaging, therefore in this section we seek to use the boundary layer methodology in order to compare results with the literature, as well as to provide further insight to the signal behaviour in the vicinity of microstructure boundaries.

An interesting observation is that when q_n is constant in time, the problem posed by Equations (6.20)–(6.23) becomes self-similar under the following re-scaling of variables:

$$U_{\text{new}}(x,t) = U\left(q_n\, x_{\text{in}}, q_n^2 t\right). \quad (6.29)$$

The inner problem for the boundary layer then reduces to a single problem, which effectively corresponds to setting $q_n = 1$:

$$\frac{\partial U}{\partial t} - \left(\frac{\partial}{\partial x} - i\right)\left(\frac{\partial}{\partial x} - i\right) U = 0, \quad (6.30)$$

$$\left(\frac{\partial}{\partial x} - i\right) U = 0 \text{ at } x = 0, \quad (6.31)$$

$$U \to \exp(-t) \text{ at } x \to \infty, \quad (6.32)$$

$$U = 1 \text{ at } t = 0. \quad (6.33)$$

In other words, the behaviours of all boundary layers at different (time-unvarying) q-vectors correspond to the same solution for $q_n = 1$, except that they are scaled according to the strength of the normal component of the q-vector. In terms of dimensional quantities, the invariant set of variables in (6.29) correspond to $(x, t) = (q_n x_{\text{dim}}, q_n^2 D t_{\text{dim}})$. However, the self-similarity does not apply when $q_n(t)$ depends on time, as Equations (6.20)–(6.22) are no longer invariant under the transformation of variables.

Two analytical methods, one based on the Laplace transform and the other using a series expansion, will be used to investigate the normalised boundary layer problem. The inner problem will also be solved numerically using a finite difference method as well. The results will then be discussed and compared with existing literature.

6.3.1 Laplace transform method

Firstly, the solution to the inner boundary layer for a constant q-vector is analysed using a Laplace transform method. By applying a Laplace transform on the time variable in Equations (6.20)–(6.23), the time-dependent behaviour is encapsulated in the Laplace variable s, resulting to an ODE for the Laplace transformed function $U_{\mathcal{L}}(x_{\text{in}}, s) = \mathcal{L}\{U(x_{\text{in}}, t)\}$:

$$\frac{\mathrm{d}^2 U_{\mathcal{L}}(x_{\text{in}}, s)}{\mathrm{d}x_{\text{in}}^2} - 2iq_n \frac{\mathrm{d}U_{\mathcal{L}}(x_{\text{in}}, s)}{\mathrm{d}x_{\text{in}}} - (s + q_n^2) U_{\mathcal{L}}(x_{\text{in}}, s) + 1 = 0, \tag{6.34}$$

$$\frac{\mathrm{d}U_{\mathcal{L}}(x_{\text{in}}, s)}{\mathrm{d}x_{\text{in}}} - iq_n U_{\mathcal{L}}(x_{\text{in}}, s) = 0 \text{ at } x_{\text{in}} = 0, \tag{6.35}$$

$$U_{\mathcal{L}}(x_{\text{in}}, s) \to \frac{1}{s + q_n^2} \text{ at } x_{\text{in}} \to \infty. \tag{6.36}$$

In these transformed equations, the time derivatives are replaced by the multiplication by s, whereas the spatial partial derivatives become ordinary derivatives, as the Laplace variable s acts as a parameter instead. The solution to this second-order ODE is given by:

$$U_{\mathcal{L}}(x_{\text{in}}, s) = -\frac{iq_n}{\sqrt{s}(s + q_n^2)} e^{(iq_n - \sqrt{s})x} + \frac{1}{s + q_n^2}. \tag{6.37}$$

It can be noted that the particular integral corresponds to the far-field solution $U_\infty(t)$, hence the deficit between the inner and outer solutions is:

$$\Delta U_{\mathcal{L}}(x_{\text{in}}, s) = \mathcal{L}\left\{U(x_{\text{in}}, t) - U_\infty(t)\right\} = -\frac{iq_n}{\sqrt{s}(s+q_n^2)} e^{(iq_n - \sqrt{s})x_{\text{in}}}. \tag{6.38}$$

The $\sqrt{s}$ within the exponent of the above expression makes it difficult to find an analytical expression for its inverse Laplace transform. Therefore, rather than solving for the exact solution for $\Delta U(x_{\text{in}}, t)$ point-wise in the spatial variable, its integral over the entire inner boundary layer can be investigated instead. The deficit of $U_{\mathcal{L}}(x_{\text{in}}, s)$ in the boundary layer is given by the following expression:

$$\int_0^\infty \Delta U_{\mathcal{L}}(x_{\text{in}}, s)\mathrm{d}x_{\text{in}} = \frac{iq_n}{\sqrt{s}(iq_n - \sqrt{s})(s+q_n^2)}. \tag{6.39}$$

The inverse Laplace transform of this expression yields:

$$\begin{aligned}\int_0^\infty (U - U_\infty)\mathrm{d}x_{\text{in}} &= \frac{1}{2i\pi} \lim_{R\to+\infty} \int_{\sigma-iR}^{\sigma+iR} \frac{iq_n}{\sqrt{s}(iq_n - \sqrt{s})(s+q_n^2)} e^{st}\mathrm{d}s \quad (\sigma > 0) \\ &= -\sqrt{\frac{t}{\pi}} + \exp(-q_n^2 t)\left(\frac{1+2q_n^2 t}{2q_n}\operatorname{erfi}(q_n\sqrt{t}) - iq_n t\right), \end{aligned} \tag{6.40}$$

where $\operatorname{erfi}(w) = -i\operatorname{erf}(iw)$ is the imaginary error function. It is noted that if the above expression is multiplied by q_n, then it becomes a function solely of the variable $q_n\sqrt{t}$. This is due to the self-similar nature of the problem with the re-scaled variables of $(q_n x_{\text{in}}, q_n^2 t)$, as demonstrated in (6.30)–(6.33), hence the signal deficit integral over the boundary layer can be expressed as a one-parameter function of $q_n\sqrt{t}$:

$$\int_0^\infty (U - U_\infty)\mathrm{d}(q_n x_{\text{in}}) = F(q_n\sqrt{t}), \tag{6.41}$$

$$\text{where } F(w) = -\frac{w}{\sqrt{\pi}} + \exp(-w^2)\left(\frac{1+2w^2}{2}\operatorname{erfi}(w) - iw^2\right). \tag{6.42}$$

The signal deficit at the $t \to \infty$ limit can be evaluated using the final value theorem as well:

$$\lim_{t\to\infty}(U - U_\infty) = \lim_{s\to 0} s\,\Delta U_{\mathcal{L}}(x_{\text{in}}, s) = 0. \tag{6.43}$$

Hence, the signal deficit indeed should decay to zero at infinite time. However, it must be noted that the boundary layer analysis is based on the short time assumption, as described from the definition of the characteristic time in (6.4). In other words, it is assumed that the far-field is unperturbed from any boundary conditions and behave as if it is an infinite medium. In reality, at the $t \to \infty$ limit, the far-field will also be influenced by other surfaces; in this case the behaviour can no longer be analysed as an isolated boundary layer.

Finally, returning to dimensional quantities, the dimensional signal deficit per unit surface area is given by:

$$\begin{aligned} \varepsilon d \int_0^\infty (U - U_\infty) \mathrm{d}x_{\mathrm{in}} &= q_c^{-1} q_n^{-1} \int_0^\infty (U - U_\infty) \mathrm{d}(q_n x_{\mathrm{in}}) \\ &= q_n^{*-1} F(q_n^* \sqrt{Dt^*}), \end{aligned} \tag{6.44}$$

where the asterisks denote the dimensional variables. The total voxel-integrated signal is therefore:

$$S_{\mathrm{voxel}} = \int_\Omega M_0 \exp\left(-|\mathbf{q}^*|^2 Dt^*\right) \mathrm{d}V + \int_{\partial\Omega} M_0 \exp\left(-q_t^{*2} Dt^*\right) q_n^{*-1} F(q_n^* \sqrt{Dt^*}) \mathrm{d}A. \tag{6.45}$$

In this expression, the subsequent image formation process is assumed to be ideal, in other words all locations within the voxel region has an equal weighting in the integration. If the image formation process is imperfect, then the above equation can be easily modified by including the local signal amplification and phase shift in the integrals.

A particular observation from the above expression is that if the initial condition is constant and image formation process is ideal, the imaginary part of the signal integrates exactly to zero. This is because the imaginary part of the voxel signal is given by:

$$\begin{aligned} \mathrm{Im}(S_{\mathrm{voxel}}) &= \int_{\partial\Omega} M_0 \exp\left(-q_t^{*2} Dt^*\right) q_n^{*-1} \exp\left(-q_n^{*2} Dt^*\right) (-q_n^{*2} Dt^*) \mathrm{d}A \\ &= \int_{\partial\Omega} -M_0 \exp\left(-|\mathbf{q}^*|^2 Dt^*\right) |\mathbf{q}^*| Dt^* \cos\theta \mathrm{d}A = 0. \end{aligned} \tag{6.46}$$

Here θ is the angle between the surface normal and the applied q-vector at each point on $\partial\Omega$. Since $\partial\Omega$ is formed by a union of volume-bounding surfaces, the integral $\int_{\partial\Omega} \cos\theta \mathrm{d}A$

is constantly zero. This result can be extended to when the voxel signal encounters a coherent phase shift at all locations, however when there is incoherent phase shift of spins when the voxel signal is formed, the imaginary part does not necessarily integrate to zero.

6.3.2 Series expansion method for small q-value

A number of modelling results in literature, such as those of [43, 77, 87], are based on the asymptotic limit of bringing $t \to 0$ whilst keeping q constant. These models therefore correspond to going horizontally left on the q-t parameter space (Fig 6.3), instead of keeping the signal decay rate at $O(1)$ and going diagonally to the top-left. Indeed, these models are interested in a different regime and the resulting signal decay rate is therefore much smaller than $O(1)$ too. Nevertheless it is possible to take a further $q \to 0$ limit from the above boundary layer analysis to relate the two regimes. On the q-t parameter space, this corresponds to first taking the diagonal asymptotic limit, and then going downwards by taking a small q limit, thus ending up in the same region, as shown in the dashed arrows in Figure 6.3.

Their are two motivations for this analysis; firstly it allows the comparison with existing models derived using a random walk or a diffusion propagator approach. Secondly, existing models, as well as the Laplace transform method described in Section 6.3.1, only give the signal deficit integral. Here we attempt to go into further details, and aim to derive equations (at least at the q $\to$ 0 limit) describing the spatial variation of the signal within the boundary layer.

Starting from Equations (6.20)–(6.23), we can investigate the behaviour at the $q_n \to 0$ limit by assuming a series expansion ansatz:

$$U = U^{(0)} + q_n U^{(1)} + q_n^2 U^{(2)} + \cdots \tag{6.47}$$

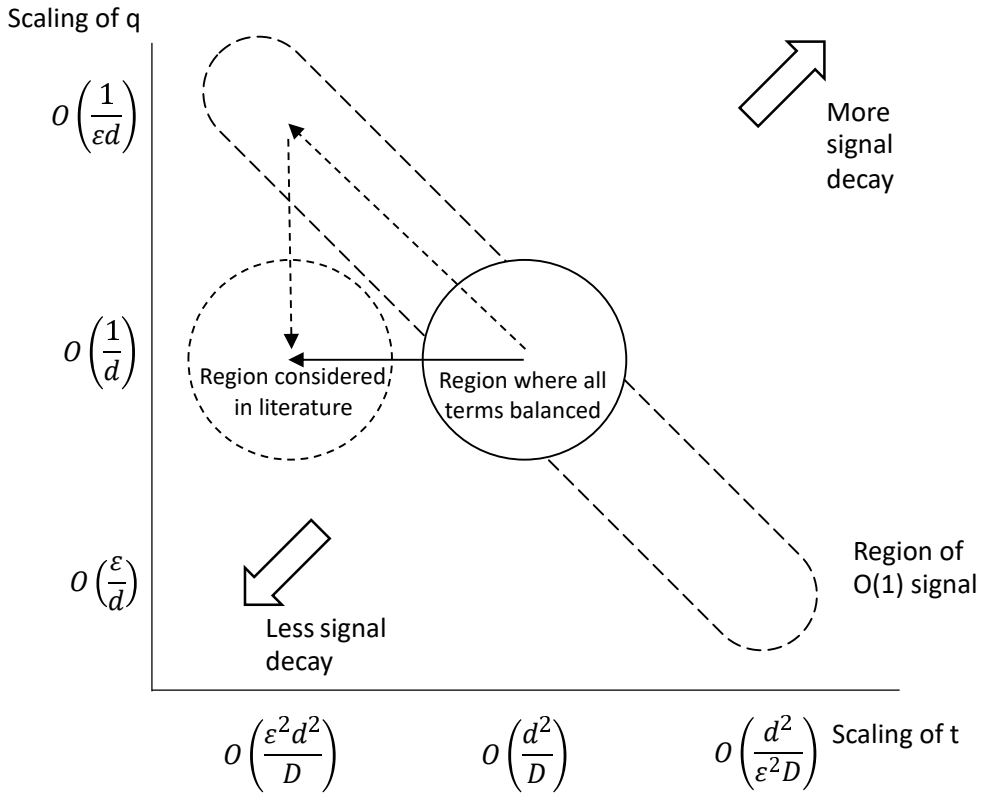

Figure 6.3: In the q-t parameter space, existing surface-to-volume models, such as [43, 77, 87], are derived by considering $t \to 0$ whilst keeping q constant (solid arrow). The same regime can also be reached by first doing the boundary layer analysis, and then taking a further $q \to 0$ limit (dashed arrows).

Substituting this into (6.20)–(6.23), the leading order problem gives:

$$\frac{\partial U^{(0)}}{\partial t} - \frac{\partial^2 U^{(0)}}{\partial x_{\text{in}}^2} = 0, \tag{6.48}$$

$$\frac{\partial U^{(0)}}{\partial x_{\text{in}}} = 0 \text{ at } x_{\text{in}} = 0, \tag{6.49}$$

$$U^{(0)} \to 1 \text{ at } x_{\text{in}} \to \infty, \tag{6.50}$$

$$U^{(0)} = 1 \text{ at } t = 0, \tag{6.51}$$

which has a trivial solution of $U^{(0)} = 1$.

Proceeding to the next order gives the $U^{(1)}$ problem:

$$\frac{\partial U^{(1)}}{\partial t} - \frac{\partial^2 U^{(1)}}{\partial x_{\text{in}}^2} = 0, \tag{6.52}$$

$$\frac{\partial U^{(1)}}{\partial x_{\text{in}}} = iU^{(0)} \text{ at } x_{\text{in}} = 0, \tag{6.53}$$

$$U^{(1)} \to 0 \text{ at } x_{\text{in}} \to \infty, \tag{6.54}$$

$$U^{(1)} = 0 \text{ at } t = 0. \tag{6.55}$$

This represents a diffusion problem in the $x_{\text{in}} > 0$ semi-infinite domain, with a constant flux condition at $x_{\text{in}} = 0$ as $U^{(0)}$ is constant. The exact solution for $U^{(1)}$ can be derived using a similarity solution. This is because if $U^{(1)}(x_{\text{in}}, t)$ is a solution, then for any arbitrary constant a, the function $a^{-1}U^{(1)}(ax_{\text{in}}, a^2 t)$ also satisfies (6.52)–(6.55). Therefore, by letting $\eta = x_{\text{in}}/\sqrt{t}$ be the similarity group, the solution takes the form of $U^{(1)} = t^{1/2}F^{(1)}(\eta)$, and the PDE problem in (6.52)–(6.55) reduces to a second-order ODE with boundary conditions:

$$\frac{\mathrm{d}^2 F^{(1)}}{\mathrm{d}\eta^2} + \frac{\eta}{2}\frac{\mathrm{d}F^{(1)}}{\mathrm{d}\eta} - \frac{F^{(1)}}{2} = 0, \tag{6.56}$$

$$\frac{\mathrm{d}F^{(1)}}{\mathrm{d}\eta} = i \text{ at } \eta = 0, \tag{6.57}$$

$$F^{(1)} \to 0 \text{ at } \eta \to \infty. \tag{6.58}$$

We note that the $t = 0$ initial condition (6.55) is automatically satisfied by the form of the similarity solution, so it does not appear in the ODE problem. The general solution to equations of the form of (6.56) can be expressed in terms of repeated integrals of the error function [1]:

$$F^{(1)} = A\,\mathrm{i}_1\,\mathrm{erfc}\left(\frac{\eta}{2}\right) + B\,\mathrm{i}_1\,\mathrm{erfc}\left(-\frac{\eta}{2}\right), \tag{6.59}$$

where A and B are constants and $\mathrm{i}_1\,\mathrm{erfc}(z) = \int_z^\infty \mathrm{erfc}(w)\,\mathrm{d}w$ is the first integral of the error function. Substituting in the boundary conditions yields $A = -2i$ and $B = 0$, hence the solution for $U^{(1)}$ is:

$$U^{(1)} = -2i\sqrt{t}\;\mathrm{i}_1\,\mathrm{erfc}\left(\frac{x_{\text{in}}}{2\sqrt{t}}\right). \tag{6.60}$$

We can see that while the leading order solution $U^{(0)}$ is real, the first order correction $U^{(1)}$ is a purely imaginary function. This means that when we consider the magnitude, the contribution from $U^{(1)}$ is in fact of order $O(q_n^2)$, as it is effectively a purely quadrature component of correction. Therefore, in order to evaluate the leading order correction to the magnitude signal, the second order problem needs to be considered as well. The $U^{(2)}$ problem is given by:

$$\frac{\partial U^{(2)}}{\partial t} - \frac{\partial^2 U^{(2)}}{\partial x_{\text{in}}^2} + 2i\frac{\partial U^{(1)}}{\partial x_{\text{in}}} + U^{(0)} = 0, \tag{6.61}$$

$$\frac{\partial U^{(2)}}{\partial x_{\text{in}}} = iU^{(1)} \text{ at } x_{\text{in}} = 0, \tag{6.62}$$

$$U^{(2)} \to -t \text{ at } x_{\text{in}} \to \infty, \tag{6.63}$$

$$U^{(2)} = 0 \text{ at } t = 0. \tag{6.64}$$

Substituting in the solutions for $U^{(0)}$ and $U^{(1)}$ leads to:

$$\frac{\partial U^{(2)}}{\partial t} - \frac{\partial^2 U^{(2)}}{\partial x_{\text{in}}^2} - 2\,\text{erfc}\left(\frac{z}{2\sqrt{t}}\right) + 1 = 0, \tag{6.65}$$

$$\frac{\partial U^{(2)}}{\partial x_{\text{in}}} = \frac{2}{\sqrt{\pi}}\sqrt{t} \text{ at } x_{\text{in}} = 0, \tag{6.66}$$

$$U^{(2)} \to -t \text{ at } x_{\text{in}} \to \infty, \tag{6.67}$$

$$U^{(2)} = 0 \text{ at } t = 0. \tag{6.68}$$

This problem can again be solved using the similarity group $\eta = x_{\text{in}}/\sqrt{t}$, but here the solution takes the form $U^{(2)} = tF^{(2)}(\eta)$ to match the $\sqrt{t}$-varying flux boundary condition in (6.66). The corresponding ODE problem and boundary conditions are given by:

$$\frac{\mathrm{d}^2 F^{(2)}}{\mathrm{d}\eta^2} + \frac{\eta}{2}\frac{\mathrm{d}F^{(2)}}{\mathrm{d}\eta} - F^{(2)} + 2\,\text{erfc}\left(\frac{\eta}{2}\right) - 1 = 0, \tag{6.69}$$

$$\frac{\mathrm{d}F^{(2)}}{\mathrm{d}\eta} = \frac{2}{\sqrt{\pi}} \text{ at } \eta = 0, \tag{6.70}$$

$$F^{(2)} \to -1 \text{ at } \eta \to \infty. \tag{6.71}$$

The complimentary function for (6.69) again can be represented by repeated integrals of

the complementary error function:

$$F_{\mathrm{cf}}^{(2)} = A\,\mathrm{i}_2\,\mathrm{erfc}\left(\frac{\eta}{2}\right) + B\,\mathrm{i}_2\,\mathrm{erfc}\left(-\frac{\eta}{2}\right), \tag{6.72}$$

where $\mathrm{i}_2\,\mathrm{erfc}(z) = \int_z^\infty \mathrm{i}_1\,\mathrm{erfc}(w)\,\mathrm{d}w$ is the second repeated integral of the error function. Meanwhile, the particular integral for the ODE problem is given by:

$$F_{\mathrm{pi}}^{(2)} = 2\,\mathrm{erfc}\left(\frac{\eta}{2}\right) - 1. \tag{6.73}$$

Applying the boundary conditions yields $A = -8$ and $B = 0$, therefore:

$$U^{(2)} = t\left(-8\ \mathrm{i}_2\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right) + 2\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right) - 1\right), \tag{6.74}$$

which gives the second order correction to the inner solution.

By gathering the zeroth, first and second order terms, an approximate solution accurate up to $O(q_n^3)$ is therefore given by:

$$U = 1 - iq_n\sqrt{t}\,\mathrm{i}_1\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right) + q_n^2 t\left(-8\,\mathrm{i}_2\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right) + 2\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right) - 1\right) + O(q_n^3). \tag{6.75}$$

Note that the above series expression for U is always a function of the following two variables: $q_n\sqrt{t}$ and $x_{\mathrm{in}}/\sqrt{t}$; this is again consistent with the fact that when the q-vector is constant in time, the inner problem is self-similar under the variable transformation described in (6.29).

As discussed previously, the existence of the inner boundary layer is due to mismatch between the boundary condition and the far-field solution. The presence of the non-conforming boundary condition propagates over time through the growing boundary layer, and this phenomenon can be shown by considering the difference between the inter and outer solutions:

$$\begin{aligned} U - U_\infty &= U - (1 - q_n^2 t) + O(q_n^3) \\ &= -iq_n\sqrt{t}\,\mathrm{i}_1\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right) + q_n^2 t\left(-8\,\mathrm{i}_2\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right) + 2\,\mathrm{erfc}\left(\frac{x_{\mathrm{in}}}{2\sqrt{t}}\right)\right) + O(q_n^3). \end{aligned} \tag{6.76}$$

In terms of the width of the boundary layer, it can be seen that it is proportional to $\sqrt{t}$. Physically, this can be explained by the fact that information propagates spatially, through diffusion processes, to a distance scaling as $\sqrt{t}$, so the presence of the boundary condition causes the boundary layer to widen at a corresponding rate.

Considering the deficit between the inner and outer solutions over the boundary layer, we can integrate the approximation for $(U - U_\infty)$ over x_{in}:

$$\int_0^\infty (U - U_\infty) \mathrm{d}x_{\text{in}} = -iq_n t + \frac{4}{3\sqrt{\pi}} q_n^2 \, t^{3/2} + O(q_n^3). \tag{6.77}$$

Then, including back the tangential component q_t and multiplying the length scale rescaling factor for the boundary layer:

$$\begin{aligned} \varepsilon \int_0^\infty (S_{\text{in}} - S_{\text{out}}) \mathrm{d}x_{\text{in}} &= \varepsilon \left(1 - q_t^2 t + O(q^4)\right) \left(-iq_n t + \frac{4}{3\sqrt{\pi}} q_n^2 \, t^{3/2} + O(q^3)\right) \\ &= \varepsilon \left(-iq_n t + \frac{4}{3\sqrt{\pi}} q_n^2 \, t^{3/2}\right) + O(\varepsilon q^3). \end{aligned} \tag{6.78}$$

The eventual voxel-integrated signal is hence:

$$S_{\text{voxel}} = \int_\Omega S_{\text{out}} \, \mathrm{d}V + \int_{\partial\Omega} \varepsilon e^{-\int_0^t q_t^2 \mathrm{d}\tau} \left(-iq_n t + \frac{4}{3\sqrt{\pi}} q_n^2 \, t^{3/2}\right) + O(\varepsilon q^3). \tag{6.79}$$

6.3.3 Comparison between exact and series solutions

To compare the results derived using the two different methods, a series expansion in q_n can be performed to the exact solution for the signal deficit (6.40):

$$\int_0^\infty (U - U_\infty) \mathrm{d}x = -iq_n t + \frac{4}{3\sqrt{\pi}} q_n^2 \, t^{3/2} + iq_n^3 t^2 - \frac{16}{15\sqrt{\pi}} q_n^4 \, t^{5/2} + \cdots, \tag{6.80}$$

the first two terms of which are identical to the result of (6.77). This means the two results are consistent with each other at the $q_n \to 0$ limit.

For larger q_n values, the comparison between the series and exact solutions are shown in Figure 6.4, which plots the time evolution of the total signal deficit. It can be clearly seen that the series solution approximation is only good for small values of q_n; when $q_n^2 Dt = 0.2$ the real and imaginary parts of the series solution are over-estimated by

17.2% and 22.1% already. Meanwhile, for $q_n^2 Dt > 0.2$, the two solutions diverge very quickly. In fact, even though the exact solution is bounded and slowly decay to 0 at infinity, the series solution grows infinitely over time. Therefore, care must be taken when considering the validity of the series solution, depending on the magnitude of the key parameter of $q_n^2 Dt$.

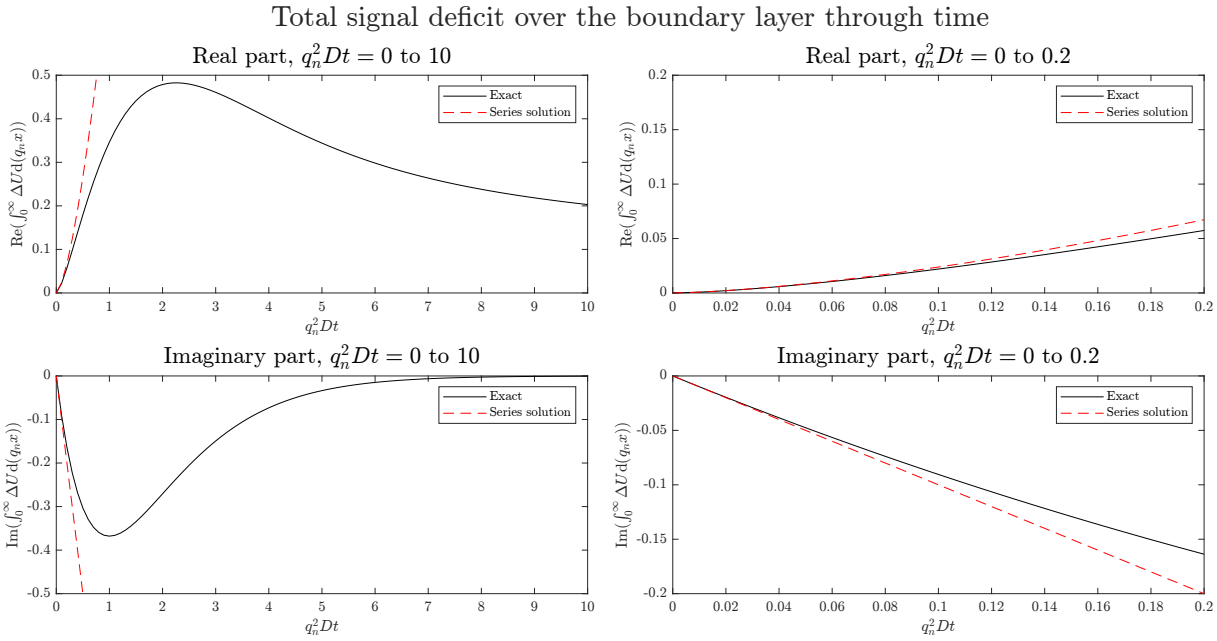

Figure 6.4: Comparison between the exact solution, derived using Laplace transform, and the series solution to the inner problem for constant q_n.

6.3.4 Numerical solution using a finite difference method

In this section, a numerical solution to the inner problem when q_n is constant is presented. The motivation for this is that although the Laplace transform method in Section 6.3.1 gives an analytical expression for the total deficit integrated over the boundary layer, the spatial profile of the solution is difficult to obtain analytically. Meanwhile, the series solution in Section 6.3.2 gives the spatial profile, but is only valid at the $q_n \to 0$ limit. Therefore, in order to compute how the boundary layer evolves over time at $q_n = O(1)$, the inner problem is solved numerically.

As explained above, regardless of the value of q_n, the inner problem is self-similar and reduces to a single problem in the variables of $(q_n x_{\mathrm{in}}, q_n^2 t)$, as given by (6.30)–(6.33).

For clarity of presentation, this set of normalised variables are denoted as (x, t) in this section, and again we note that their relation to the dimensional quantities are $(x, t) = (q_n x, q_n^2 D t)_{\text{dim}}$. To simplify the far-field condition, Equations (6.30)–(6.33) can be rewritten using the difference between U and the far-field solution, i.e. $\Delta U(x, t) = U - \exp(-t)$:

$$\frac{\partial \Delta U}{\partial t} - \frac{\partial^2 \Delta U}{\partial x^2} + 2i\frac{\partial \Delta U}{\partial x} + \Delta U = 0, \tag{6.81}$$

$$\frac{\partial \Delta U}{\partial x} - i\Delta U - i\exp(-t) = 0 \text{ at } x = 0, \tag{6.82}$$

$$\Delta U \to 0 \text{ at } x \to \infty, \tag{6.83}$$

$$\Delta U = 0 \text{ at } t = 0. \tag{6.84}$$

The spatial domain for the above equation is $x \in [0, \infty)$, with a far-field decay condition at infinity. However, to implement the discretisation approximation using finite difference, it is much more convenient to have a finite interval instead as the spatial domain. Therefore, the far-field decay condition is replaced by a Dirichlet boundary condition at $x = x_{\max}$, which is justified as long as $x_{\max}$ is set to be far away enough from the non-zero behaviour local to the $x = 0$ boundary. Since the boundary layer should widen at a diffusive rate of $\sqrt{t}$, for a time window $t \in [0, t_{\max}]$ in which the numerical solution is sought, the Dirichlet boundary condition is justified if $x_{\max}/\sqrt{t_{\max}} \gg 1$.

The solution domain, given by $(x, t) \in [0, x_{\max}] \times [0, t_{\max}]$, can then be discretised into $(m + 1) \times (n + 1)$ equally spaced grid points, with the grid spacing given by $\delta x = x_{\max}/m$ and $\delta t = t_{\max}/n$. We therefore seek to solve for ΔU_j^k at all $(j, k) \in [0, m] \times [0, n]$, which approximates the value of $\Delta U(j\delta x, k\delta t)$ in the continuous solution. Using the backward Euler scheme, which is unconditionally stable, the approximations to (6.81)–(6.84) are

given by:

$$\frac{\Delta U_j^k - \Delta U_j^{k-1}}{\delta t} - \frac{\Delta U_{j+1}^k - 2\Delta U_j^k + \Delta U_{j-1}^k}{\delta x^2} + 2i\frac{\Delta U_{j+1}^k - \Delta U_{j-1}^k}{2\delta x} + \Delta U_j^k = 0 \quad \text{for } (j,k) \in [1, m-1] \times [1, n], \tag{6.85}$$

$$\frac{\Delta U_1^k - \Delta U_0^k}{\delta x} - i\Delta U_0^k - i\exp(-k\delta t) = 0 \text{ for } k \in [1, n], \tag{6.86}$$

$$\Delta U_m^k = 0 \text{ for } k \in [0, n], \tag{6.87}$$

$$\Delta U_j^0 = 0 \text{ for } j \in [0, m]. \tag{6.88}$$

The linear equations (6.85)–(6.88) can then be iteratively solved by incrementing the time step k. At each time step ($k \geq 1$), Equations (6.85) and (6.86) stipulate a matrix equation in the following form:

$$\mathbf{A}\left[\Delta U_0^k \ \cdots \ \Delta U_{m-1}^k\right]^T = \mathbf{b}.$$

In particular, since most of the entries in the $m \times m$ matrix $\mathbf{A}$ are zero, the above matrix equation can be solved efficiently using a sparse algorithm. In this implementation, the in-built `sparse` functionality in MATLAB (ver. R2022a, Mathworks Inc) is used.

For this implementation, the chosen discretisation parameters and its resulting spatial and temporal resolutions are listed in Table 6.1. The computed solution based on the discretisation parameters of Set 1 is plotted in Figure 6.5, where the far-field solution is added back on to show the evolution of $U(x,t)$ over time, with the axes of the plot being the relevant dimensionless groups of $(q_n x, q_n^2 Dt)$ as explained in (6.26) and (6.27).

In terms of accuracy of the numerical solution, a comparison of the total deficit integrated over x with the analytical expression (6.40) is shown in Figure 6.6(a-b). It can be seen that the two solutions agree very closely. The magnitude of the relative error for

Table 6.1: Discretisation parameters used for the finite difference method in (6.85)–(6.88).

Parameter	$x_{\max}$	$t_{\max}$	m	n	δx	δt	$x_{\max}/\sqrt{t_{\max}}$
Set 1	50	10	5,000	5,000	0.01	0.002	15.81
Set 2	10	0.4	5,000	5,000	2e-3	8e-5	15.81

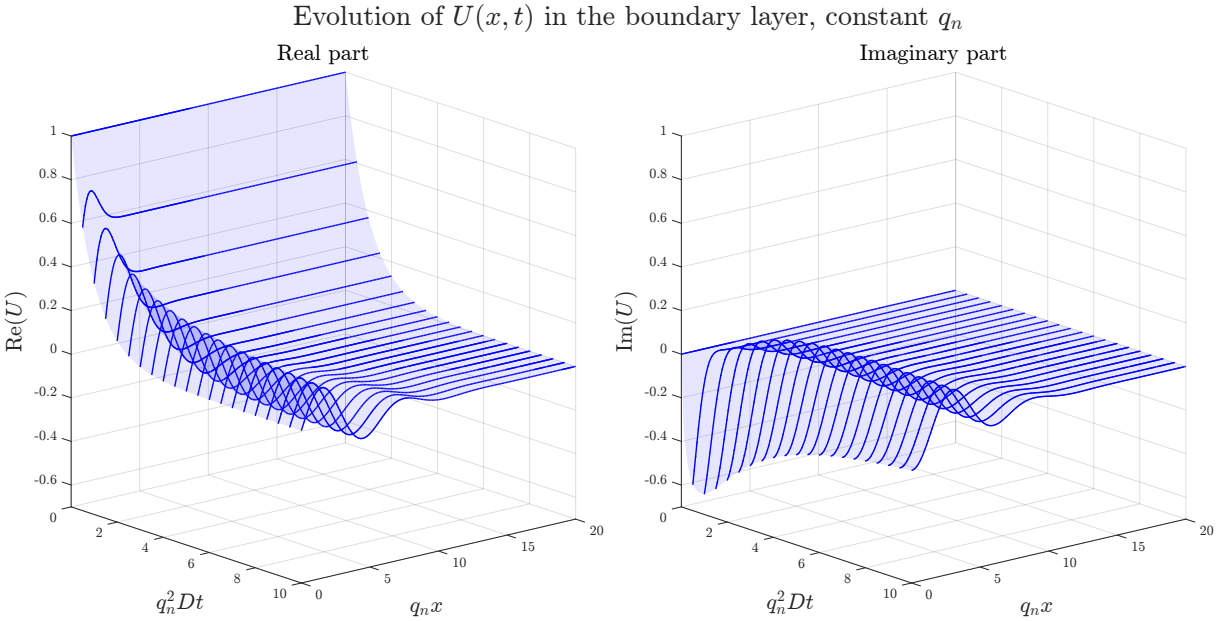

Figure 6.5: The evolution of the normalised signal $U(x,t)$ near a boundary over time, when the normal component of the q-vector (q_n) is constant. Note that the spatial axis is clipped to 20 for ease of visualisation, the numerical solution is computed with $(q_n x)_{\max} = 50$.

each case is also computed and plotted in Figure 6.6(c). For the baseline grid spacing $(m,n) = (5000, 5000)$, the relative error is within 5% for $t \geq 0.05$. However, for small $q_n^2 Dt$, the relative error is larger due to the small initial magnitude of ΔU. A grid refinement test, where the number of grid points is doubled in the spatial and temporal axes respectively, has also been performed for the discretisation parameters of Set 1. As shown in Figure 6.6(c), the error is insensitive to the change in temporal resolution, but it is halved when the spatial resolution is increased by a factor of two.

As we can see in Figure (6.5), the signal $U(x,t)$ is perturbed from its far-field behaviour near the boundary at $x = 0$. The perturbation, in the form of ripples, exist in both real and imaginary components of the signal, and it spreads into the far field over time. The region in which such perturbations are significant is the boundary layer. It can be seen that the peak amplitude of the perturbation initially increases, but subsequently slowly reduces as the ripples spread over space. The same trend can also be observed in the total signal deficit integrated over space, as shown in Figure 6.6(a-b).

Meanwhile, to compare the numerical solution to the series solution in Section 6.3.2,

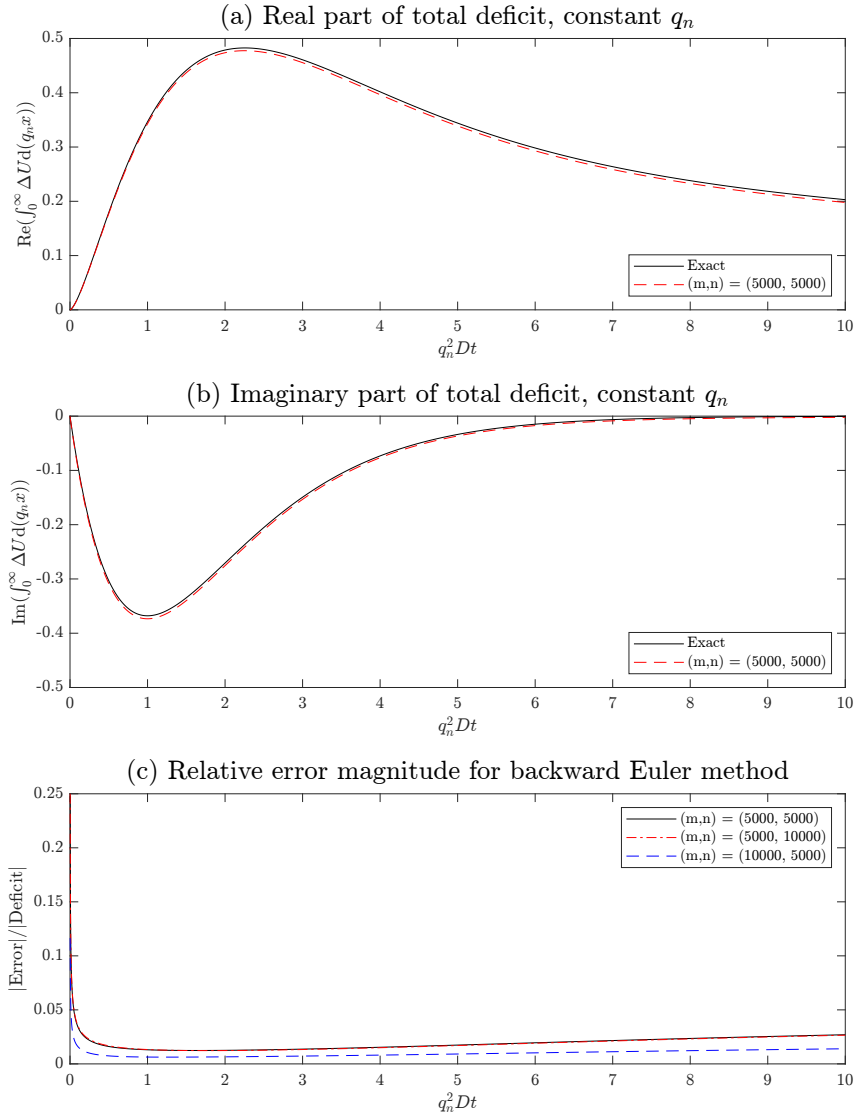

Figure 6.6: Comparison of the total deficit $\int_0^\infty \Delta U \, \mathrm{d}(q_n x)$ between the numerical solution (Parameters Set 1) and exact solution from Eq. (6.40). The magnitudes of relative error are plotted for different grid spacings, with (m, n) being the number of grid points in space and time respectively.

a second set of discretisation parameters (Set 2) is chosen to solve for smaller values of $q_n^2 Dt$. This set of parameters are chosen to improve the solution accuracy for such small $q_n^2 Dt$, with a higher spatial and temporal resolution but for smaller ranges of $x_{\max} = 10$ and $t_{\max} = 0.4$. The performance and relative error plots for Set 2 is shown in Figure 6.7.

The comparison between the time evolutions of the signal $U(x, t)$ for $q_n^2 Dt \leq 0.4$, derived using the numerical method and the series method, is plotted in Figure 6.8. As we can see, at this time range the series solution does replicate the shape of the signal

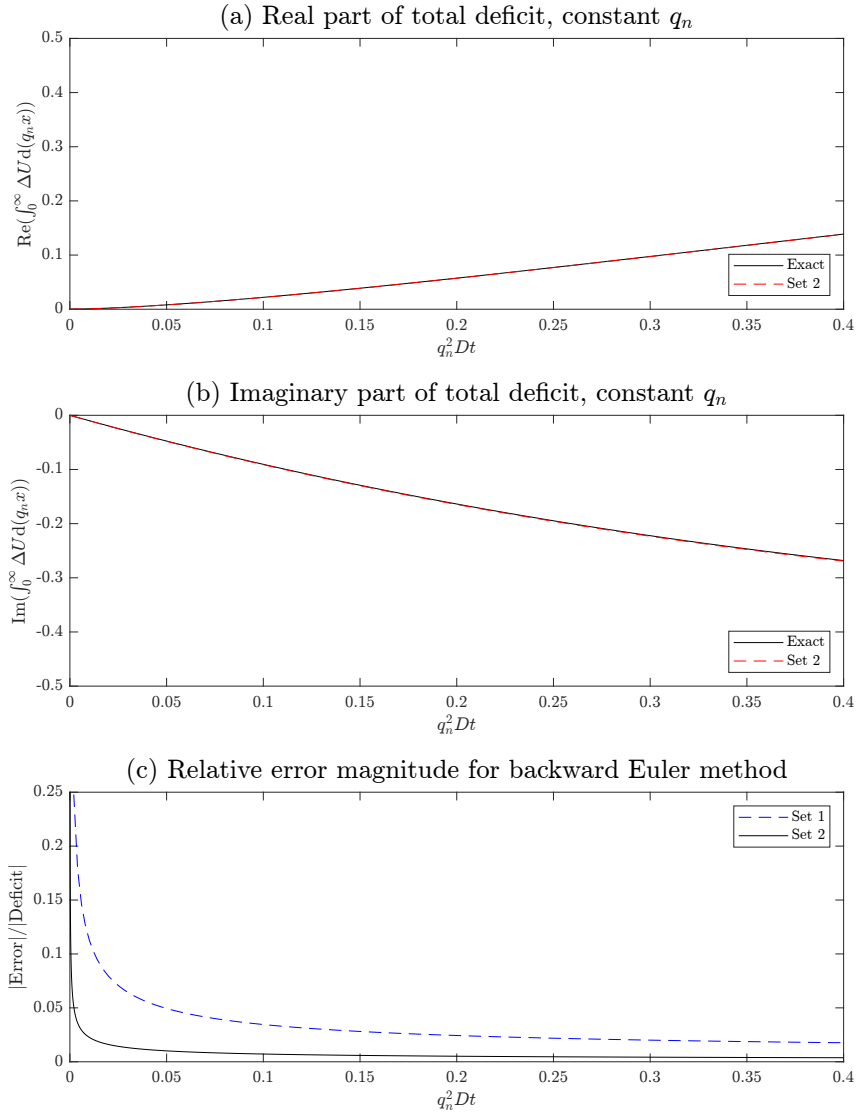

Figure 6.7: Comparison of the total deficit $\int_0^\infty \Delta U \, \mathrm{d}(q_n x)$ between the numerical solution (Parameters Set 2) and exact solution from Eq. (6.40). The accuracy of the solution for small $q_n^2 Dt$ is much improved by increasing resolution, albeit for a smaller range of $x_{\max}$ and $t_{\max}$.

profile, however the overestimation of the signal deficit gets progressively worse over time.

6.3.5 Summary and discussion

In this section, the solution to the inner boundary layer problem has been considered for the constant q-vector case. This has been approached using three different methods, using Laplace transforms, a small q-value series expansion, and by numerical solution respectively. The results from all three methods have been shown to be consistent with

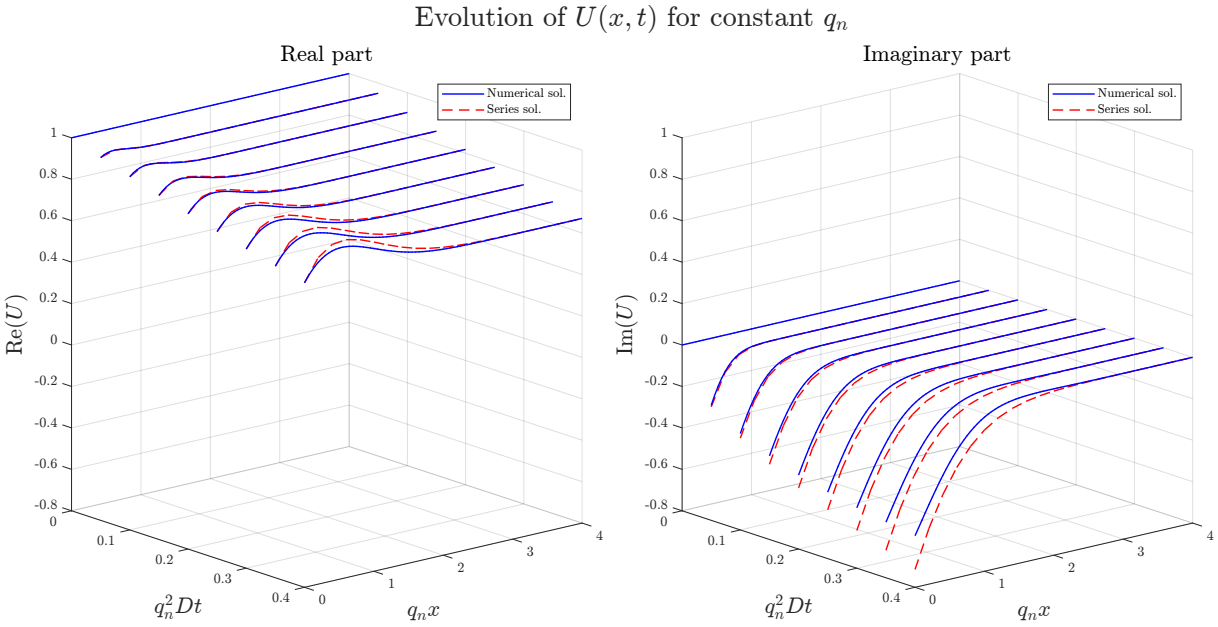

Figure 6.8: Comparison of between the numerical solution (Parameters Set 2) and series solution for $q_n^2 Dt \leq 0.4$.

each other, however each method has its own strengths and weaknesses. The Laplace transform method gives an exact and analytic derivation for the total signal deficit over the boundary layer, however it does not reveal the spatial profile of the boundary layer. Meanwhile, the small q-value expansion yields an analytic expression for the boundary layer signal as a function of both space and time, however it is valid only asymptotically for small q-values. Finally, the numerical solution is the most versatile and describes the spatial and temporal behaviour of the boundary layer at all q-values, however there may not be an analytic expression to describe it.

Even though the constant q-vector condition corresponds to an unphysical assumption, which is the phase of spins being modulated instantaneously by an infinitely strong but short pulse of gradient field, the main reason for considering this case is to compare with existing literature (see Chapter 2, Section 2.4.6). Existing literature has focussed on the magnitude of the complex signal averaged over the entire voxel, whereas the derivations in this section reveal both real and imaginary contributions to the complex signal. If the imaginary parts are ignored, the results from the above derivations are consistent with existing relations. For example, in the small q-value limit, the real part of expression

(6.79) is the same as reported in literature [43, 77, 87]. Meanwhile, for the $O(1)$ q-value case, the real part of (6.40) is also consistent with the relation derived using a Green's function methodology on the diffusion equation [40].

The significance of the additional imaginary part of the boundary layer contribution is dependent on the image acquisition and formation process from the diffusion-modified magnetisation signal. For each image acquisition, the resulting signal at each imaging voxel can be expressed as a convolution between the underlying magnetisation signal and an imaging kernel. So far in an analysis, the acquired image signal in the voxel has been idealised as a uniform integration over the voxel volume, as described in Equations (6.25), (6.79) and (6.40). Under these assumptions, it is shown that the imaginary part of the expressions integrates constantly to zero. The same result applies when all spins encounter a coherent phase shift when the signal is formed as well. However, in reality, the image acquisition methods used for diffusion imaging, typically single-shot echo-planar imaging (EPI), are susceptible to off-resonance artefacts caused by magnetic heterogeneities [104]. This leads to EPI images being distorted, blurred, or encountering phase shift due to such effects [13]. Therefore, the convolution kernel for an imaging voxel is no longer a perfect top-hat function, but instead is spatially varying and complex-valued. In this case, the newly revealed imaginary part of the boundary layer solution will play a part and will need to be considered in the modelling of image signals as well.

6.4 Varying q-vector – rectangular gradient pulse

As mentioned in the previous section, the constant q-vector model corresponds to an unphysical assumption of instantaneous modulation of spins, as this requires an infinite magnetic field strength. In reality, the finite pulse of gradient magnetic field occurs over a considerable amount of time, causing the q-vector to ramp up and down instead.

Figure 6.9 shows the diagram of a diffusion-weighted spin-echo sequence with rectangular gradient pulses. For such pulse sequences, the modulating and demodulating

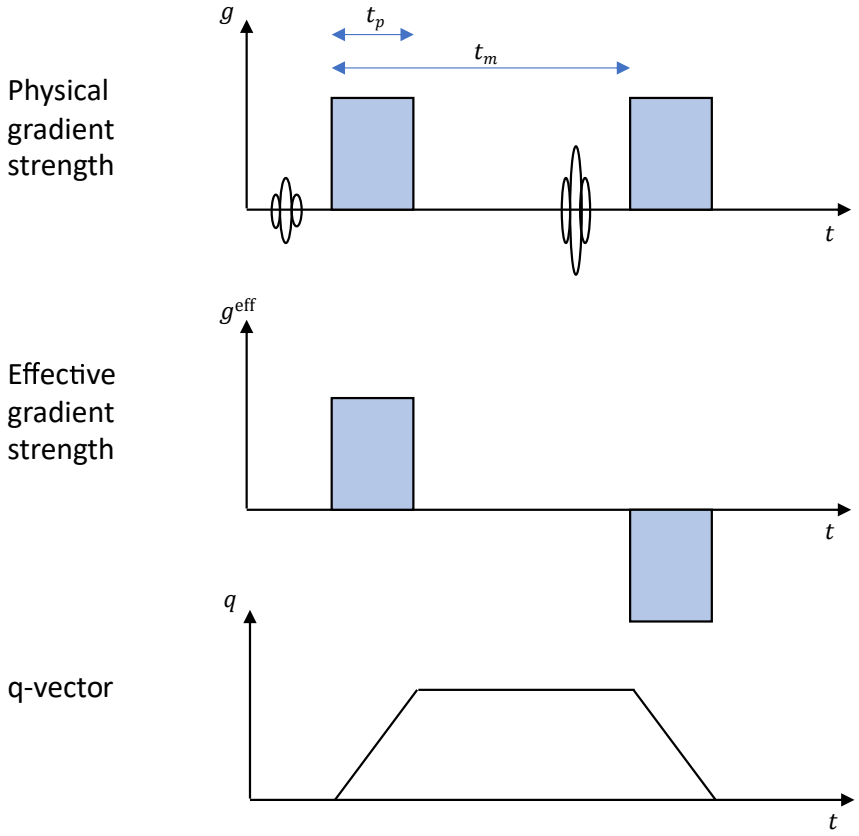

Figure 6.9: Time profile of a DW-SE pulse sequence with rectangular gradient pulses, and its associated q-vector profile.

gradient pulses are constant and have the same duration, hence the time profile of the q-vector, being the time integral of the gradient strength, takes the form of a symmetric trapezium. The timings of such pulse sequences are characterised by the pulse duration and the spacing between them, also known as the mixing time. In clinical MRI literature, the pulse and mixing durations are commonly notated as δ and Δ respectively. However, for clarity of presentation (as δ, Δ denotes discretisation intervals and difference quantities respectively), here we instead notate them as t_p and t_m, as labelled in Figure 6.9.

In this section, the effect of this time response of the q-vector to the short-time regime is investigated through numerically solving the boundary layer inner problem with such q-vector profiles.

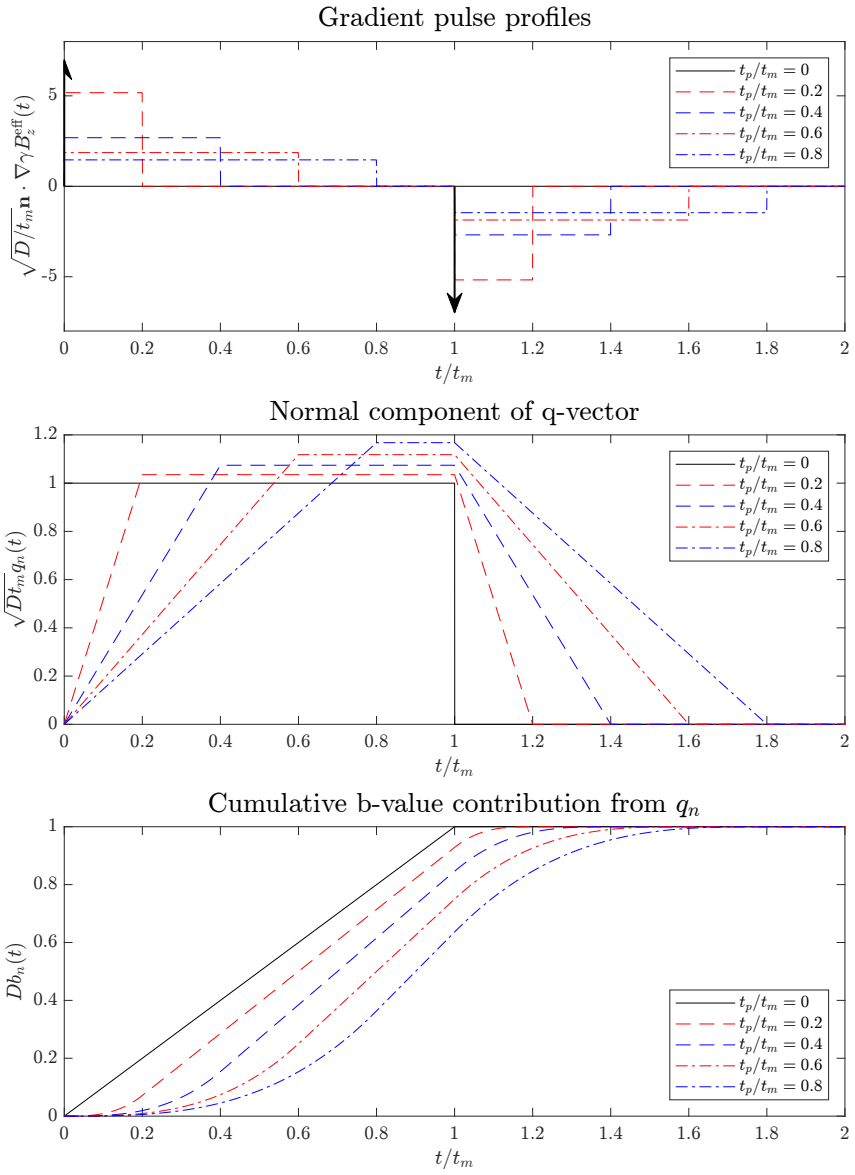

Figure 6.10: Time profiles for the gradient pulses, the normal component of q-vector, and its contribution to the cumulative b-value. Nominal q-value = 1.0. Note the final b-value (i.e. the nominal q-value), rather than the peak q-value, is set to be constant across the profiles.

6.4.1 Numerical study definition

In this study, we consider the solution to the inner problem for the boundary layer for a number of such q-vector profiles. In particular, we are interested in the effect of a non-zero pulse duration t_p given a fixed mixing time t_m, and how the behaviour differs from the constant q-vector case, which corresponds to $t_p = 0$. The inner problem, in terms of

non-dimensional variables is defined by Equations (6.20)–(6.23), which is presented here again for ease of reference:

$$\frac{\partial U}{\partial t} - \left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right)\left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right) U = 0, \tag{6.89}$$

$$\left(\frac{\partial}{\partial x_{\text{in}}} - iq_n\right) U = 0 \text{ at } x_{\text{in}} = 0, \tag{6.90}$$

$$U \to \exp\left(-\int_0^t q_n^2 \, \mathrm{d}\tau\right) \text{ at } x_{\text{in}} \to \infty, \tag{6.91}$$

$$U = 1 \text{ at } t = 0. \tag{6.92}$$

The non-dimensional variables relate to dimensional quantities via the relations of (6.26) and (6.27), normalising with respective to some characteristic quantities. For this problem, the time scale of interest is the mixing time t_m, hence the associated characteristic q-value is given by $q_c = (Dt_m)^{-1/2}$. The normalised q-vector component normal to the surface hence depends on the diffusivity and the orientation of the surface relative to the q-vector. The numerical study therefore investigates the solution to the inner problem for pulse profiles with a range of pulse durations ($t_p/t_m = 0$ to 0.95 at intervals of 0.5), at a number of nominal q-values ($q_{n,\text{nom}} = 0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.5$ and 2.0).

As we can see in (6.91), the outer bulk signal is the exponential of $-\int_0^t q_n^2 \mathrm{d}\tau$, which in effect is the non-dimensional version of the b-value in the direction normal to the surface. Therefore, when comparing between time profiles with different t_p, rather than keeping the q-vector constant, it is more relevant to fix the b-value, and hence the bulk far-field signal. The time profiles for pulse sequences with a nominal q-value of $q_n = 1.0$ are plotted in Figure 6.10. As shown in the plots, the peak q-value increases slightly as t_p increases, so that the final b-value remains constant across the profiles.

6.4.2 Discretisation of inner problem

Similar to the constant q-vector case, the inner problem equations are re-written in terms of the deficit variable $\Delta U = U - U_\infty$, and then discretised using a backward Euler scheme:

$$\frac{\Delta U_j^k - \Delta U_j^{k-1}}{\delta t} - \frac{\Delta U_{j+1}^k - 2\Delta U_j^k + \Delta U_{j-1}^k}{\delta x^2} + 2iq_n^k \frac{\Delta U_{j+1}^k - \Delta U_{j-1}^k}{2\delta x} + (q_n^k)^2 \Delta U_j^k = 0 \quad \text{for } (j,k) \in [1, m-1] \times [1, n], \tag{6.93}$$

$$\frac{\Delta U_1^k - \Delta U_0^k}{\delta x} - iq_n^k \Delta U_0^k - iq_n^k U_\infty^k = 0 \text{ for } k \in [1, n], \tag{6.94}$$

$$\Delta U_m^k = 0 \text{ for } k \in [0, n], \tag{6.95}$$

$$\Delta U_j^0 = 0 \text{ for } j \in [0, m]. \tag{6.96}$$

In these equations, $U_\infty^k = \exp(-\int_0^{k\delta t} q_n^2 \mathrm{d}\tau)$ is the outer bulk signal sampled at the k-th time step. These equations can again be solved by incrementing through the time steps using a sparse matrix inversion algorithm at each time step. The discretisation parameters used are listed in Table 6.2:

Table 6.2: Discretisation parameters used for the finite difference method in (6.93)–(6.96).

Parameter	$x_{\max}$	$t_{\max}$	m	n	δx	δt	$x_{\max}/\sqrt{t_{\max}}$
Value	20	2.0	5,000	5,000	4e-3	4e-4	14.14

6.4.3 Results of numerical study

Figure 6.11 shows the dependence of the total signal deficit over the boundary layer on the pulse duration at different nominal q-vector components normal to the boundary surface. At all nominal q-values, the total signal deficit increases with the pulse duration. The relationship between the signal deficit and t_p/t_m, however, is not perfectly linear. This is expected due to q_n being part of the differential operator in the differential equation itself, leading to a non-linear dependence.

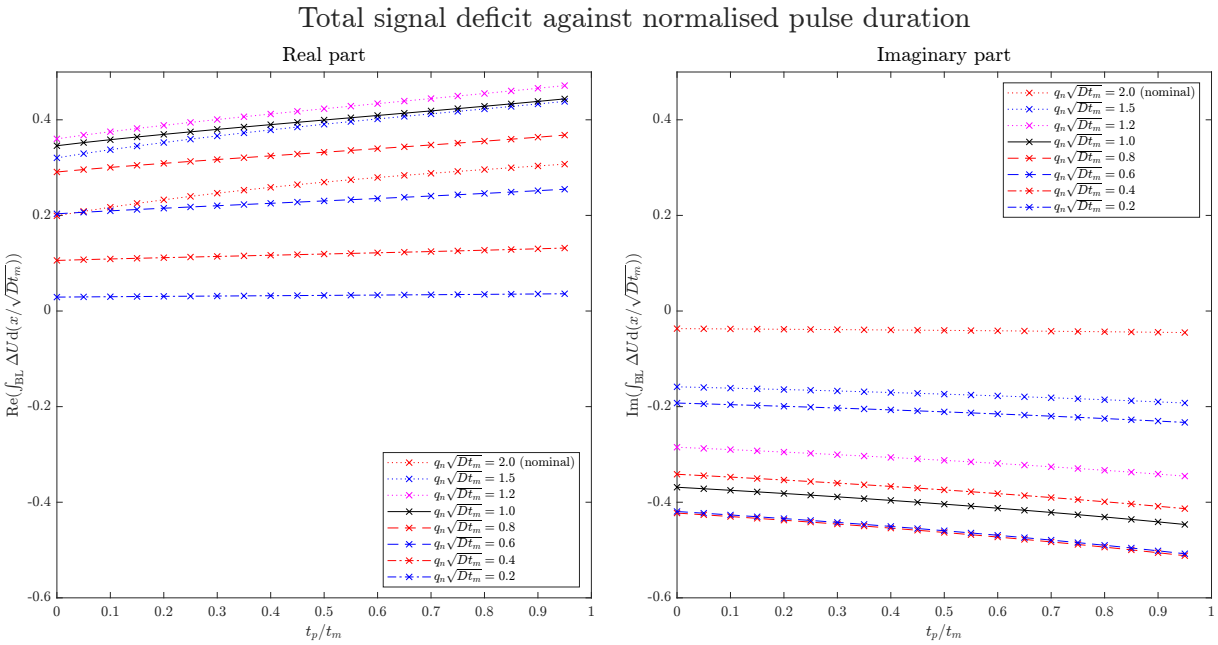

Figure 6.11: Total signal deficit integrated through the boundary layer thickness for different pulse durations and nominal q-vector components normal to boundary surface.

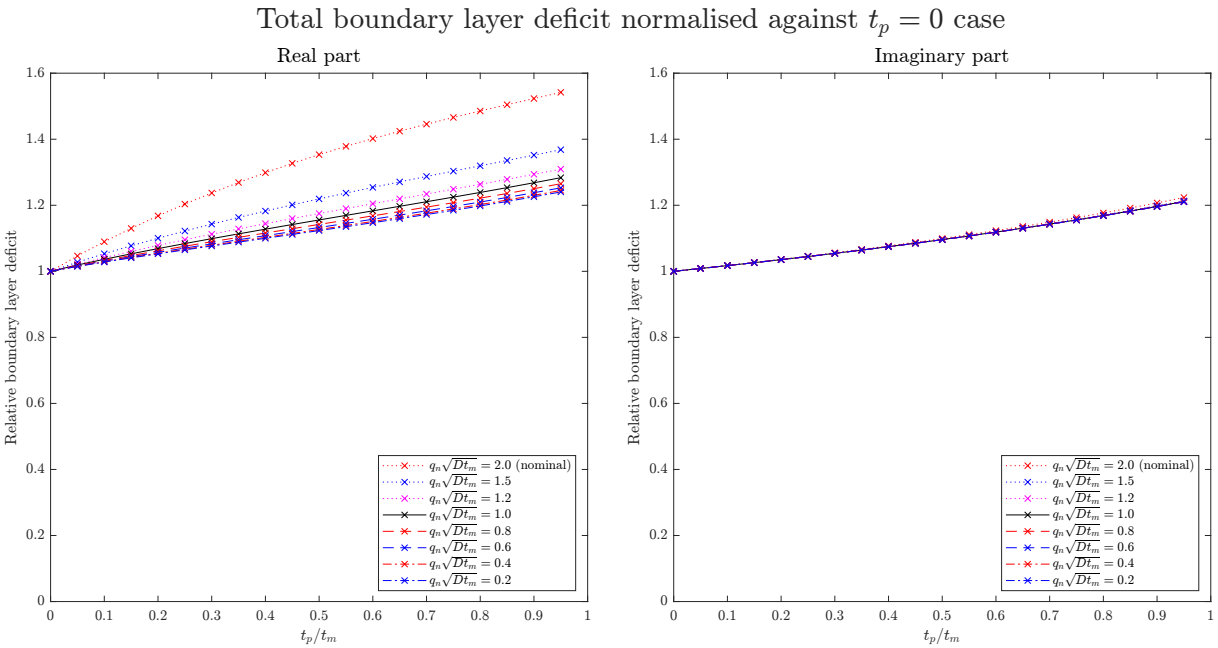

Figure 6.12: Total boundary deficit normalised against the instantaneous pulse case for a range of nominal normal q-values.

To compare with the instantaneous pulse (i.e. constant q-vector) case, the total boundary signal deficit is normalised against the $t_p = 0$ case and plotted in Figure 6.12. It becomes apparent that the signal deficit in the boundary layer is quite sensitive to the non-zero pulse duration. For a non-dimensionalised q-value of $q_n = 1.0$, the real signal deficit at $t_p = 0.3t_m$ is already 9.9% higher than the $t_p = 0$ case; this increases to 15.6% at $t_p = 0.5t_m$, which is a substantial percentage. The imaginary part of the signal deficit is also affected, with the value being 5.4% and 9.6% higher at $t_p = 0.3t_m$ and $t_p = 0.5t_m$ respectively.

The strength of this effect from the non-zero pulse duration also has a dependence on the q-value as well, particularly for the real part of the signal. As depicted in Figure 6.12, the stronger the q-value, the more sensitive the signal deficit is to the pulse duration, and the dependence is much more pronounced for the real part compared to the imaginary part. For a normalised q-value of $q_n = 2.0$, the signal deficit can be as much as 23.7% and 35.3% higher at $t_p = 0.3t_m$ and $t_p = 0.5t_m$ respectively. It is also noted that the non-linearity also increases as q-value increases; for a non-dimensionalised q_n of up to 1.2 the deficit-pulse duration relation is still more or less linear, however the concavity of the curve becomes remarkably pronounced for higher q_n. This means that if the normalised q-value stays within a reasonable range, the relation between the signal deficit and pulse duration can be reasonably approximated by a linear relationship, leading to a possible two-point extrapolation correction method. The details of such a correction method will be discussed in the next section.

6.5 Correcting for non-zero pulse duration

In the previous section, the effect of time-varying q-vector profiles, which is an inevitable consequence of the limited gradient strength available from hardware construction, to the boundary layer signal deficit has been analysed. However, much of the existing literature, either on the biophysical modelling of tissue under diffusion MRI, or models and algo-

rithms to infer medium properties from image signals, is based on the constant q-vector assumption. A natural question therefore arises: is it possible to correct for the effect of an unphysical assumption of a constant q-vector, i.e. $t_p = 0$, such that we can estimate the hypothetical image signal from an imagined diffusion MRI using an unphysical constant q-vector? If such a method or protocol can be found, then it would extend the range of physical diffusion MRI experiments which could be analysed using existing modelling methods found in the literature, in particular ones where t_p is comparable to t_m when the diffusion time is short, for example z[5, 47].

6.5.1 Concept of two-point extrapolation correction

The results from the numerical study, presented in Section 6.4.3, have shown that the relation between the signal deficit and pulse duration is roughly linear, as long as the non-dimensional q-vector (normal component) is not too large. Although the gradients of these linear relationships depend on the exact parameter of q_n, the fact that it is linear presents an opportunity of extrapolating between two sample points, and hence estimating the voxel signal if the modulation and demodulation had happened instantly. If we recall the expression for the voxel-integrated image signal, it is given by:

$$S_{\text{voxel}} = M_0 \int_\Omega \exp\left(-D \int_0^{t^*} |\mathbf{q}^*|^2 \, \mathrm{d}\tau\right) \mathrm{d}V + M_0 \int_{\partial\Omega} \sqrt{Dt_m} \exp\left(-D \sin^2\theta \int_0^{t^*} |\mathbf{q}^*|^2 \mathrm{d}\tau\right) \left(\int_0^\infty (U - U_\infty) \mathrm{d}x_{\text{in}}\right) \mathrm{d}A, \quad (6.97)$$

where θ is the angle between the q-vector and the normal of the surface element $\mathrm{d}A$. If the time profiles of the q-vector are chosen to have different pulse durations but conserve the final b-value b^*, as shown previously in Figure 6.10, then the $\int_0^{t^*} |\mathbf{q}^*|^2 \mathrm{d}\tau$ terms can simply be replaced by the fixed quantity b^* of the pulse sequence. Meanwhile, the signal contribution from the boundary layer term, as shown in the analysis in the previous section, is dependent on the nominal q-vector normal component ($q_{n,\text{nom}}$) and the pulse duration (t_p). This relationship can be approximated by a linear model, in which the

slope and intercepts are functions of the nominal q-vector normal component:

$$\int_0^\infty (U - U_\infty) \mathrm{d}x_{\mathrm{in}} = F\left(q_{n,\mathrm{nom}}, t_p\right) \approx m(q_{n,\mathrm{nom}}) t_p + c(q_{n,\mathrm{nom}}), \tag{6.98}$$

where $q_{n,\mathrm{nom}} = \sqrt{Dt_m}|\mathbf{q}^*_{\mathrm{nom}}| \cos\theta = \sqrt{b^* D}\cos\theta$. Due to the linearity of integration over the boundary surface $\mathrm{d}A$, a linear model can hence be assembled for the dependence of the voxel-integrated image signal on the pulse duration:

$$S_{\mathrm{voxel}} \approx M t_p + C, \tag{6.99}$$

where:

$$M = M_0 \int_{\partial\Omega} \sqrt{Dt_m} \exp\left(-b^* D \sin^2\theta\right) m(q_{n,\mathrm{nom}}) \mathrm{d}A, \tag{6.100}$$

$$C = M_0 \int_\Omega \exp\left(-b^* D\right) \mathrm{d}V + M_0 \int_{\partial\Omega} \sqrt{Dt_m} \exp\left(-b^* D \sin^2\theta\right) c(q_{n,\mathrm{nom}}) \mathrm{d}A. \tag{6.101}$$

In the above expressions, it can be seen that the slope and intercept for the linear voxel signal model depends on the b-value of the sequence, the diffusivity of the medium, as well as the geometry of the boundary surface $\partial\Omega$ due to the integration of θ-dependent terms. In practice, the exact diffusivity and microstructure geometry of the imaged medium is unknown, hence it is not possible to deduce the values of M and C through modelling alone. Nevertheless, the linear model allows an easy empirical characterisation. As long as the mixing duration t_m and the b-value are held constant, by measuring the value of S_{voxel} at two different pulse durations, the above linear model can be applied. The signal model can then be extrapolated to the $t_p = 0$ point to provide an estimation for the signal in this unphysical scenario.

6.5.2 Graphical interpretation of two-point extrapolation

Figure 6.13 shows a graphical interpretation of the two-point extrapolation method. In the plots, the horizontal axis is the ratio between the pulse and mixing durations, whereas the vertical axis is the boundary layer signal deficit normalised against the $t_p = 0$ case. The family of curves in dotted lines show the relation between the signal and the pulse

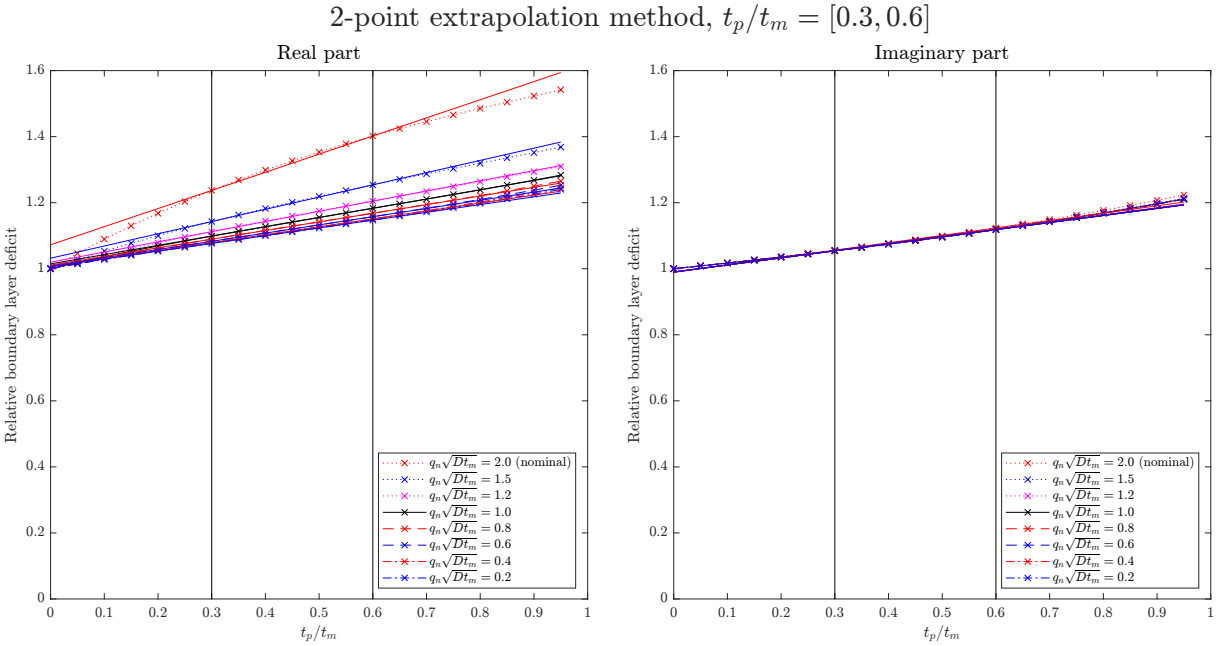

Figure 6.13: Example of a two-point extrapolation to estimate signal deficit with constant q-vector, using two finite pulses with different durations. Dotted lines: boundary layer model results; solid lines: linear extrapolation model.

duration, which depends on the dimensionless number of the nominal normal q-value:

$$q_{n,\mathrm{nom}} = q^*_{n,\mathrm{nom}}\sqrt{Dt_m} = |\mathbf{q}^*_{\mathrm{nom}}|\sqrt{Dt_m}\cos\theta = \sqrt{b^*D}\cos\theta. \tag{6.102}$$

In an actual imaging voxel, the boundary layer signal deficit at each surface element $\mathrm{d}A$ on the boundary surface $\partial\Omega$ is represented by an instance from the family of dotted curves, depending on the local value of $q_{n,\mathrm{nom}}$. Graphically, the extrapolation method is hence the estimation of each of these curves with a straight line, so that rather than integrating instances of the curves, instances of linear relations are integrated instead over all surface elements to get the voxel image signal. As an example, the extrapolation is performed from the two points of $t_p = 0.3\,t_m$ and $t_p = 0.6\,t_m$ in Figure 6.13. The estimation of the $t_p = 0$ signal, which should be 1 on the graph, is therefore given by the y-intercept of the respective straight lines joined from the two extrapolation points.

6.5.3 Accuracy of the two-point extrapolation method

In terms of using the two-point extrapolation method in practice, one of the practical considerations is the accuracy of such estimation. As we can see in Figure 6.13, the y-intercepts of the extrapolation lines do not coincide unity exactly, with the error coming from the non-linearity of the signal-t_p relations themselves. The non-linearity is dependent on the nominal non-dimensional q_n, which in turn is dependent on not only the geometric orientation of the surface element $\mathrm{d}A$ in question, but also the physical quantities of the unknown diffusion coefficient D relative to the parameters $\mathbf{q}$ and t_m used for the imaging experiment (see Eq. 6.102). The relative extrapolation error for the example in Section 6.5.2 is plotted against the nominal non-dimensional q_n in Figure 6.14. For the real part of the signal deficit, it can be seen that the extrapolation error increases substantially at higher values of q_n, whereas the imaginary part is affected much less.

In general, for an imaging voxel, since the boundary surface contains surface elements of different angles, the overall extrapolation error after the surface integration will be somewhere between zero and the maximum value of $q^*_{n,\mathrm{nom}}\sqrt{Dt_m}$ in the medium. A conservative estimation of the extrapolation error can therefore be evaluated by assuming the worst case for all surface elements, which is when the surface normal is aligned with the q-vector, and using a conservative overestimation of the expected diffusivity D in the medium.

6.6 Varying q-vector – irregular gradient pulses

So far in the previous section, only gradient pulses with a rectangular shape have been considered in the boundary layer calculations. However, in reality due to hardware limitations, it is not possible to switch on or off the gradient magnetic field instantaneously, as the rate of change of the gradient field strength is limited by a maximum slew rate [122]. More generally, the gradient pulse lobe can have an irregular shape instead, for example to account for practical considerations such as minimising eddy currents [100].

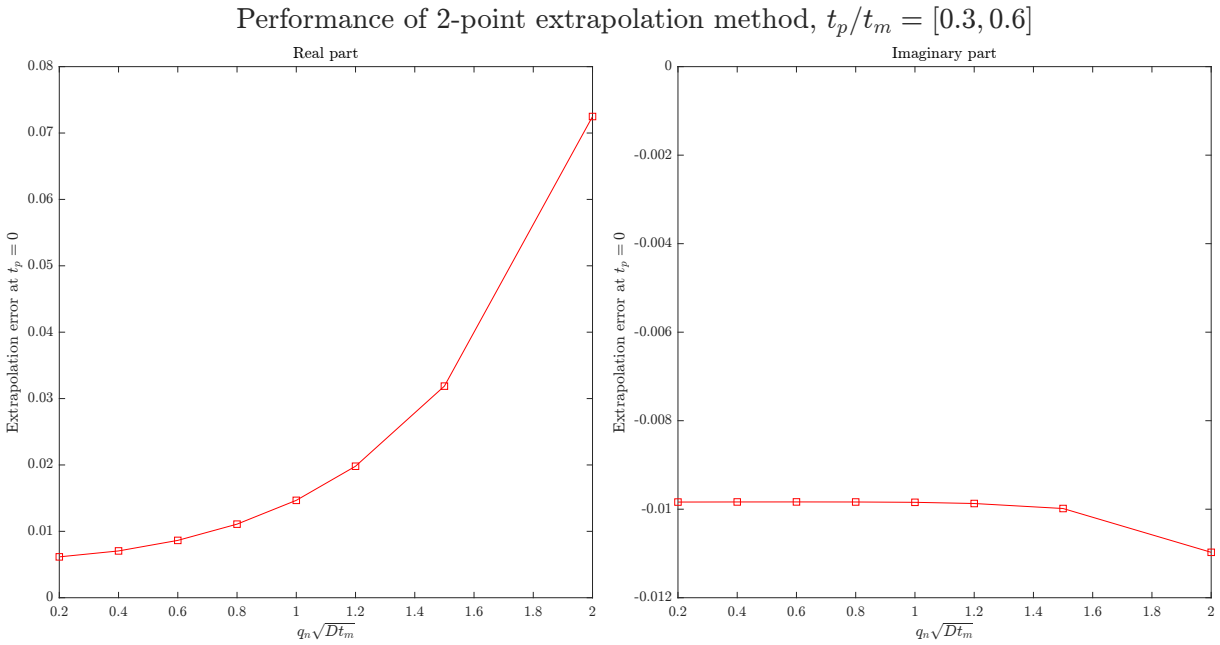

Figure 6.14: Relative extrapolation error at different q-vectors, using $t_p/t_m = 0.3$ and 0.6.

In this section, the effect of more irregular gradient pulse shapes on the boundary layer calculations is investigated through further numerical studies.

6.6.1 Aim and definition of numerical studies

Similar to the previous section, the aim of the numerical studies is to characterise the effect of irregular gradient pulse shapes, and to evaluate the possibility of approximating them with equivalent models using simpler pulse shapes with known calculations. The following numerical studies therefore investigate the following hypothesis: given a certain irregular gradient pulse shape, is it possible to relate the boundary layer calculations to an equivalent rectangular pulse? In particular, since in the imaged medium, the boundary surface orients at an unknown distribution of angles relative to the applied q-vector, thus leading to a distribution of nominal q_n values too. Therefore, it is essential that such an equivalent rectangular pulse timing is consistent across all nominal q_n values.

In the following numerical studies, four different gradient pulse shapes are used, as shown in Figure 6.15. For these gradient profiles, the pulse and mixing durations are defined in the same manner as the rectangular pulse profiles in the previous section. In

terms of symmetry, the half-wave sinusoidal pulse and triangular pulse are symmetric themselves, therefore there is always also a symmetry between the modulating and demodulating pulses. Meanwhile, the quarter-wave pulse is not symmetric, therefore the translation and reflection cases for the demodulating pulse are different and are considered separately. Similar to the previous numerical study, in order to keep the far-field bulk signal constant over the different cases, the final b-values of the profiles are held constant across the four different pulse profiles, based on the nominal q-value if the q-vector is constant for the same b-value. This thus leads to a slightly higher actual peak q-value, compared to the nominal q-value.

For each of the pulse shapes, four different pulse durations ($t_p/t_m = [0.2, 0.4, 0.6, 0.8]$) are considered at the nominal q-values of $q_{n,\mathrm{nom}} = [0.2, 0.4, 0.6, 0.8, 1.0, 1.2, 1.5, 2.0]$. In terms of the implementation of the numerical algorithm, the same equations (6.93)–(6.96) and discretisation parameters (Table 6.2) as defined for rectangular pulses are used.

6.6.2 Results of numerical studies

In order to evaluate the possibility of finding an equivalent rectangular pulse to approximate the signal from an irregular pulse, the total signal deficit over the boundary layer after the whole irregular pulse sequence is evaluated and compared against the curves in Figure 6.11, which show the signal deficit for rectangular pulses at different pulse durations. This set of curves can then be used for interpolation to deduce an equivalent rectangular pulse duration, given a certain nominal q-value.

The interpolation results for the half-wave and triangular pulses are shown in Figure 6.16. It can be seen that for both shapes of gradient pulses, there are consistent sets of equivalent rectangular pulse durations, regardless of the nominal normal q-value. Meanwhile, for the interpolation results for the quarter-wave pulses (Figure 6.17), consistent sets of equivalent rectangular pulse durations could not be found in either of the translation or reflection cases. An interesting observation is that the inconsistencies over the nominal q-value mainly appear in the imaginary part for the translated demodulation

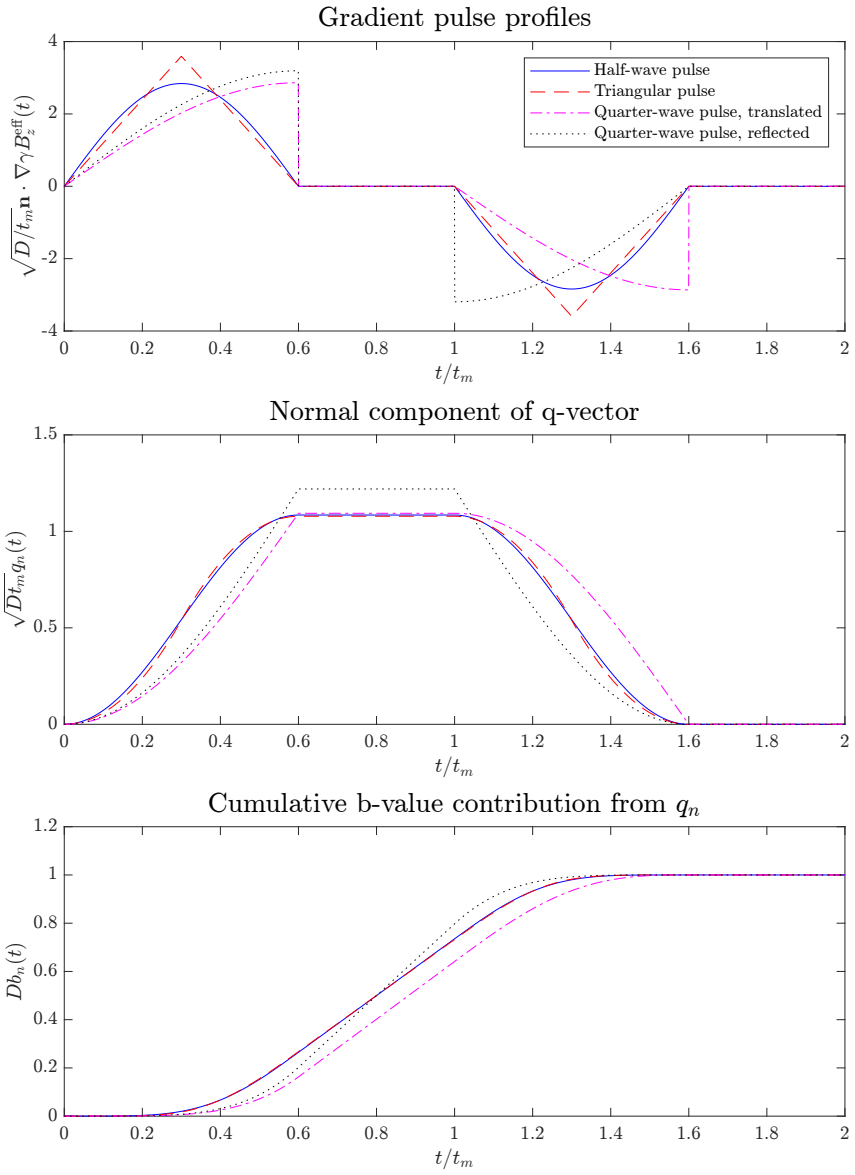

Figure 6.15: The four non-rectangular gradient pulse profiles considered in the numerical study, shown here with parameters $t_p/t_m = 0.6$ and $q_{n,\text{nom}} = 1.0$. Note the final b-value (i.e. the nominal q-value), rather than the peak q-value, is set to be constant across the profiles.

case, whereas for the reflected demodulation case they appear in the real part instead. These interpolation results seem to suggest that the condition for an equivalent rectangular pulse is that the demodulation pulse needs to be simultaneously a translation and a reflection of the modulation pulse, in other words the pulses need to be symmetric themselves.

For the symmetric pulses, an additional observation is that the equivalent rectangular

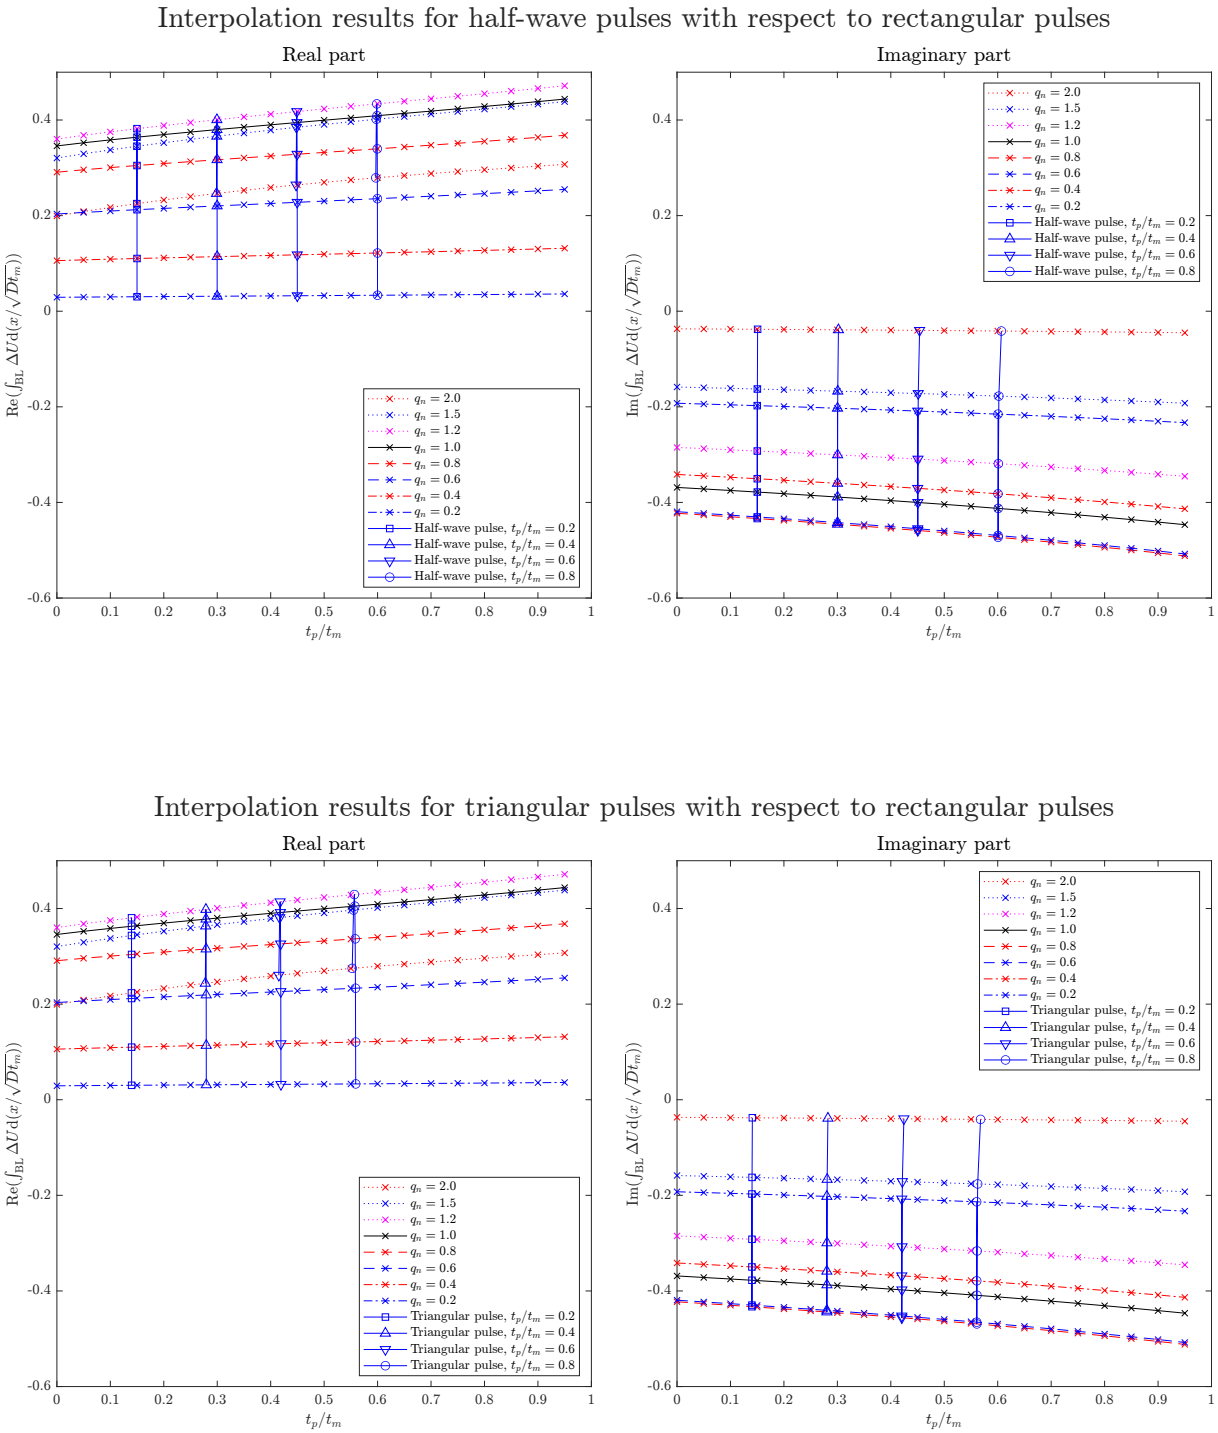

Figure 6.16: Total signal deficit over boundary layer plotted against normalised pulse duration for rectangular pulses, with results from half-wave (top) and triangular (bottom) pulses interpolated onto the curves. Note the equivalent pulse durations are associated to scaling factors of 0.75 and 0.7 for the two cases respectively.

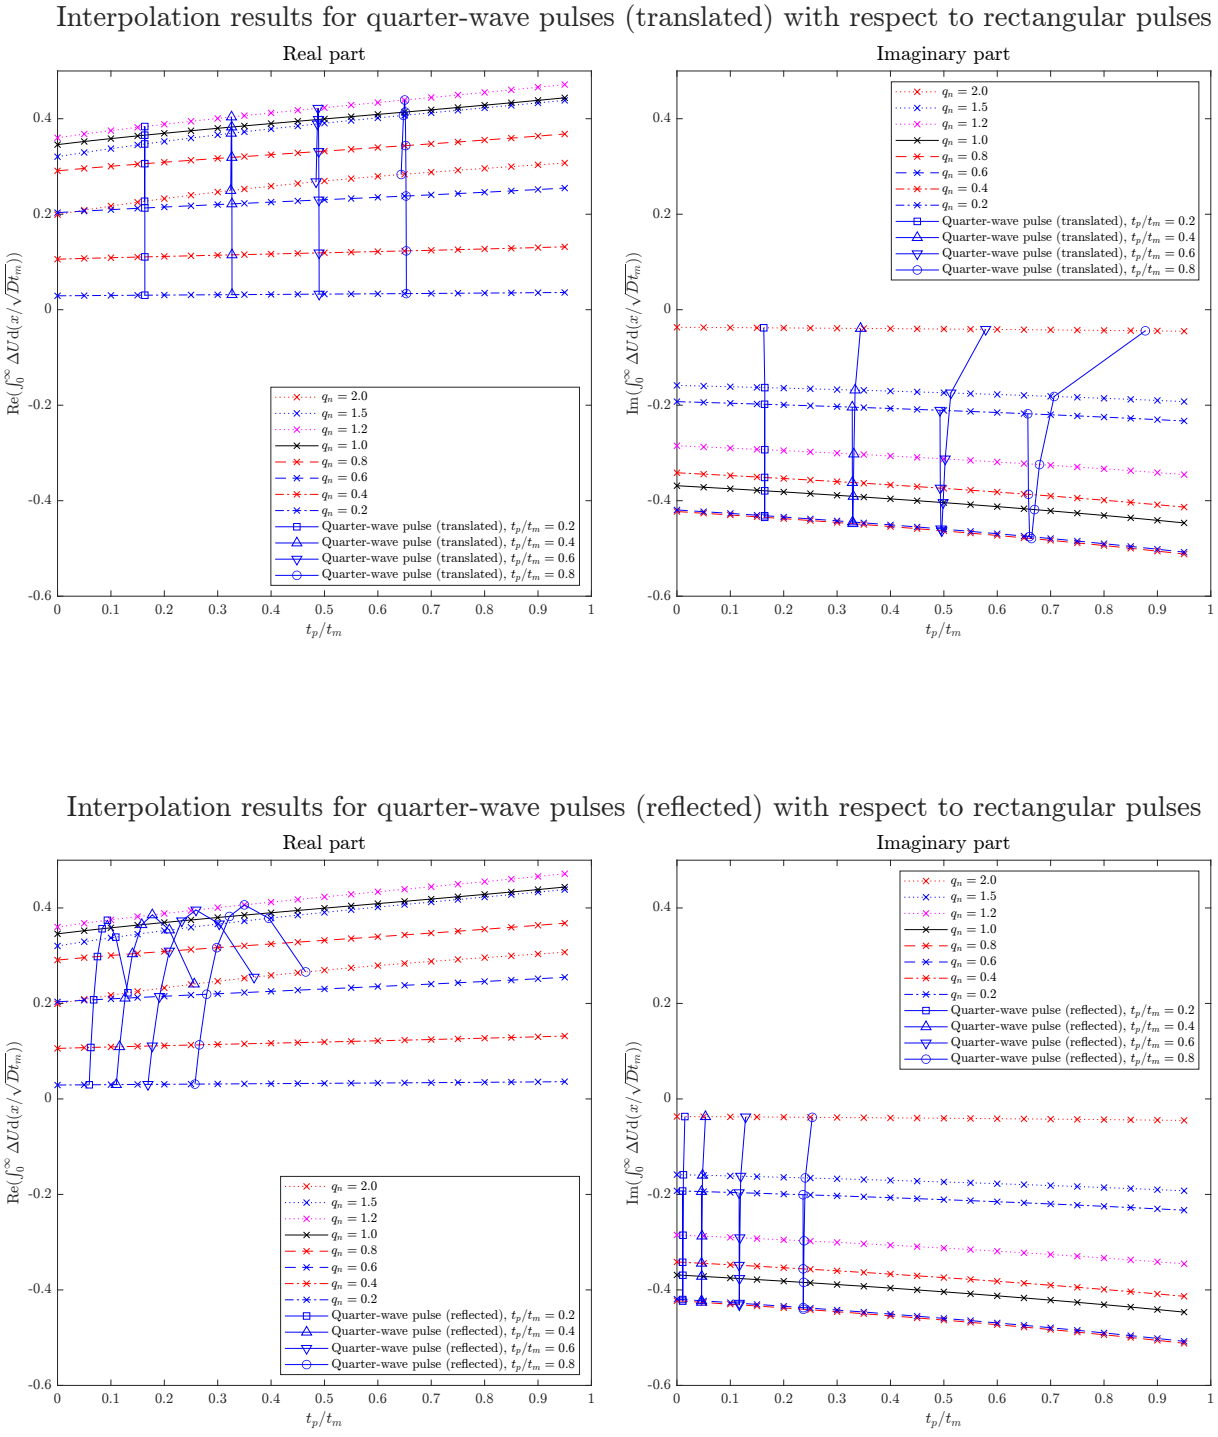

Figure 6.17: Total signal deficit over boundary layer plotted against normalised pulse duration for rectangular pulses, with results from quarter-wave pulses (top: translated demodulation, bottom: reflected demodulation) interpolated onto the curves.

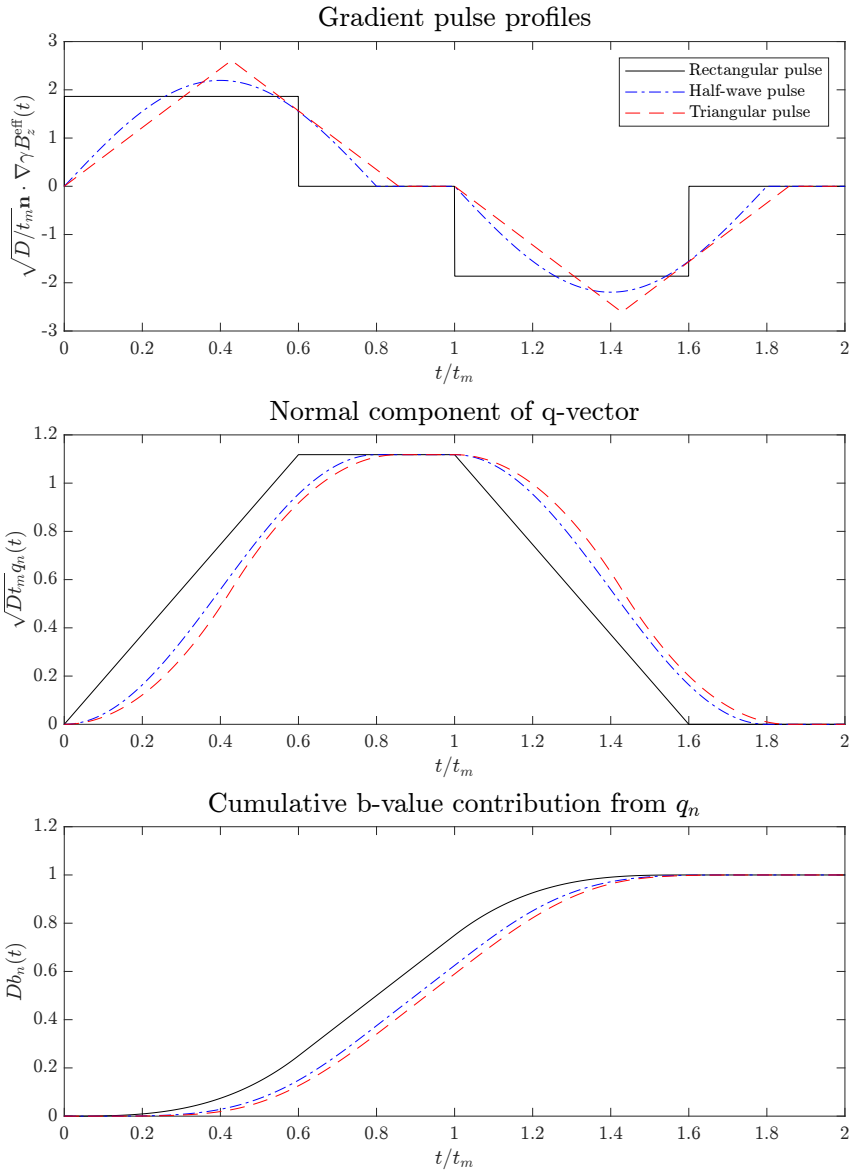

Figure 6.18: Comparison between three equivalent pulse profiles that produce approximately the same signal from the boundary layer model.

pulse durations can be described by an approximate scaling factor of about $t_p' = 0.75t_p$ for the half-wave pulses, and $t_p' = 0.7t_p$ for the triangular pulse. In order to examine the possible origin of such scaling factors, a comparison of the equivalent half-wave and triangular pulse profiles for a particular rectangular pulse is shown in Figure 6.18. The key common characteristics between these profiles are that they share the same peak q-value as well as the final b-value. In fact, it can be shown that for a given pulse shape, these two conditions uniquely define the pulse amplitude and duration required to obtain the

equivalent pulse. For example, the half-wave sinusoidal pulse profiles can be parametrised by the amplitude A and pulse duration t_p as:

$$|\gamma \mathbf{g}(t)| = \begin{cases} A\sin(\pi t/t_p), & 0 \leq t < t_p \\ -A\sin(\pi(t-t_m)/t_p), & t_m \leq t < t_m + t_p, \\ 0, & \text{otherwise.} \end{cases}$$

The associated peak q-value and final b-value for this profile are given by:

$$\begin{aligned} q_{\max} &= \int_0^{t_p} |\gamma \mathbf{g}(t)| \mathrm{d}t = \frac{2}{\pi} A t_p, \\ b_{\text{final}} &= \int_0^{t_m+t_p} \left(\int_0^t |\gamma \mathbf{g}(\tau)| \mathrm{d}\tau \right)^2 \mathrm{d}t = \left(\frac{2}{\pi} A t_p \right)^2 \left(t_m - \frac{t_p}{4} \right). \end{aligned}$$

Meanwhile in a similar fashion, for a rectangular pulse with amplitude A' and pulse duration t'_p, these two quantities are:

$$\begin{aligned} q_{\max} &= A' t'_p, \\ b_{\text{final}} &= (A' t'_p)^2 \left(t_m - \frac{t'_p}{3} \right). \end{aligned}$$

By equating the two sets of expressions, one can then deduce that the equivalent rectangular pulse parameters are given by:

$$t'_p = \frac{3}{4} t_p, \;\; A' = \frac{8}{3\pi} A.$$

Similarly, for the triangular pulse profile, the peak q-value and final b-value are:

$$\begin{aligned} q_{\max} &= \frac{A t_p}{2}, \\ b_{\text{final}} &= \left(\frac{A t_p}{2} \right)^2 \left(t_m - \frac{7}{30} t_p \right), \end{aligned}$$

which leads to the equivalent parameters:

$$t'_p = \frac{7}{10} t_p, \;\; A' = \frac{5}{7} A.$$

These relations therefore suggest that the consistent linear scaling for equivalent rectan-

gular pulses may relate to the conditions of holding the peak q-vale and final b-value over the pulse sequence constant. However, these criteria for deducing the equivalent rectangular pulses are only empirical findings from the interpolation studies. This means that they are only approximate, and may not apply to all general pulse shapes or profiles. In reality, due to the way the q_n terms appear in the differential operator of the PDE problem, one might expect the dependence of the solution on the parameters of the q-vector profile to be non-trivial, nevertheless the empirical findings in this section might provide some insight to how an approximate equivalent model may be deduced with a simple pen-and-paper calculation, without the need of the full solution of the PDE problem.

6.6.3 Extending to multiple pairs of modulation-demodulation pulses

So far, in the numerical studies only single pairs of modulation-demodulation pulses have been considered. However, in some DW-SE experiments, multiple lobes of modulation or demodulation pulses may be chosen due to various considerations. Therefore, in this section, two additional pulse sequences are considered, where the modulation-demodulation pulses each consist of two lobes rather than one lobe. The profiles of the pulse sequences are shown in Figure 6.19, and are chosen such that they satisfy the symmetry condition discovered above. The purpose of simulating these two pulse sequences is to test if the empirical findings from the previous section also apply to more general cases.

Similar to the other irregular pulse profiles, the simulation results for the two multiple-lobed sequences are interpolated onto the rectangular pulse data curves, resulting in the plots in Figure 6.20. According to the empirical findings in the previous section, if the same criteria of fixing the maximum q-value and final b-value constant are used to deduce an equivalent rectangular pulse duration, the expected effective pulse durations for the two pulse sequences should scale respectively as:

$$t_p' = \frac{15}{16} t_p \text{ (sequence 1), and } t_p' = \left(\frac{6}{\pi} - \frac{3}{4}\right) t_p \text{ (sequence 2).}$$

These two expected scaling factors are plotted as black dotted lines in Figure 6.20 as well. These plots show that the empirical findings from the section above still largely apply to these two more complicated pulse profiles, albeit with a decreased accuracy. In particular, the performance worsens at longer pulse durations, and at higher nominal q-values in the normal direction to the surface. The discrepancy between the empirical model and the numerical results also appear mainly in the real part, rather than the imaginary part, of the signal solution. Pulse sequence 1 also performs significantly better than pulse sequence 2, in terms of being in line with the empirical model; this could be because the shape of pulse sequence 1 is less irregular itself and resembles the rectangular pulses more.

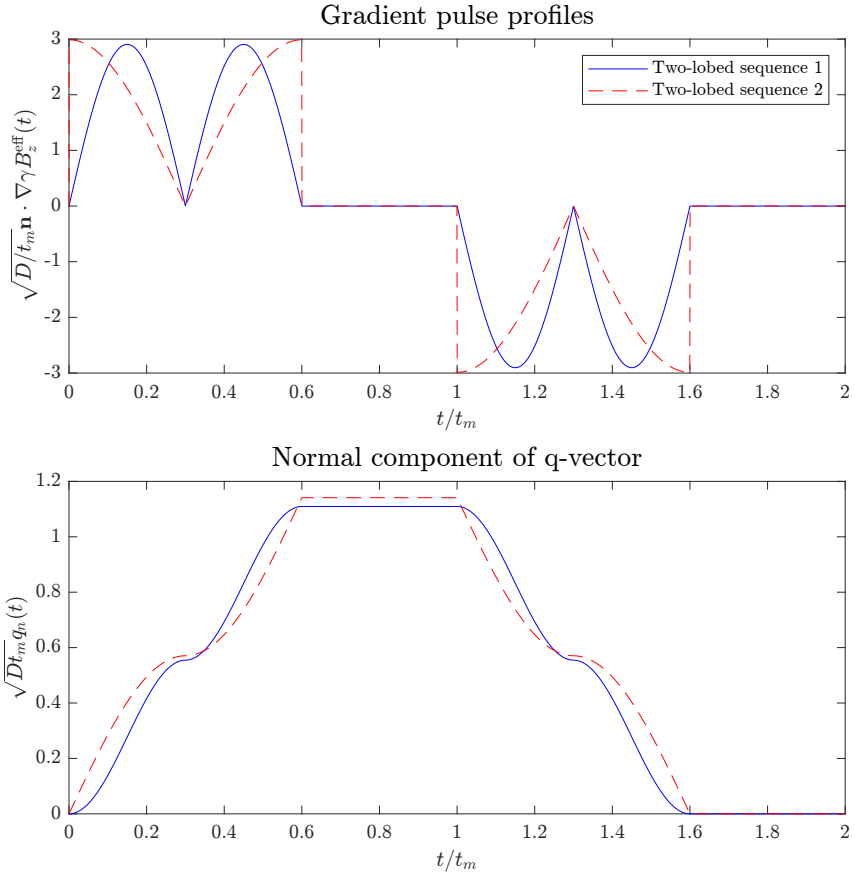

Figure 6.19: Two pulse sequences with two pairs of modulation-demodulation gradient pulses, and its associate q-vector profile.

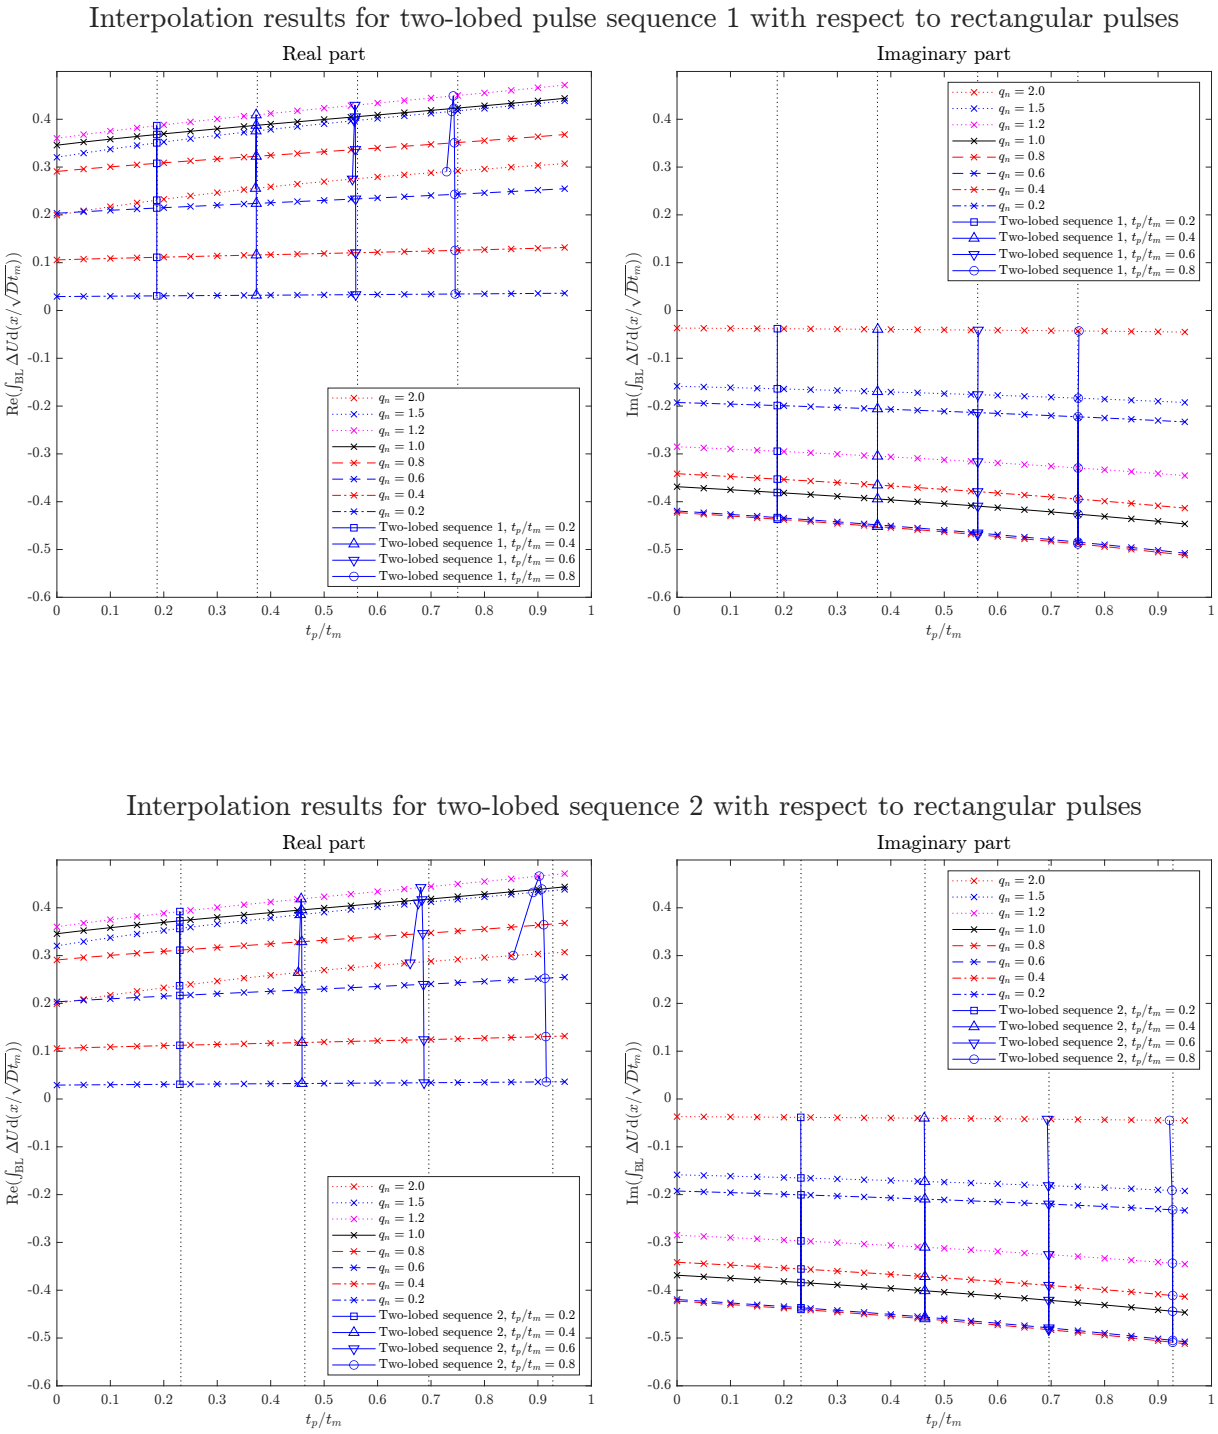

Figure 6.20: Total signal deficit over boundary layer for two-lobed sequences defined in Figure 6.19, interpolated onto results for rectangular pulses. Blacked dotted lines: expected equivalent rectangular pulse durations based on findings from previous section.

6.6.4 Discussion

In this section, the effect of gradient pulse profiles which are irregular in shape, rather than just being rectangular, has been considered in the boundary layer calculations. A number of different irregular profiles have been used to numerically solve for the boundary layer signal. The results are then interpolated back onto the signal-pulse duration curves for rectangular pulses, in an attempt to find equivalent rectangular pulses that would result in the same signal, no matter what the nominal q-value is.

The empirical findings from the numerical studies seem to suggest that as long as both the modulation and demodulation pulses are symmetric themselves, then an approximate equivalent rectangular pulse does exist. The parameters of this equivalent rectangular pulse may also be found by equating maximum q-value and final b-value for these two pulse profiles. This empirical relation was also tested on pulse sequences which contain multiple pairs of modulation-demodulation pulses, and it still seems to apply, albeit with slightly poorer accuracy. In particular, it seem to diverge more from the actual solution at higher q-values and at longer pulse durations, as well as if the shape of the pulse becomes more irregular and diverges more from a simple rectangular shape.

6.7 Conclusion and implication of results

To conclude this chapter, the behaviour of DW-SE in the short-time regime has been investigated by the asymptotic analysis of the Bloch-Torrey equation. Under this particular scale separation between the time scale of the pulse sequence and the characteristic time needed for diffusion over the microstructure length scale, DW-SE can be described by a boundary layer model. For the bulk of the signal away from the boundary surfaces, the signal decays as if there are no diffusion barriers; meanwhile in the close vicinity to the boundary surface, the signal can be characterised by an inner PDE problem. The contribution of the tangential component of the q-vector can be further factored away, reducing the inner problems into a single family of problems (6.20)–(6.23), which are of one spatial

and one time variable, parametrised only by the normal component of the q-vector locally at the boundary surface element. The boundary condition, which represents the diffusion barrier in physical terms, thus causes a deficit in signal decay which can be deduced by solving the inner problem. The overall signal observed over a voxel is thus an integration of such decay deficit over all boundary surface elements within the voxel.

When the q-vector is constant, the family of inner problems in fact is self-similar and can be reduced to a single problem by a rescaling of variables. Effectively, the boundary layer behaviour depends solely on the values of the two non-dimensional groups of $q_n x$ and $q_n^2 Dt$, which essentially are the relevant rescaled spatial and temporal variables. The inner problem is then solved using three different methods in Section 6.3, namely a Laplace transform method, a series expansion method and a finite difference method. The Laplace transform method yields an analytic expression for the total signal decay deficit over the boundary layer, meanwhile the series expansion method resolves the spatial profile of the boundary layer signal evolving over time, albeit at the small q limit. The finite difference method shows both the spatial and time profile of the boundary layer signal, although only in a numerical sense rather than as an analytic expression. The results of the analysis all agree with existing literature, and more importantly reveal an imaginary part contribution that has been ignored in literature so far. If the subsequent image formation protocol of the DW-SE experiment is perfect, in the sense that the all spins experience a coherent phase shift before being integrated to give the final voxel image signal value, then the imaginary part indeed integrates to zero. However if this is not the case the imaginary part will have a contribution of comparable magnitude to the real part.

In a realistic DW-SE experiment, unfortunately the constant q-vector assumption is not physical, as the strength of the gradient pulse has to be finite. Therefore, the profile of the q-vector varies over time. The effect of this on the boundary layer model has been investigated by using such time profiles when applying the finite difference scheme to numerically solve the inner problem. In Section 6.4, the simpler case of the gradient pulses being rectangular is studied. In this case, given a nominal q-value, the pulse is

characterised only by the pulse duration t_p relative to the spacing between the pulses t_m, also known as the mixing duration. Whilst keeping the integral $\int_0^t q_n^2 D \mathrm{d}\tau$, in other words the b-value contribution from the normal component, constant over the different profiles, it is found that the boundary layer signal deficit increases approximately linearly with the pulse duration. Although it means that if a constant q-vector model is naively applied it will not be accurate, the approximately linear relationship does allow for a two-point extrapolation method proposed in Section 6.5. This extrapolation method estimates a hypothetical, unphysical image signal for a constant q-vector case, based on two physical measurements at two different pulse durations.

In Section 6.6, the above analysis is extended from rectangular pulses to more general, irregular gradient pulse profiles. It is found empirically from these numerical results that if the shape of the gradient pulse is symmetric, then there exists an equivalent rectangular pulse across that would lead to approximately the same boundary layer signal deficit across different q_n values. This is a particularly useful property, as in an imaged medium the boundary surface orientates in a distribution of angles relative to the q-vector, hence the normal components of the q-vectors vary depending on the local orientation. A coherent equivalent pulse duration means that even after integrating the signal deficits over an unknown distribution of angles, the equivalency is still valid. However, for asymmetric pulse shapes, such a consistent equivalent timing does not exist. Furthermore, for the symmetric pulses, it is found that the equivalent pulse timing can be reliably deduced by equating the maximum q-value and final b-value between the equivalent pulses. These two simple relations thus allow the equivalent timings to be deduced by a simple pen-and-paper calculation, without the need of any numerical simulation.

In practical terms, the above findings can be useful in allowing for a better way to design DW-SE experiments and analysing the data. In particular, even though the majority of DW-SE modelling in the literature involve the unphysical and often inaccurate assumption of a constant q-vector, these models are not rendered useless, as the hypothetical corresponding signal can be estimated from measurements using two different DW-SE

pulse sequences and performing the extrapolation, as long as the gradient pulses used are symmetric. The proposed pipeline is as follows: imagine if we are interested in a hypothetical signal when an unknown medium is imaged at a certain constant q-vector. The hardware limitations, such as the maximum gradient strength and slew rates, may mean that there are certain restrictions on possible shapes of the gradient pulses. Nevertheless, as long as two different symmetric pulse profiles with the same final b-value are used, their respective image signal measurements can be related to an equivalent rectangular pulse using the relations described in Section 6.6. Subsequently, the two-point extrapolation method can be performed using these two equivalent rectangular pulse timings to deduce the hypothetical image signal. The result will be a better estimation of what would be the image signal, compared to naively using either image signal measurement and applying the constant q-vector assumption straight away. The hypothetical image signal can then serve as a better input to existing models, so that the subsequent analysis can lead to inferred medium properties with higher fidelity.

Chapter 7

Numerical simulation of DW-SE for the intermediate-time regime

7.1 Introduction

In this chapter, the intermediate-time regime of DW-SE is considered. Previously in Chapters 5 and 6, two asymptotic models for diffusion MRI in porous media have been derived for the long and short time regimes of DW-SE respectively. However, as all the terms in the Bloch-Torrey equation are balanced in magnitude in the intermediate-time regime, there is no scale separation to allow for asymptotic analysis. Therefore, numerical solution of the full equation is sought instead. A finite element method will be applied to solve the Bloch-Torrey equation over two simple model geometries of white matter, as well as a realistic microstructure derived from microscopy data.

In order to evaluate the behaviour of DW-SE in the regime bridging the two asymptotic models, simulations using constant q-vectors will firstly be considered. Although they correspond to unphysical assumptions, they simplify the definition of simulation cases and allow for easier comparison. Subsequently, the effects of time-varying q-vectors due to realistic gradient pulses will be simulated, with an aim to validate and extend the correction methods proposed previously in Chapter 6 for the boundary layer regime.

7.2 Finite element formulation

In this section, the finite element formulation for the Bloch-Torrey equation is described. As per the previous chapters, we start with the co-rotating, co-decaying form of the Bloch-Torrey equation:

$$\frac{\partial S}{\partial t} + (\nabla - i\mathbf{q}) \cdot (-\mathbf{D}(\nabla - i\mathbf{q})S) = 0 \quad \text{in } \Omega, \tag{7.1}$$

$$\mathbf{n} \cdot (-\mathbf{D}(\nabla - i\mathbf{q})S) = 0 \quad \text{on } \partial\Omega, \tag{7.2}$$

$$S = M_0 \quad \text{at } t = 0. \tag{7.3}$$

The co-rotating, co-decaying formulation provides two major advantages over the original equation in terms of implementing the finite element method. The first one relates to the periodic nature of the domain, in the sense that the transformed variable S is also periodic given periodic initial conditions. Whereas in the original formulation, the applied magnetic field B_z contains a component which varies linearly in space, leading to some kind of jump conditions required to bridge the opposite sides of a periodic cell. It is therefore much simpler to consider the co-rotating, co-decaying form, as one simply needs to pose the equation over a periodic domain and enforce a simple periodic boundary condition to obtain the solution.

The second advantage is the spatial resolution of the mesh needed to resolve the solution, particularly at high q-values. The original Bloch-Torrey formulation uses a stationary frame of reference for phase, so the modulation of spins leads to a spatially fast-varying signal. Hence, the mesh needs to be even finer than the modulation wavelength to spatially resolve the solution. Meanwhile, the modulation behaviour is already accounted for in the co-rotating formulation and encapsulated in the q-vector, so the mesh only needs to resolve the regions where there are significant spatial variations in the co-rotating signal, i.e. in the boundary layer for high q-values. This thus greatly reduces the computational expense of the finite element implementation.

The first step of discretising the parabolic problem is to use a backward Euler scheme to discretise the time variable. This involves approximating the time derivative at a given time step k with the finite difference between the k-th and $(k-1)$-th time steps:

$$\frac{S^k - S^{k-1}}{\delta t} + (\nabla - i\mathbf{q}^k) \cdot \left(-\mathbf{D}(\nabla - i\mathbf{q}^k)S^k\right) = 0 \ \text{ in } \Omega, \tag{7.4}$$

$$\mathbf{n} \cdot \left(-\mathbf{D}(\nabla - i\mathbf{q}^k)S^k\right) = 0 \ \text{ on } \partial\Omega, \tag{7.5}$$

$$S^0 = M_0. \tag{7.6}$$

Here the superscript k denotes the sampling of a time-varying quantity at the k-th time step, and δt is the duration of each time step. This leads to an elliptic problem for S^k at each time step k, with the solution at the previous time step S^{k-1} appearing in a source term in the problem.

At each time step, the elliptic problem can then be solved numerically using a finite element method, similar to that described previously in Section 5.4. Firstly, the variational formulation of the problem can be obtained by multiplying any test function to (7.4) and integrating by parts using the boundary condition (7.5). The equivalent variational problem of (7.4)–(7.5) is therefore to solve for $S^k \in \mathcal{H}^1_{\text{per}}(\Omega)$ such that:

$$\int_\Omega \left(\mathbf{D}\nabla S^k \cdot \nabla \bar{v} + i\mathbf{q}^k \cdot \mathbf{D}\nabla S^k \bar{v} - i\mathbf{D}\mathbf{q}^k S^k \cdot \nabla \bar{v} + \left(\mathbf{q}^k \cdot \mathbf{D}\mathbf{q}^k + \frac{1}{\delta t}\right) S^k \bar{v}\right) \mathrm{d}x = \int_\Omega \frac{1}{\delta t} S^{k-1} \bar{v} \, \mathrm{d}x, \ \forall v \in \mathcal{H}^1_{\text{per}}(\Omega). \tag{7.7}$$

Since the PDE is a complex-valued problem, the function space $\mathcal{H}^1_{\text{per}}(\Omega)$, for both the solution and the test function, is the space of *complex* periodic functions over the periodic domain Ω with square-integrable first derivatives, with the cursive notation differentiating it from the *real* function space $H^1_{\text{per}}(\Omega)$. The continuous domain Ω can then be approximated by a discrete mesh Ω_{mesh}, which leads to a finite-dimensional function space $\mathcal{H}^1_{\text{per}}(\Omega_{\text{mesh}})$ over which the same variational problem can be posed, resulting to a discrete linear problem that can be solved numerically.

Similar to the numerical solution of the cell problem in Chapter 5, in this chapter the open-source FEniCS package [71] is used to implement the finite method. The choice of FEniCS is to be in line with scientific reproducibility principles through reusable codes, as well as to take advantage of its capability for bespoke variational problems.

7.3 Transitional region between asymptotic regimes

7.3.1 Overview of numerical simulations

In this section, the numerical simulations concern the case when the q-vector is constant over time. As explained previously in Chapter 6, this corresponds to the unphysical assumption of instantaneous modulation of spins. Nevertheless, this assumption reduces the number of parameters needed to define a simulation scenario, allowing for simpler and fewer simulations as well as easier comparison and analysis of the results. In particular, the main objective of the numerical studies in this section is to investigate the transition behaviour between the long-time and short-time asymptotic regimes; the effects of time-varying q-vectors due to realistic gradient pulses will be considered later in Section 7.4.

The numerical simulations are performed at a range of q-values over several orders of magnitude, either side of the region corresponding to the characteristic geometric scales. Similar to the regimes considered in Chapters 5 and 6, we are interested in the region of $O(1)$ signal and contrast, which corresponds to the $|\mathbf{q}|^2 Dt = O(1)$ diagonal strip in the q-t parameter space. For each q-value used, the maximum simulation time is set to correspond to $|\mathbf{q}|^2 Dt = 3$. This value is chosen as it corresponds to a b-value of $b = |\mathbf{q}|^2 t = 1000$ s mm^{-2} if $D = 3.0 \times 10^{-9}$ m^2 s^{-1}, which are typical values for DW-SE imaging of water at 37 °C. [1] Moreover, since the q-vector is constant throughout the simulation, the results for any shorter times can be obtained simply by looking at the corresponding time step in the simulation, hence only one run per q-value is needed to

[1] Note that this typical value of diffusion coefficient is for a homogeneous sample of water, in biological imaging the effective diffusion coefficient will be lower with the existence of diffusion barriers.

cover a range of times in the $|\mathbf{q}|^2 Dt \sim 1$ strip.

7.3.2 Domain geometries

In the numerical simulations, the geometries are assumed to be uniform and infinite in the longitudinal direction, whereas in the transverse directions they are periodic in both x- and y-directions. These are reasonable assumptions as axons in white matter are usually in the form of a long bundle of fibres. These are also necessary assumptions for the numerical implementation, both because of limited availability of data (only a small 2D cross-sectional slice is available) and to reduce computational expense. Under these assumptions, the problem reduces to only two spatial dimensions, and it only needs to be posed over a two-dimensional periodic unit cell, with periodic conditions enforced on respective edges. This is because the behaviour in the longitudinal direction is homogeneous and does not couple into the transverse directions. Even if the applied q-vector contains any longitudinal component, its effect is simply a multiplication of $\exp(-q_z^2 Dt)$ to the signal which can be factored out. Hence, in the simulations, the equations are solved over the cross-sectional periodic cells of the geometries, and only q-vectors applied along a direction in the transverse plane are considered.

Three different geometries are considered in this section. The first two are simplified models of axons in white matter, consisting of periodic cylindrical inclusions. There are two reasons for using these simple models, one is to test the codes with simpler meshes, the other is to provide easier insight to the solution behaviour. The cross-sectional periodic unit cells of the geometries are shown in Figure 7.1. Geometry A, which consists of a grid of cylinders with spacing d between the centres and radii of $0.25d$, is the same as that used in Chapter 5, Section 5.4.3. The slightly more complicated Geometry B also has a $d \times d$ periodic unit cell, but each cell contains three cylinders of different sizes. The cylindrical surfaces in both geometries are modelled as no-flux boundaries, separating each geometry into interior and exterior compartments.

A third geometry based on realistic white matter microstructure is then considered,

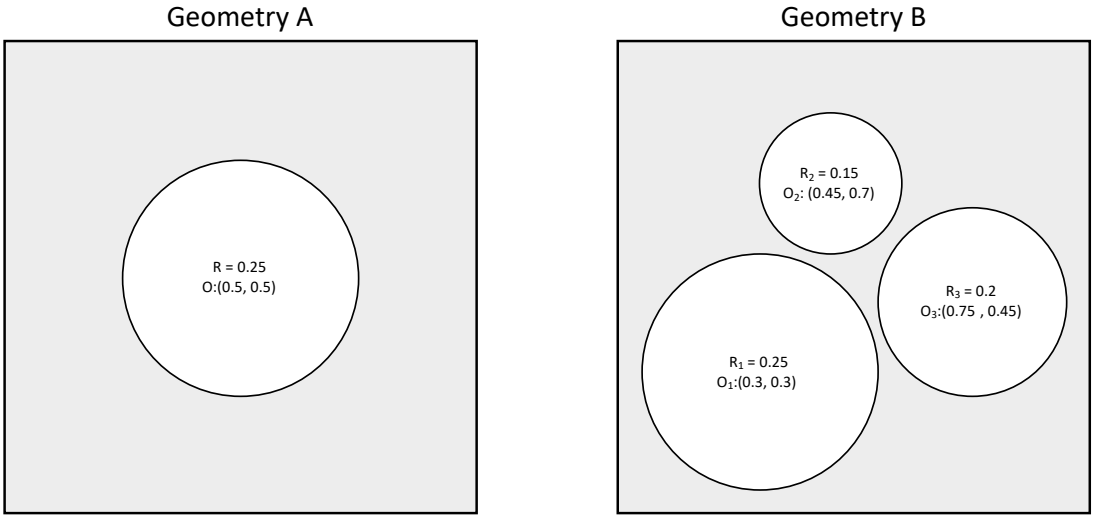

Figure 7.1: Periodic unit cells of the two test geometries. The dimensions are normalised relative to the periodicity length, which is d in both horizontal (x-) and vertical (y-) directions. The centres of circles are relative to the bottom-left corner.

as shown in Figure 7.2. It is one of the meshes used in Chapter 5, Section 5.5, which are derived from the transmission microscopy data of mouse white matter collected by Xu et al. [121]. The periodic domain consists of three compartments, namely the extra-axonal space, the myelin sheath and the intra-axonal space. In the simulations, only the extra-axonal and intra-axonal spaces are considered, with the myelin sheath being modelled as no-flux barriers between the two compartments. This is because the myelin sheath consists mainly of layers of lipid, and the image signal from the myelin sheath compartment is often nulled with the use of fat-suppression imaging techniques. The domain is augmented by appending mirror images of the image sample, so that periodic boundary conditions can be enforced to the opposite edges of the domain. For the detailed description of the meshing method and parameters, please see Appendix A.

7.3.3 Definition of simulation parameters

In all simulations, the diffusion coefficient D of the medium is assumed to be constant and uniform. As mentioned above, the maximum time for each simulation is set to be $t_{\max} = 3(|\mathbf{q}|^2 D)^{-1}$ respectively, and the number of time steps used is 1500. The initial condition is assumed to be uniform as well, without loss of generality this is set to be one.

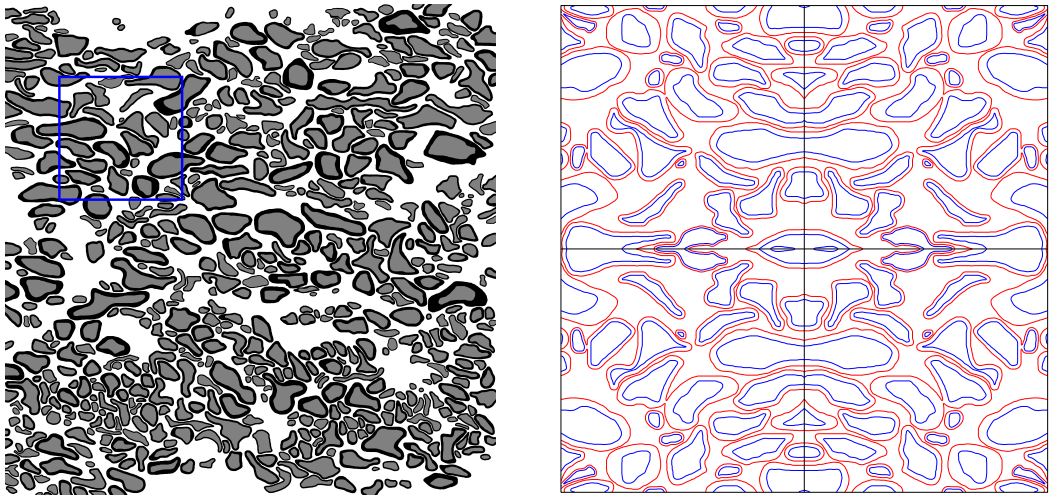

(a) Location of sample in TEM image (b) Periodic domain

Figure 7.2: Realistic microstructure geometry based on transmission electron microscopy data. A 7.1 μm $\times$7.1 μm (1000 px $\times$1000 px) sample has been cropped from the microscopy image (a), and then augmented to form the periodic domain (b).

All simulations are implemented in terms of non-dimensional quantities, as this allows the results to be easily scaled for any geometric length scale d and diffusivity D. Hence, the results are generally presented in non-dimensional axes as well, with the relevant non-dimensionalisation labelled on the axes.

For the first two simple geometries, 17 q-values in the range from $0.2d$ to $100d$ at roughly equal intervals on the log axis are considered, where d is the periodicity length. Meanwhile, for the realistic microstructure, 19 q-values in the range from 1.41×10^5 to 1.41×10^8 m^{-1} at roughly equal logarithmic intervals are used. In order to investigate the angular dependence on the q-vector, for the exterior compartment of Geometry A, q-vectors are applied at 0°, 15°, 30° and 45° relative to the x-axis. Meanwhile for all other cases, the q-vectors are only applied along the x-direction.

7.3.4 Results for simple model geometries

Qualitative description of signal behaviour

The signal behaviour of the image signal in Geometries A and B across multiple q-values and times are shown in Figures 7.3 and 7.4 respectively. From these visualisations, it

can be seen the signal behaviour transitions from a multi-compartmental homogenisation regime at low q-values (ie long time) to a boundary layer regime at high q-values (ie short time), thus agreeing with the asymptotic derivations in Chapters 5 and 6. Both the real and imaginary parts are shown in the figures. However, it is found that the imaginary parts always integrate to zero, which is a result consistent with both long-time and short-time asymptotic models[2]. Nevertheless, the spatial distribution of the complex signal intensity may still be of interest, especially if the image formation process is imperfect and introduces non-uniform phase shift in the voxel.

In the low q-value regime, the difference in signal magnitudes between the two compartments is clear. In particular, there is almost no signal decay in the interior compartment, which is a result of its isolated, disconnected nature. Meanwhile, at the other end of the spectrum when the q-value is high and the time is short, the signal only differs from a decaying background intensity within the boundary layers, which are attached to the diffusion barriers on both sides. The thickness of the boundary layers is largest when the boundary surface is normal to the q-vector applied and diminishes when they are tangent, which is consistent with the analysis in Chapter 6. The predicted ripples of the signal in the boundary layer, which is particularly significant in the imaginary part, can also be observed in the figures.

As the q-value decreases from the boundary layer regime, the thickness of the boundary layers increase as the scale separation becomes weaker. Eventually, the boundary layer reaches and interferes with another boundary layer originating from the diffusion barrier at the other side, thus bringing the signal behaviour into a transition region. In Figures 7.3 and 7.4, the $|\mathbf{q}|d = 4.47$ case, and to some extent the $|\mathbf{q}|d = 20$ case for Geometry B, show this kind of transitional signal behaviour quite clearly. In these cases, neither asymptotic models can describe the spatial distribution of signal intensity accurately. Meanwhile, comparing between the geometries, it can be seen that the transition behaviour in Geometry B is more complex. This is due to the increased number of ob-

[2]For the long-time model, this is because the effective diffusion tensor is always real, and the initial condition has a zero imaginary part. For the short-time model, see Section 6.3, Eq. (6.46).

structions in more irregular locations, leading to boundary layers interacting in a more complex manner, as shown in the $|\mathbf{q}|d = 20$ case.

Quantitative validation of asymptotic models

In terms of quantitative comparisons between the simulation results and the asymptotic models, the averaged signal intensities are considered. These correspond to observed MR image signals assuming they are formed by an ideal image acquisition process. Figures 7.5 and 7.6 show these results over multiple q-values, which are plotted with a log axis, and at two different times corresponding to $|\mathbf{q}|^2 Dt = 1$ and 3 respectively. The signals are averaged over individual compartments, as well as together in the combined signal with the compartments weighted according to their volume fractions. As mentioned previously, the imaginary parts of the signals always integrate to zero, therefore only the real parts are considered in the plots.

These results are overlaid on top of the signal values calculated using the two asymptotic models, as well as the signal if there were no diffusion barriers at all. For the low-q, long-time homogenisation model, the effective diffusion tensors for the exterior compartments are computed by numerically solving the cell problem using the method described in Section 5.4, whereas that for the interior compartments is trivially zero in both transverse directions. Since these effective diffusion tensors do not depend on the q-value, and in each plot the b-value i.e. $|\mathbf{q}|^2 Dt$ is held constant, the homogenisation model signals appear as constant horizontal lines on the plots. Meanwhile for the high-q, short-time boundary layer model, the boundary layer deficit for each surface element, which is a function of $q\sqrt{t}$ and its angle, is integrated numerically over the boundary surfaces for each case. The resulting curves converge to the no diffusion barrier line at the $|\mathbf{q}| \to \infty$ limit, as the thickness of the boundary layers tend to zero.

As shown in the plots, the simulation results validate the asymptotic models derived in the previous chapters. The transition region between the two asymptotes, however, shows some interesting behaviour. There is a trend that the transition region of Geometry

B spans across a wider range of q-values compared to that of Geometry A. A likely explanation for this is that the features in Geometry B are more complex and are of variable sizes, therefore it takes a wider range of q-values for the boundary layers to extend and completely "average" these features out to become a homogenised model. Moreover, the respective asymptotic models are not strictly an upper bound on the signal value either, with some overshooting and undershooting behaviour in various cases. However, in these two geometries at least, the overshoot in one compartment seems to partially cancel out the undershoot in the other, leading to the apparent effect that the transition region in the combined signal spans over a narrower range of q-values. However, more test geometries and theoretical analysis will be needed to justify whether this is in general the case.

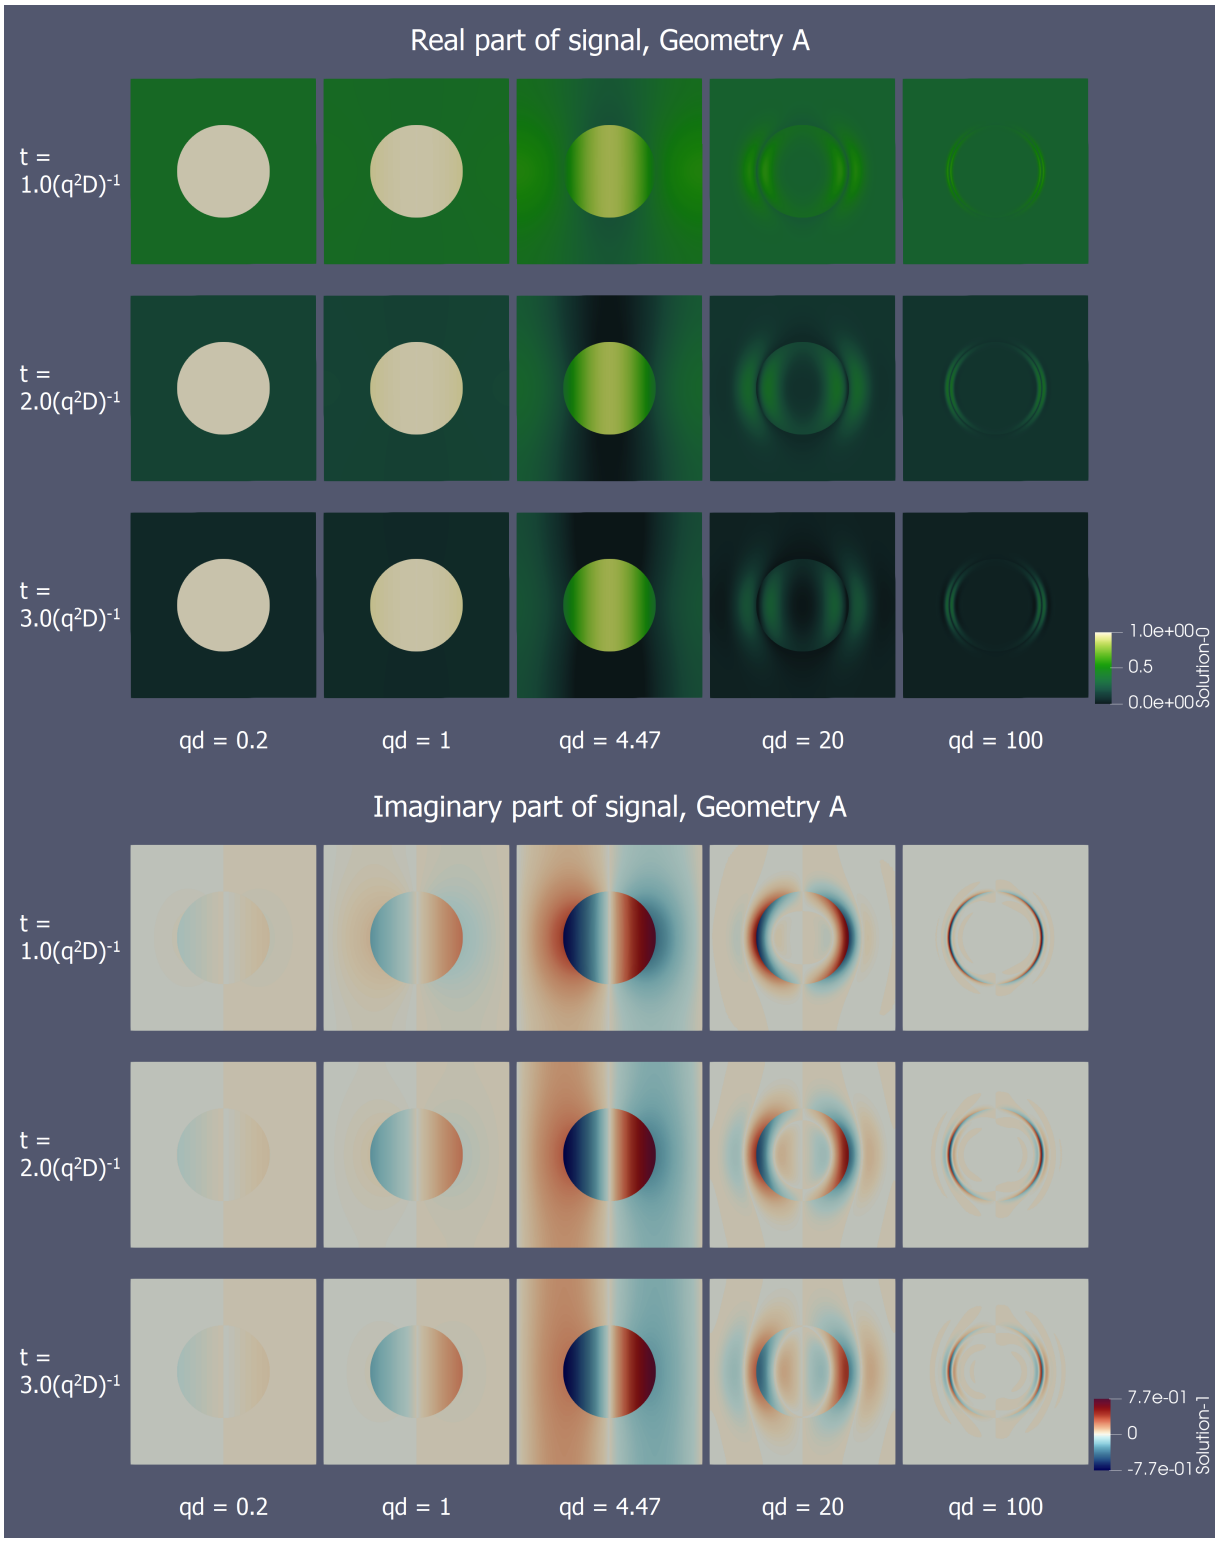

Figure 7.3: A visualisation of the image signal behaviour in Geometry A across multiple q-values and times. A separate divergent colour map is used for the imaginary part to demonstrate its antisymmetric nature. There is a clear transition between the low q-value (i.e. long time) multi-compartmental homogenisation and high q-value (i.e. short time) boundary layer regimes.

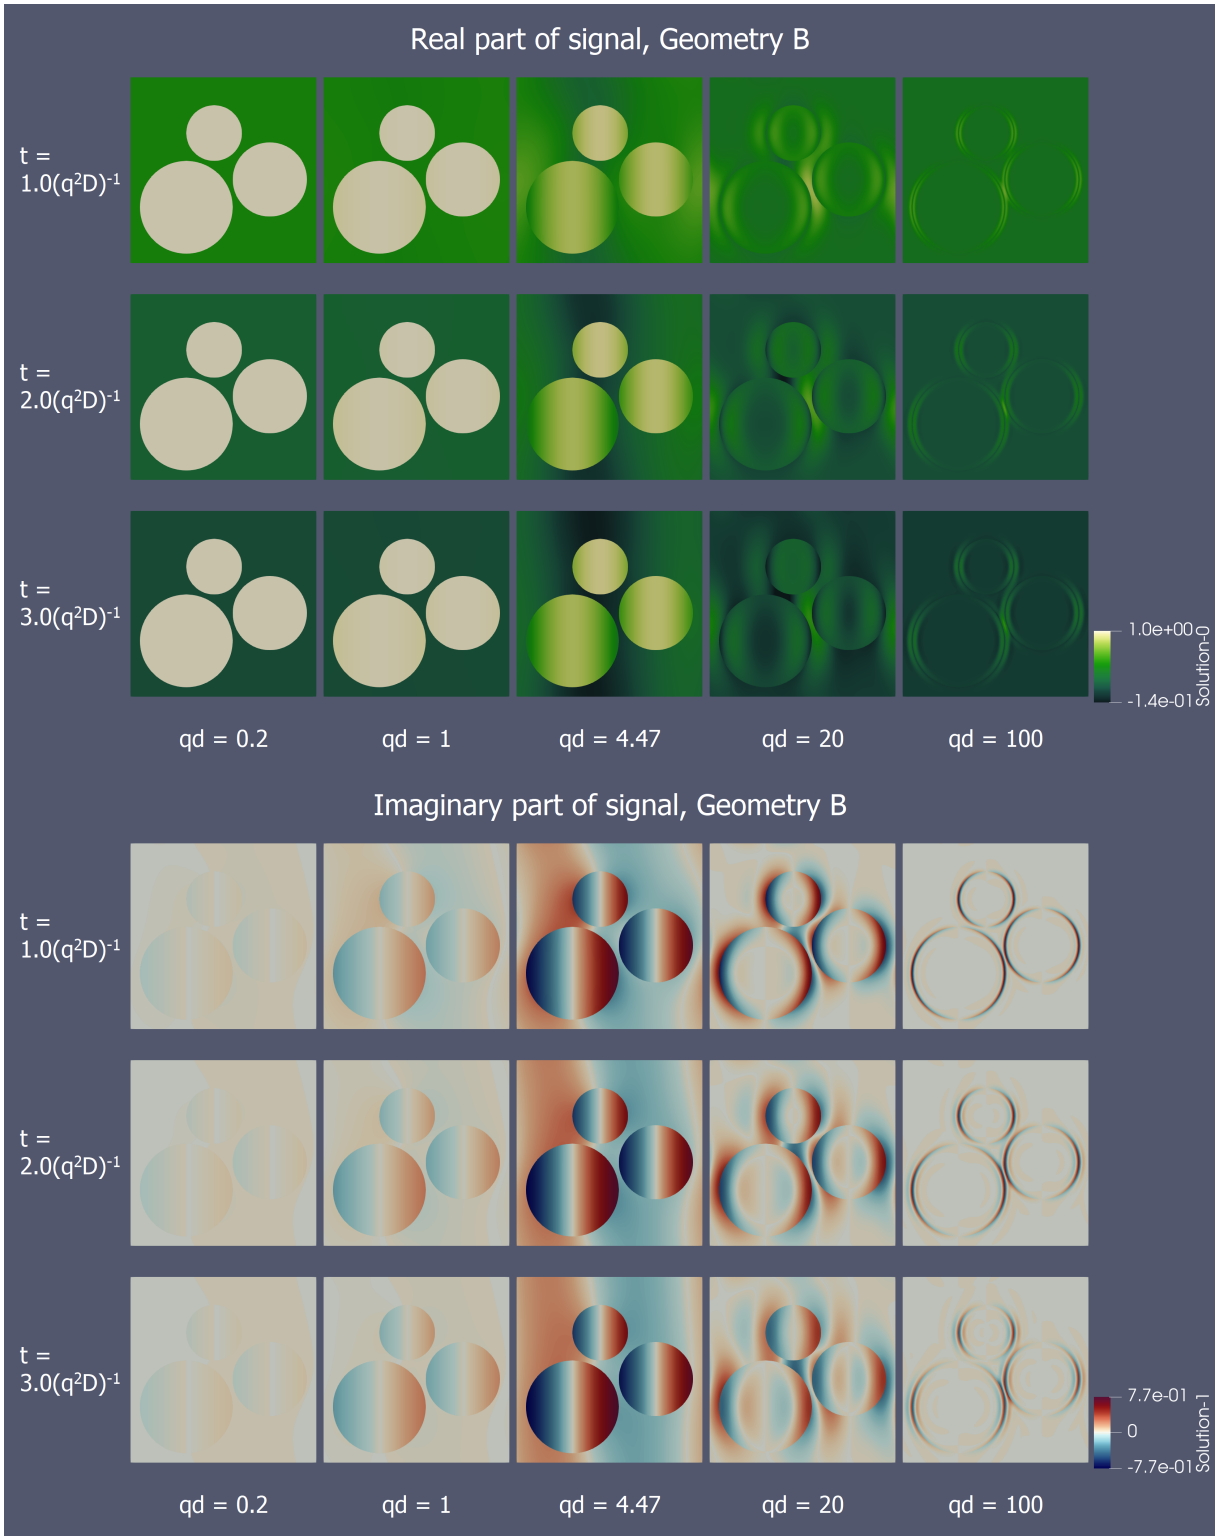

Figure 7.4: Similar visualisation to Figure 7.3 but for Geometry B instead. The signal behaviour follows a very similar trend, however the interaction of boundary layers becomes more complex in the transition region, as shown in the $|\mathbf{q}|d = 20$ figures.

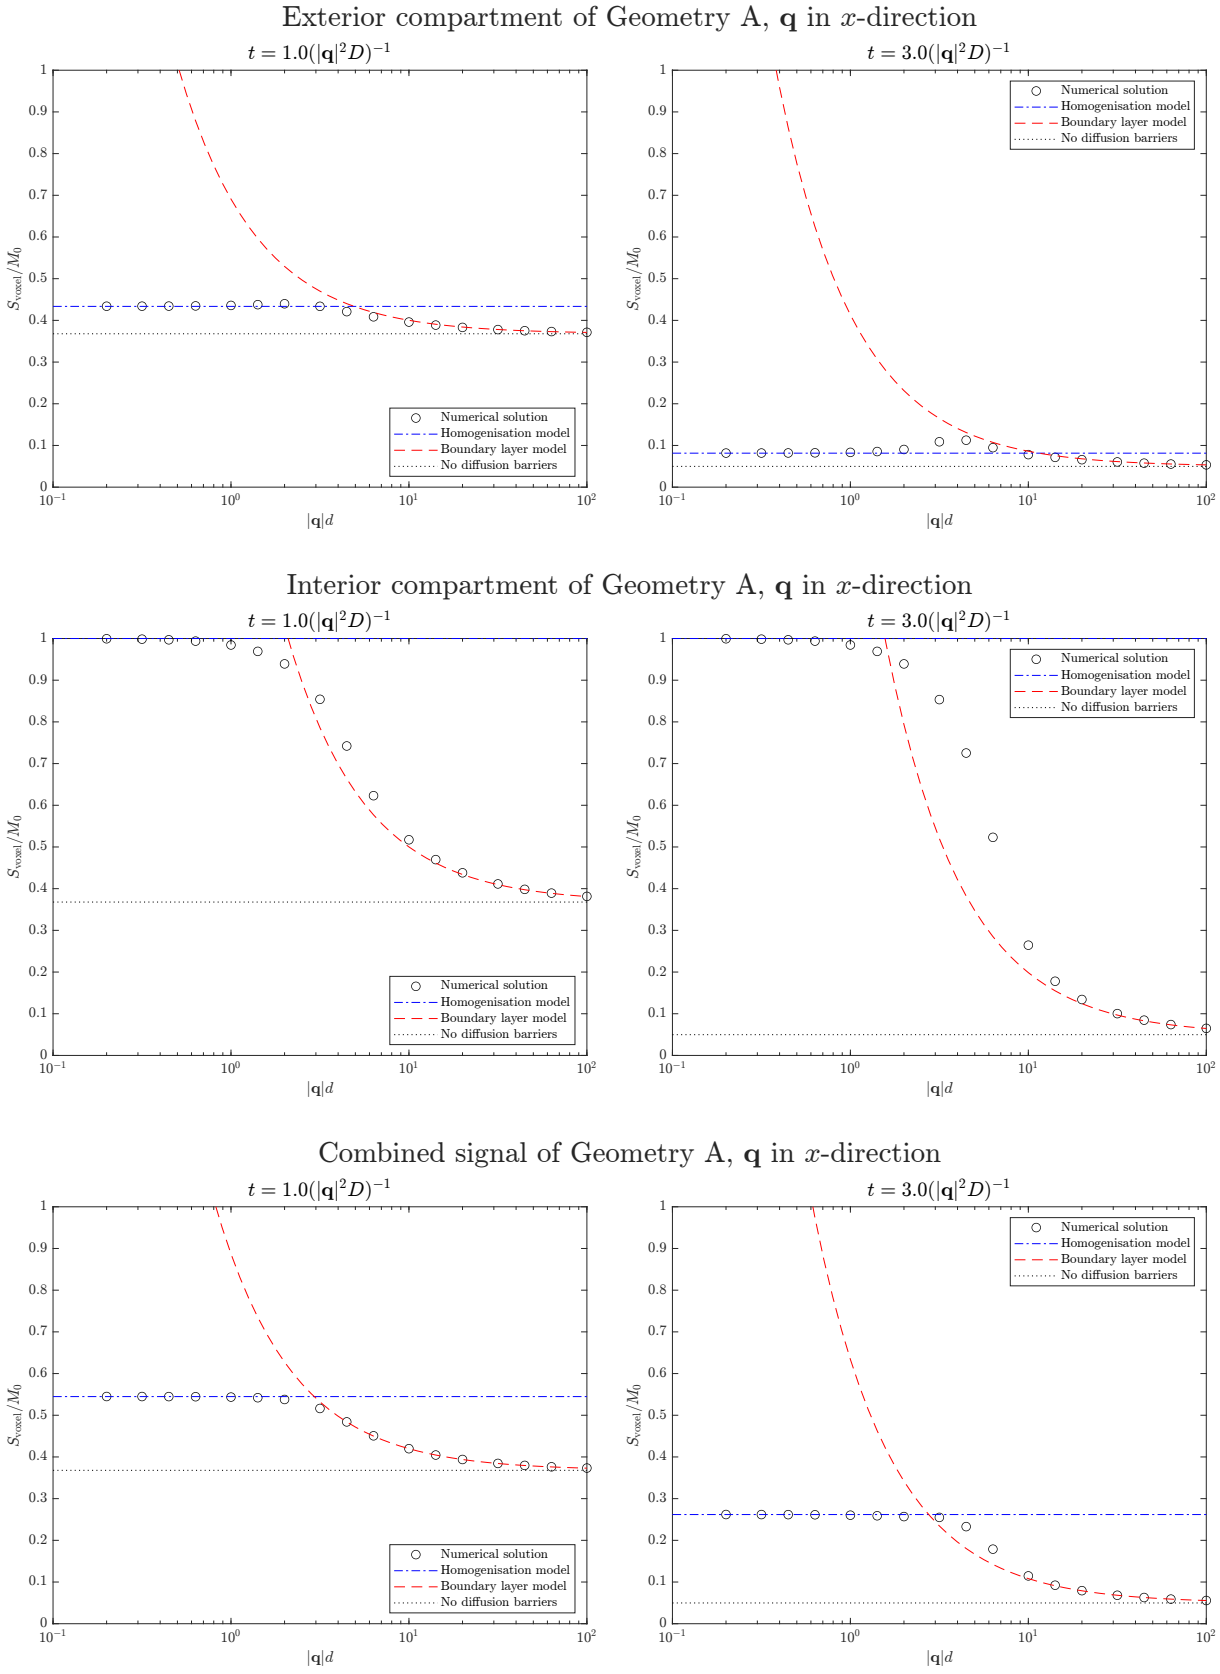

Figure 7.5: The comparison between numerical solutions and asymptotic models over a range of q-values and times for Geometry A. The signal intensities are averaged over individual compartments, as well as the two compartments combined. The plots are normalised against the periodicity length d and diffusion coefficient D.

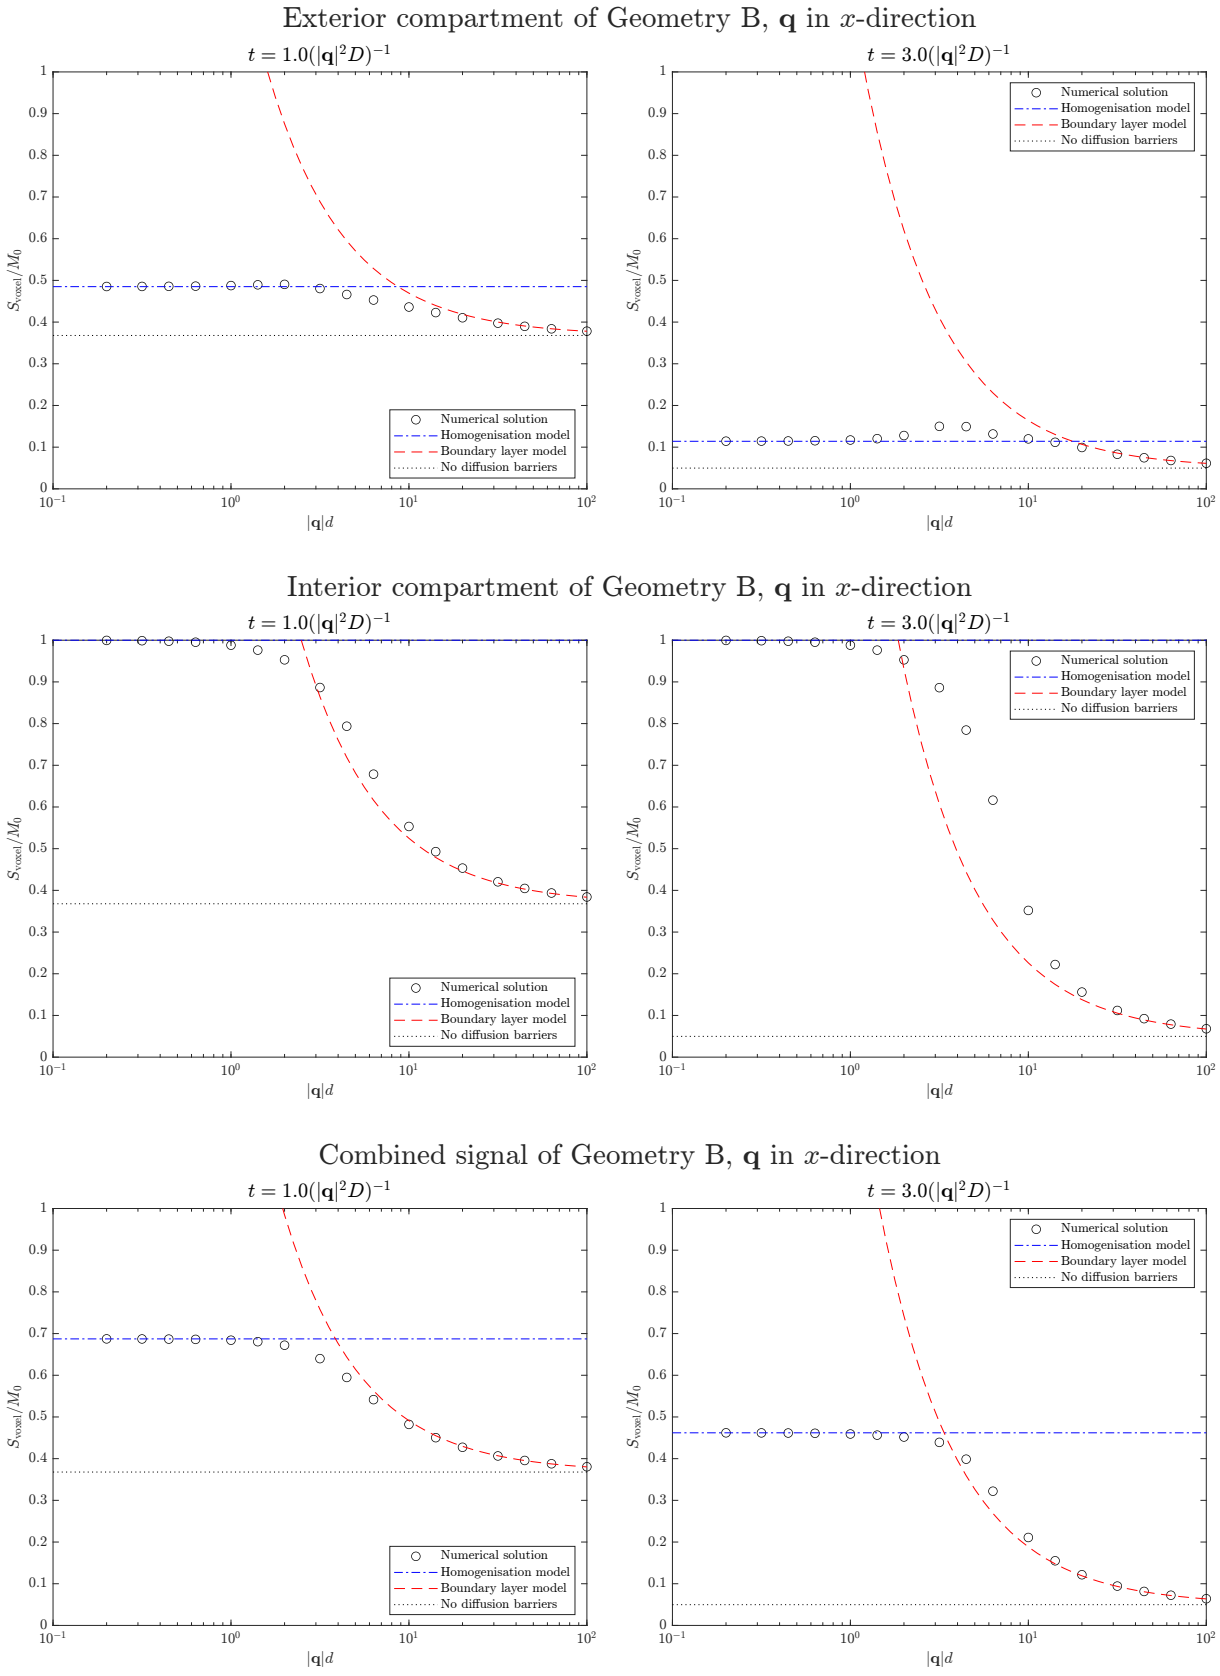

Figure 7.6: The comparison between numerical solutions and asymptotic models over a range of q-values and times for Geometry B. The transition region is broader than that for Geometry A, likely due to the increased number of different features in the geometry.

Angular dependence of transition region

In addition to the above studies where the q-vectors are all applied in the same direction, the angular dependence of the signal behaviour is studied by applying q-vectors at a number of angles in Geometry A. The specific reason for choosing Geometry A to study the angular behaviour is because of its interesting property. Its asymptotic models at both low-q, long-time and high-q, short-time regimes are both isotropic[3], yet the lattice layout of the inclusions in fact breaks the rotational symmetry in the exterior compartment. This means that regardless of the q-vector angle, the signal should converge to the same asymptotes at both high and low q-value limits, allowing for controlled comparisons between the transition region behaviours at different q-vector angles.

The simulation results are visualised and plotted in Figure 7.7. At $t = (|\mathbf{q}^2|D)^{-1}$, there is no significant differences between the different q-vector angles. However, at $t = 3(|\mathbf{q}^2|D)^{-1}$, the angular dependence becomes very pronounced, as shown in the visualisation and in the plots by the differing amounts of overshooting in the transition regions. This phenomenon is closely related to the "diffusion-diffraction" concept proposed in the literature (see Chapter 2, Section 2.4.5), which uses an analogy to wave diffraction to explain the behaviour when the modulation wavelength is of a similar magnitude to the microstructure. However, since it takes time for the eigenmodes of the system to separate by decaying at different rates, this angular sensitivity is only apparent at longer times. An inevitable consequence of this is when considering the combined signal across both compartments, the decaying signal magnitude in the exterior compartment will be dominated by the signal within the interior compartment, which decays at a lower rate due to its isolated, disconnected nature. In the particular case of Geometry A, the interior compartment is genuinely isotropic, since it is a collection of individual symmetric cylinders. Hence, in the bottom plots of Figure 7.7, it can be seen that the angular sensitivity of the transition region is somewhat masked in the combined signal.

[3]For the long-time model, the effective diffusion tensor is isotropic in the transverse directions, see Chapter 5, Section 5.4.3. For the short-time model, the boundary surfaces have an even angular distribution, hence the boundary layer model stays the same regardless of the q-vector direction.

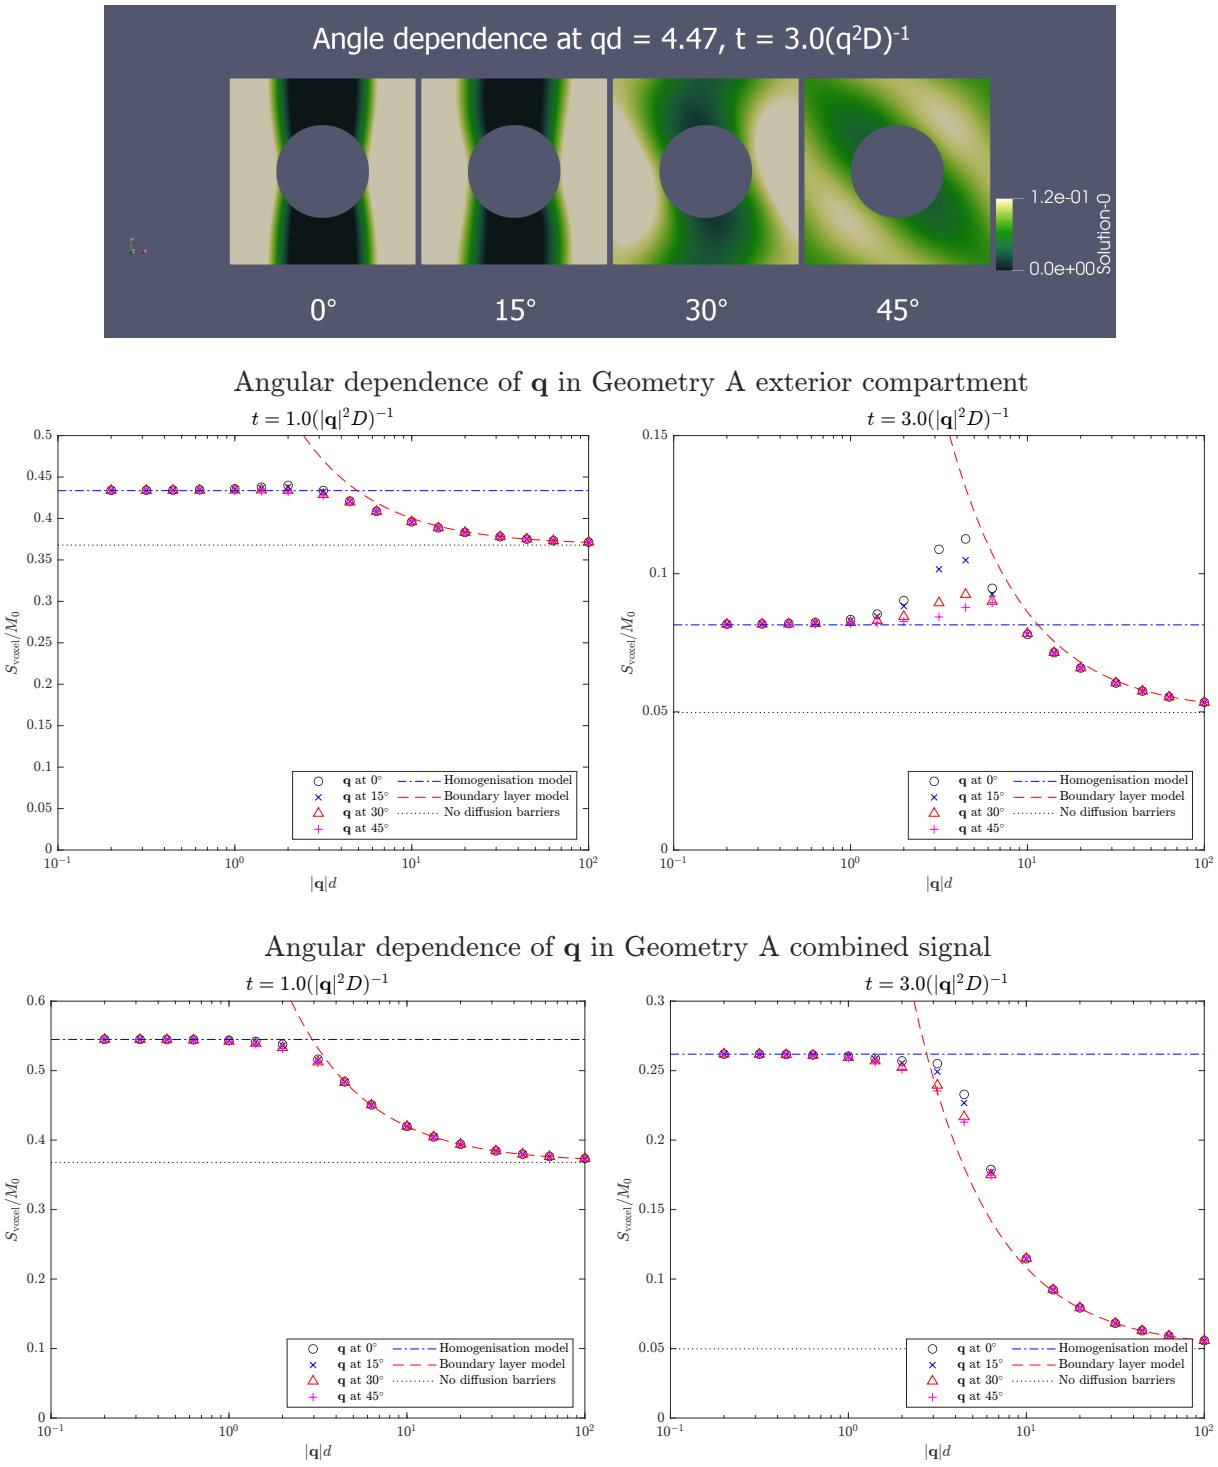

Figure 7.7: Top row: visualisation of the signal intensity distribution when the q-vector is applied at different angles. Middle row: the average signal over the exterior compartment, showing the angular sensitivity of the transition region at longer times. Bottom row: Average signal over all compartments. The angular dependence is less apparent once the isotropic interior compartment is also taken into account.

7.3.5 Results for realistic microstructure

Within the realistic microstructure, the simulations show the same transition behaviour from the multi-compartmental regime to the boundary layer regime. As with the data from model geometries, the spatial distribution of the signal intensity over multiple q-values and times are shown in Figure 7.8, whereas the averaged signals over individual and combined compartments are plotted in Figure 7.9 along with the results from asymptotic models.

It can be seen that the transition region spans roughly a decade of q-values. This is significantly broader in terms of its q-value range in the realistic microstructure than in the model geometries, which is expected due to the much larger number of irregular features present in the geometry. Moreover, it can be observed that the transition region in the extra-axonal space occurs at higher q-values (about 5 times) compared to the intra-axonal space.

There are two possible explanations or hypotheses for this phenomenon. The first one relates to the relative locations of the intersections between asymptotes. The boundary layer curves in the plots reflect the surface-to-volume ratios in the respective compartments; for this microstructure geometry they are similar in the two individual compartments since they have almost the same volume fractions (extra-axonal: 0.341, intra-axonal: 0.343). However, the homogenisation models differ significantly, as one compartment is isolated whereas the other is connected. Therefore, the model intersection happens at higher q-values for the extra-axonal space, so the transition region should be located at higher q-values accordingly. The second hypothesis relates to the length scale reflected by the signal behaviour in each compartment. The close-packing nature of the microstructure geometry (see Figure 7.2) means that the characteristic length *between* different axons in the extra-axonal space is shorter than that *within* individual axons in the intra-axonal space.

The two explanations are related to each other, as the similar volume fractions should

imply the separation in characteristic lengths between connected and isolated domains. Using a square lattice of cylinders as an example, the diameter needs to be $\sqrt{2/\pi} = 0.798$ times the periodicity for the volumes in the exterior and interior spaces to be equal, meaning that the cylinders are only 0.202 times the periodicity apart at the closest location. This ratio is roughly consistent with the findings from the simulations.

This finding regarding the different locations of transition regions may be useful in informing the design, interpretation and biophysical modelling of diffusion MRI studies. Referring to the studies reviewed in Section 3.3.3, recent MRI studies using high gradients reach q-values of around 5×10^5 m^{-1}. The simulation data in Figure 7.9 show that at this q-value, the extra-axonal space is still described by a homogenisation model, whereas the intra-axonal space has already entered the transition region. This is also shown in the spatial distribution of signal intensity at $|\mathbf{q}| = 6.30 \times 10^5$ m^{-1} (see Figure 7.8), with the signal intensity being much more variable within the axons. This also agrees with the findings of Grebenkov [42], in which it is suggested that the extra-axonal signal starts to deviate from a Gaussian behaviour at the order of $|\mathbf{q}| = 3 \times 10^5$ m^{-1}. Hence, for diffusion MRI studies within this regime, the information about the length scales of the microstructure is only contained in the intra-axonal contributions to the signal, as the homogenised extra-axonal diffusion tensor does not depend on the length scale of the feature at all, however if the gradient strength becomes even stronger non-Gaussian behaviour can start to be expected as we move into the transition regime in the extra-axonal space as well.

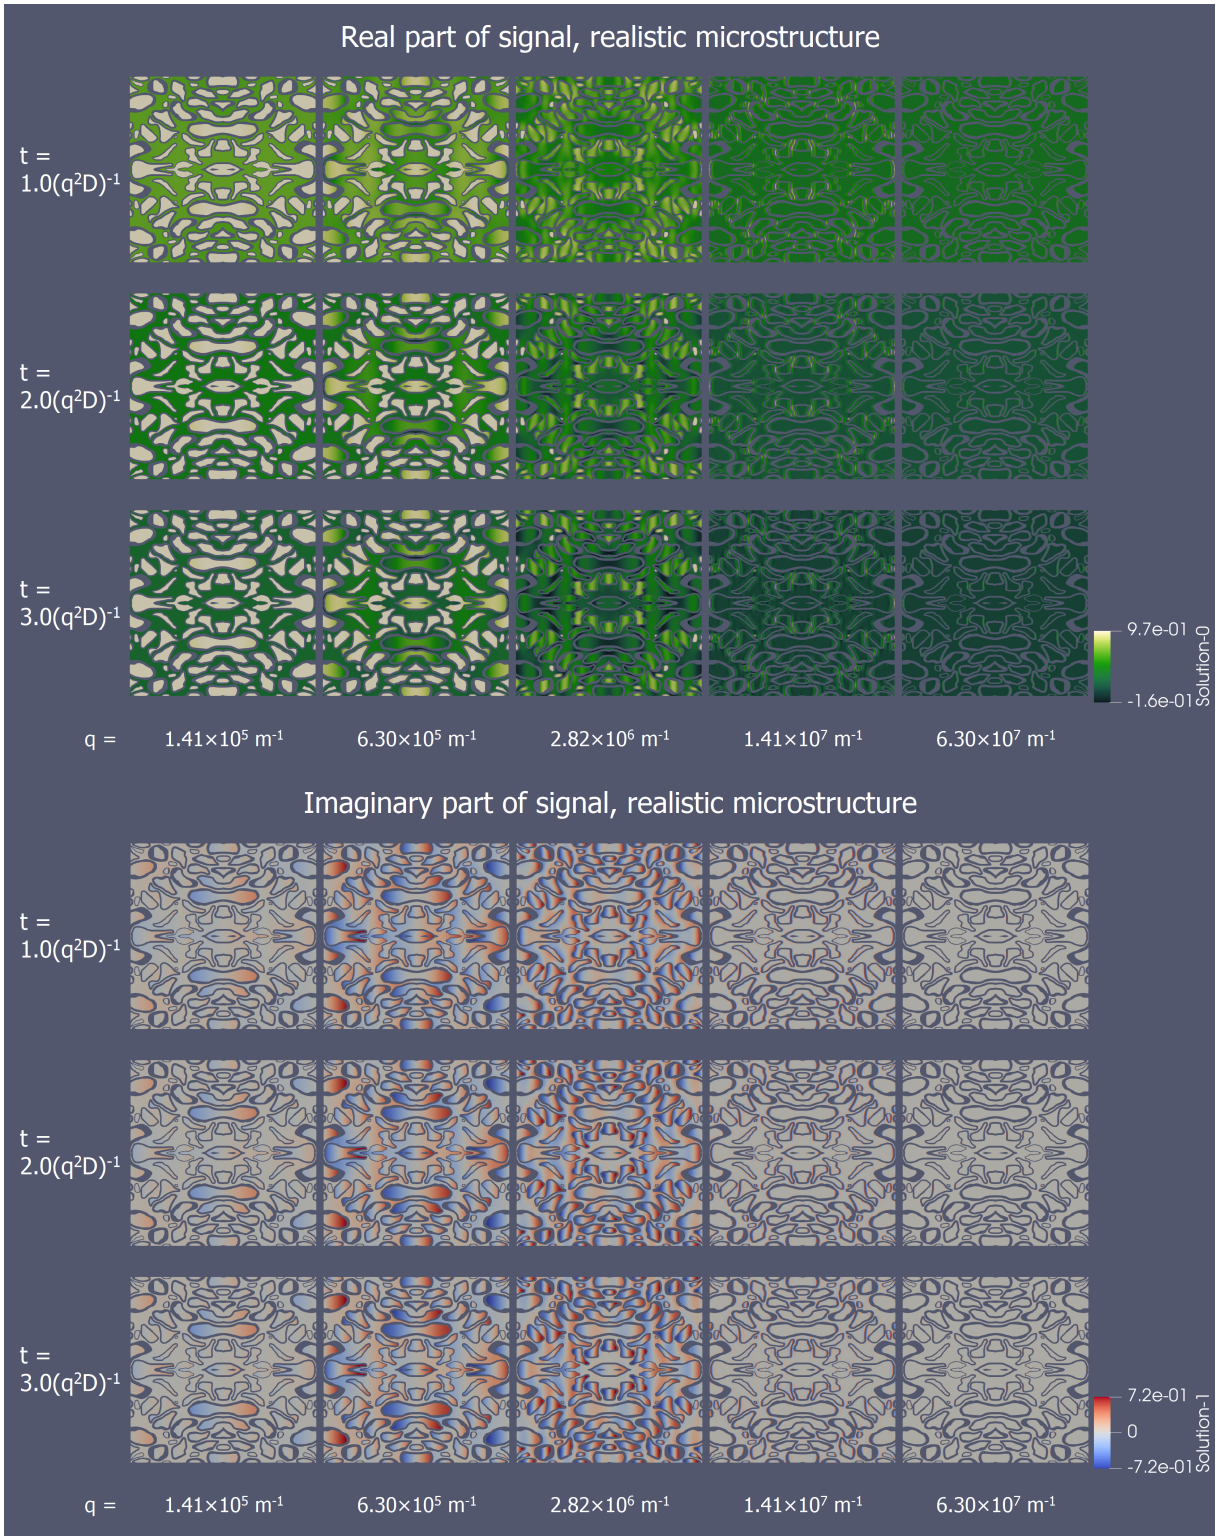

Figure 7.8: Spatial distribution of signal intensity in the realistic microstructure across multiple q-values and times, showing the transition between homogenisation and boundary layer regimes.

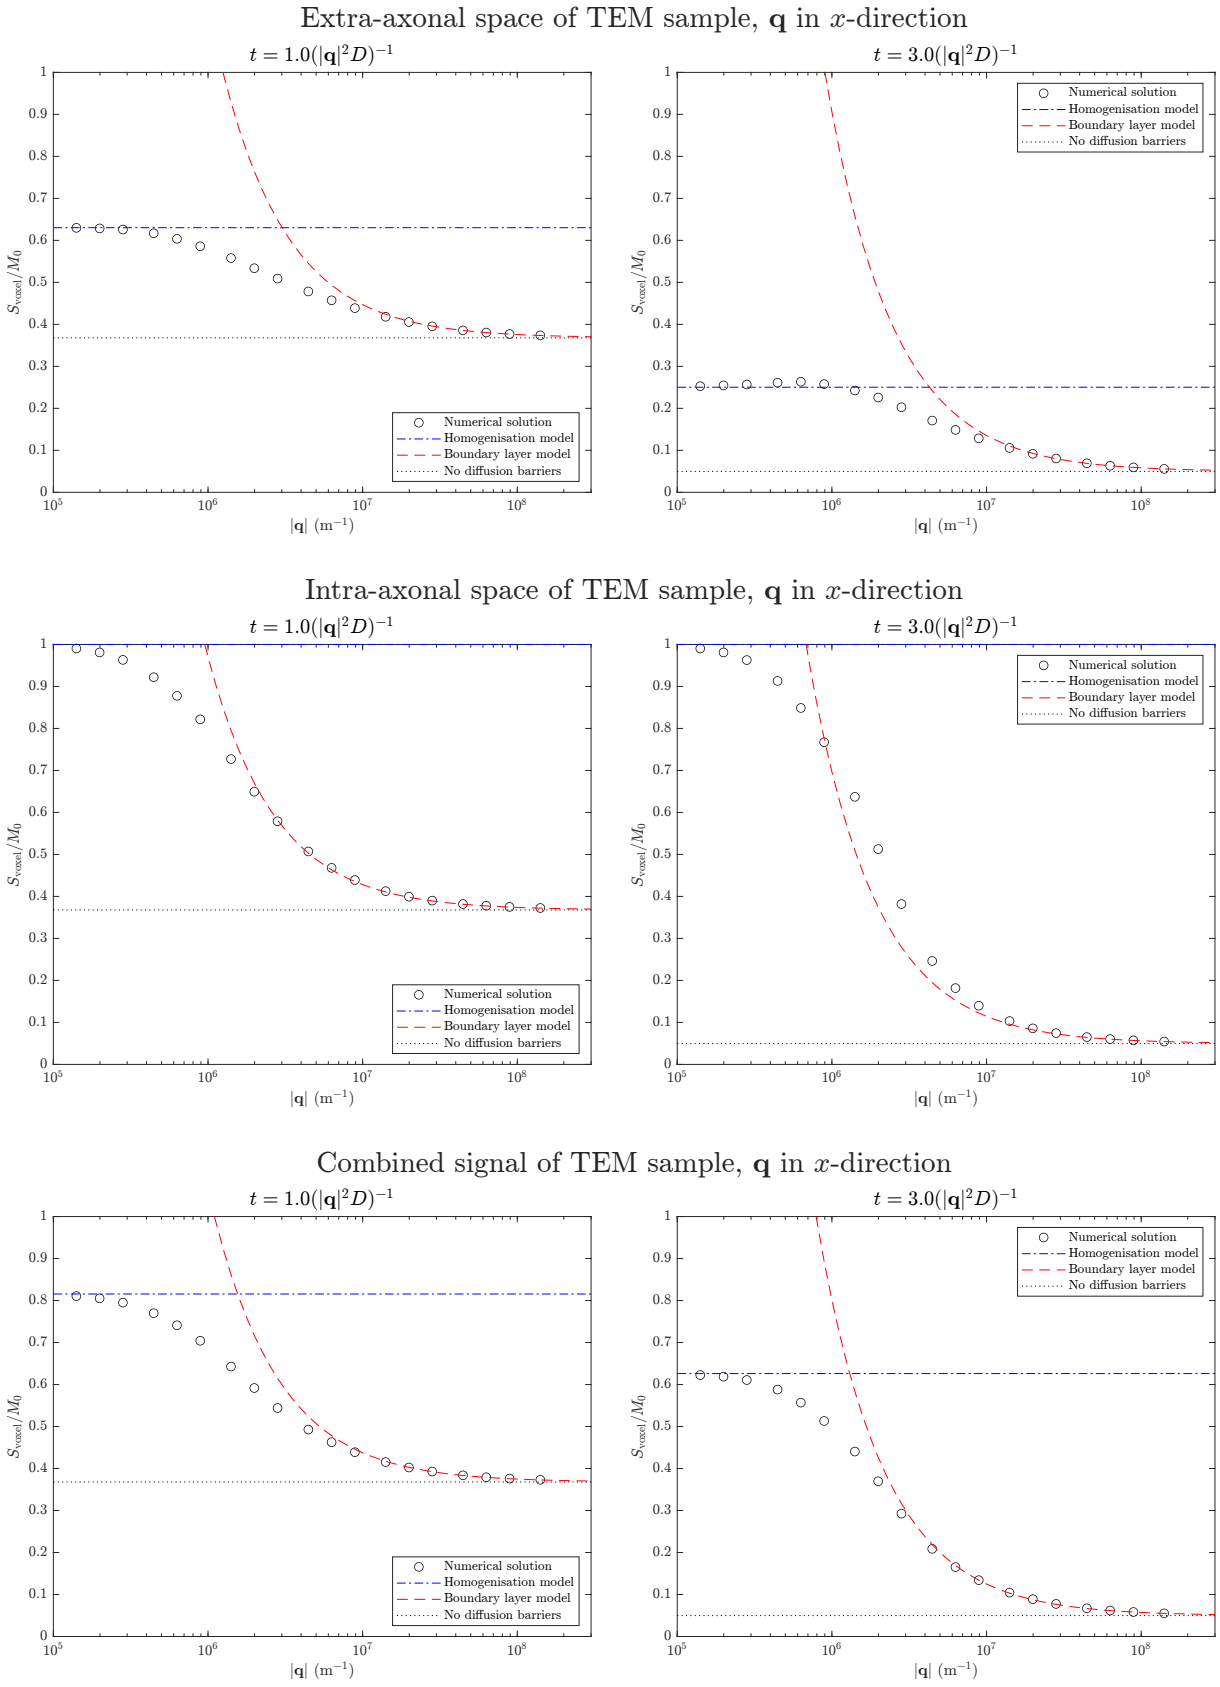

Figure 7.9: The comparison between numerical solutions and asymptotic models over a range of q-values and times for a realistic microstructure geometry. The transition region covers a range of about a decade in q-value. It also occurs at higher q-values in the extra-axonal space (roughly 5 times difference).

7.4 Effects of time-varying q-vector

7.4.1 Overview of section

In the previous section, the transition behaviour between the asymptotic regimes has been studied through simulations using a constant q-vector. However, as discussed before, this corresponds to the unphysical assumption of instantaneous modulation of spins with an infinitely strong gradient field. In the short-time regime, the effects of time-varying q-vector have been studied in Chapter 6, at least asymptotically. Meanwhile in the long-time regime, the derivation of the homogenisation model does not involve the constant q-vector assumption.

In this section, the finite element simulations are therefore utilised to bridge the gap and study the effects of time-varying q-vector in the intermediate-time transition regime. The aim of this section is to firstly validate the results and the two-point correction method proposal in Chapter 6 for the boundary layer regime. By using the finite element simulations over the entire geometry, the feasibility and accuracy of extending the same correction protocols to other regimes can then be evaluated. In particular, since the b-value is held constant across the pulses used in the two-point correction protocol, and the signal in the homogenisation model only depends on the b-value but not the time profile of the q-vector, the correction method is expected to also work in the homogenisation regime (as there is no need for correction). However, the accuracy of the correction method in the transition region remains to be seen.

7.4.2 Definition of numerical simulations

In the finite element simulations, time-varying q-vectors are applied along the x-direction to both extra-axonal and intra-axonal spaces in the realistic microstructure geometry used in the previous section. Sets of gradient pulses, which give rise to different q-vector profiles, are considered over four nominal q-values between 1.41×10^5 m^{-1} and 1.41×10^8 m^{-1} at

equal logarithmic intervals of a decade. According to the results from the previous section, these q-values should cover the whole range of different regimes. For each nominal q-value, two different mixing durations (t_m) are considered, corresponding to $|\mathbf{q}|^2 D t_m = 1$ and 2.

At each nominal q-value and time, the set of different gradient pulses is defined using the same methodology in Sections 6.5 and 6.6. Firstly, unphysical pulses in the form of delta functions, which correspond to the constant q-vector case, are considered. Secondly, rectangular pulses with pulse durations (t_p) being 0.3 and 0.6 times the mixing duration are considered. The peak q-value for these pulses are adjusted so that the overall b-values of the pulse sequences, i.e. $\int_0^t |\mathbf{q}|^2 D \, \mathrm{d}\tau$, are held constant. Finally, half-wave pulses with $t_p/t_m = 0.4,\ 0.8$ and triangular pulses with $t_p/t_m = 0.429,\ 0.857$ are considered, as the analysis in Section 6.6 suggests that they should be equivalent to the rectangular pulses. The comparison between the equivalent pulse sequences for the $t_p^{\mathrm{eq}}/t_m = 0.6$ case can be found in Figure 6.18 in the previous chapter. These gradient pulse profiles are integrated to give the q-vector profile, which are then discretised in 1000 time steps over $t \in [0, 2t_m]$ and implemented in the finite element simulations.

7.4.3 Results

The results of the simulations are plotted in Figure 7.10. Firstly, the signals from the suggested equivalent pulses are compared, which are shown as differently shaped data points on the plots. The results show excellent agreement in all cases, with the relative differences within the expected discretisation error from the backward Euler method. The performances of the two-point correction method at different q-values are then analysed. The extrapolations are plotted as lines in the plots, which allow for a comparison between the extrapolated x-intercept and the actual signal from a hypothetical delta function pulse. The comparison between the extrapolation error and the estimation error using a $t_p/t_m = 0.3$ pulse alone is shown in Table 7.1. As shown in the figure, there is as expected little need for correction in the $|\mathbf{q}|_{\mathrm{nom}} = 1.41 \times 10^5$ m^{-1} case, which is in the homogenisation regime. Meanwhile, for $|\mathbf{q}|_{\mathrm{nom}} = 1.41 \times 10^7$ m^{-1} and 1.41×10^8 m^{-1}, which

Table 7.1: Comparisons between the error in predicting the constant q-vector signal by using $t_p/t_m = 0.3$ pulses alone and extrapolation using $t_p/t_m = [0.3, 0.6]$ pulses. Ext / Int = extra-axonal / intra-axonal spaces.

$t_m = 1.0(|\mathbf{q}|^2 D)^{-1}$:

$\lvert\mathbf{q}\rvert_{\text{nom}}$ (m^{-1})	1.41×10^5		1.41×10^6		1.41×10^7		1.41×10^8	
	Ext	Int	Ext	Int	Ext	Int	Ext	Int
$t_p/t_m = 0.3$ only	0.37%	0.72%	2.01%	3.97%	1.06%	1.06%	0.20%	0.17%
Extrapolation	0.33%	0.63%	0.90%	1.54%	0.30%	0.18%	0.08%	0.08%

$t_m = 2.0(|\mathbf{q}|^2 D)^{-1}$:

$\lvert\mathbf{q}\rvert_{\text{nom}}$ (m^{-1})	1.41×10^5		1.41×10^6		1.41×10^7		1.41×10^8	
	Ext	Int	Ext	Int	Ext	Int	Ext	Int
$t_p/t_m = 0.3$ only	0.46%	0.83%	4.68%	7.89%	3.38%	3.68%	0.68%	0.56%
Extrapolation	0.45%	0.77%	2.40%	3.93%	1.12%	0.72%	0.23%	0.21%

are both well into the boundary layer regime, the corrections show good agreement with the delta function pulse signals. The performance of the correction method, however, is not as good in the $|\mathbf{q}|_{\text{nom}} = 1.41 \times 10^6$ m^{-1} case, which is in the transition region between homogenisation and boundary layer regimes. However, of all the q-values considered, the error before extrapolations is the greatest in this case. The extrapolation at least offers some improvement to the accuracy of predicting the hypothetical constant q-vector signal, with the error about halved comparing to if only $t_p/t_m = 0.3$ pulses are used.

7.4.4 Limitations and future work

The simulations in this section have validated the use of equivalent rectangular pulses to model the signal from irregular but symmetric gradient pulses over q-values ranging from the homogenisation regime to the boundary layer regime. Furthermore, even though only four q-values each a decade apart were considered, the simulations provided some measurement of the performance and usefulness of the two-point extrapolation correction method proposed in the previous chapter.

However, as shown in the plots in Figure 7.10, there seems to be a big gap in the transition regime between $|\mathbf{q}|_{\text{nom}} = 1.41 \times 10^6$ m^{-1} and 1.41×10^7 m^{-1}, where the extrapolation correction would be significant and hence of great interest. The range of normalised b-

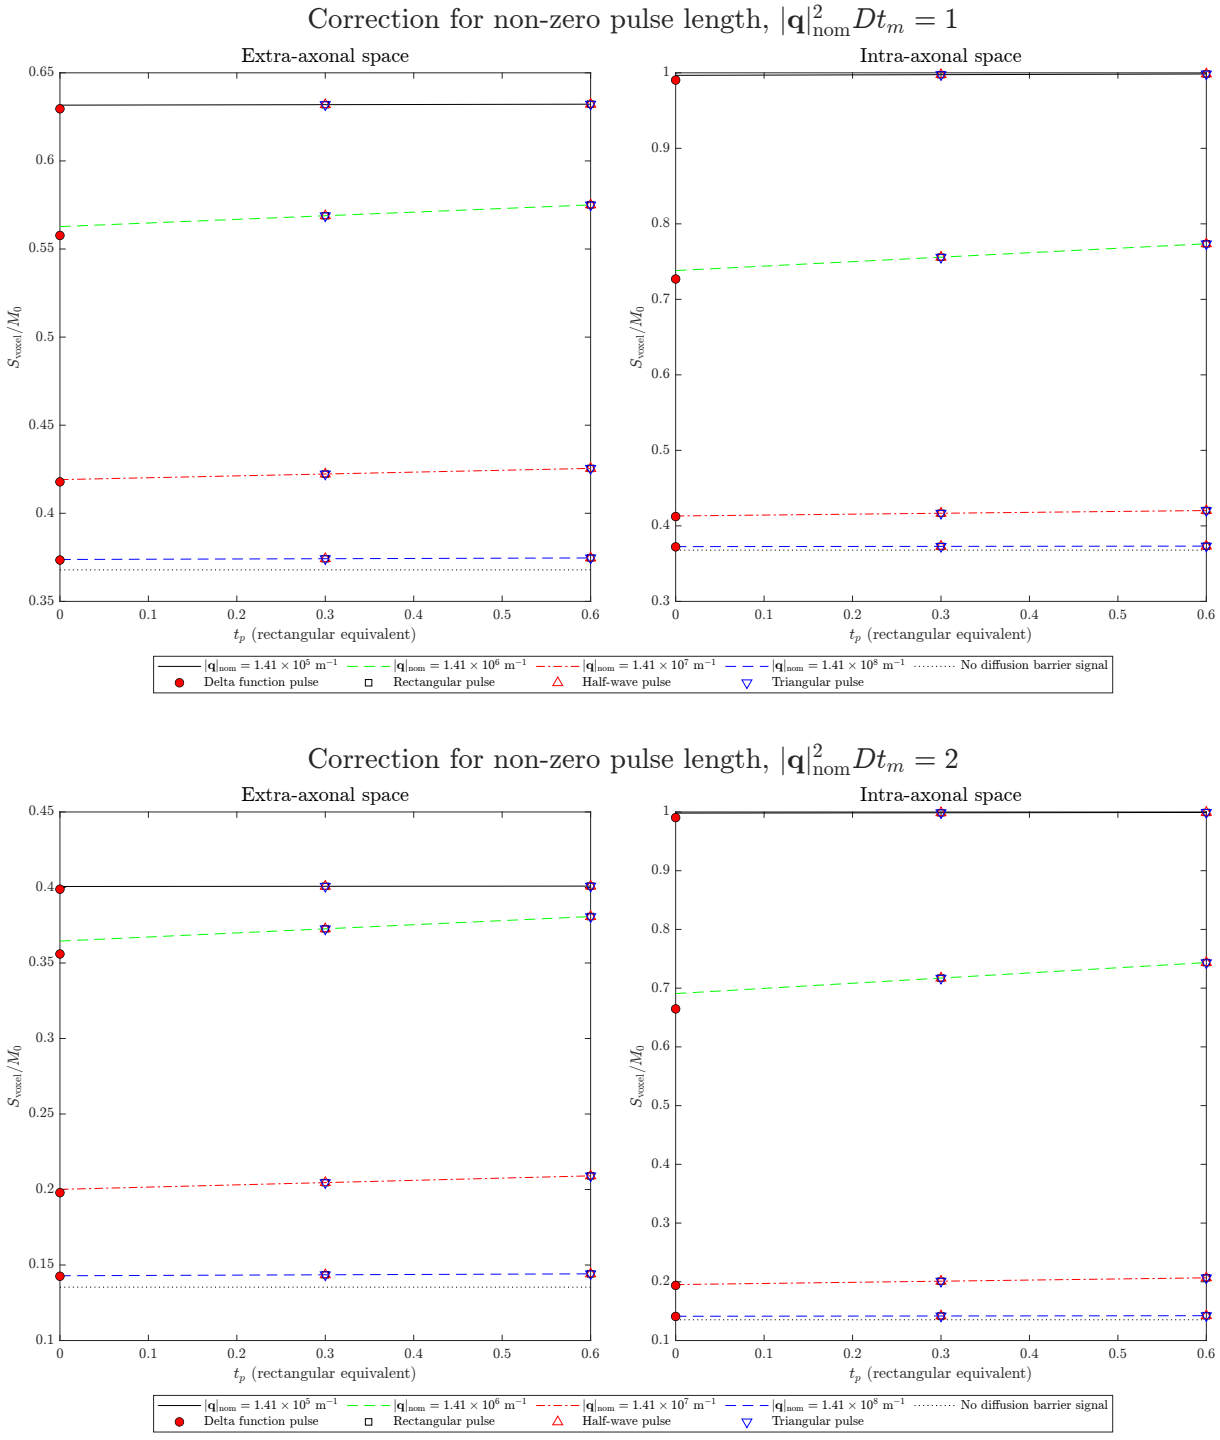

Figure 7.10: Effect of time-varying q-vectors due to finite gradient pulses at a range of nominal q-values. The lines represent the two-point extrapolation correction method proposed in Section 6.5 to estimate the hypothetical constant q-vector signal, which requires delta function pulses effectively. For the irregular pulses, the equivalent rectangular pulse lengths are calculated using the method proposed in Section 6.6.

values, i.e. $|\mathbf{q}|^2 D t_m$, has been limited to only 2 in the study as well; ideally simulations should cover a larger range to reflect the higher b-values used in recent diffusion MRI research experiments.

One major downside of these numerical studies is their computational complexity, as the lack of time invariance in the q-vector means that each data point comes from a separate simulation. The simulations took nearly three hours each to run on a six-core Intel® Xeon® CPU to produce the results presented, thus only four q-values were considered in this section. This does also highlight, however, the value of the asymptotic models derived in Chapters 5 and 6 in reducing the computational complexity by orders of magnitude.

In the future, the analysis and results in this section can be extended to more simulation scenarios to address the gaps in current simulation data. In particular, the finite element method presented is implemented using parallel computing, hence it is easily scalable to high-performance distributed computing clusters. Access to such computing resources will speed up the simulation time significantly, as well as allowing for finer time steps and bigger meshes to achieve better simulation results.

7.5 Summary

In this chapter, the intermediate-time regime of DW-SE, where all the terms in the Bloch-Torrey equation are similar in magnitude, is studied using numerical simulations. This is achieved by discretising the co-rotating, co-decaying form of the Bloch-Torrey equation using a finite element method. Two simple model geometries of white matter and a realistic microstructure based on microscopy data are considered, with simulations being conducted over a range of q-values covering three orders of magnitude, extending into the two asymptotic regimes covered in Chapters 5 and 6. The results show good agreement with the asymptotic models in their respective limits, whereas a transition region is observed between them. The width of the transition region increases with the complexity of

the geometry, and it is found to span across about a decade in q-value for the realistic microstructure. The transition region is also found to be sensitive to the angle of the applied q-vector in a regular lattice structure. Meanwhile, for the realistic microstructure geometry, it is found that the transition region occurs at a higher q-value in the extra-axonal space than in the intra-axonal space. This finding might be relevant to recent diffusion MRI studies, as in their operating regimes the intra-axonal signal contains information about the microstructure length scale, while the extra-axonal space still behaves as an effective homogeneous medium.

The effects of time-varying q-vectors, which is a result of finite gradient fields used in realistic diffusion MRI experiments, are then analysed. The finite element method allows the evaluation of both the two-point correction method and the equivalent pulse method, which are proposed originally in the boundary layer regime in Chapter 6, in the intermediate-time regime as well. It is found that the equivalent pulse method produces almost identical signals at a wide range of q-values, covering from the homogenisation regime to the boundary regime. Meanwhile, in the intermediate regime, the two-point correction method roughly halves the prediction error of a hypothetical constant q-vector image signal. However, in the future, more data points will be needed to further evaluate the performance of the two-point correction method, especially in the transition region between the two asymptotic regimes where the correction has the greatest effect.

Chapter 8

Conclusion and future directions

8.1 Summary of research

In this thesis, the behaviour of diffusion MRI in heterogeneous media has been modelled systematically using multi-scale asymptotic methods. Both magnetic heterogeneities in the form of magnetic susceptibility variations and molecular transport heterogeneities in the form of no-flux barriers have been considered simultaneously in the modelling.

In Chapter 3, a novel distribution function formulation of diffusion-weighted spin-echo (DW-SE) has been proposed, which allows for easier interpretation of the signal behaviour using an augmented dimension diffusion-advection analogy. With a change of frame of reference in the augmented dimension, the local spatial modulation wave-vector, i.e. the q-vector, and the time scale were found to be the two relevant scales to characterise a DW-SE experiment. This is similar in concept to recently proposed q-t characterisations of diffusion MRI, although the novelty of the change of frame of reference within the distribution formulation leads to the notion of *local* variations of q-vectors, which is essential in the treatment of micro-scale magnetic heterogeneities.

Chapter 4 then considered the local q-vector variations due to magnetic heterogeneities microstructure. Multi-scale analysis showed that the local q-vector variations are of comparable magnitudes to the macro-scale applied q-vector, if current typical research hard-

ware is used in the long-time regime. This thus led to the simultaneous modelling of magnetic and transport heterogeneities for the long-time regime of DW-SE in Chapter 5. The result is a multi-compartmental effective homogeneous model, where the effective diffusion tensors can be computed by solving the cell problem over a representative periodic microstructure cell. However, the effects of microstructure scale q-vector variations have been found to integrate exactly to zero.

Chapter 6 then considered the short-time regime of DW-SE using boundary layer analysis. In this regime, the effects of transport heterogeneities are only significant in the boundary layer near the boundary surfaces. The results when the q-vector is constant over time agree with the literature, and they also reveal a spatially varying imaginary part of the complex signal intensity that has not been reported in the literature before. This imaginary part integrates to zero over the boundary surface, however if the image acquisition process is non-uniform the resulting signal will have a resultant imaginary part. The effects of time-varying q-vectors to the boundary layer model have then been analysed. These correspond to finite gradient field strengths, rather than unphysical infinite fields needed for the constant q-vector assumption. The results led to the proposal of a correction method for irregular symmetric gradient pulses, firstly by calculating an equivalent rectangular pulse by two simple relations, and secondly by a two-point extrapolation method to estimate the hypothetical constant q-vector signal.

Finally, in Chapter 7 the intermediate-time regime of DW-SE has been considered through the use of numerical simulations, both on simple model geometries and a realistic white matter microstructure. This chapter validated the two asymptotic models derived in previous chapters, and provided some evaluation of the transition region between the two asymptotic regimes. The transition region was found to span roughly a decade in q-value in the white matter sample, and was found to occur at higher q-values in the extra-axonal space compared to the intra-axonal space. The validity of the correction method for time-varying q-vectors proposed in Chapter 6 when being extended to the transition region has also been evaluated.

8.2 Future research directions

There are some avenues for future research which are opened up by the work of this thesis. For the asymptotic analysis, there are some modifications to the models which could be considered to account for different operating conditions of diffusion MRI. These will be briefly outlined in the following paragraphs.

Consideration of stronger advection

The multi-scale homogenisation model in Chapter 5 has assumed advection to be weak in the compartment, which is a reasonable assumption when perfusion is not significant. However, in media where advection is stronger, for instance in vascular tumours, the homogenisation model can be modified to consider higher Péclet numbers. By drawing comparison to the homogenisation of the diffusion-advection equation, for Péclet numbers up to $O(\varepsilon^{-1})$ it is anticipated that the effective diffusion model will transition to an effective advection-driven dispersion model instead.

This regime of stronger advection will likely be more relevant in heavily vascularised tissue, for example vascular tumours, than in the white matter of the brain where the volume fraction of blood is much lower. Multi-scale homogenisation approaches have been employed in these tissues for the purpose of modelling blood flow and drug delivery [102], so it will be of much interest to use a similar approach to investigate the behaviour of diffusion MRI, and whether diffusion MRI can be a feasible tool for quantifying vascularisation, and hence a biomarker of the pathological condition of these tumours.

Effects of ultra-high static field

In Chapter 5, the local spatial variation of the q-vector has been considered to be of comparable magnitude to the macro-scale applied q-vector, based on the hardware used in typical diffusion MRI experiments. However, there is a trend of using significantly higher static fields for MRI experiments, with 7 T systems already in use for human whole-body

imaging and even higher fields available for animal systems. As shown in the analysis of Chapter 4, this would increase the magnitude of local q-vector variations. The effects of this can be investigated by modifying the asymptotic formulations correspondingly.

The modelling of this regime may have important implications for future development of diffusion MRI. Ultra-high field systems have the benefit of higher signal-to-noise ratio (SNR), which therefore facilitate more complex sensitisation sequences whilst keeping acceptable signal. 7 T MRI systems has already found its way into diffusion MRI usages [45], with the increased off-resonance effects at macro-scale tissue interfaces being observed and corrected for using distortion maps [96]. The modelling of off-resonance at the micro-scale, which is beyond the spatial resolution of distortion maps, will therefore be an interesting topic for further research.

Extending to other boundary conditions

So far in this thesis, only no-flux boundary conditions were considered for the heterogeneous molecular transport. However, in some pathological conditions such as stroke, there has been suggestions that water transport between extra-cellular and intra-cellular spaces is increased. Such effects may be formulated as leaky boundary conditions, and the behaviour of this coupling can be investigated in both short-time and long-time asymptotic models.

The modelling of leaky boundary conditions may offer some further insight into the pathological pathways of stroke, especially immediately after onset when subtle changes to the microstructure happen, before the bulk changes to relaxation parameters become visible. It is therefore important to understand how water might move between vascularisation, the extra-cellular space, and within brain cells, and the use of diffusion MRI to sensitise leaky boundaries may offer some progress.

8.3 Final remarks

The work in this thesis has demonstrated the value of multi-scale modelling in understanding and interpreting diffusion MRI. Coming from a theoretical modelling perspective, this thesis complements other research work coming from data-driven approaches. Signal models are necessary to simplify the inverse problem of interpreting imaging data, and the systematic asymptotic analysis presented in the thesis provides the theoretical basis for them. It is hoped that these will inform future diffusion MRI developments, both in terms of protocol design and data interpretation, and eventually leading to advancements in biomedical research and improvements in clinical care.

Appendix A

Automated mesh generation from segmented microscopy images

In this appendix, the automated process of generating meshes of porous domains from a segmented microscopy image is presented. The transmission electron microscopy image of a mouse white matter sample, as well as its segmentation, are obtained from Xu et al. [121]. The image data were collected at an isotropic resolution of 7.1nm $\times$7.1nm in a 4000 $\times$4000 matrix, with each pixel labelled as one of the three compartments, namely the extra-axonal space, the myelin sheath, and the intra-axonal space. A pipeline is therefore created to use the segmentation map information to create a finite element mesh representing the three domains.

As described in Section 5.5, 100 smaller 1000 $\times$1000 samples were used for the evaluation of the effective diffusion tensors. Therefore, 100 locations are picked at random, with the x and y co-ordinates of the top-left corner sampled from a uniform distribution between 1 and 3000, to crop out the 1000 $\times$1000 segmentation maps. In order to create a periodic domain which is representative of the microstructure information in these samples, they are augmented by appending their mirror images along both directions, as shown in Figure A.1. The details of the definition of the contours in the periodic domain, as well as the generation of periodic meshes based on the contours, is described below.

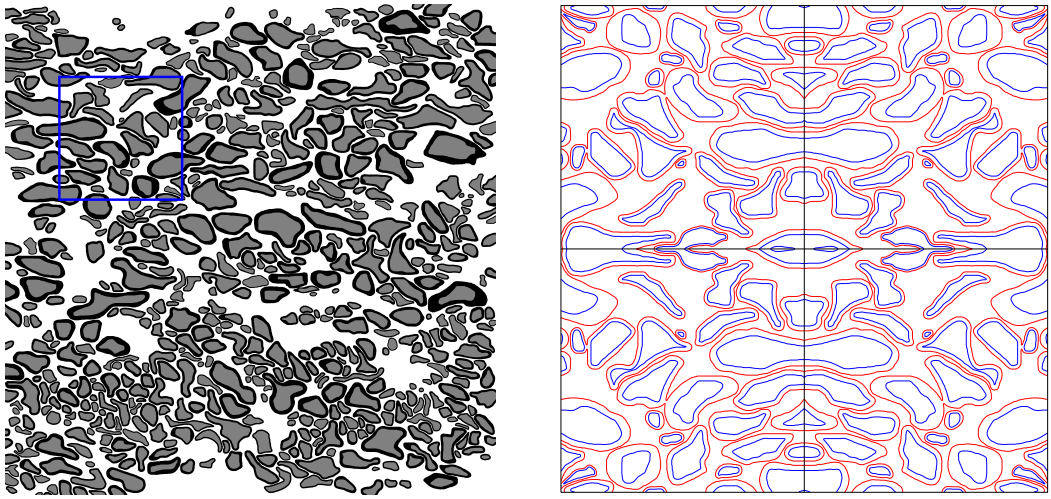

(a) Location of sample in TEM image (b) Periodic domain

Figure A.1: Realistic microstructure geometry based on transmission electron microscopy data. A 7.1 μm $\times$7.1 μm (1000 px $\times$1000 px) sample has been cropped from the microscopy image (a), and then augmented to form the periodic domain (b).

Since the segmentation map consists only of labels of individual pixels, this means the edges between the compartments are pixelated and ragged, rather than a smooth boundary which would be expected in reality. Therefore, the first step of processing is the smoothing of the edges, and then a threshold can be applied to define contours which are smoother. By labelling the intra-axonal space as 1 and everywhere else as 0, a threshold of 0.75 was used to define the interior boundaries, which are shown in blue in Figure A.1. Similarly, a threshold of 0.75 between 0 as extra-axonal space and 1 as everywhere else was used to define the exterior boundaries. The threshold process was implemented with the `contour` function in MATLAB (R2022a, MATLAB Inc.), with the output being a set of points on the edges between pixels, where the interpolated signal values cross the threshold value. The contours are therefore the curves which join these points. The reason of the choice of 0.75 is because of the segmentation appears to favour an interior label rather than an exterior label for pixels along the boundary. Due to the closely packed nature of the cells, if a 0.5 threshold was used, many cells would end up merging together, leading to unrealistic, disconnected extra-axonal space geometries.

With the set of points defining the contours defined, they can be used to define the geometry of the meshing domain. This is implemented using the OpenCASCADE geome-

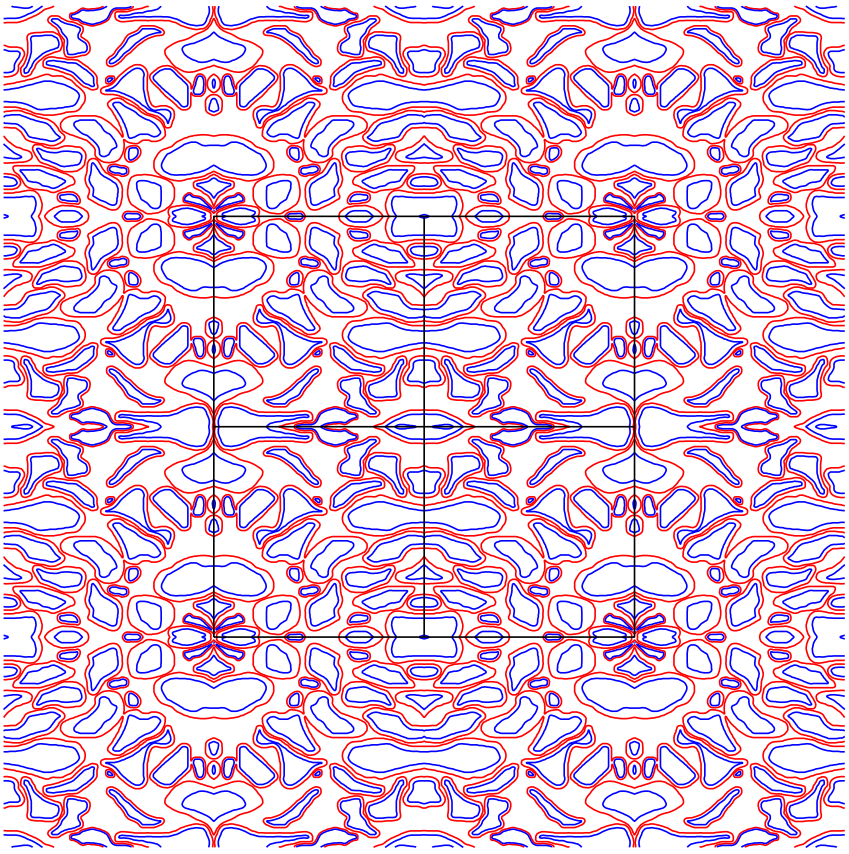

Figure A.2: Drawing contours on a 4000 ×4000 domain, so the cubic spline loops define intra-axonal and myelin sheath surface entities. The intersection of these entities with the 2000 ×2000 square then creates the required entities of the periodic domain.

try backend of the open-source `gmsh` software [41]. For each contour curve, a cubic spline is used to join together the points defined by the threshold process. These curves, along with the boundary edges of the 2000px ×2000px domain, form the bounding curves for the surface entities in the mesh.

To simplify the geometry definition process, as well as to ensure that the straight edge segments either side of the 2000px ×2000px domain match each other to enforce the periodicity condition, instead of drawing the boundary contours in a 2000px ×2000px domain, they are drawn on a 4000px ×4000px domain instead, again by tessellation of mirror images. The intra-axonal and myelin sheath entities can then be easily defined by using the cubic spline loops as bounding curves. The intersection of these entities with the 2000px ×2000px square thus define the required entities for the periodic domain with periodicity guaranteed. An example of this process is shown in Figure A.2.

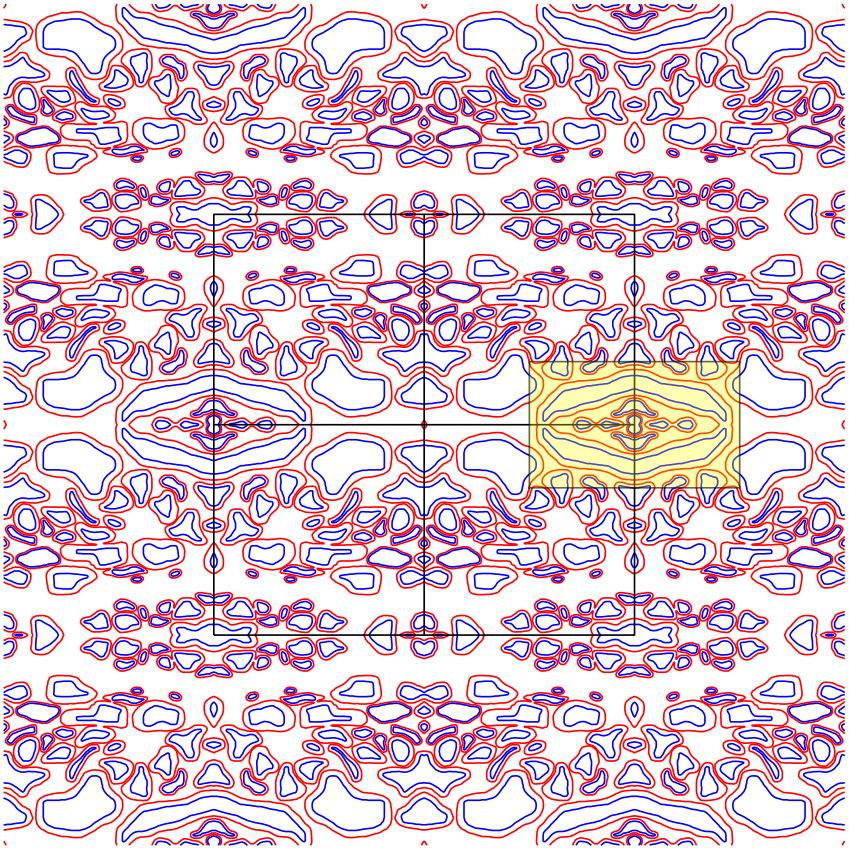

Figure A.3: Problematic cellular topology resulting from domain augmentation with mirror images of samples, as shown in shaded region. The domain augmentation is necessary to enforce periodicity conditions, however it may result in geometries where some cells are completely contained in another cell, as shown in this figure. In theses cases, the samples are discarded.

One problem of the augmentation of the domain using mirror images is that sometimes the straight edges cropping the domain can cut across a cell multiple times, so when the mirror images are appended unrealistic cell topologies can be formed. This is an unavoidable result, but this domain augmentation is necessary due to the need for periodicity to be enforced, as well as for the microstructure information to be kept as best as possible in the periodic meshes. If the geometry contains cells completely contained in another cell, then it is discarded. 29 out of the 100 random samples were therefore discarded, with one example shown in Figure A.3. Meanwhile, if there are instances when extra-axonal space is encapsulated by a cell, as shown in Figure A.4, due to their isolated nature, they are simply treated as part of the myelin sheath domain instead.

Finally, with the geometry defined, the meshes can be generated by triangulation. In

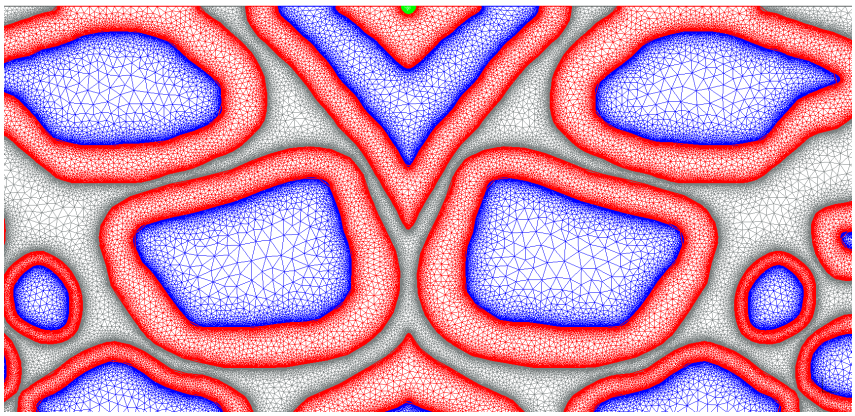

Figure A.4: In instances when the extra-axonal space is encapsulated by another cell, such as the example in this figure (small region at the top of figure highlighted in green), it is treated as part of the myelin sheath compartment instead.

order to resolve the anticipated boundary layer behaviour in the numerical simulations in Chapter 7, the mesh is refined to resolve the features near the boundaries at a pixel scale. Therefore, in the regions within one pixel length either side of a boundary curve, the local mesh size is set to 0.5px. The target mesh size varies linearly as the distance from a boundary increases, eventually reaching a background mesh size of 10px if a certain point is at least 30px away from any boundary curve. With these mesh controls in place, the meshes are generated using Delaunay triangulation. An example of the mesh control results is shown in Figure A.5. In the resulting meshes, the number of elements in the extra-axonal and intra-axonal spaces are in the order of 1 to 2 million. The statistics over the 79 meshes created, along with the associated area fractions, are presented in Figure A.6.

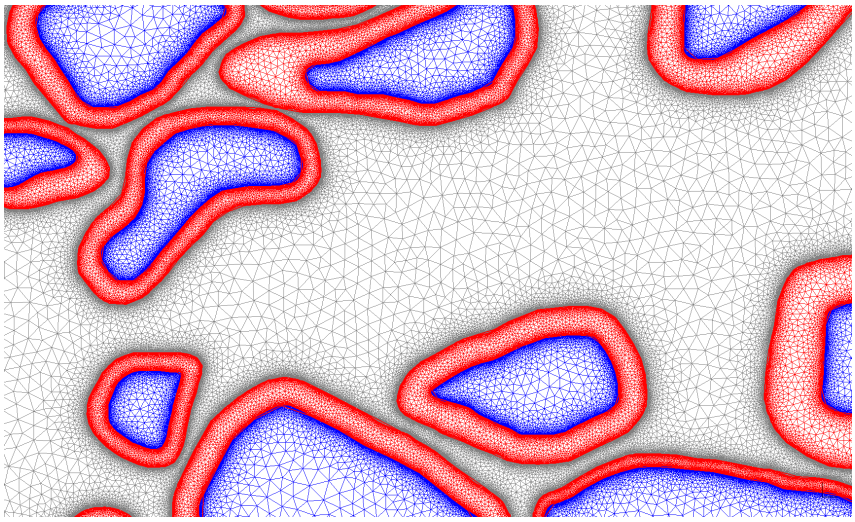

Figure A.5: Example of the mesh size control in mesh generation.

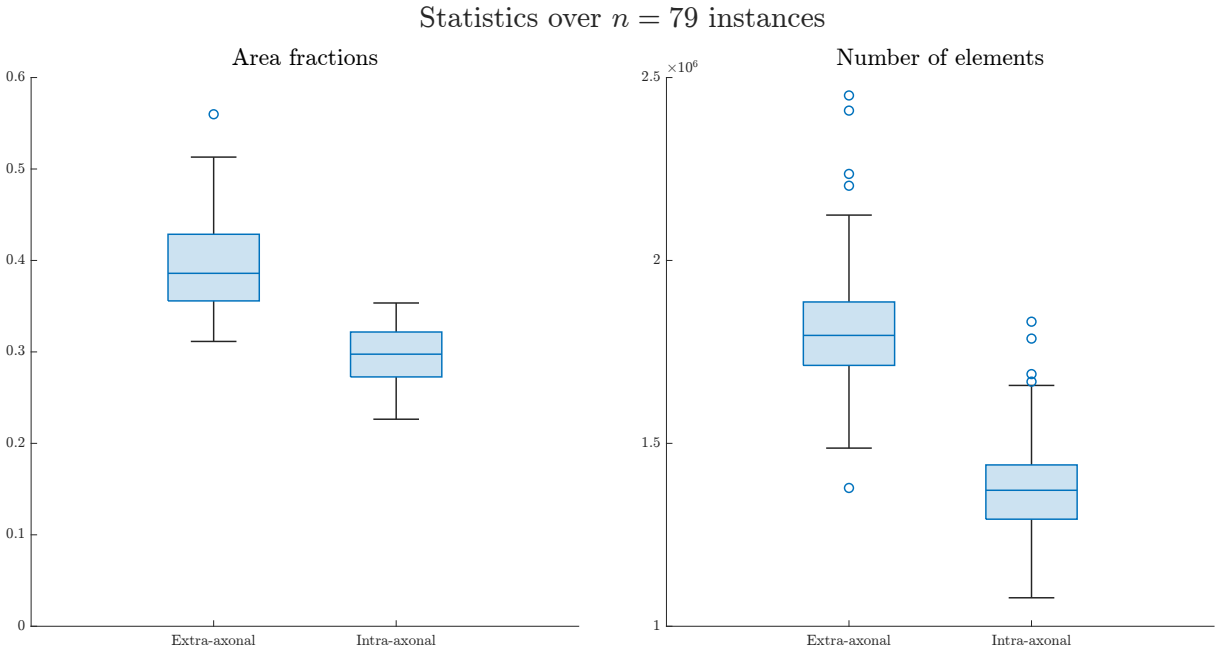

Figure A.6: Statistics of the area fractions and number of elements over the 79 meshes generated. The boxes show the range between the lower quartile to the upper quartile, whereas the whiskers show the range of data excluding the outliers, defined as data points more than 1.5 times the inter-quartile range away from the box.

Bibliography

[1] Abramowitz, M. and Stegun, I. A. "Repeated Integrals of the Error Function". In: *Handbook of mathematical functions : with formulas, graphs, and mathematical tables*. 10th printing, with corrections. 1972.

[2] Ackerman, J. J. and Neil, J. J. "Biophysics of Diffusion in Cells". In: *Diffusion MRI: Theory, Methods, and Applications*. Ed. by Jones, D. K. Oxford University Press, 2010.

[3] Allaire, G. "Homogenization and Two-Scale Convergence". In: *SIAM J. Math. Anal.* 23.6 (1992), pp. 1482–1518. DOI: 10.1137/0523084.

[4] Assaf, Y. and Cohen, Y. "Structural information in neuronal tissue as revealed by q-space diffusion NMR spectroscopy of metabolites in bovine optic nerve". In: *NMR Biomed.* (1999), p. 10.

[5] Assaf, Y. et al. "Axcaliber: A method for measuring axon diameter distribution from diffusion MRI". In: *Magn. Reson. Med.* 59.6 (2008), pp. 1347–1354. DOI: 10.1002/mrm.21577.

[6] Auriault, J.-L. "Heterogeneous periodic and random media. Are the equivalent macroscopic descriptions similar?" In: *Int. J. Eng. Sci.* 49.8 (2011), pp. 806–808. DOI: 10.1016/j.ijengsci.2011.01.005.

[7] Auriault, J.-L., Boutin, C., and Geindreau, C. *Homogenization of Coupled Phenomena in Heterogenous Media*. ISTE & Wiley, 2009.

[8] Bass, R. F. *Diffusions and Elliptic Operators*. Springer, 1998.

[9] Basser, P. J. "Relationships between diffusion tensor and q-space MRI". In: *Magn. Reson. Med.* 47.2 (2002), pp. 392–397. DOI: 10.1002/mrm.10052.

[10] Basser, P. J., Mattiello, J., and Le Bihan, D. "MR diffusion tensor spectroscopy and imaging". In: *Biophys. J.* 66.1 (1994), pp. 259–267. DOI: 10.1016/S0006-3495(94)80775-1.

[11] Basser, P. J. and Pierpaoli, C. "Microstructural and Physiological Features of Tissues Elucidated by Quantitative-Diffusion-Tensor MRI". In: *J. Magn. Reson. B* 111.3 (1996), pp. 209–219. DOI: https://doi.org/10.1006/jmrb.1996.0086.

[12] Basser, P. J. and Pierpaoli, C. "A simplified method to measure the diffusion tensor from seven MR images". In: *Magn. Reson. Med.* 39.6 (1998), pp. 928–934. DOI: 10.1002/mrm.1910390610.

[13] Bernstein, M. A., King, K. F., and Zhou, X. J. “Echo Train Pulse Sequences”. In: *Handbook of MRI Pulse Sequences*. Ed. by Bernstein, M. A., King, K. F., and Zhou, X. J. Elsevier, 2004, pp. 702–801. DOI: https://doi.org/10.1016/B978-012092861-3/50023-6.

[14] Bloch, F. “Nuclear Induction”. In: *Phys. Rev.* 70.7-8 (1946), pp. 460–474. DOI: 10.1103/PhysRev.70.460.

[15] Bloch, F., Hansen, W. W., and Packard, M. “The Nuclear Induction Experiment”. In: *Phys. Rev.* 70.7-8 (1946), pp. 474–485. DOI: 10.1103/PhysRev.70.474.

[16] Brenner, S. C. and Scott, L. R. *The Mathematical Theory of Finite Element Methods*. 3rd ed. Springer, 2008.

[17] Callaghan, P. T. “Pulsed-Gradient Spin-Echo NMR for Planar, Cylindrical, and Spherical Pores under Conditions of Wall Relaxation”. In: *J. Magn. Reson. A* 113.1 (1995), pp. 53–59. DOI: https://doi.org/10.1006/jmra.1995.1055.

[18] Callaghan, P. T. “A Simple Matrix Formalism for Spin Echo Analysis of Restricted Diffusion under Generalized Gradient Waveforms”. In: *J. Magn. Reson.* 129.1 (1997), pp. 74–84. DOI: 10.1006/jmre.1997.1233.

[19] Callaghan, P. T. “Physics of Diffusion”. In: *Diffusion MRI: Theory, Methods, and Applications*. Oxford University Press, 2010.

[20] Callaghan, P. T. “Flow and dispersion”. In: *Translational Dynamics and Magnetic Resonance: Principles of Pulsed Gradient Spin Echo NMR*. Oxford University Press, 2011.

[21] Callaghan, P. T. et al. “High-resolution q-space imaging in porous structures”. In: *J. Magn. Reson.* 90.1 (1990), pp. 177–182. DOI: 10.1016/0022-2364(90)90376-K.

[22] Cantor, D. and Jonas, J. “Automated measurement of self-diffusion coefficients by the spin-echo method”. In: *J. Magn. Reson.* 28.2 (1977), pp. 157–162. DOI: 10.1016/0022-2364(77)90141-X.

[23] Carr, H. Y. and Purcell, E. M. “Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments”. In: *Phys. Rev.* 94.3 (1954), pp. 630–638. DOI: 10.1103/PhysRev.94.630.

[24] Cassot, F. et al. “A Novel Three-Dimensional Computer-Assisted Method for a Quantitative Study of Microvascular Networks of the Human Cerebral Cortex”. In: *Microcirculation* 13.1 (2006), pp. 1–18. DOI: 10.1080/10739680500383407.

[25] Cheung, M. M. et al. “Does diffusion kurtosis imaging lead to better neural tissue characterization? A rodent brain maturation study”. In: *NeuroImage* 45.2 (2009), pp. 386–392. DOI: 10.1016/j.neuroimage.2008.12.018.

[26] Chizhik, V. I. et al. *Magnetic Resonance and Its Applications*. Springer, 2014. DOI: 10.1007/978-3-319-05299-1.

[27] Ciarlet, P. G. *The Finite Element Method for Elliptic Problems*. Society for Industrial and Applied Mathematics, 2002.

[28] Cioranescu, D., Damlamian, A., and Griso, G. “The Periodic Unfolding Method in Homogenization”. In: *SIAM J. Math. Anal.* 40.4 (2008), pp. 1585–1620. DOI: 10.1137/080713148.

[29] Clark, C., Barker, G., and Tofts, P. “An in Vivo Evaluation of the Effects of Local Magnetic Susceptibility-Induced Gradients on Water Diffusion Measurements in Human Brain”. In: *Journal of Magnetic Resonance* 141.1 (1999), pp. 52–61. DOI: 10.1006/jmre.1999.1872.

[30] Clark, C. A. and Le Bihan, D. “Water diffusion compartmentation and anisotropy at high b values in the human brain”. In: *Magn. Reson. Med.* 44.6 (2000), pp. 852–859. DOI: 10.1002/1522-2594(200012)44:6<852::AID-MRM5>3.0.CO;2-A.

[31] Cohen, Y. and Assaf, Y. “Extracting Geometric Properties of White Matter with q-Space Diffusion MRI (QSI)”. In: *Diffusion MRI: Theory, Methods, and Applications*. Ed. by Jones, D. K. Oxford University Press, 2010.

[32] Cory, D. G. and Garroway, A. N. “Measurement of translational displacement probabilities by NMR: An indicator of compartmentation”. In: *Magn. Reson. Med.* 14.3 (1990), pp. 435–444. DOI: 10.1002/mrm.1910140303.

[33] Damlamian, A. “An Introduction to Periodic Homogenization”. In: *Multiscale Problems: Theory, Numerical Approximation and Applications*. Higher Education Press & World Scientific, 2011, pp. 1–27.

[34] Duong, T. Q. et al. “Evaluation of extra- and intracellular apparent diffusion in normal and globally ischemic rat brain via 19F NMR”. In: *Magn. Reson. Med.* 40.1 (1998), pp. 1–13. DOI: 10.1002/mrm.1910400102.

[35] Einstein, A. “Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen”. In: *Ann. Phys.* 17 (1905), pp. 549–560.

[36] El-Bouri, W. K. and Payne, S. J. “Multi-scale homogenization of blood flow in 3-dimensional human cerebral microvascular networks”. In: *J. Theor. Biol* 380 (2015), pp. 40–47. DOI: 10.1016/j.jtbi.2015.05.011.

[37] Ferizi, U. et al. “White matter compartment models for in vivo diffusion MRI at 300 mT/m”. In: *NeuroImage* 118 (2015), pp. 468–483. DOI: 10.1016/j.neuroimage.2015.06.027.

[38] Fick, R. H. J. “Advanced dMRI Signal Modeling for Tissue Microstructure Characterization”. PhD thesis. Universite Cote d’Azur, 2017.

[39] Fick, R. H. et al. “Non-parametric graphnet-regularized representation of dMRI in space and time”. In: *Medical Image Analysis* 43 (2018), pp. 37–53. DOI: 10.1016/j.media.2017.09.002.

[40] Frøhlich, A. F., Østergaard, L., and Kiselev, V. G. “Effect of impermeable boundaries on diffusion-attenuated MR signal”. In: *Journal of Magnetic Resonance* 179.2 (2006), pp. 223–233. DOI: 10.1016/j.jmr.2005.12.005.

[41] Geuzaine, C. and Remacle, J.-F. “Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities”. In: *Int. J. Numer. Methods Eng.* 79.11 (2009), pp. 1309–1331. DOI: 10.1002/nme.2579.

[42] Grebenkov, D. S. “Diffusion MRI/NMR at high gradients: Challenges and perspectives”. In: *Microporous and Mesoporous Materials* 269 (2018), pp. 79–82. DOI: 10.1016/j.micromeso.2017.02.002.

[43] Haddar, H., Li, J.-R., and Schiavi, S. “A Macroscopic Model for the Diffusion MRI Signal Accounting for Time-Dependent Diffusivity”. In: *SIAM J. Appl. Math.* 76.3 (2016), pp. 930–949. DOI: 10.1137/15M1019398.

[44] Hahn, E. L. “Spin Echoes”. In: *Phys. Rev.* 80.4 (1950), pp. 580–594. DOI: 10.1103/PhysRev.80.580.

[45] Heidemann, R. M. et al. “Diffusion imaging in humans at 7T using readout-segmented EPI and GRAPPA”. In: *Magnetic Resonance in Medicine* 64.1 (2010), pp. 9–14. DOI: https://doi.org/10.1002/mrm.22480.

[46] Howard, A. F. et al. *The BigMac dataset: an open resource combining multi-contrast MRI and microscopy in the macaque brain [preprint].* 2022. DOI: 10.1101/2022.09.08.506363.

[47] Huang, S. Y. et al. “The impact of gradient strength on in vivo diffusion MRI estimates of axon diameter”. In: *NeuroImage* 106 (2015), pp. 464–472. DOI: 10.1016/j.neuroimage.2014.12.008.

[48] Jaeschke, S. H. F. “Motion Estimation Using Radiofrequency Scattering of a Parallel Transmit Coil for Cardiac Magnetic Resonance Imaging at 7T”. DPhil Thesis. University of Oxford, 2019.

[49] Jenkinson, M., Wilson, J. L., and Jezzard, P. “Perturbation method for magnetic field calculations of nonconductive objects”. In: *Magn. Reson. Med.* 52.3 (2004), pp. 471–477. DOI: 10.1002/mrm.20194.

[50] Jensen, J. H. and Chandra, R. “NMR relaxation in tissues with weak magnetic inhomogeneities”. In: *Magn. Reson. Med.* 44.1 (2000), pp. 144–156.

[51] Jensen, J. H. et al. “Diffusional kurtosis imaging: The quantification of non-gaussian water diffusion by means of magnetic resonance imaging”. In: *Magn. Reson. Med.* 53.6 (2005), pp. 1432–1440. DOI: 10.1002/mrm.20508.

[52] Jezzard, P. and Clare, S. “Principles of nuclear magnetic resonance and MRI”. In: *Functional MRI: an introduction to methods*. Ed. by Jezzard, P., Matthews, P. M., and Smith, S. M. Oxford University Press, 2001.

[53] Kärger, J. and Heink, W. “The propagator representation of molecular transport in microporous crystallites”. In: *J. Magn. Reson.* 51.1 (1983), pp. 1–7. DOI: 10.1016/0022-2364(83)90094-X.

[54] King, M. D. et al. “q-Space imaging of the brain”. In: *Magn. Reson. Med.* 32.6 (1994), pp. 707–713. DOI: 10.1002/mrm.1910320605.

[55] Kiselev, V. G. and Posse, S. “Analytical Theory of Susceptibility Induced NMR Signal Dephasing in a Cerebrovascular Network”. In: *Phys. Rev. Lett.* 81.25 (1998), pp. 5696–5699. DOI: 10.1103/PhysRevLett.81.5696.

[56] Kiselev, V. G. “The Cumulant Expansion: An Overarching Mathematical Framework For Understanding Diffusion NMR”. In: *Diffusion MRI: Theory, Methods, and Applications*. Ed. by Jones, D. K. Oxford University Press, 2010.

[57] Lam, W. W., Jbabdi, S., and Miller, K. L. “A model for extra-axonal diffusion spectra with frequency-dependent restriction”. In: *Magn. Reson. Med.* 73.6 (2015), pp. 2306–2320. DOI: 10.1002/mrm.25363.

[58] Lazar, M. et al. “Estimation of the orientation distribution function from diffusional kurtosis imaging”. In: *Magn. Reson. Med.* 60.4 (2008), pp. 774–781. DOI: 10.1002/mrm.21725.

[59] Le Bihan, D. “Intravoxel incoherent motion imaging using steady-state free precession”. In: *Magn. Reson. Med.* 7.3 (1988), pp. 346–351. DOI: 10.1002/mrm.1910070312.

[60] Le Bihan, D. “Molecular diffusion, tissue microdynamics and microstructure”. In: *NMR Biomed.* 8.7 (1995), pp. 375–386. DOI: 10.1002/nbm.1940080711.

[61] Le Bihan, D. “Magnetic Resonance Diffusion Imaging: Introduction and Concepts”. In: *Diffusion MRI: Theory, Methods, and Applications*. Ed. by Jones, D. K. Oxford University Press, 2010.

[62] Le Bihan, D. “What can we see with IVIM MRI?” In: *NeuroImage* 187 (2019), pp. 56–67. DOI: 10.1016/j.neuroimage.2017.12.062.

[63] Le Bihan, D., Turner, R., and Macfall, J. R. “Effects of intravoxel incoherent motions (IVIM) in steady-state free precession (SSFP) imaging: application to molecular diffusion imaging”. In: *Magn. Reson. Med.* 10.3 (1989), pp. 324–337. DOI: 10.1002/mrm.1910100305.

[64] Le Bihan, D. and Turner, R. “The capillary network: a link between IVIM and classical perfusion”. In: *Magn. Reson. Med.* 27.1 (1992), pp. 171–178. DOI: 10.1002/mrm.1910270116.

[65] Le Bihan, D. et al. “MR imaging of intravoxel incoherent motions: application to diffusion and perfusion in neurologic disorders.” In: *Radiology* 161.2 (1986), pp. 401–407. DOI: 10.1148/radiology.161.2.3763909.

[66] Le Bihan, D. et al. “Separation of diffusion and perfusion in intravoxel incoherent motion MR imaging”. In: *Radiology* 168.2 (1988), pp. 497–505. DOI: 10.1148/radiology.168.2.3393671.

[67] Le Bihan, D. et al. “Diffusion tensor imaging: Concepts and applications”. In: *J. Magn. Reson. Imaging* 13.4 (2001), pp. 534–546. DOI: 10.1002/jmri.1076.

[68] Lee, H.-H., Fieremans, E., and Novikov, D. S. “What dominates the time dependence of diffusion transverse to axons: Intra- or extra-axonal water?” In: *NeuroImage* 182 (2018), pp. 500–510. DOI: 10.1016/j.neuroimage.2017.12.038.

[69] Levitt, M. H. “Relaxation”. In: *Spin Dynamics: Basics of Nuclear Magnetic Resonance*. 2nd ed. Wiley, 2007, pp. 543–595.

[70] Levitt, M. H. *Spin Dynamics: Basics of Nuclear Magnetic Resonance*. 2nd ed. Wiley, 2007.

[71] Logg, A., Mardal, K.-A., Wells, G. N., et al. *Automated Solution of Differential Equations by the Finite Element Method*. Ed. by Logg, A., Mardal, K.-A., and Wells, G. N. Springer, 2012. DOI: 10.1007/978-3-642-23099-8.

[72] Lövblad, K. O. et al. “Diffusion-weighted MRI in cortical ischaemia”. In: *Neuroradiology* 46.3 (2004), pp. 175–182. DOI: 10.1007/s00234-003-1133-7.

[73] Manz, B., Alexander, P., and Gladden, L. F. “Correlations between dispersion and structure in porous media probed by nuclear magnetic resonance”. In: *Phys. Fluids* 11.2 (1999), pp. 259–267. DOI: 10.1063/1.869876.

[74] Merboldt, K.-D., Hänicke, W., and Frahm, J. “Self-diffusion NMR imaging using stimulated echoes”. In: *J. Magn. Reson.* 64.3 (1985), pp. 479–486. DOI: 10.1016/0022-2364(85)90111-8.

[75] Merboldt, K.-D., Hänicke, W., and Frahm, J. “Diffusion imaging using stimulated echoes”. In: *Magn. Reson. Med.* 19.2 (1991), pp. 233–239. DOI: 10.1002/mrm.1910190208.

[76] Merboldt, K.-D. et al. “Rapid NMR imaging of molecular self-diffusion using a modified CE-FAST sequence”. In: *J. Magn. Reson.* 82.1 (1989), pp. 115–121. DOI: 10.1016/0022-2364(89)90170-4.

[77] Mitra, P. P. et al. “Diffusion propagator as a probe of the structure of porous media”. In: *Phys. Rev. Lett.* 68.24 (1992), pp. 3555–3558. DOI: 10.1103/PhysRevLett.68.3555.

[78] Mulkern, R. V. et al. “Multi-component apparent diffusion coefficients in human brain”. In: *NMR Biomed.* 12.1 (1999), pp. 51–62. DOI: 10.1002/(SICI)1099-1492(199902)12:1<51::AID-NBM546>3.0.CO;2-E.

[79] Nicholson, C. and Phillips, J. M. “Ion diffusion modified by tortuosity and volume fraction in the extracellular microenvironment of the rat cerebellum.” In: *J. Physiol.* 321.1 (1981), pp. 225–257. DOI: 10.1113/jphysiol.1981.sp013981.

[80] Niendorf, T. et al. “Biexponential diffusion attenuation in various states of brain tissue: Implications for diffusion-weighted imaging”. In: *Magn. Reson. Med.* 36.6 (1996), pp. 847–857. DOI: 10.1002/mrm.1910360607.

[81] Norris, D. G., Niendorf, T., and Leibfritz, D. “Healthy and infarcted brain tissues studied at short diffusion times: The origins of apparent restriction and the reduction in apparent diffusion coefficient”. In: *NMR Biomed.* 7.7 (1994), pp. 304–310. DOI: 10.1002/nbm.1940070703.

[82] Novikov, D. S. et al. “Revealing mesoscopic structural universality with diffusion”. In: *Proc. Natl. Acad. Sci.* 111.14 (2014), pp. 5088–5093. DOI: 10.1073/pnas.1316944111.

[83] Novikov, D. and Kiselev, V. “Transverse NMR relaxation in magnetically heterogeneous media”. In: *J. Magn. Reson.* 195.1 (2008), pp. 33–39. DOI: 10.1016/j.jmr.2008.08.005.

[84] Novikov, D. S. and Kiselev, V. G. “Effective medium theory of a diffusion-weighted signal”. In: *NMR Biomed.* 23.7 (2010), pp. 682–697. DOI: 10.1002/nbm.1584.

[85] Novikov, D. S. and Kiselev, V. G. “Surface-to-volume ratio with oscillating gradients”. In: *Journal of Magnetic Resonance* 210.1 (2011), pp. 141–145. DOI: 10.1016/j.jmr.2011.02.011.

[86] Novikov, D. S., Kiselev, V. G., and Jespersen, S. N. “On modeling”. In: *Magn. Reson. Med* 79.6 (2018), pp. 3172–3193. DOI: 10.1002/mrm.27101.

[87] Novikov, D. S. et al. “Random walks with barriers”. In: *Nature Phys* 7.6 (2011), pp. 508–514. DOI: 10.1038/nphys1936.

[88] Ockendon, J. R. et al. *Applied Partial Differential Equations*. Rev. ed. Oxford University Press, 2003.

[89] Ogawa, S. et al. “Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model”. In: *Biophys. J.* 64.3 (1993), pp. 803–812. DOI: 10.1016/S0006-3495(93)81441-3.

[90] Okell, T. W. “Assessment of Collateral Blood Flow in the Brain using Magnetic Resonance Imaging”. DPhil Thesis. University of Oxford, 2011.

[91] Özarslan, E. and Basser, P. J. “MR diffusion–“diffraction” phenomenon in multi-pulse-field-gradient experiments”. In: *J. Magn. Reson.* 188.2 (2007), pp. 285–294. DOI: 10.1016/j.jmr.2007.08.002.

[92] Pfeuffer, J. et al. “Water signal attenuation in diffusion-weighted 1H NMR experiments during cerebral ischemia: influence of intracellular restrictions, extracellular tortuosity, and exchange”. In: *Magn. Reson. Med.* 16.9 (1998), pp. 1023–1032. DOI: 10.1016/S0730-725X(98)00107-6.

[93] Pierpaoli, C. and Basser, P. J. “Toward a quantitative assessment of diffusion anisotropy”. In: *Magn. Reson. Med.* 36.6 (1996), pp. 893–906. DOI: 10.1002/mrm.1910360612.

[94] Pierpaoli, C. et al. “Diffusion tensor MR imaging of the human brain.” In: *Radiology* 201.3 (1996), pp. 637–648. DOI: 10.1148/radiology.201.3.8939209.

[95] Purcell, E. M., Torrey, H. C., and Pound, R. V. “Resonance Absorption by Nuclear Magnetic Moments in a Solid”. In: *Phys. Rev.* 69.1-2 (1946), pp. 37–38. DOI: 10.1103/PhysRev.69.37.

[96] Rijssel, M. J. van et al. “Reducing distortions in echo-planar breast imaging at ultrahigh field with high-resolution off-resonance maps”. In: *Magn. Reson. Med.* 82.1 (2019), pp. 425–435. DOI: https://doi.org/10.1002/mrm.27701.

[97] Ryland, B. N. and Callaghan, P. T. “Spin Echo Analysis of Restricted Diffusion under Generalized Gradient Waveforms for Spherical Pores with Relaxivity and Interconnections”. In: *Isr. J. Chem.* 43.1-2 (2010), pp. 1–7. DOI: 10.1560/JF3Q-URL3-5U20-WHLY.

[98] Sato, K.-i. and Ueno, T. “Multi-dimensional diffusion and the Markov process on the boundary”. In: *J. Math. Kyoto Univ.* 4.3 (1965), pp. 529–605. DOI: 10.1215/kjm/1250524605.

[99] Seymour, J. D. and Callaghan, P. T. “Generalized approach to NMR analysis of flow and dispersion in porous media”. In: *AIChE J.* 43.8 (1997), pp. 2096–2111. DOI: 10.1002/aic.690430817.

[100] Shao, X. et al. “Mapping water exchange across the blood-brain barrier using 3D diffusion-prepared arterial spin labeled perfusion MRI”. In: *Magn. Reson. Med.* 81.5 (2019), pp. 3065–3079. DOI: 10.1002/mrm.27632.

[101] Shinmoto, H. et al. “Biexponential apparent diffusion coefficients in prostate cancer”. In: *Magn. Reson. Imaging* 27.3 (2009), pp. 355–359. DOI: 10.1016/j.mri.2008.07.008.

[102] Shipley, R. J. and Chapman, S. J. “Multiscale Modelling of Fluid and Drug Transport in Vascular Tumours”. In: *Bull. Math. Biol.* 72.6 (2010), pp. 1464–1491. DOI: 10.1007/s11538-010-9504-9.

[103] Silva, M. D. et al. “Separating changes in the intra- and extracellular water apparent diffusion coefficient following focal cerebral ischemia in the rat brain”. In: *Magn. Reson. Med.* 48.5 (2002), pp. 826–837. DOI: 10.1002/mrm.10296.

[104] Skare, S. T. and Bammer, R. “EPI-Based Pulse Sequences for Diffusion Tensor MRI”. In: *Diffusion MRI: Theory, Methods, and Applications*. Ed. by Jones, D. K. Oxford University Press, 2010.

[105] Skorokhod, A. V. “Stochastic Equations for Diffusion Processes in a Bounded Region”. Trans. by Greenleaf, N. In: *Theory Probab. Its Appl.* 6.3 (1961), pp. 264–274. DOI: 10.1137/1106035.

[106] Smith, A. F. et al. “Transmural Variation and Anisotropy of Microvascular Flow Conductivity in the Rat Myocardium”. In: *Ann Biomed Eng* 42.9 (2014), pp. 1966–1977. DOI: 10.1007/s10439-014-1028-2.

[107] Stejskal, E. O. “Use of Spin Echoes in a Pulsed Magnetic-Field Gradient to Study Anisotropic, Restricted Diffusion and Flow”. In: *J. Chem. Phys.* 43.10 (1965), pp. 3597–3603. DOI: 10.1063/1.1696526.

[108] Stejskal, E. O. and Tanner, J. E. “Spin Diffusion Measurements: Spin Echoes in the Presence of a Time-Dependent Field Gradient”. In: *J. Chem. Phys.* 42.1 (1965), pp. 288–292. DOI: 10.1063/1.1695690.

[109] Stepišnik, J. “Analysis of NMR self-diffusion measurements by a density matrix calculation”. In: *Physica B* 104.3 (1981), pp. 350–364. DOI: 10.1016/0378-4363(81)90182-0.

[110] Stilbs, P. “Fourier transform pulsed-gradient spin-echo studies of molecular diffusion”. In: *Prog. Nucl. Magn. Reson. Spectrosc.* 19.1 (1987), pp. 1–45. DOI: 10.1016/0079-6565(87)80007-9.

[111] Tanner, J. E. “Use of the Stimulated Echo in NMR Diffusion Studies”. In: *J. Chem. Phys.* 52.5 (1970), pp. 2523–2526. DOI: 10.1063/1.1673336.

[112] Tong, D. C. et al. “Relationship Between Apparent Diffusion Coefficient and Subsequent Hemorrhagic Transformation Following Acute Ischemic Stroke”. In: *Stroke* 31.10 (2000), pp. 2378–2384. DOI: 10.1161/01.STR.31.10.2378.

[113] Torrey, H. C. “Bloch Equations with Diffusion Terms”. In: *Phys. Rev.* 104.3 (1956), pp. 563–565. DOI: 10.1103/PhysRev.104.563.

[114] Turner, R. et al. “Echo-planar imaging of intravoxel incoherent motion.” In: *Radiology* 177.2 (1990), pp. 407–414. DOI: 10.1148/radiology.177.2.2217777.

[115] Van, A. T., Holdsworth, S. J., and Bammer, R. “In vivo investigation of restricted diffusion in the human brain with optimized oscillating diffusion gradient encoding: In Vivo Oscillating Gradient Diffusion in the Human Brain”. In: *Magn. Reson. Med.* 71.1 (2014), pp. 83–94. DOI: 10.1002/mrm.24632.

[116] Warach, S. et al. “Acute human stroke studied by whole brain echo planar diffusion-weighted magnetic resonance imaging”. In: *Ann Neurol.* 37.2 (1995), pp. 231–241. DOI: 10.1002/ana.410370214.

[117] Watanabe, S. “On stochastic differential equations for multi-dimensional diffusion processes with boundary conditions”. In: *J. Math. Kyoto Univ.* 11.1 (1971), pp. 169–180. DOI: 10.1215/kjm/1250523692.

[118] Wharton, S. and Bowtell, R. “Fiber orientation-dependent white matter contrast in gradient echo MRI”. In: *Proceedings of the National Academy of Sciences* 109.45 (2012), pp. 18559–18564. DOI: 10.1073/pnas.1211075109.

[119] Xu, J., Does, M. D., and Gore, J. C. “Dependence of temporal diffusion spectra on microstructural properties of biological tissues”. In: *Magn. Reson. Imaging* 29.3 (2011), pp. 380–390. DOI: 10.1016/j.mri.2010.10.002.

[120] Xu, J. et al. “Mapping mean axon diameter and axonal volume fraction by MRI using temporal diffusion spectroscopy”. In: *NeuroImage* 103 (2014), pp. 10–19. DOI: 10.1016/j.neuroimage.2014.09.006.

[121] Xu, T. et al. “The effect of realistic geometries on the susceptibility-weighted MR signal in white matter: Effect of Realistic Geometries on Susceptibility-Weighted MR Signal in White Matter”. In: *Magn. Reson. Med.* 79.1 (2018), pp. 489–500. DOI: 10.1002/mrm.26689.

[122] Yadav, N. N., Stait-Gardner, T., and Price, W. S. “Hardware Considerations for Diffusion MRI”. In: *Diffusion MRI: Theory, Methods, and Applications*. Ed. by Jones, D. K. Oxford University Press, 2010.