

Convergence and Parameter Choice for Monte-Carlo Simulations of Diffusion MRI

Matt G. Hall* and Daniel C. Alexander, *Member, IEEE*

Abstract—This paper describes a general and flexible Monte-Carlo simulation framework for diffusing spins that generates realistic synthetic data for diffusion magnetic resonance imaging. Similar systems in the literature consider only simple substrates and their authors do not consider convergence and parameter optimization. We show how to run Monte-Carlo simulations within complex irregular substrates. We compare the results of the Monte-Carlo simulation to an analytical model of restricted diffusion to assess precision and accuracy of the generated results. We obtain an optimal combination of spins and updates for a given run time by trading off number of updates in favor of number of spins such that precision and accuracy of synthesized data are both optimized. Further experiments demonstrate the system using a tissue environment that current analytic models cannot capture. This tissue model incorporates swelling, abutting, and deformation. Swelling-induced restriction in the extracellular space due to the effects of abutting cylinders leads to large departures from the predictions of the analytical model, which does not capture these effects. This swelling-induced restriction may be an important mechanism in explaining the changes in apparent diffusion constant observed in the aftermath of acute ischemic stroke.

Index Terms—Data synthesis, diffusion magnetic resonance imaging (MRI), diffusion tensor imaging, Monte-Carlo simulation, validation.

I. INTRODUCTION

DIFFUSION-WEIGHTED magnetic resonance imaging (MRI) measures the diffusive motion of water *in vivo* in the direction of an applied magnetic field gradient. Particle scattering through diffusion depends on the tissue microstructure and diffusion MRI thus supports inferences about the underlying microstructure. For example, diffusion tensor imaging (DTI) fits a symmetric, positive-definite tensor to six or more diffusion-weighted images on the assumption of Gaussian-distributed water-molecule displacements. The principal direction of the diffusion tensor provides an estimate of white-matter fibre orientations. Tractography techniques follow the principle direction point-to-point through the image to reconstruct macroscopic fibre trajectories and infer brain

connectivity. These techniques have emerged as powerful tools for probing both anatomy [1], [2] and pathology [3].

Diffusion MRI potentially provides more detailed information about tissues than merely the principal direction of white matter fibres since features like cell size, density and permeability all affect the dispersion of water in tissues. One approach is to fit a model to a set of diffusion-weighted measurements with parameters that describe how the microstructure affects diffusion. For example, the CHARMED model [4] of white matter is impermeable, nonabutting parallel cylinders with a gamma distribution of radii. Alexander [5] simplifies the model by assuming a single radius and cylindrically-symmetric extracellular diffusion. Sen and Basser [6] add thickness to cylinder walls (like myelin sheaths) and show that different parameter settings reflect the observations from brain tissue with various pathological conditions. These models potentially provide new biomarkers for diseases and new anatomical insights. However, significant questions remain as to how well fitted parameters reflect the actual tissue microstructure. This is a difficult question to answer, since ground-truth information from living tissue is hard to obtain.

Synthetic data provides a cheap and powerful tool to investigate the precision and accuracy of imaging techniques. With synthetic data, we can compare estimated parameters to the known parameters of the simulation and investigate the effects of noise and of breaking modeling assumptions in a controlled way. It is vital, however, that synthetic models are representative of *in vivo* systems.

Diffusion MR data synthesis falls into two broad categories: analytical models based on modeling the diffusion process in some geometry to obtain exact or approximate solutions to the diffusion equation, which then provide estimates of diffusion-weighted measurements, and numerical models which simulate the diffusion process either by numerical solution of the diffusion equation in a known environment (e.g., [7], [8]) or directly via a Monte-Carlo approach (e.g., [9], [10]). Although numerical solutions do not provide the same level of mathematical insight that analytic models do, numerical techniques are capable of investigating a much wider class of systems for which no analytical solutions exist and thus capture more biological complexity than analytical models.

Furthermore, Monte-Carlo simulations provide the ability to address quite subtle questions about which aspects of the measured dynamics have an impact on the measurement process. Within the Monte-Carlo framework, we are able to study not only diffusion in an arbitrary environment, but also different spin-barrier interaction mechanisms or step beyond the Brownian approximation and introduce spin-spin interactions. The extent to which either of these mechanisms effect the measurements is an open question that only Monte-Carlo modeling

Manuscript received November 26, 2008; revised February 09, 2009. First published March 04, 2009; current version published August 26, 2009. The work of M. G. Hall was supported by EPSRC under Grant EP/E064590/1. The work of D. C. Alexander was supported by EP/E007748. *Asterisk indicates corresponding author.*

*M. G. Hall is with the Centre for Medical Image Computing (CMIC), University College London, WC1E 6BT London, U.K. (e-mail: m.hall@cs.ucl.ac.uk).

D. C. Alexander is with the Centre for Medical Image Computing (CMIC), University College London, WC1E 6BT London, U.K..

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMI.2009.2015756

can address. For reasons of flexibility and generality, we pursue the Monte-Carlo modeling framework here. Although several authors have used the approach previously, as we discuss in Section II, existing systems use simple regular geometries for the simulation. Here, we describe and demonstrate a general and flexible system that is simple to adapt for arbitrary geometry and dynamics.

This paper performs a detailed optimization of simulation parameters to maximize accuracy and precision by comparison with an analytical model for a simple geometry. It then considers more complex irregular geometry. We provide algorithms for constructing such environments and show how to maintain integrity of the simulation within them. Specifically, we use the Monte-Carlo model to investigate tissue swelling which is a common effect in diseases such as stroke. We present a new model that includes abutting and deforming cylinders. Current analytical models cannot capture these kinds of effect. We observe that these mechanisms introduce significant differences between synthetic measurements from the Monte-Carlo and approximate analytical models. The experiment demonstrates the need for Monte-Carlo simulation and also provides a plausible explanation for well-known diffusion MRI signal changes that occur in the aftermath of stroke.

The rest of this paper is structured as follows. Section II contains background on diffusion-weighted MRI and the measurement process we emulate within the simulation. Section III describes the simulation procedure and its parameters. The analytical model is described in Sections IV. Section V describes two experiments. To find optimal settings within the simulation (Section V-A) we measure the standard deviation of the synthetic signal and the differences between the Monte-Carlo and analytical models while varying numbers of spins and number of updates. We then use the optimal settings to investigate the model of swelling tissue (Section V-B). All results are discussed in Section VI.

II. BACKGROUND

Diffusion-weighted MRI measures the local properties of water diffusion, often via a pulsed gradient spin-echo (PGSE) Stejskal–Tanner sequence. Fig. 1 illustrates a PGSE sequence and the timescales associated with it. An initial 90° pulse aligns spin phases perpendicular to the applied static field B_0 . During the first half of the pulse sequence up to the 180° and time τ , spin j will experience a phase shift

$$\phi_j^{(1)}(\mathbf{G}, \delta) = \gamma B_0 \tau + \gamma \int_{t_1}^{t_2} \mathbf{G} \cdot \mathbf{x}_j(t) dt. \quad (1)$$

The first term is the phase shift due to the static field B_0 for echo time 2τ and the second due to the gradient pulse switched on at time t_1 and off at time $t_2 = t_1 + \delta$ with gradient strength and orientation \mathbf{G} , γ is the gyromagnetic ratio for protons in water and $\mathbf{x}_j(t)$ is the displacement of spin j at time t . If we assume the gradient is constant during the pulse, we can take it out of the integral to give

$$\phi_j^{(1)}(\mathbf{G}, \delta) = \gamma B_0 \tau + \gamma \mathbf{G} \cdot \int_{t_1}^{t_2} \mathbf{x}_j(t) dt. \quad (2)$$

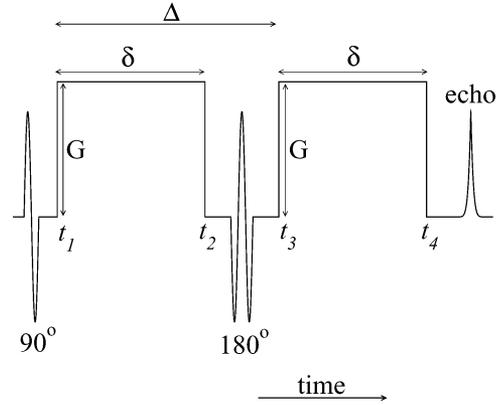


Fig. 1. Schematic representation of a PGSE MR pulse sequence. The first gradient block begins at time t_1 and ends at time $t_2 = t_1 + \delta$. The second gradient block begins at time $t_3 = t_1 + \Delta$ and ends at $t_4 = t_3 + \delta$.

A 180° pulse occurs after the end of the first gradient pulse. This flips the sign of the phase of all spins in the sample. At time $t_3 = t_1 + \Delta$ a second gradient pulse, identical to the first, starts. The second pulse is turned off at time $t_4 = t_3 + \delta$. The phase shift experienced by each spin in the second half of the pulse sequence is

$$\phi_j^{(2)}(\mathbf{G}, \Delta, \delta) = \gamma B_0 \tau + \gamma \mathbf{G} \cdot \int_{t_3}^{t_4} \mathbf{x}_j(t) dt. \quad (3)$$

The net phase shift experienced by a spin is the difference between the two phase shifts

$$\phi_j(\mathbf{G}, \Delta, \delta) = \phi_j^{(1)}(\mathbf{G}, \delta) - \phi_j^{(2)}(\mathbf{G}, \delta, \Delta). \quad (4)$$

For a stationary spin, the second pulse cancels out the first leaving no net effect on the phase. A moving spin, however, retains a residual phase offset that depends on its path during the sequence, i.e.,

$$\phi_j(\mathbf{G}, \Delta, \delta) = \gamma \mathbf{G} \cdot \left(\int_{t_1}^{t_2} \mathbf{x}_j(t) dt - \int_{t_3}^{t_4} \mathbf{x}_j(t') dt' \right). \quad (5)$$

The system contains an ensemble of spins, so the signal

$$S(\mathbf{G}, \Delta, \delta) = S(\mathbf{0}, \Delta, \delta) \int_{-\infty}^{+\infty} P(\phi) \exp(i\phi) d\phi \quad (6)$$

where P is the spin phase distribution and $S(\mathbf{0}, \Delta, \delta)$ is the “unweighted signal” acquired with zero diffusion-weighting gradients but the same echo time. The distribution of particle displacements p and the scan parameters \mathbf{G} , Δ and δ control the phase distribution P [11].

Considering only the real component of this signal gives

$$S(\mathbf{G}, \Delta, \delta) = S(\mathbf{0}, \Delta, \delta) \int_{-\infty}^{+\infty} P(\phi) \cos(\phi) d\phi. \quad (7)$$

This is the quantity of interest in most diffusion MRI analysis techniques.

A. Data Synthesis in Diffusion MRI

Validation and calibration techniques often make use of synthetic data. Two classes of technique generate synthetic diffusion MRI data. The first uses analytic models of p and the mea-

surement process to approximate the signal. The second simulates particle scattering within geometric models of material microstructure and emulates the measurement process.

1) *Analytic Models:* The simplest and most common model for generating synthetic data is the Gaussian. For example, Alexander and Barker [12] run simulations with synthetic data to determine optimal pulse-sequence settings for DT-MRI. They also synthesize data from a weighted sum of Gaussians to optimize acquisition for fitting mixture models. Tuch [13], Tournier *et al.* [14], and Alexander [15] (amongst others) use Gaussian-mixture models to test and compare multiple-fibre reconstruction algorithms.

The literature also contains analytic models for the scatter pattern within simple restricting geometries, such as spheres, cylinders and parallel planes [16]–[18], as well as the MR signal from particles scattering within these geometries. Various researchers construct simple geometric models from these primitives, so that p has analytic form. For example, Szafer *et al.* [19] construct a brain-tissue model consisting of nonabutting semi-permeable square cylinders. Stanisiz constructs a model of diffusion in various restricting intracellular geometries that utilizes a tortuosity approximation in the extracellular space [20]. Peled *et al.* [21] fit models of diffusion-weighted signal attenuation in parallel solid cylinders and in annular myelinated compartments to data acquired from bullfrog sciatic nerve. Other models, e.g., [22], use similar primitives to construct models for p in red blood cells.

2) *Simulations:* The models above improve on the Gaussian, but still grossly oversimplify real tissue microstructures. Real white-matter axons bend, abut and have a range of orientations, radii and wall thicknesses. Moreover, synthetic measurements from these models usually account only approximately for finite pulse widths. Monte-Carlo simulations of particle motion within geometric models avoid these limitations. The approach maintains a population of “spins,” each executing a random walk within a tissue model of arbitrary complexity. Models may contain semi-permeable barriers and material interfaces with which the spins interact. Material properties, such as diffusivity, T_1 and T_2 , may vary arbitrarily over space. In diffusion MRI simulations, each spin has a magnetization that evolves depending on position and magnetic-field gradient at each timestep. This approach also incorporates finite- δ effects and other pulse-sequence variables naturally.

Szafer *et al.* [19], Stanisiz *et al.* [23], and Duh *et al.* [24] use Monte-Carlo simulations to validate analytic models of diffusion in brain tissue. Regan and Kuchel [25] use Monte-Carlo simulations of diffusion within red-blood-cells to investigate the potential for measuring red-blood-cell-wall permeability. Liu *et al.* [10] generate synthetic data to test their multiple-fibre reconstruction algorithm using Monte-Carlo simulations in the intracellular compartment of various configurations of impermeable cylinders. Ford and Hackney [9] model the rat spinal-cord with close-packed cylinders and use Monte-Carlo simulations to predict diffusivity changes after trauma. In later work [7], [8], they construct 3-D tissue models by extruding a segmentation of a 2-D light-microscope image slice and obtain synthetic data by solving the diffusion equation within the model using a finite-difference approach.

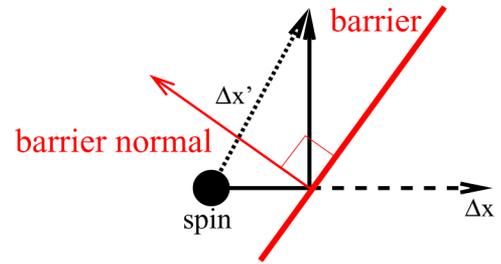


Fig. 2. If an update step takes a walker across a barrier it is elastically reflected. The length of the reflected step is such that the sum of the lengths of its components is the same as the length of the unreflected step.

Lipinski [26] uses a 2-D Monte-Carlo diffusion simulation to investigate the effects of cell swelling on the measured diffusivity. Here, a model of tissue is constructed from a light microscope image slice. Each pixel in the slice is classified as intracellular or extracellular using a simple thresholding procedure and after removing isolated intracellular voxels from the image, diffusion is simulated in the extracellular space only, and in 2-D. Although this is the most complex tissue model in the literature, the comparative crudity of the tissue classification and the lack of an intracellular compartment mean that the model is still greatly simplified compared to the underlying tissue.

The flexibility of Monte-Carlo modeling allows us to investigate the effect of mechanisms that are not included in analytical models by comparing Monte-Carlo data with analytical models. Moreover, because they are closer to the dynamics of diffusion than finite-element approaches, Monte-Carlo modeling can address questions that are difficult to capture even with a finite-element approach. For example, the question of the importance of the precise mechanism of interaction between spins and barriers can be readily addressed in a Monte-Carlo simulation, but would be difficult to incorporate correctly into other numerical approaches. This combination of flexibility and ability to capture very detailed dynamics motivates the present work.

III. MONTE-CARLO DIFFUSION SIMULATION

This section describes the Monte-Carlo model. We model the population of spins as random walkers in a 3-D environment and model the PGSE sequence to track phases over the trajectories of the spins and thus derive measurements. The environment contains a static configuration of barriers, which impede the motion of spins. We will refer to this configuration of barriers as the “substrate.” The simulation has a population of spins contained within a single voxel which contains the substrate. The simulation is used to synthesize a set of noise-free measurements from diffusing spins on a specified substrate to which noise can then be added.

A. Diffusion Simulation

Spins are initially randomly distributed across the substrate and then updated according to the following rule (see Fig. 2).

- 1) Generate a step vector $\Delta\mathbf{x}$. Steps are of constant length ℓ and random orientation.
- 2) Check if the step crosses a barrier.
 - If no barrier is crossed, the walker executes the step; walker position $\mathbf{x} \rightarrow \mathbf{x} + \Delta\mathbf{x}$.

- If the step would take the walker across a barrier, the barrier elastically reflects by the spin. The spin's position $\mathbf{x} \rightarrow \mathbf{x} + \Delta\mathbf{x}'$ where $\Delta\mathbf{x}'$ is the vector of the reflected step (see Fig. 2).
- Make the step to the barrier and repeat the above with the reflected portion of the step until no further barrier intersections are detected.

The check for multiple reflections in the update step last step is essential to avoid significant errors.

Since step length ℓ defines the effective diffusion constant for the population, we set it to

$$\ell = \sqrt{6 \frac{Dt_s}{T}} \quad (8)$$

where D is the diffusion constant, $t_s = 0.1$ s is the duration of the diffusion simulation and T is the number of timesteps in the simulation.

Previous simulations, e.g., [10], draw steps from a Gaussian distribution. Here, we use fixed-length steps with random orientation. This choice is a numerical optimization that is mathematically justified by the statistical properties of random walks. Although a Gaussian distribution of velocities occurs in a physical system, it is not necessary to generate steps with Gaussian distributed lengths to approximate them. The Central Limit Theorem [27] ensures that values obtained summing repeatedly over random values drawn from *any* distribution will converge on a Gaussian distribution. Since the positions of spins in our model are obtained by summing over steps in a random walk, we can approximate the positions expected from summing randomly oriented steps with Gaussian distributed lengths by summing instead over randomly oriented steps with uniform length and still observe the Gaussian distributed displacements that we would expect in the absence of restricting geometry. Fixed-length steps are simpler to generate than those with Gaussian distributed lengths, leading to an improvement in speed of execution. This also reduces fluctuations in the mean-squared displacement of spins as a function of time, improving convergence in the model.

The diffusion simulation environment is periodic. Spins are free to diffuse anywhere in space, subject to interactions with barriers. The particular configuration of barriers considered in a simulation (such as, for example, a small number of cylinders defining a regular packing of parallel cylinders or a larger region containing randomly positioned cylinders) can be thought of as defining a unit cell. When a spin's position is updated, its position with respect to the barriers for testing intersections is obtained by taking its position in space modulo the unit cell size.

By considering all substrates as defining a periodic unit cell, spins are guaranteed to encounter an environment consisting of barriers to diffusion no matter how far they diffuse. Spin positions are initially uniformly distributed across the unit cell, so that the allocation of a specific spin to either intra or extracellular space is statistical chance, with impermeable barriers ensuring that intracellular spins remain intracellular and vice-versa. The combination of impermeable barriers and the fact that the intracellular and extracellular diffusivities are the same in our model means that the equilibrium concentrations of spins in the two compartments are equal. This means that the

equilibrium distribution of spin positions is of equal concentration in the intra- and extracellular compartments with numbers of spins in each compartment proportional to each compartment's volume fraction, i.e., a uniform distribution across the unit-cell.

B. Synthetic Measurement

In the simulation, we approximate (7) by accumulating a phase shift for each spin. In each update, a spin receives a phase change $d\Phi$ from the simulated scan given the current time and its position. Following (2), this phase change is given by

$$d\Phi = a\gamma\mathbf{G}(t) \cdot \mathbf{x}(t)dt \quad (9)$$

where $\mathbf{G}(t)$ is the gradient vector, assumed constant over the timestep, and $\mathbf{x}(t)$ the spin's displacement at time t , dt is the duration of the timestep and $a = +1$ during the first gradient block and -1 during the second. All phase changes are zero outside the gradient blocks since $\mathbf{G}(t) = \mathbf{0}$.

To avoid numerical error, we take the origin for a spin to be its initial position. Because the gradient is linear, the choice of origin is arbitrary and calculating the phase shift in this way does not impact the results. Without this choice of origin, numerical errors due to repeated summation in voxels of realistic sizes dominate the contributions from particle net-displacements.

Spins accumulate a phase change for each gradient block during the course of the simulation, and synthetic signals S are generated by summing the contributions from all spins at echo time

$$S = \sum_j e^{i\Phi_j} \quad (10)$$

which is analogous to (7).

This scheme is readily generalizable to other diffusion-weighted pulse sequences such as twice-refocused [28] or STEAM [29].

IV. ANALYTICAL MODEL

In the next section, we validate and test the simulation using a simple substrate for which we can obtain a close analytic approximation to the signal. The model treats the diffusion environment as a two-compartment system: intracellular water is restricted within cylinders and extra cellular water diffuses around the cylinders, which are non-touching. We use the Gaussian phase distribution approximation [30], [16] to generate the intracellular signal. Cylinder walls are of zero thickness and permeability is set to zero. We assume free diffusion parallel to cylinders. Diffusion in the extracellular compartment perpendicular to cylinders is modeled using a tortuosity approximation. The model uses Szafer *et al.*'s tortuosity approximation for regularly packed cylinders [19] which sets an effective diffusion constant

$$D_E = \frac{D}{\lambda^2} \quad (11)$$

where D is the free diffusivity and λ^2 is the tortuosity given by

$$\lambda^2 = \begin{cases} \left(1 + V_I^{\frac{3}{2}}\right), & \text{perpendicular to fibre direction} \\ 1, & \text{parallel to fibre direction} \end{cases} \quad (12)$$

TABLE I
EXPLANATION OF PARAMETERS USED IN SIMULATIONS. PARAMETERS ARE GROUPED BY THEIR SOURCE: SCAN PARAMETERS, BIOPHYSICAL PARAMETERS, AND INTERNAL SIMULATION SPECIFIC PARAMETERS

Parameter	Value	Meaning	Source
Δ	0.01-0.1s	gradient block interval	feasible scan sequence range
δ	0.005-0.08s	gradient block duration	feasible scan sequence range
G	0.01Tm^{-1} - 0.1Tm^{-1}	gradient strength	feasible scan sequence range
D	$2.0 \times 10^{-9}\text{m}^2\text{s}^{-1}$	diffusivity	biophysical measurement
r	$1.0 \times 10^{-6}\text{m}$	cylinder radius in regular packing experiments	fit to histological data [36]
R	$2.1 \times 10^{-6}\text{m}$	cylinder separation in regular packing experiment	fit to histological data [36]
θ	6.91×10^{-7}	scale parameter in radius gamma distribution	fit to histological data [36]
k	2.331	shape parameter in radius gamma distribution	fit to histological data [36]
N	(various)	number of spins in simulation	experiment set 1
T	(various)	number of updates in simulation	experiment set 1
ℓ	(various)	length of steps in random walk	d and T using Eq. 8
U	10^6 - 10^8	complexity measure	estimated from computational constraints

and V_I is the intracellular volume fraction. The total analytic synthetic signal is

$$\mathcal{S} = (1 - V_I)S_E + V_I S_I \quad (13)$$

where

$$S_E = \exp(-bD_E) \quad (14)$$

is the signal from the external compartment, S_I is the signal from inside the cylinders and $b = (\gamma\delta G)^2(\Delta - \delta)/(3)$.

V. EXPERIMENTS

This section outlines two sets of experiments. In the first set, we investigate the trade-off between number of spins and number of updates within a fixed runtime for simulations. The second experiment explores a more complex substrate which provides a model for the observed changes in the diffusion MRI signal in the aftermath of ischaemic stroke. Following ischaemic insult, the apparent diffusion constant (ADC) in the region of the lesion decreases [3]. We construct a geometric model of microstructural tissue changes during stroke and perform Monte Carlo simulations using the optimal parameter settings determined in the first set of experiments.

A. Simulation Parameter Optimization

We want to ensure the best possible accuracy and precision of data synthesized from Monte-Carlo simulation. The computational requirement to run a simulation depends on both the number of updates in the simulation and the number of spins simulated. This experiment investigates trading off number of spins and number of updates. We hypothesize that by trading off one against the other it is possible to improve precision (i.e., reduce variation in signals synthesized by a set of models that have different initial spin positions and trajectories but identical substrates and scan parameters) and accuracy.

To quantify accuracy we measure the difference between synthetic measurements from the Monte-Carlo model and an approximating analytical model. It is important to recognize that the analytical model is itself an approximation and should not be regarded as a gold-standard for comparison. However, the approximation is good enough to reveal parameter combinations for which the simulation behaves poorly (see [31] for a partial evaluation).

To generate data from simulation, it is necessary to choose various parameters. In addition to physical quantities such as the diffusion constant, substrate size, cylinder radii and position, and scan-specific quantities such as diffusion time and gradient strength and duration, simulation-specific quantities must also be chosen. Diffusivity and cylinder radii settings come from knowledge of the microstructure we are modeling and scan parameters from the pulse sequence we wish to emulate. Table I summarizes the parameter settings used in these experiments. The physical quantities are chosen to be representative of parameters reported in the literature for *in vivo* brain tissue.

The central loop of the simulation updates every spin's position in every update, meaning that the total number of operations is proportional to the product $U = NT$ where N is the number of spins and T is the number of updates. We use U as a measure of the complexity of the simulation. We hypothesize that varying T and N has a significant effect on the precision of results at fixed U and aim to find the optimal balance. In this set of experiments we perform simulations with different pairs of N and T for $U = 10^6$, $U = 10^7$, and $U = 10^8$ and seek the configuration that minimizes the standard deviation of signals synthesized from sets of simulations while maintaining approximate agreement with the analytical model.

Methods: Simulations are performed on hexagonally-packed parallel cylinder substrates. Cylinders have radius 10^{-6} m and separation (defined as the distance between the central axes of neighboring cylinders) 2.1×10^{-6} m. These parameters ensure that cylinders do not abut and lead to an intracellular volume fraction $V_I = 0.71$, which is roughly representative of human white matter tissue. Simulations with each combination of N and T are repeated 30 times with different initial spin positions. Synthetic signals are generated over the parameter ranges given in Table I with seven increments over the range of each variable. Data were synthesized for all combinations of scan parameters for which $\delta < \Delta$. In each case, N and T are varied in pairs with $N = 10$ initially, and thereafter is increased in powers of 10 until $T = (U)/(N) = 10$.

All simulations have a duration of 0.1 s regardless of pulse sequence. In this way we can unambiguously identify a duration dt with a particular value of T , namely

$$dt = \frac{0.1}{T}. \quad (15)$$

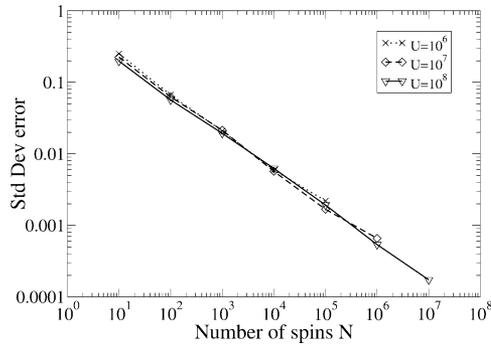


Fig. 3. Convergence of synthetic diffusion-weighted signals. Worst-case standard deviations are shown for sets of 30 simulations with different pairs of N and T for several constant $U = NT$. Results shown $U = 10^6$, $U = 10^7$, and $U = 10^8$ for sets of hexagonally packed cylinders.

Synthetic diffusion-weighted measurements are synthesized with gradient directions parallel and perpendicular to the principal cylinder axes. Standard deviations of normalized signals (i.e., $S(G, \Delta, \delta)/S(\mathbf{0}, \Delta, \delta)$) are obtained over 30 repeats with different random seeds and signals generated for one diffusion-weighted value in each direction and one unweighted measurement.

Unweighted measurements for all repeats always have the same value since they are unrelated to the diffusion process being simulated. In the absence of a diffusion-encoding gradient, the phase shift of all spins is zero, hence from (10) the magnitude of the unweighted signal is equal to the number of spins in the simulation.

Results: Standard deviations are calculated for each set of 30 repeats of the simulation. The worst case (i.e., the highest standard deviation) over all combinations of scan parameters provides a measure of precision in the Monte-Carlo synthesized measurements. Worst-case standard deviation of synthetic signals as a function of number of spins N for $U = 10^6$, $U = 10^7$, and $U = 10^8$ is shown in Fig. 3. Increasing N asymptotically reduces the standard deviations for all choices of U .

Accuracy of synthesized measurements is explored by considering the mean squared-difference between the measurements synthesized from the analytical and Monte-Carlo models. As in the precision study, the worst-case mean squared-difference between the mean Monte-Carlo synthesized signal and the analytical model over all scan parameter combinations and repeats is plotted in Fig. 4 as a function of N for $U = 10^6$, $U = 10^7$ and $U = 10^8$. Increasing N shows minima in each curve around $N = U/1000$, with an increase thereafter. Below the turning point in the curve, very little difference is observed between the curves for different U .

Conclusion: We observe that for fixed U increasing number of spins leads to an asymptotic decrease in worst-case standard deviation of signals generated between repeats. This is principally observed in the direction perpendicular to cylinder axes. Parallel to cylinder axes diffusion is free and thus much easier to simulate with a smaller number of spins.

Bias shows a more complicated behavior with increasing N . We observe that increasing number of spins leads to a decreased bias for all U , but this trend does not continue indefinitely. For each U we see a minimum and subsequent increase in bias once

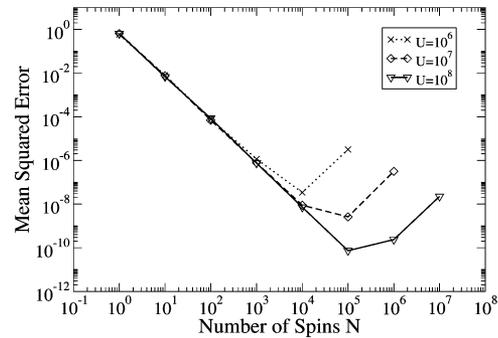


Fig. 4. Bias of synthetic diffusion-weighted signals. Worst-case mean squared-differences between signals from the Monte-Carlo and analytical models are shown for sets of 30 simulations with different pairs of N and T for several constant $U = NT$. Results shown $U = 10^6$, $U = 10^7$, and $U = 10^8$ for sets of hexagonally packed cylinders.

the number of timesteps T drops below 1000, corresponding to $dt = 10^{-4}$ s. Extremely long timesteps are a source of bias as the diffusion approximation breaks down. An extreme case would be a situation in which there was only a single timestep in the simulation: the contributions from the two gradient blocks would cancel each other perfectly and no diffusion-weighting would be observed. We are interested in the combination of U and N for which the standard deviation and the bias are both minimized. To minimize variation in the results of future experiments, we choose $U = 10^8$, $N = 10^5$, and $T = 10^3$.

B. Diffusion-Weighted Measurements in the Aftermath of Ischaemic Stroke

Diffusion-weighted images acquired in the immediate aftermath of ischaemic stroke exhibit a decrease in observed ADC in the region of the insult [3]. The decrease disappears after a few hours and is replaced by an overall increase in ADC which persists for a longer time [32]. The reason for the initial decrease is not known, but it is likely a result of cell swelling. Spins in the extracellular compartment are hindered by cell-walls which, in the absence of exchange, they must walk around. An increase in cell radius due to swelling leads to an increase in tortuosity and a concomitant decrease in ADC in this compartment. In addition to this mechanism, swelling also decreases the extracellular volume fraction. This decreases the contribution of extracellular diffusion to the overall diffusion-weighted signal in favor of more restricted diffusion in the intracellular compartment.

This increase in contribution to the signal from the intracellular space has been proposed as an explanation for the observed reduction in ADC after stroke. Here we investigate the idea that swelling induces restriction in the extracellular space due to a loss of percolation. Other authors have speculated on this idea, e.g., [22], [33]–[35].

We hypothesize that the increase in V_I and λ^2 in the simple model of Section IV alone cannot explain the level of change in the *in vivo* signal that accompanies cell swelling. However, loss of percolation in the extracellular space can produce such changes. Loss of percolation introduces additional restricted diffusion from cutoff regions in the extracellular space and hence a reduction in the diffusion attenuation in the normalized signal.

To investigate this idea, we construct a substrate to model diffusion in an environment where fibre radius increases due to

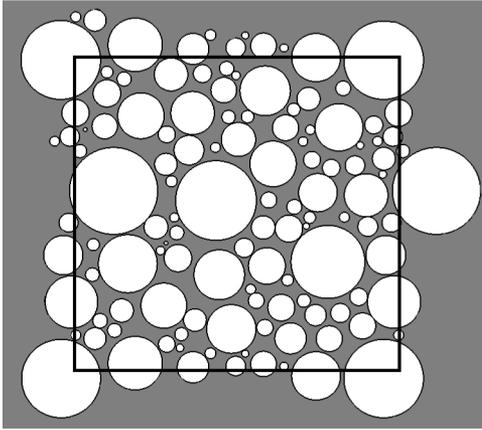


Fig. 5. Cross section of nonoverlapping parallel cylinders with gamma distributed radii. Additional cylinders are added to ensure periodic boundaries to the voxel. The square illustrates the boundaries of the voxel. Within the square, the pattern of cylinders tiles exactly in the plane.

tissue swelling. The initial, unswollen, substrate contains 100 cylinders with radii drawn from a gamma distribution

$$P(x; k, \theta) = \frac{x^{k-1} e^{-x/\theta}}{\Gamma(k) \theta^k} \quad (16)$$

with scale parameter $k = 6.91 \times 10^{-6}$ m and shape parameter $\theta = 2.331$. These values are obtained by fitting a gamma distribution to the histograms of axon radii in Aboitiz *et al.* [36] (Fig. 4, midbody). The cylinders are positioned within a cubic region of space 3.65×10^{-5} m on each side, giving an initial intracellular volume fraction $V_I = 0.65$. Initially, (i.e., prior to swelling) cylinders are guaranteed not to touch. In order to achieve the desired positioning, the algorithm for placing cylinders proceeds as follows:

- 1) Draw 100 radii from the gamma distribution.
- 2) Sort them into descending order.
- 3) Starting with the largest, for each cylinder:
 - a) choose a random position for the cylinder on the substrate;
 - b) if the cylinder position overlaps the edges of the voxel, create copies that overlap the opposing edge(s) to ensure substrate periodicity (see Fig. 5);
 - c) if the new cylinder or any of the copies touches or overlaps any other cylinder that has been already placed, discard and return to 3(a);
 - d) add new cylinder plus copies to the substrate.

Using this procedure, it is possible to obtain intracellular volume fractions similar to those observed in healthy white matter tissue, although some degree of experimentation is needed to find combinations of substrate size, gamma distribution parameters and number of cylinders in order to maximize it.

Tissue swelling is simulated by iteratively enlarging the cylinders on the substrate. Cylinder radii are enlarged by an amount proportional to their volume (in line with Boyle's Law). Since their positions are not altered, larger radii will inevitably lead to overlaps between cylinders. Overlapping cylinders are deformed along the chord of intersection in the plane of cross section so that they abut. The substrate is updated in each swelling iteration as follows.

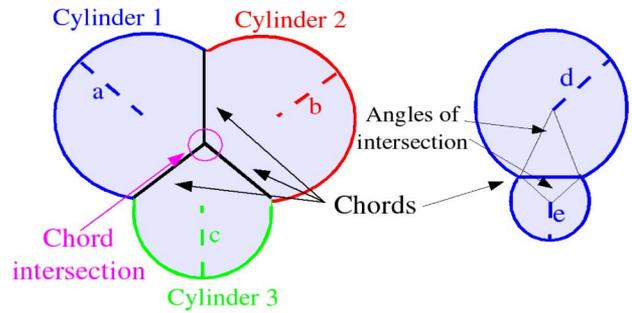


Fig. 6. Cylinder deformation scheme. Cylinders that would otherwise intersect each other are deformed along the chord of intersection (left). Chords that intersect each other are shortened accordingly. Cylinder swelling is limited by the angle of intersection (right). If either angle is greater than 90° both cylinders stop expanding.

- 1) Generate a radius increment Δr_j for each cylinder proportional to initial volume.
- 2) For each cylinder position j :
 - a) using the cylinder positions obtained in the initial placement phase, replace the cylinders with initial radius r_{0j} at each position with a new one of radius $r_{ij} = r_{0j} + i\Delta r_j$ where i is the swelling iteration number, $i = 0, \dots, n$ where n is the number of inflammation increments, or $r_{(i-1)j}$ if the cylinder has stopped expanding;
 - b) check if the new cylinder intersects any that are already placed. If yes, find the chord of intersection and use it as the deformed surface of the two cylinders as illustrated in Fig. 6;
 - c) if the angle of intersection (i.e., the angle subtended by the chord of intersection see Fig. 6) between two cylinders is greater than 90° in either one, stop both from expanding in subsequent iterations.

Fig. 7 shows the cross section of a substrate with cylinders positioned using this method, and the changes that occur as the cylinders swell. The extracellular space rapidly becomes divided by a continuous barrier of deformed circles. This is known in the statistical mechanics literature as the percolation threshold [37]. As swelling continues, the extracellular space becomes increasingly segmented into smaller and smaller regions in which diffusion is restricted rather than hindered.

Methods: Swelling of white matter tissue is simulated by incrementally increasing the radius of cylinders on the substrate. At each increment we perform 30 simulations with different cylinder and initial spin positions. This is performed for 15 volume increments. Volume fraction is estimated numerically by placing 10 000 points in a square grid in the plane of cross-section of the substrate and checking whether each of them is inside or outside of a cylinder.

This set of experiments uses a similar numerical model to the last for comparison. Since the cylinders are randomly rather than regularly packed we replace the tortuosity perpendicular to cylinder axes with

$$\lambda^2 = (1 - V_I)^{-1} \quad (17)$$

to model the randomly packed cylinders used here [19].

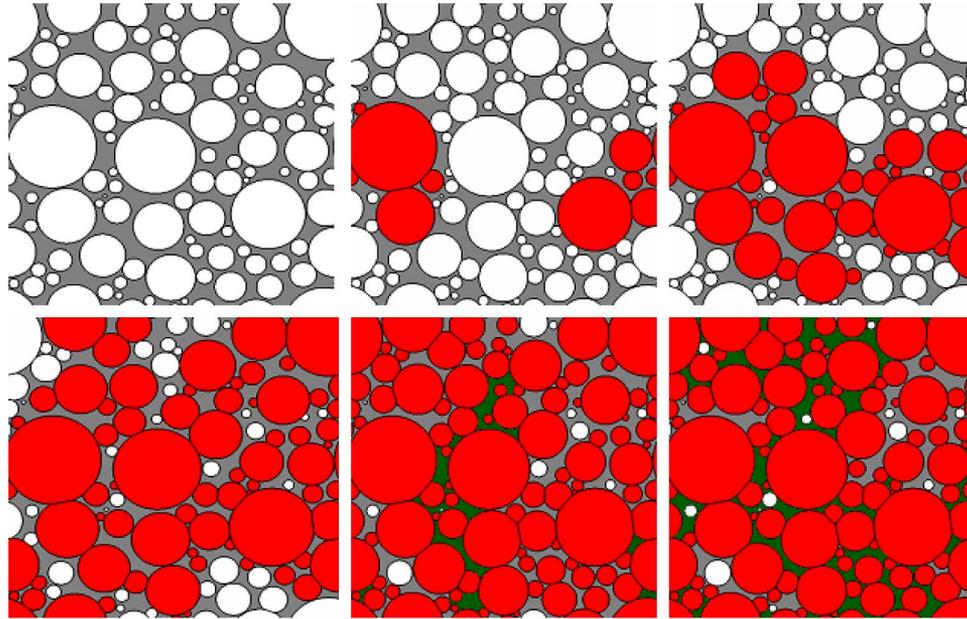


Fig. 7. Cross section of the substrate at various swelling iterations. The degree of swelling increases across rows (starting top left, iteration 3, 4, 5, 6, 9, 14). As cylinders swell, the number of cylinders deformed due to overlaps increases. These abutting cylinders are shown in red. At the third iteration the intersections form a complete barrier, segmenting the extracellular space. As swelling continues to increase, the extracellular space and some regions become completely cut off (green). Eventually green regions become very widespread.

Sets of 30 simulations with $U = 10^8$, $N = 10^5$, $T = 10^3$ were performed at each swelling iteration and the mean with error-bars showing standard deviation in each set are plotted as a function of V_I and compared to the analytical model.

Results: Fig. 8 shows the synthetic measurements from Monte-Carlo simulation and the analytical model as a function of intracellular volume fraction V_I . When cylinders are nonoverlapping ($V_I < 0.68$), the predictions of the analytical and Monte-Carlo models agree closely. As the degree of swelling increases, and the cylinders begin to touch, the predicted signal from the numerical model increases more quickly than that from the analytical model (i.e., ADC decreases more quickly). Fig. 8 also provides a plot of the number of overlapping cylinders. The divergence between the models, seen in both normalized diffusion signal and ADC, increases as the number of overlaps increases and we infer that it is the loss of percolation that is responsible for the divergence.

Conclusion: The main difference between the analytical and Monte-Carlo models is the inclusion of abutting cylinders and as a consequence, the formation of clusters of abutting cylinders and extracellular restriction. This effect turns out to have a large impact on the signal. We observe that the differences between the predictions of the two models increases as the number of intersections increases.

Two types of effect occur due to abutting cylinders. In the early stages of swelling, cylinders form small, island-like clusters. These smaller clusters have the effect of locally increasing tortuosity above the level that would be expected if clustering is ignored and as such causes a deviation between the two models. A more important mechanism, however, is the formation of system-spanning clusters that introduce extracellular restriction. Once the diffusion environment has been segmented by a system-spanning cluster, the tortuosity rises dramatically.

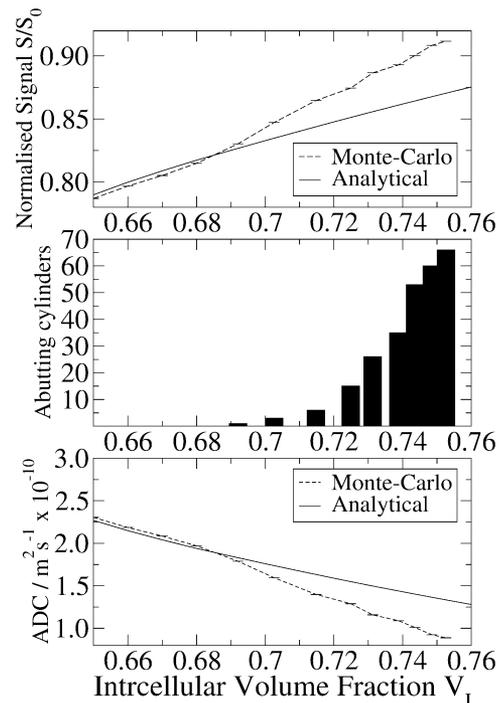


Fig. 8. Synthetic signals perpendicular to fibre directions from tissue undergoing swelling. Diffusion-weighted measurements from Monte-Carlo simulation (top) and measured ADC (bottom) along with the number of abutting cylinders (middle) in each iteration of the tissue model as a function of intracellular volume fraction.

When the extracellular space has been completely divided into restricted regions, the probability of a spin moving from one side of the diffusion environment to the other is reduced to zero and tortuosity over large length scales is effectively infinite.

In terms of cluster formation and the statistics of percolation, the system of 100 cylinders considered here is quite small. In an *in vivo* voxel we would expect a vastly greater number of axons, on the order of hundreds of thousands, but this increase in number of axons would mean that the trend observed here was smoother and less susceptible to fluctuations. The probability of forming a cluster of a given size at a given volume fraction is independent of system size [37], and larger systems have the effect of suppressing disorder-induced fluctuations, rather than changing the magnitude of the effect itself.

We conclude that it is crucial to use a Monte-Carlo model that is capable of capturing subtle effects when considering more complex biophysical effects on the diffusion-weighted signal, and that such models are capable of describing a broader class of phenomena than analytical models whilst still being computationally feasible to run. The algorithms we have presented for placing distributions of objects and maintaining periodicity generalize to a wide variety of similar substrates. The code is flexible enough to run simulations within any such substrate.

VI. DISCUSSION

We have described a Monte-Carlo model of diffusion in a restricting environment and investigated the simulation parameter settings that optimize accuracy and precision of data synthesized from the Monte-Carlo model, comparing it with the predictions of an analytical model. We then used the optimal parameter settings to investigate diffusion in an environment in which cylinders swell up, abutting and deforming due to the swelling. This abutting and deformation cannot be modeled analytically and we observe significant differences between the Monte-Carlo and Analytical models as abutting becomes more widespread. This illustrates that the Monte-Carlo model is capable of including effects that are too subtle to model analytically, and furthermore that such effects can be important when synthesizing diffusion-weighted signals. We have also seen that care should be taken when choosing simulation-specific parameters such as number of spins and number of timesteps. A poor choice of these can significantly reduce the accuracy and precision of results.

The results suggest that, broadly, more spins is more important than more updates. This reflects anecdotal wisdom that more spins, rather than more updates, leads to better statistics, but the literature contains no previous empirical validation. Clearly, the exact optimal values of number of spins and number of updates is a function of the complexity of the diffusion environment. It is not the purpose of these experiments to indicate definitive numbers, but to provide a rule of thumb that may be taken into consideration when performing similar simulations. Strictly, a version of the first experiment in this publication should be included as a necessary calibration step in any work that makes use of a Monte-Carlo simulation to generate diffusion-weighted measurements to assure optimal parameter choice but for similar kinds of experiment the ratio is unlikely to vary considerably.

The model of cell swelling we propose here could be refined in a number of ways. For example, it is likely that cells displace as well as deform as they swell and abut. These ideas are beyond the scope of this work, which aims mainly to demonstrate use

of our general simulation system for studying complex models of this type. It is straightforward to implement more complex models and test them within the existing simulator.

Implementation of a Monte-Carlo system is a nontrivial task. During the course of the development of the code, several issues have emerged as being important to a successful implementation of Monte-Carlo diffusion simulation for diffusion MRI. They are listed here in no particular order of importance.

- 1) *Fixed step lengths*—The central limit theorem ensures this results in Gaussian displacements in the absence of restricting barriers. Convergence is faster than for Gaussian distributed step lengths. Step directions are uniformly distributed on the sphere.
- 2) *Gradient block/dt mismatch*—Since, in general, the length of a gradient block is not an integer multiple of dt this mismatch can cause the effective gradient block duration to differ from the assumed value. This introduces bias into the synthesized signals. The effect of gradient duration mismatch is avoided here by including partial contributions from updates that overlap the ends of gradient blocks.
- 3) *Multiple barrier reflections*—In complex restricting geometries, it is often the case that a step will reflect from more than one barrier, bounce between barriers or reflect off the same barrier several times in a single step. Although it is tempting to approximate by ignoring multiple reflections, it is critical that they are handled correctly and explicitly or significant errors can occur.
- 4) *Substrate edges*—Care must be taken when considering what happens when a spin leaves the region of space occupied by the substrate. In this work we use periodic boundaries, making the restricting environment effectively infinite. In this case we use a global set of coordinates for spin position in all situations (including calculating phase shifts) except checking for barrier intersection. For barrier intersection a spin's position may be taken modulo the system size.

Where substrates are more complicated this issue becomes more important. For example, the swelling substrate requires cylinders that overlap substrate edges to be cloned so that they appear on both sides of the periodic tile. This further means that information about whether a cylinder is swelling or not must be propagated to all clones.

The simulation code is implemented in Java 1.5. Although a minor performance increase is possible by choosing a different language (such as C/C++), a Java implementation means that the code is easily portable and platform-independent. A typical run in the first set of experiments requires about 90 min and a typical run in the second set of experiments takes about 15 times longer (since one simulation per swelling iteration must be run) and hence require 15 times as much time to run (between 15 and 24 h depending on the machine). This means that considerable computational effort must be expended to obtain the several sets of results required to assess variance in the results presented (the experiments shown were run in batches on the UCL CMIC grid computing cluster).

Although the computational load for the current experiments is large, requiring a computing cluster to run in a feasible amount of time, by choosing simulation parameters N and T

correctly it is possible to reduce variation between runs to the point where an ensemble of simulations is not necessary. This means that the simulations as they currently stand can be run in a few hours on a single desktop PC.

The current implementation still leaves room for considerable optimization. A spatial optimization scheme such as a BSP-tree [38] or regular spatial grid [39], for example, would enable considerably more complicated substrates to be considered, including tissue geometries derived from microscope images, to add increased realism. The system itself is modular and so including different pulse sequences, substrates and spin dynamics is straightforward.

Permeability is not considered here, although the system does implement it. Stanisiz [20] reports that realistic values for permeability are very small and so the effect of realistic permeability on the current set of results would also likely be small.

Although the effects of a small permeability should not affect the results greatly, it is possible that it is important elsewhere. For example, in system with different diffusivities in different compartments permeability complicates simulation work, since the simulation must be allowed to settle into equilibrium concentrations in each compartment before being used to synthesize data.

The precise mechanism of interaction between barriers and spins can be chosen freely and the system implements changes in mechanism. In previous work (e.g., [6], [9]) reflecting barriers are frequently assumed, and we have made the same choice here.

Future applications are numerous. The question of to what extent different spin-barrier interaction mechanisms affect diffusion-weighted measurements [40] is currently an open question and would make an interesting avenue of research. Preliminary results suggest it is important. In addition to testing reconstruction mechanisms, the model may also be able to contribute to investigating the wider question of novel biomarkers for various neurological pathologies.

ACKNOWLEDGMENT

The simulation code used to generate these results is available online as part of the Camino diffusion MR toolkit (<http://www.camino.org.uk/>).

REFERENCES

- [1] S. Mori, B. J. Crain, and P. J. Chackoand van Zijl, "PCM three dimensional tracking of axonal projections in the brain by magnetic resonance imaging," *Ann. Neurol.*, vol. 45, pp. 265–269, 1999.
- [2] M. Catani, R. J. Howard, S. Pajevic, and D. K. Jones, "Virtual *in vivo* interactive dissection of white matter fasciculi in the human brain," *NeuroImage*, vol. 17, pp. 77–94, 2002.
- [3] J. Kucharczyk, J. Mintorovitch, H. Asgari, and M. Moseley, "Diffusion/perfusion MR imaging of acute cerebral ischemia," *Magn. Reson. Med.*, vol. 19, pp. 311–315, 1991.
- [4] Y. Assaf and P. J. Basser, "Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain," *NeuroImage*, vol. 30, pp. 1100–1111, 2005.
- [5] D. C. Alexander, "A general framework for experimental design and its application in measuring direct tissue-microstructure features," *Magn. Reson. Med.*, vol. 60, pp. 439–448, 2008.
- [6] P. N. Sen and P. J. Basser, "A model for diffusion in white matter," *Biophys. J.*, vol. 89, pp. 2927–2938, 2005.
- [7] C.-L. Chin, F. W. Wehrli, S. N. Hwang, M. Takahashi, and D. B. Hackney, "Biexponential diffusion in the rat spinal cord: Computer simulations based on anatomical images of axonal architecture," *Magn. Reson. Med.*, vol. 47, pp. 455–460, 2002.
- [8] S. N. Hwang, C.-L. Chin, F. W. Wehrli, and D. B. Hackney, "An image-based difference model for simulating restricted diffusion," *Magn. Reson. Med.*, vol. 50, pp. 373–382, 2003.
- [9] J. C. Ford and D. B. Hackney, "Numerical model for calculation of apparent diffusion coefficients (ADC) in permeable cylinders—Comparison with measured ADC in rat spinal cord white matter," *Magn. Reson. Med.*, vol. 37, pp. 387–394, 1997.
- [10] C. Liu, R. Bammer, M. Acar, and M. E. Moseley, "Characterizing non-Gaussian diffusion by using generalized diffusion tensors," *Magn. Reson. Med.*, vol. 51, pp. 924–937, 2004.
- [11] W. S. Price, "Pulsed-field gradient nuclear magnetic resonance as a tool for studying translational diffusion: Part I. Basic theory," *Concepts Magn. Reson.*, vol. 9, pp. 299–336, 1997.
- [12] D. C. Alexander and G. J. Barker, "Optimal imaging parameters for fibre-orientation estimation in diffusion MRI," *NeuroImage*, vol. 27, pp. 357–367, 2005.
- [13] D. S. Tuch, T. G. Reese, M. R. Wiegell, and V. J. Wedeen, "Diffusion MRI of complex neural architecture," *Neuron*, vol. 40, pp. 885–895, 2003.
- [14] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly, "Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution," *NeuroImage*, vol. 23, pp. 1176–1185, 2004.
- [15] D. C. Alexander, "Multiple fibre reconstructions for diffusion MRI," *Ann. NY Acad. Sci.*, vol. 1046, pp. 113–133, 2005.
- [16] C. H. Neuman, "Spin echo of spins diffusing in a bounded medium," *J. Chem. Phys.*, vol. 60, pp. 4508–4511, 1974.
- [17] P. T. Callaghan, "Pulsed-gradient spin-echo NMR for planar, cylindrical and spherical pores under conditions of wall relaxation," *Magn. Reson. Med.*, vol. 113, pp. 53–59, 1995.
- [18] O. Soderman and B. Jonsson, "Restricted diffusion in cylindrical geometry," *J. Magn. Reson. A*, vol. 117, pp. 94–97, 1995.
- [19] A. Szafer, J. Zhong, and J. C. Gore, "Theoretical model for water diffusion in tissues," *Magn. Reson. Med.*, vol. 33, pp. 697–712, 1995.
- [20] G. J. Stanisiz, "Diffusion MR in biological systems: Tissue compartments and exchange," *Israel J. Chem.*, vol. 43, pp. 33–44, 2003.
- [21] S. Peled, D. G. Cory, S. A. Raymond, D. A. Kirschner, and F. A. Jolesz, "Water diffusion, T_2 and compartmentation in frog sciatic nerve," *Magn. Reson. Med.*, vol. 42, pp. 911–918, 1999.
- [22] L. L. Latour, K. Svoboda, P. P. Mitra, and C. H. Sotak, "Time-dependent diffusion of water in a biological model system," *Proc. Nat. Acad. Sci.*, vol. 91, pp. 1229–1233, 1994.
- [23] G. J. Stanisiz, A. Szafer, G. A. Wright, and M. Henkelman, "An analytical model of restricted diffusion in bovine optic nerve," *Magn. Reson. Med.*, vol. 37, pp. 103–111, 1997.
- [24] A. Duh, A. Mohoric, and J. Stepisnik, "Computer simulation of the spin-echo spatial distribution in the case of restricted self-diffusion," *J. Magn. Reson.*, vol. 148, pp. 257–266, 2001.
- [25] D. G. Regan and P. W. Kuchel, "Simulations of NMR-detected diffusion in suspension of red cells; The signatures of q-space plots of various lattice arrangements," *Eur. Biophys. J.*, vol. 31, pp. 563–574, 2003.
- [26] H.-G. Lipinski, "Monte-Carlo simulation of extracellular diffusion in brain tissues," *Phys. Med. Biol.*, vol. 35, pp. 441–447, 1990.
- [27] O. Kallenberg, *Foundations of Modern Probability*. New York: Springer-Verlag, 1997.
- [28] T. J. Reese, O. Heid, R. M. Weisskoff, and V. J. Wedeen, "Reduction of eddy-current-induced distortion using a twice refocused spin-echo," *Magn. Reson. Med.*, vol. 49, pp. 177–182, 2003.
- [29] J. Frahm, K. D. Merboldt, W. Haenicke, and A. Haase, "Stimulated echo imaging," *J. Magn. Reson.*, vol. 64, no. 1, pp. 81–93, 1985.
- [30] J. S. Murday and R. M. Cotts, "Self-diffusion coefficient of liquid lithium," *J. Chem. Phys.*, vol. 48, pp. 4938–4945, 1984.
- [31] B. Balinov, Jonsson, P. Linse, and Soderman, "The NMR self-diffusion method applied to restricted diffusion. Simulation of echo attenuation from molecules in spheres and between planes," *J. Magn. Reson. Series A*, vol. 104, pp. 17–25, 1993.
- [32] C. H. Sotak, "The role of diffusion tensor imaging in the evaluation of ischemic brain injury—A review," *NMR Biomed.*, vol. 15, pp. 561–569, 2002.
- [33] D. G. Norris, T. Niendorf, and D. Leibfritz, "A theory of diffusion contrast in healthy infarcted tissue," in *Proc. SMRM 12th Annu. Meeting*, 1993, vol. 579.
- [34] E. Syková, J. Svoboda, J. Polákand, and A. Chvátal, "Extracellular volume fraction and diffusion characteristics during progressive ischemia and terminal anoxia in the spinal cord of the rat," *J. Cereb. Blood Flow*, vol. 14, pp. 301–311, 1994.

- [35] A. W. Anderson, J. Zhong, O. A. Petroff, A. Szafer, B. R. Ransom, J. W. Pritchard, and J. C. Gore, "Effects of osmotically driven cell volume changes on diffusion-weighted imaging of the rat optic nerve," *Magn. Reson. Med.*, vol. 35, pp. 162–167, 1996.
- [36] F. Aboitiz, A. B. Scheibel, R. S. Fisher, and E. Zaidel, "Fiber composition of the human corpus callosum," *Brain Res.*, vol. 598, pp. 143–153, 1992.
- [37] D. Ben-Avraham and S. Havlin, *Diffusion and Reactions in Fractals and Disordered Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2000.
- [38] , D. Kirk, Ed., *Graphics Gems III*. San Diego, CA: AP Professional, 1994.
- [39] M. Slater, Y. Chrysanthou, and A. Steed, *Computer Graphics and Virtual Environments: From Realism to Real Time*. Reading, MA: Addison Wesley, 2001.
- [40] D. LeBihan, S.-I. Urayama, T. Aso, T. Hanakawa, and H. Fukuyama, "Direct and fast detection of activation in the human brain with diffusion MRI," *Proc. Nat. Acad. Sci.*, vol. 103, no. 21, pp. 8263–8268, 2006.