Monte-Carlo simulation of diffusion MRI with realistic voxel sizes

Matt G Hall[1], Gemma Nedjati-Gilani[1], and Daniel C Alexander[1]
[1]*Centre for Medical Image Computing, University College London, London, United Kingdom*

Target Audience
Researchers developing novel microstructure imaging methods that use diffusion MRI, researchers using or developing biophysical simulations for diffusion MR data synthesis.

Purpose
Monte-Carlo simulation of spins executing random walks in environments that restrict their motion has established itself as an indispensable tool for the development of microstructural imaging methods [1][2][3]. Detailed models of the tissue environment experienced by diffusing spins provide a ground truth against which new methods can be calibrated and tested. A common choice as a model of white matter is randomly-packed parallel cylinders with radii drawn from a gamma distribution. In previous work computational constraints have meant that the size of the region simulated has typically been orders-of-magnitude smaller than a typical scan voxel, limiting realism and statistic power of tissue models. This work describes a technique that allows the size of this region to be greatly increased, allowing simulation regions to be of comparable size to voxels in a human diffusion MRI acquisition. The increased region size and concomitant increase in the number of samples from the underlying cylinder radius distribution leads to reduced variation in synthesized signals without increasing the number of spins simulated.

Methods
Diffusion is simulated using the approach of [1]. We sub-divide the region simulated into a regular 3D grid of "subvoxels" of a given size and construct a map of subvoxels to the set of cylinders intersecting each one, built dynamically during cylinder placement. The map means that a new object can be identified with the small set of subvoxels it intersects, and so only need be explicitly tested against other cylinders intersecting the same subvoxels, greatly reducing the number of intersection checks required. This map is used to place subsequent cylinders during substrate construction and then used again to aid spin intersection checking during dynamics simulation.

Map construction and cylinder placement
Cylinders are placed sequentially in descending order of size by picking a random candidate location which is accepted only if the new cylinder intersects no others. Cylinder radius and position allow the set of intersecting subvoxels to be calculated easily. The map then provides set of intersection candidates to check against. If no intersections are detected the new cylinder is added to the subvoxel map, otherwise a new location is generated and the check repeated.

Spin-cylinder intersection
The subvoxel-to-cylinder map generated during cylinder placement is used during spin dynamics simulation to test for intersection between spins and cylinders. Each step in a spin's random walk can be mapped to a set of subvoxels [4]. Again, intersection is only possible between the step and cylinders intersecting these subvoxels – the rest may be safely ignored.

Simulation parameters
We perform simulations in environments containing 1,000,000 non-overlapping parallel cylinders with radii drawn from a Gamma distribution with parameters from [5]. Diffusion is simulated using 100,000 spins and 5000 spin updates.

Experiments
We investigate optimal subvoxel grid size by comparing runtimes of simulations with different subvoxel sizes and gamma distribution parameters. We then investigate signal variance by synthesizing diffusion-weighted measurements using a PGSE sequence over a range of pulse sequence parameters ($\Delta \in (0.01\text{-}0.1)$s, $|\mathbf{G}| \in (0.01\text{-}1)\text{Tm}^{-1}$, δ=0.01s) with 30 repeats of each simulation. For comparison, we also perform the same experiment with 100 cylinders drawn from the same distribution packed into a smaller spatial region which preserves intracellular volume fraction.

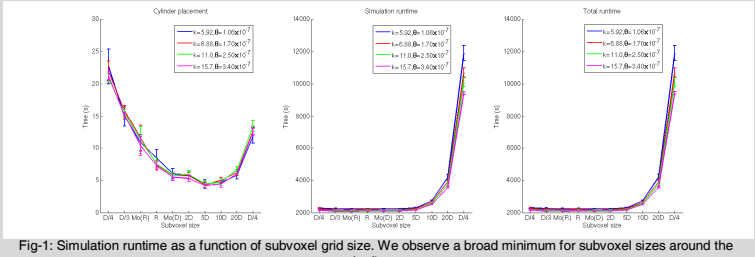

Fig-1: Simulation runtime as a function of subvoxel grid size. We observe a broad minimum for subvoxel sizes around the mode diameter.

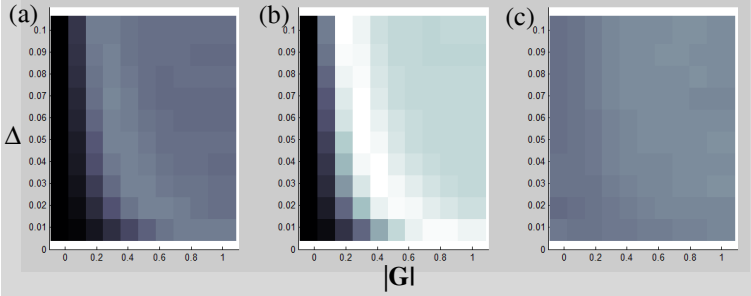

Fig-2: Std devs of syntheitic diffusion signal as a function of |G| and Δ for (a) 1,000,000 and (b) 100 cylinders. Ratio of std devs shown in (c).

Results
Fig-1 shows the time taken to construct substrates and execute spin dynamics for simulations with different subvoxel grid sizes. We observe a broad minimum with total runtime effectively flat over a range of subvoxel sizes from half the mean radius to eight times the mean radius. The minimum is particularly pronounced in the cylinder construction phase of the simulation, although the dynamics phase dominates the total runtime. Fig-1 also illustrates the performance gains from the subdivision algorithm: the difference between the best and worst cases is more than an order of magnitude, and both represent a significant improvement over the naïve case. Fig-2 shows the standard deviation in synthetic diffusion-weighted measurements using substrates with (a) 1,000,000 and (b) 100 cylinders as a function of $|\mathbf{G}|$ and Δ. In both cases we observe a contour of maximum variance which follows a line of constant b-value, indicating there is an optimal b-value which is sensitive to cylinder radius distribution. The same pattern is observed for other values of δ (not shown), where the change in pulse length acts to shift the pattern in the $|\mathbf{G}|$-Δ plane, but does not qualitatively change it. Fig-2(c) shows the ratio of std devs in (b) to those in (a). We observe an approximately consistent factor of 2.1, which increases by a small amount with increasing diffusion weighting. Std devs in (a) are on the order of 2-3% of the mean, those in (b) are 6-7% of the mean.

Discussion & Conclusions
Uniform spatial subdivision allows Monte-Carlo diffusion simulations to be performed in much larger spatial regions than previously. We have shown how spatial subdivision subvoxel size may be tuned to achieve an optimal (and significant) performance increase. Our total region size in these experiments is a little over 1.4mm isotropic. Although a little smaller than a typical clinical scan this is nevertheless larger than the 1.2mm voxel size acquired in the Human Connectome Project [6]. Larger tissue regions also mean larger samples from the underlying distribution of axon radii, which we believe is the source of the reduced variation in synthetic diffusion-weighted measurements. The broad minimum observed in total runtime indicates that the technique is relatively insensitive to the particular size of the subvoxel grid over a range of values, making it easy to choose a good size for a given substrate. We also note that there is no reason to limit the application of this technique to cylinder-based tissue models. The same procedure can be applied to any geometric tissue description, including triangle meshes or other primitives such as spheres or tetrahedra. The limiting factor in substrate size is memory – more memory allows a larger substrate. Other, non-uniform, spatial subdivisions exist [4] and under some circumstances are known to have lower memory requirements. The closely-packed substrates considered here, however, contain very little empty space and would not significantly benefit from non-uniform techniques.

References
[1] Hall & Alexander *IEEE TMI* 28, 9 (2009) [2] Panagiotaki et al *NeuroImage* 59, 3 (2011) [3] Drobnjak et al *J Magn Reson* 212(2) (2011) [4] Slater et al, *Computer Graphics & Virtual Environments* (2002) [5] Aboitiz et al, *Brain Res* 598 (1992) [6] van Essen et al *NeuroImage* 62, 4 (2012)