A Multi-Scale Model of Brain White Matter Structure and Its Solution from Diffusion MRI

PhD Thesis Proposal

Candidate: Jadrian Miles

Department of Computer Science

Brown University, Providence, Rhode Island

Second revision: 2009-10-15

First revision: 2009-06-22

Summary of Changes to the 2009-06-22 Revision

The author thanks the reviewers—David Laidlaw, Peter Basser, John Hughes, and Ben Raphael—for their comments on the revision of this thesis proposal prepared 2009-06-22, and has made the following changes.

- An introduction has been added to §1, Contributions, to clarify the relationship among the components of the proposed work. The sub-sections for individual contributions have also been clarified and streamlined.
- Hypotheses of uniqueness of the solution of optimization systems have been clarified, in particular in §2.3.1. Uniqueness claims apply only to the solution to the regularized objective function for the microstructure parameters at a single position.
- Implementation details have been added throughout to the descriptions of the macrostructure and microstructure models. A new section (§2.5.2) has been added to discuss known brain structures that will be challenging to model.
- The role of clinical applications, §2.9, has been reduced in response to concerns that the applications were impractical to undertake in the thesis period. Clinical applications will primarily be used for informal validation of the various development versions of the modeling system.
- §3.2, Established Collaborations, has been expanded with recent collaborations, in particular Drs. Gregory Balls and Lawrence Frank, the developers of the DifSim Monte Carlo diffusion simulation software [6].
- §4, Research Design and Methods, has been reorganized to match the structure of the Contributions, §1.
- Entries in §4.8, the Timeline, have been expanded to clarify their meaning, with references to the relevant prose descriptions. Potential conference and journal submissions have also been specified. Note that ISMRM, MICCAI, and Vis are conferences. ISMRM accepts only abstracts, either for a 20-minute talk or a poster. MICCAI and Vis submissions may be posters or conference papers. NeuroImage and MRM are journals relevant to the field of diffusion MRI of the brain.
- Animal MRI and histology data have been included throughout the proposal for validation of the model and its solution technique.
- Connections to other sub-fields of computer science have been emphasized. The proposed research depends upon machine learning techniques (§2.6.1) and will contribute to the machine learning and computer science theory communities (§3.1).

This page intentionally left blank.

Introduction

The intended contributions of the thesis work proposed herein are a mathematical model of the macroscopic and microscopic structure of the white matter of the human brain, a technique to compute the parameters of such a model from magnetic resonance imaging scans of a given individual's brain and to evaluate the goodness of fit of the resulting model instance to the available data, and applications of this model to medical research.

The white matter of the brain is made up of axons that connect different regions of the grey matter to each other and to the body. Axons that connect nearby sites tend to do so along the same path. This gives rise to a higher-level structure of so-called "tracts", and similar tracts group together into "fascicles". The mathematical model will comprise a gross segmentation of the white matter into non-exclusive closed regions of similar structure, approximating fascicles, each region accompanied by statistical descriptions of the axons within it: their trajectories, diameters, spacing, physical properties, etc. By computing the agreement between the model and medical images of a real brain, numerical optimization techniques may be applied to adjust the parameters of the model so that it best describes the observations (see Figure 1). At each step of its development, the model and its solution technique will be tested with diffusion and histological datasets, both simulated and real, to establish their stability over variation due to noise and imaging parameters, their accuracy across anatomical variation, and their efficacy in confirming known white matter pathologies.

The candidate hypothesizes that two significant advantages over current techniques for modeling the white matter of the brain will result from this novel combination of large- and small-scale structure in one model. One hypothesized advantage is that it will allow the computation of a unique optimal solution at each point for a more detailed model of microscopic structure than currently exists (see discussion in §2.3.1). The other is that it will provide a more accurate and precise reconstruction of macroscopic structure for the brain overall. These advantages will result from the optimization-based solution process, in which information from large- and small-scale structures may each inform the solution of the other. For example, characteristics of the microscopic structure may be assumed to change smoothly over the volume of a given white matter fascicle, constraining the local solution of these characteristics, while discontinuities in the input images that would violate this smoothness constraint may instead be accounted for by the presence of a fascicle boundary.

In addition to the immediate advantages of the proposed modeling system, it has potential for broader impact through future applications in computational medical imaging and clinical neuroscience research. The multi-scale model would allow for a holistic and heirarchical view of the tissue structure not only for the purposes of fitting the model parameters to a specific set of images, but also for further post-processing and novel analysis of diffusion data; for interactive software for visualization, selection, and inspection of white matter structure; and for the definition and measurement of statistics within and between subjects. Some of these applications will be demonstrated as part of the proposed work but many more are possible as future work by other researchers.

Background and Terminology. *"Diffusion MRI"* is a magnetic resonance imaging technique that remotely measures water self-diffusion and thereby enables the noninvasive observation of the effects of white matter structure on this diffusion. The output of a diffusion MRI scan is a set of so-called *"diffusion-weighted images"* or *"DWIs"*, which will be the principal input to the proposed techniques. Each DWI is a raster volume—that is, a regular three-dimensional grid of *"voxels"*, the volumetric analogue of pixels. The value in each voxel in a DWI is a noisy measurement of water self-diffusion at that position in the brain subject to the parameters of that DWI.

Throughout this proposal, certain terms will be used with specific intended meanings. *"Images"* primarily indicates DWIs, but in several places could also include data from other types of MRI, in particular T1- or T2-weighted images. *"Macrostructure"* denotes white matter fiber trajectories at the millimeter scale as well as the morphology of the gross segments of brain tissue: coherent white matter structures and contiguous volumes of grey matter or cerebrospinal fluid. *"WM"*, *"GM"*, and *"CSF"* respectively refer to these tissue and fluid types. *"Microstructure"* denotes the characteristics of the white matter on a sub-millimeter level, including details of fluid exchange, fluid self-diffusion, and cell geometry. *"Model"* denotes a mathematical model—a set of parameters and a consistent interpretation thereof—while an *"instance"* of that model is a particular set of values assigned to the parameters. *"Inverse solving"* denotes the application of numerical optimization techniques to fit a global model of a system to input observations; this is in contrast to so-called *"forward modeling"* in which the global model results as the aggregation of smaller, independent models that are fit to subsets of the total observations. A *"modeling system"* is the combination of a model and a technique used to fit it to observations. The terms *"axon"*, *"fiber"*, *"tract"*, and *"fascicle"* will refer to the levels of the anatomical heirarchy of the white matter of a real brain, in order of increasing size, while *"curves"* and *"bundles"* are the geometrical representations that will attempt to approximate them in the model.

1 Contributions

This section briefly describes the principal components of the proposed brain modeling system, which may be described and developed independently of each other. The system comprises three parts: a macrostructure model that decomposes the brain volume into large structures, a microstructure model that describes the small-scale tissue properties within each structure, and a solution technique to fit the parameters of the models to the input images. Following a "spiral" development paradigm, the complete system will go through several revisions. Each revision will be evaluated before being refined further.

1. First, the complete macrostructure model, containing a simplistic microstructure model, will be developed, as well as a process to initialize this model from input images (§1.1). In order to evaluate goodness of fit for this model and to drive the later solution process, a process to generate synthetic images from the model will be developed (§1.2).
2. Next, an iterative optimization process will be developed to refine the initialized model instance from the previous step to better fit the input images (§1.3).
3. Next, a detailed microstructure model will be developed and studied, using physics simulations, independently of the other components (§1.4).
4. Finally, the detailed microstructure model will be incorporated into the macrostructure model, and an optimization process will be developed for this combined model (§1.5).

Each of these revisions constitutes a distinct novel research contribution. Clinical applications of the modeling system will primarily be used informally to provide validation and guidance during the development stages listed above. More significant clinical contributions (§4.7) may be pursued as time allows.

1.1 A geometric model of whole-brain macrostructure and white matter fascicles. The proposed model will represent volumes of contiguous grey matter ("GM"), volumes of cerebrospinal fluid ("CSF"), and coherent fascicles of white matter ("WM") tracts by their enclosing surfaces. The model for a WM fascicle will also compactly represent the trajectories of its constituent tracts, such that a curve representing a tract may be reconstructed from the model without explicitly storing its trajectory. A process, based primarily on existing techniques, will be developed to generate an instance of the model from images in a forward-modeling fashion (see §2.4,2.5). A model instance generated in this way will form the initial candidate solution for the inverse-solving systems described in §1.3 and §1.5. The model and forward-modeling process will be evaluated by constructing a model instance from images, and then quantitatively comparing reconstructed WM curves to the forward-modeled curves from which the instance was constructed. For evaluation details, see §4.1.4.

1.2 A method for generating synthetic MR images from a geometric brain model. A process will be developed to convert an instance of the macrostructure model (§1.1) into a volume of instances of an existing, simple voxelwise tissue microstructure model. From this, it will be possible to generate synthetic images from the macrostructure model (see §2.6). The model will be evaluated by constructing an instance from

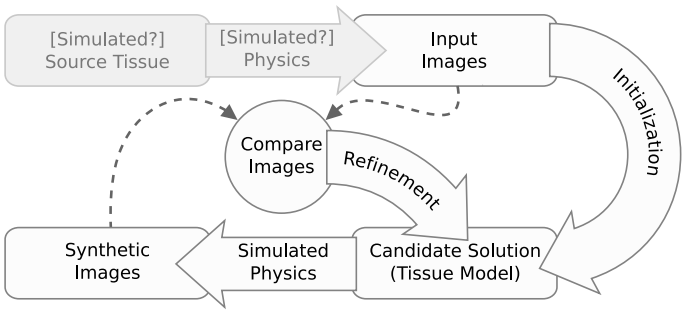

Figure 1: A flowchart of the proposed inverse solving technique, described in §1.3 and §1.5. The source of the input to the solver is indicated in darker boxes, while the solver itself is illustrated by the lighter boxes. Starting from a set of input images (which may be real images from an MRI scanner or synthetic images), an initialization step creates a tissue model instance that serves as the initial candidate solution to approximate the unknown source tissue. Synthetic images are generated from the candidate and then compared to the input images. Based on the differences between the sets of images, the candidate is refined and then new synthetic images are generated.

images, generating synthetic images from the model, and comparing these to the original images (see §4.2.1). The accuracy of the "reconstructed" images is expected to be at least 75% but with many localized artifacts that motivate the optimization-based fitting technique (see similar results in [35]).

1.3 A method for inverse solving a geometric brain model. The inverse solving system depicted in Figure 1 will be implemented by initializing the candidate solution model instance from input images as in §1.1, generating synthetic images as in §1.2, and refining the model based on biologically-motivated regularizations and the differences between the synthetic and input images (see §4.3. The system will be evaluated in turn with two sources of input. First, the input will be synthetic images of computational phantoms, constructed as macrostructure model instances, so that the final candidate solution may be directly compared to the phantom "tissue". Second, the input will be real images, and the synthetic images of the initial and final candidates will be compared to the real input images.

1.4 A demonstration of ambiguous diffusion response in detailed microstructure. A microstructure model will be developed that represents more microstructural characteristics of the WM than the state of the art (see §2.2). This will involve almost certainly too many microstructure parameters to derive a closed-form solution for the diffusion response. Instead, synthetic images will be generated using 3D diffusion simulations on examples that cover the microstructure model's parameter space (§2.6). Machine learning techniques will be applied to the simulation output to learn an approximation of the diffusion response as a function over the parameters of the microstructure model. Ambiguities with respect to the diffusion imaging response will be demonstrated among sets of parameters, motivating the combination of macrostructure and microstructure in one model to constrain the microstructure solution (see §2.3).

1.5 A method for inverse solving a geometric brain model with detailed microstructure. The model will be refined to incorporate the detailed microstructure model of §1.4, and the inverse solving method from §1.3 will be similarly updated. To demonstrate the improvement over the original model (which uses a simplistic microstructure model), the new model and the original model will be solved for the same set of input images, and then synthetic images of the final candidate model instances will be compared to the input images. The goal of this step is 95% accuracy for the new model, in addition to stability over imaging variations.

1.6 High resolution MRI data. Validation of the various development steps will require images of healthy brains acquired under a variety of conditions; in particular, DWIs at multiple gradient magnitudes and high spatial and angular resolution will be required (see §2.7,2.8). Publicly available data will be used when possible, in particular for high-resolution diffusion MRI and histology data of animals. Otherwise, acquisitions will be performed locally and the data made available online.

1.7 Applications. At various points throughout the progress of the proposed research, the systems developed will be applied to other quantitative diffusion imaging techniques and to clinical topics (see §2.9).

2 Background and Significance

2.1 Advantages. The proposed modeling and solution framework could result in improved precision and accuracy of brain tissue reconstruction over current techniques at both the macro- and micro-structural levels, and thereby be useful as a tool for neuroscience researchers investigating white matter. Within the field of computational diffusion MR brain tissue modeling, it is believed to be:

- the first optimization-based technique for whole-brain macrostructure modeling that works directly from diffusion and structural images, rather than a higher-level diffusion forward model (c.f. §2.5.3);
- the first technique that models both macrostructure and microstructure and fits them simultaneously;
- one of the only techniques that explicitly quantifies and optimizes its goodness of fit to the input images; and
- a significant advance beyond the state of the art in terms of the number of microstructure characteristics incorporated into a model (c.f. §2.2,2.3).

The solution technique as proposed would require no human intervention to fit the model to a set of input images. The formulation of a high-level mathematical model for white matter macrostructure itself has potential as a basis for future work by other researchers on problems in the tissue modeling field unrelated to those investigated by the proposed research.

2.2 Background: Microstructure Modeling. Among the earliest applications of diffusion MRI in neuroscience was the investigation of WM microstructure properties at the voxel level. The earliest and still most widespread technique for this is diffusion tensor imaging ("DTI"), which fits a second-order tensor to the orientation distribution of apparent diffusion coefficients in each voxel [7]. A variety of measures of the diffusion tensor were developed, including the popular tensor invariant called "fractional anisotropy" (FA) [9]. Especially in clinically-oriented research, voxelwise analysis with DTI, and particularly FA, continues to be the most common use of the technology; e.g., [33, 36, 32, 24]. However, the diffusion tensor model has widely acknowledged disadvantages, such as poor fit in partial volume voxels (where distinct fiber populations or tissue types coexist in the same voxel) and sensitivity to image resolution and other acquisition parameters [1, 48, 62, 2]. FA and other anisotropy measures also suffer from ambiguous interpretation in terms of actual tissue properties; demyelination, axon dropout, and mixed fiber populations, for example, can all cause a decrease in anisotropy [42].

One response to these disadvantages has been the development of mathematical models of the microstructure itself that can be solved directly from the DWIs, e.g., [60, 5, 4, 3]. Parameters of these models include orientation of axon populations, diffusion coefficients of the intra- and extracellular fluid, volume fractions of the various fluid compartments, distribution of axon diameters or "calibers", and, more recently, fluid exchange between compartments [69]. In order to solve these models, however, various assumptions have been made for each—most often that the fibers are oriented in a single, known direction—and none solve for all the parameters listed.

2.3 Challenge: Microstructure Models May Be Ambiguous. In the general case of full-brain diffusion MRI, however, fiber orientation at each point is not known. It has been argued that fluid exchange is also a necessary component of any microstructure model that hopes to accurately approximate the real tissue [44]. Axon caliber affects nerve signal conduction speed [52], making it an important property to be modeled as well. Axon calibers are known to vary widely within the volume of a single voxel in most regions of the WM, implying a need to model a distribution of calibers, not just a single value [4]. Furthermore, not only do many regions exist where fiber tracts cross, axons even within a single tract exhibit some variability in orientation [31], a property that has as yet not been incorporated into any microstructure model.

As the list of desirable parameters in a microstructure model grows, the system to be solved grows more underdetermined relative to the available diffusion MRI measurements. The candidate hypothesizes, in fact, that even if the full four-dimensional probability density function of diffusion were known for a single voxel, a sufficiently detailed but biologically realistic microstructure model could still be underdetermined; that is, that there would exist distinct instances of the model that resulted in the same diffusion profile.

2.3.1 Approach: Microstructure Model Regularization. The remedy proposed here to these theoretical limitations is the principled application of non-local regularization based on generic, known properties of white matter tissue. WM axons that terminate inside the brain do so, in healthy brains, only within grey matter [56, 49], and therefore the number of axons in the cross-section of a non-branching fascicle is conserved along its length within the brain. In a given fascicle cross-section with area A, axon volume fraction V, and axon radius distribution $N(r)$, where $\int_0^\infty N(r)\,\mathrm{d}r$ is the total number of axons in the cross-section, then this conservation principle implies $\int_0^\infty N(r)\pi r^2\,\mathrm{d}r = VA$. Since the same axons are members of a fascicle along its entire length, it is also reasonable to assume a fixed (or at least smoothly varying) distribution of radii and intracellular diffusion coefficients, and perhaps even of exchange rates, within a bundle. Whereas the diffusion response in a single voxel may be insufficient to give a unique solution to the microstructure parameters for that voxel in isolation, the candidate hypothesizes that by incorporating these and other regularizing assumptions into the objective function for fitting microstructure parameters to the diffusion response, a unique optimal solution to the local microstructure model will arise.

The computational tractability of the optimization process for the whole model may be improved by modeling the smoothly changing parameters as a sparse sampling or simple function over the length of each bundle, thereby reducing the number of free variables in the model, rather than storing parameters for a full microstructure model for each voxel. This may also provide robustness against variation in resolution.

2.3.2 Related Work: Microstructure Model Regularization. Many examples exist of previous work on regularizing the solution of diffusion models, both for the diffusion tensor model [19, 65, 46] and for higher-order models [14, 55, 20]. All of these, however, assume equal smoothness of the diffusion field in all directions; that is, when fitting the model in a given voxel in one fascicle, data from a voxel in a different fascicle affect the fit as much as those from a voxel an equal distance away in the same fascicle do. This necessarily leads to blurring of the boundaries between fascicles and between the WM and GM, and therefore exacerbates partial

volume issues and the limited resolution of diffusion MRI. The proposed model, however, incorporates both the macrostructure and the microstructure, and the proposed solution technique regularizes microstructure only within bundles, leaving their boundaries sharp at sub-voxel resolution.

2.4 Background: Macrostructure Modeling with Tractography. Large-scale and even full-brain macrostructure have also been a subject of investigation with diffusion MRI. The most widely used technique for modeling the WM macrostructure is known as "tractography", in which space curves representing likely paths of fibers through the WM volume are computed from the voxelwise diffusion model instances [8, 41]. Full-brain sets of tractography curves have been used in fundamental anatomical research of cataloguing the regular structure of the WM fibers [45] and to study cortical connectivity [25]. Manually-selected subsets of curves have also been used as approximations to known WM fascicles in order to study the effects of diseases on specific structures [29]. There is a great deal of variation between subjects in brain morphology [54, 53], indicating that atlases may be insufficient for identifying common structures [70]. Accurate macrostructure reconstruction for each subject individually may therefore offer a better alternative for the comparison of WM structures between subjects.

Once again, the best-established diffusion model underlying tractography algorithms is DTI, but it has known disadvantages. The poor fit of the diffusion tensor in partial volume voxels and at fiber crossings can lead simple tractography algorithms to terminate curves prematurely within the WM or to generate curves with spurious redirections into other fiber bundles [34]. In response to the failure conditions of the diffusion tensor model, a number of higher-level diffusion models have also been proposed [11, 27, 62, 26, 61], and research is active on tractography algorithms that take advantage of their higher angular resolution [66].

Though more sophisticated tractography algorithms that run on DTI or higher-order models generally perform better than simple ones, they can all result in unrealistic reconstructions. A common failure case is that when fiber tracts cross, a significant number of curves for one tract either terminate at the crossing region, redirect into the other tract, or otherwise lose their directional coherence, while only the remaining fraction successfully pass through to the same tract on the other side of the crossing; e.g., results in [47, 55]. Other techniques appear more successful; e.g., [51, 21]. All of these algorithms improve over traditional DTI-based tractography, however, and any could be incorporated into the proposed work as a step in the model initialization. (There is also another class of techniques called *probabilistic* tractography that shows promising results for connectivity applications (e.g., [10, 12]) but does not appear to be applicable to the proposed model.)

2.5 Background: Macrostructure Modeling with Curve Clusters. The proposed solution technique initializes the macrostructure model by generating a set of curves with tractography and then identifying volumes of coherent WM macrostructure by clustering the curves together. Clustering in general is a fundamental problem in computer science theory [43]. Many options for curve clustering specifically for brain modeling have been investigated in the literature (e.g., Moberts, et al.'s review [40]). The choice of curve-to-curve similarity measure and clustering algorithm have a strong effect on the resulting clustering.

2.5.1 Approach: Geometric Macrostructure Modeling. The candidate hypothesizes that a properly designed curve similarity measure and clustering algorithm, run on a whole-brain set of curves, can automatically recover structure from collections of "broken" curves (those that fail to run the length of fiber they approximate) that no single region of interest could capture. A curve similarity measure and clustering algorithm will be designed specifically for this goal. An additional design challenge for the clustering algorithm will be the large number of curves to be clustered, typically numbering in the hundreds of thousands.

Once curves are clustered, the bounding volumes of the clusters themselves will become the initial values of the model's WM bundles. As these are treated as approximations of WM fascicles, the aforementioned known properties of fascicles may be enforced on them. Rather than store all the curves contained within a cluster, which by construction follow similar trajectories, a sparse representation of the trajectories will be generated and included in the bundle model.

In order to assert that bundles terminate only within the GM (or at the lower edge of the modeled volume, where the spinal cord projects to the body), the model must also have a reckoning of the GM volume. Another property to be enforced is that any "empty space" in the volume of the brain is occupied by cerebrospinal fluid, which can also be identified with MRI and therefore included in the model. Volumes of GM and CSF will be represented in the model by a triangle mesh of their boundaries, which may be generated by established tissue classification techniques and isosurface mesh-generation algorithms; specifically, the classification schemes described in [58] and [53], and the advancing front algorithm of Schreiner, et al. [57] will be used. Another property that is widely accepted is that the curvature of WM tracts is limited [41, 49, 15], particularly in the

interior of the WM; a related but less widely asserted property is that WM fibers are normal to the WM/GM surface where they insert into the GM [63, 38]. In the initialization phase of the inverse solving process, the macrostructure model instance generated by the clustering will be directly adjusted so that all WM bundles terminate at the GM and the entire volume of the WM is covered by bundles, further correcting for errors made during tractography. During the actual optimization, these properties will define the space of feasible solutions and serve as constraints that may be relaxed depending on the optimization technique. The objective function itself will include terms to favor low bundle curvature and perpendicular insertion of curves into the GM.

2.5.2 Challenge: Modeling Brain Structures with Difficult Shapes. The optic chiasm and the cingulum provide well-known challenges to macrostructure models. The optic chiasm contains four comingled tracts, two of which cross (or *decussate*), while the other two remain on their respective sides. These mixed fiber populations confound DTI tractography but have been resolved with tractography over higher-order diffusion models [66]. Provided that it is initialized with a good tractography instance, the proposed macrostructure model would model each of the four tracts as distinct bundles that overlap in the volume of the optic chiasm.

The cingulum is only a single tract, but its path lies along the inner surface of the cortical GM from the cingulate gyrus to the temporal lobe. At each end, fibers within the tract turn sharply to insert into the GM while others continue past to insert further along the structure, resulting in a flattened fanning shape. This difficult case constrains the design of the macrostructure model, as any potential model must be able to express the structure of the cingulum. One possibility is illustrated in figure 2.

2.5.3 Related Work: Inverse Solving for Regularized Macrostructure Models. Inverse solving has been applied to macrostructure models, specifically tractography, in the past with encouraging results. The "spin-glass" or "spaghetti-plate" tractography model attempts to find curves that terminate only at the WM/GM interface and have low curvature, while still fitting the data as well as possible [37]. It does this by globally minimizing a configuration energy that includes tract curvature and local alignment to the diffusion tensor field and requiring that tracts terminate at the boundary. Recent work has extended this concept to higher-order local diffusion models [23]. The proposed technique is distinct from these in that it solves for macrostructure and sophisticated microstructure simultaneously, that it regularizes the macrostructure at the level of bundles rather than individual curves, and that it uses the DWIs directly as input rather than working from a higher-level diffusion model.

2.6 Related Work: Synthetic DWIs. The final component required for the proposed inverse solving technique is a procedure to generate synthetic DWIs from the combined macro- and microstructure model. This problem has already been studied for the case of tractography curves and diffusion tensors as the respective models by Leemans et al. [35] and Close et al. [15]. Leemans et al. show especially compelling results by generating tractography curves from DTI and then comparing the reconstructed tensor images to the originals. The partial success of this simple approach suggests feasibility of the more complex proposed system, specifically showing that the initial candidate solution for the optimization process may be quite close to the optimal.

The solution techniques for the microstructure models described above also require that the diffusion MRI signal for a given model instance be defined [60, 5, 4, 3, 69]. All of these therefore could generate synthetic DWIs for a volume of voxelwise model instances.

Figure 2: A cartoon of the axonal structure of the cingulum (left) and a possible representation of such a structure as a single bundle in a macrostructure model (right). The cingulum lies along the inner surface of the cortical GM and each fiber turns sharply to insert into the GM over a relatively long region at either end, in the cingulate gyrus and the temporal lobe (the distance between these GM sites is not to scale). In the equivalent macrostructure model, two surfaces on the bounding volume (dark lines) are specified as the terminating surfaces for curves, while the other surfaces of the volume show where the outermost curves lie. In the interior of the bundle, a center line determines the orientation of normal planes (grey lines) that define the cross-sections of the bundle.

2.6.1 Approach: Diffusion Simulation. It may not be possible to write a closed-form definition of the diffusion MR response for the detailed microstructure model proposed above. However, the diffusion response for a given microstructure may also be computed directly through physics simulation, as in Balls and Frank's DifSim software [6]. This software will be extended to accommodate the detailed microstructure model proposed above, and the diffusion MR response will be simulated and recorded for many microstructure configurations. These simulated values will be used either to validate the closed-form expression of the diffusion response, or to determine the function for the response using machine learning techniques.

2.7 Background: Data Acquisition. The diffusion tensor model requires a minimum of only seven DWIs for reconstruction [7], though gathering more samples improves the signal-to-noise ratio (SNR) [42]. Higher-level models and detailed microstructure models require more DWIs, an extreme case being diffusion spectrum imaging [2]. Acquiring more DWIs and at higher spatial resolution increases scanning time, which is a practical limitation in clinical applications [42]. While lower-resolution data are more realistic inputs to software for clinical use, the development of a new modeling technique may be simplified by working from high-quality data before attempting to accommodate more realistic datasets.

2.8 Challenge: Sensitivity to Data Variation. Reproducibility of theoretical results for diffusion MRI processing with data different from those used in the original procedure is often difficult to achieve, for several reasons. SNR is often relatively low in diffusion MRI studies, so noise can lead to significant variation even between two otherwise identical acquisitions of the same subject. Variation in brain morphology between subjects, as explained above, leads to further reproducibility issues. Lastly, since diffusion MRI is still a young field, diverse pulse sequences and scanning protocols have been developed for acquiring DWIs, none of which will emerge as a universal standard in the forseeable future, and quirks of individual scanners may also give slightly different imaging results for the same subject. Furthermore, diffusion MRI suffers from a number of sources of noise, artifacts, distortions, and mis-registration [2], and many options exist for pre-processing the data to correct these. Ultimately, each choice of scanning protocol and image correction results in a unique set of image statistics, once again providing a barrier to reproducibility [30].

2.8.1 Approach: Multi-Faceted Validation. Validation of robustness across several modes of data variation is therefore essential. The first implementation of the inverse solving system will use several computational phantoms to test for robustness to morphological changes, possibly generated by the software described in [15]. To test for stability, several orientations of the phantom may be used relative to the imaging coordinate system, and varying levels of noise may also be introduced into the synthetic DWIs of the phantoms used as input.

For later implementations of the inverse solver, real data will be required as input. A few databases of diffusion MRI scans are publicly available; among them is the high angular resolution database described in [50]. This research group's local collaborators in Providence have also collected a large number of scans at lower angular resolution and with only one non-zero diffusion weighting b-value. The proposed work will make use of available data to the greatest extent possible, but three types of high-resolution, multi-b-value data are foreseen to be necessary for validation:

- Multiple acquisitions of the same subject using the same protocol, to test robustness to noise and artifacts
- Multiple acquisitions of the same subject, each with a different protocol, to test robustness to protocol-induced image variations
- Acquisitions of multiple subjects, all using the same protocol, to test robustness to inter-subject variability

Any necessary acquisitions will be performed locally on a 3-Tesla MRI scanner.

While comparing synthetic DWIs of a model instance to input DWIs allows the computation of goodness of fit in the form of the image reconstruction accuracy, a true validation requires ground truth. Paired diffusion MRI and histology of an animal brain would allow a comparison of a model instance directly to the true brain structure of a specific individual. If such data are publicly available, they will be used to evaluate the initial macrostructure model, the simple modeling system, and the final modeling system with detailed microstructure. If not, it may be possible to collect histology data from collaborators (see §3.2).

2.9 Approach: Clinical Applications. Intermediate forms of the solution system will be applied to clinical neuroscience subjects for informal validation. The model should be designed so that it can describe pathological as well as healthy brains, and the solution system should reflect differences between healthy and diseased brains for well-understood brain pathologies such as CADASIL (§2.9.1). It is not within the scope of the proposed research to perform formal clinical studies using this model, and most of the applications proposed

below will be investigated to the level of a pilot study, if not skipped entirely. Those applications that are pursued will be in collaboration with the established clinical collaborators listed in §3.2.

It is important to note that neuroscientists' conception of the anatomy of the brain is informed by functional insights and centuries of detailed dissection and histology, and therefore it is unlikely that an automatic system whose only inputs are diffusion MRI data would generate a model of the brain organized in the same fashion. It is also unknown whether the models generated for two different scans would be organized in a similar way. Therefore most clinical applications based on results of the proposed model and inverse solving method would be expected to require some manual refinement by splitting and joining macrostructure bundles.

2.9.1 Application: Bundlewise Statistics. Though there have been some successes in medical research comparing DTI measures voxelwise between subjects (e.g., [32]), the known high degree of variation in human brain morphology suggests that greater precision and statistical power may be available by comparing tissue statistics between selected WM structures identified individually for each subject. Additionally, medical hypotheses about brain disease are often constructed in terms of known structures of the brain, so tools allowing for the semi-automatic segmentation and quantification of these structures could be helpful for medical research.

Since the size and shape of WM structures may vary non-pathologically between subjects, some degree of normalization is desirable. Correia et al. [18] identify a number of values measured over manually selected tractography curves of interest that are normalized for tract size and/or intracranial volume, but they suffer from numerical instability. Furthermore, since the values are measured over space curves, their normalization with respect to a three-dimensional value is difficult to justify. The proposed macrostructure model, however, directly represents WM tracts by their volumes and representative trajectories, and it may therefore be more straightforward to define diffusion measures with respect to this model.

After the inverse solving method has been developed to the point that it is reliable on real data using a simple microstructure model, bundlewise statstics similar to those presented in [18] will be defined. Sensitivity testing as described in §2.8.1 will be performed to assure that the newly defined statistical measures are stable.

Pre-existing diffusion MRI data from healthy controls and patients with CADASIL (Cerebral Autosomal Dominant Arteriopathy with Subcortical Infarcts and Leukoencephalopathy) will be processed by the inverse-solving system. Measures will be computed on WM bundles manually selected, and possibly manually refined, from the resulting model instances. CADASIL is known to uniformly degrade the WM, so statistically significant group differences are expected.

2.9.2 Application: Bundlewise Statistics and HIV. Having validated the newly defined tractwise statistics with CADASIL patients, an experiment will be undertaken to determine focused WM damage as the result of HIV infection. Again, pre-existing diffusion MRI data from healthy controls and HIV patients will be processed by the inverse-solving system. Measures will be computed on several manually selected and refined WM bundles, and the groups will be compared for each bundle to identify localized effects of infection.

2.9.3 Application: Multiple Sclerosis Fibers at Risk. So-called "fibers at risk" ("FAR") are WM fibers that pass through a region of localized damage in the brain, such as a multiple sclerosis (MS) lesion. It has been observed that localized lesions affect brain functioning in remote regions, and it is hypothesized that WM tracts passing through lesions are more likely to degrade as the disease progresses. A pilot study that used DTI tractography to build a FAR map on the mid-sagittal section of the corpus callosum of a single MS patient was previously published [59]. Since then, longitudinal data for a larger population of MS patients have been collected. These scans will be processed by the inverse-solving system to generate instances of the macrostructure model with simple microstructure, and FAR will be re-formulated as a continuous projection of the lesion volume onto a chosen plane of interest, distorted according to the bundlewise fiber trajectory model. FAR maps will be generated for all scans on the mid-sagittal section of the corpus callosum, and the hypothesis will be tested by observing the progression of the disease within subjects. WM degradation in the plane of interest, observed in the form of decreased FA or some other measure, is expected to correlate through time with areas of identified high risk in earlier stages.

3 Preparation

While this proposal describes plans for work that has not yet been undertaken, several examples of preparation indicate that the candidate is capable of completing the proposed research.

3.1 Preliminary Work: Curve Clustering. A review of existing curve-to-curve similarity measures [17, 13, 67, 22] determined that these were insufficient for the goal of exploiting broken tractography curves to recover more complete tract clusters in a full-brain clustering. A new measure was developed that mimics the "corresponding segment" measure of [22] for arbitrary pairs of curves, and weights the similarity by the evenness of the correspondence to bias against skewed curve pairs. A parallel algorithm for efficient computation of the sparse similarity matrix was designed and implemented. Well-known problems with the overzealous clustering behavior of agglomerative methods such as that described in [67] were anticipated, and a combined strategy that includes spectral clustering is planned [43]. This combined clustering algorithm would be a significant contribution to machine learning and computer science theory, as it would exploit the efficiency of agglomerative clustering for large datasets while still affording the sensitive discrimination of spectral clustering.

3.2 Feasibility: Established Collaborations. Established collaborations exist between this research group and medical researchers around the world, in which this group provides DTI data processing and implements new computational techniques, while the outside collaborators provide raw data and pursue their own research agendas. The candidate has also collaborated directly with researchers at other institutions on work relevant to the proposed research.

- **Stephen Correia, Ronald Cohen, and Stephen Salloway: Providence, Rhode Island.** The teams led by Drs. Correia, Cohen, and Salloway have contributed more than 100 diffusion MRI scans at 1.5T and 3T to this research group over several years, and have collected more than 100 new scans at 3T and high angular resolution for a current study of HIV. In turn, this research group has collaborated with them on several clinical studies and the development of new analysis techniques. Since joining this group, the candidate has worked closely with Dr. Correia's team to provide software tools for DWI processing and to troubleshoot mathematical issues with data analysis.
- **Edward Walsh: Providence, Rhode Island.** The candidate has worked with Dr. Walsh to design a high angular resolution diffusion MRI acquisition protocol. He has also trained with the technical staff at the medical imaging center at Brown to participate in planning and carrying out MRI acquisitions of human volunteers.
- **Gregory Balls and Lawrence Frank, La Jolla, California.** The candidate negotiated with Drs. Balls and Frank to extend the functionality of their Monte Carlo diffusion simulator, DifSim [6], to work with more detailed microstructure.
- **Jason Kaufman, Pasadena, California.** Dr. Kaufman has collaborated with this research group to provide diffusion MRI and histology data for a macaque brain. These paired data may be useful in the validation of the proposed system.
- **Bruce Spottiswoode: Cape Town, South Africa.** The candidate has been the primary contact at Brown on data processing issues for a Cape Town research team studying HIV and has developed new software to accommodate their data. This project is led by PI Robert Paul of St. Louis, Missouri and was recently funded by the NIH.
- **Mark Bastin: Edinburgh, Scotland.** The candidate has coordinated processing of brain tumor data for this collaborator.
- **Jack Simon: Portland, Oregon.** The candidate met with Dr. Simon to coordinate further research on the Fibers-At-Risk project.
- **Robert Paul and Tom Conturo: St. Louis, Missouri.** The candidate has worked directly with the St. Louis-based PIs of a multi-site study of cognitive aging and has developed new software to accommodate their data.

Through these collaborations the candidate has gained familiarity with the medical context of the proposed research, including the realities of data acquisition, anatomical and biological knowledge of the brain, and the questions that are investigated with diffusion MRI. Several collaborations afford the re-use for this research of data and software that have already benefitted from substantial effort on the part of the collaborators. The involvement of established medical researchers in this project also encourages clinical applications for the proposed research.

3.3 Related Work: DTI Pipeline. Data received from all of the above-listed collaborators are processed by the same "DTI pipeline": conversion of the DWIs to an in-house format, followed by resampling, nonlinear tensor-fitting, tractography, and visualization. Since joining this group, the candidate has expanded the variety

of data formats handled by this software pipeline; prior to this work, data only from the Providence-based team could be automatically processed. He has also coordinated the growing list of collaborator data handled by this group and has organized both the data and the software tools involved.

Familiarity with this data pipeline has given the candidate first-hand experience with the variety of data formats and scanning protocols that must be managed, as well as other caveats of real medical data including noise, distortions, metadata ambiguities, and the real-world computational scale of theoretical algorithms.

3.4 Related Work: DTI Interpolation. The candidate presented a study on tradeoffs in the interpolation of DTI-derived tensor measures at the ISMRM '09 poster session [39]. This study investigated the practical results of the formal properties of diffusion tensors, and motivates the pursuit of resolution-independent WM models.

3.5 Related Work: Probabilistic Index of Connectivity. Previous research by the candidate into diffusion simulation over diffusion MRI datasets covered a somewhat different area of the diffusion processing literature, including closed-form techniques based on heat diffusion and simulation techniques that fall under the heading of probabilistic tractography. This broader knowledge of the field will assist the candidate in responding flexibly to unanticipated difficulties in the proposed research.

3.6 Related Work: DTI Protocols. The candidate has assisted Dr. Correia's research group in studying systematic variation in data sets due to different acquisition protocols for DTI. Familiarity with sources of noise and error, as well as established contact with Dr. Walsh and others at Brown's imaging facility, will assist in the design of thorough stability tests for the proposed techniques.

4 Research Design and Methods

4.1 Macrostructure Model. Much of the work to develop the geometric macrostructure model initialization process will be based either on novel implementations of established algorithms or on existing software.

4.1.1 Q-Ball Tractography. The macrostructure model will be initialized from tractography curves. The FLIRT component of FSL [28] will be used for registration and distortion correction of DWIs, and Camino [16] or DTK [64] will be used for Q-ball tractography (see §2.4,2.5.2). **Expected Duration: 0.5–1 month**

4.1.2 Curve Clustering. Work on the definition and implementation of a curve similarity measure and a clustering algorithm is in progress but incomplete (see §3.1). An informal evaluation will follow development. **Expected Duration: 2 months**

4.1.3 Macrostructure Model Initialization. The work of defining the macrostructure model and of implementing an algorithm to derive an instance of it from a clustering of tractography curves will proceed in parallel. Voxelwise segmentation masks for WM, GM, and CSF will be generated either by the FAST component of FSL [68] or by implementations of other algorithms [58, 53]. The voxelwise segmentation masks will be converted into triangle meshes using an isosurface mesh extraction algorithm [57]. Conversion of curve clusters into triangle meshes or other geometric representations of their convex hulls will require novel software development. **Expected Duration: 2 months**

4.1.4 Evaluation: Macrostructure from Tractography. The macrostructure model will not explicitly store most (or perhaps any) of the tractography curves from which it was constructed. However, it must be able to recover all directional information present in the input to within acceptable bounds. To evaluate the sensitivity of the construction process, dense sets of tractography curves will be generated from images already collected from approximately 10 human subjects. Random subsets of these curves will be selected, from 100% to 10% in 10% increments. From each subset, a macrostructure model will be generated and then used to reconstruct the original set of curves. The similarity measure used in §4.1.2 will be used to measure reconstruction error. **Expected Duration: 0.5–1 month**

4.1.5 Macrostructure Forward-Modeling Adjustment. An algorithm will be designed and implemented to adjust an instance of the macrostructure model to conform to biologically-based constraints including total coverage of the brain volume and WM termination conditions. This will involve concatenating bundles that terminate in the interior of the WM and expanding those adjacent to empty spaces. **Expected Duration: 1 month**

4.1.6 Macrostructure Rendering. In order to produce figures for papers and talks, and also for future development of interactive software tools, a method to render two-dimensional images of the constituent volumes of a macrostructure model instance will be necessary. **Expected Duration: 2 months**

4.2 Synthetic Images from the Macrostructure Model. A procedure will be defined to generate voxelwise instances of a ball-and-stick diffusion model [11] (or some other simple diffusion model established in the literature) from an instance of the adjusted macrostructure model. Every diffusion model implies an equation for the diffusion response and therefore generating images from a volume of such voxelwise model instances will be a straightforward process. **Expected Duration: 1.5 months**

4.2.1 Evaluation: Macrostructure from Images. The generic image-based evaluation for the model will be as follows. Given a set of input images, a macrostructure model and synthetic images of this model will be generated. The per-image, per-voxel error will be the difference in the values of the input and synthetic images in that voxel, normalized relative to a conservative Rician noise prior derived from the original image. The total error will be the mean error over all voxels in the brain volume, over all images.

The macrostructure model (§4.1.3) and forward-modeling adjusted macrostructure model (§4.1.5) will be evaluated on real images according to the 'full battery" of sensitivity tests described in §2.8.1 as well as on animal data with paired histology, to which the model itself may be compared. **Expected Duration: 2 months**

4.3 Inverse Solving: Macrostructure and Simple Microstructure Each iteration of the inverse solving procedure will involve computing the voxelwise error between input and synthetic images and refining the model for those bundles that intersect regions of high error. The nature of the errors will inform heuristics for choosing good refinements. Possible refinement operations on the candidate solution will include splitting or merging bundles, distorting the trajectory of an entire bundle while keeping its cross-sectional area fixed, locally distorting the boundary of a bundle while compensating for the new cross-section at the microstructure level, and uniformly scaling the cross-section of an entire bundle. The objective function will include both regularization terms for curvature, brain volume coverage, and termination criteria, as well as the image reconstruction error. **Expected Duration: 1.5 months**

4.3.1 Evaluation: Inverse Solving with Simple Microstructure The inverse solving system of §4.3 will be evaluated with three sources of data: synthetic images of computational phantoms, real images according to the 'full battery" of sensitivity tests described in §2.8.1, and real images of an animal brain with paired histology. The computational phantoms will be instances of the macrostructure model and may therefore be compared directly to the model output of the optimization. **Expected Duration: 2 months**

4.4 Microstructure Model. The microstructure model will include all the physical characteristics listed in §2.2 and §2.3, supporting multiple axon populations at different orientations, distributions of axon orientations within each population, different volume fractions for each population, distributions of axon caliber within each population, different exchange rates for each population, and distributions of diffusion coefficients within each population. If possible, a closed-form expression should be derived for the diffusion MR response in a chosen direction for a given instance of the model. The DifSim software [6] will be modified to simulate the diffusion response for arbitrary instances of this model, and will be used either to validate the closed-form diffusion response expression, or to learn an approximation to it using machine learning techniques. **Expected Duration: 3 months**

4.4.1 Demonstration of Microstructure Model Ambiguity. The candidate hypothesizes that there will exist different instances of the microstructure model that result in the same diffusion response for all diffusion gradients. An attempt will be made to demonstrate examples of ambiguous configurations, either using simulation or, preferably, a closed-form proof. In the event that connected equivalence classes of ambiguous configurations can be identified, equations describing these ambiguities will be derived. **Expected Duration: 1–2 months**

4.5 Inverse Solving: Macrostructure and Detailed Microstructure. The inverse solving implementation from §4.3 will be modified to incorporate the detailed microstructure model of §4.4. In the initialization phase, a single set of microstructure parameters (excluding axon orientation) will be assigned to each bundle in the macrostructure model to best fit the images for all voxels contained within the volume of only that bundle. The axon orientation at each position will be determined by the local trajectory fit by the macrostructure model. In addition to the macrostructure refinements listed above, microstructure refinements will include adjustments to any parameters in isolation and simultaneous adjustments of ambiguous parameters that result in no change in the diffusion response, as described in §4.4.1. The objective function will include the components specified above as well as regularization terms for smoothness of microstructure parameters along the length of each bundle and the geometrical constraints implied by conservation of axons along a bundle, as discussed in §2.3.1. **Expected Duration: 3 months**

4.5.1 Evaluation: Inverse Solving with Detailed Microstructure. The formal evaluation of the inverse solving system for the model with combined macrostructure and detailed microstructure will cover the full battery of sensitivity-testing datasets; one final candidate model will be generated for each input dataset with the system described in §4.5. In addition to comparing the final candidate models directly to each other for sensitivity tests, the utility of the detailed microstructure model will be quantified by comparing the images of the final candidate models from this evaluation to images of final candidates for the same datasets using the simple-microstructure inverse solver of §4.3 relative to the input images. Goodness of fit may be quantified by differences between a given set of images and the input set, and analyzed either as distributions or with visual demonstration of artifacts. **Expected Duration: 3 months**

4.6 Data Acquisition. As explained above, three sets of data will be required to isolate and test sensitivity to different sources of data variation at various stages throughout the research. While publicly available data may suffice for some of these needs, it is likely that custom acquisitions will need to be performed.

- To test for stability over variation due to image noise, approximately 5 scans repeated on the same subject with the same protocol, at high angular resolution and multiple b-values.
- To test for stability over variation due to acquisition parameters, approximately 5 scans repeated on the same subject with different protocols, varying the number of b-values, the angular and spatial resolution, and pulse sequence.
- To test for stability over inter-subject variability, approximately 10 scans of different subjects with a fixed protocol, at high angular resolution and multiple b-values.

Due to the large number of scans that may be necessary, plans will be initiated immediately to acquire data while other goals are being pursued. **Expected Duration: 9 months**

4.7 Optional Clinical Applications. The following studies applying the proposed model to medical research are optional and sufficiently time-consuming that they cannot all be included in a reasonable research plan. Informal versions of some, especially §4.7.3, may be useful for guidance during the main research activity. All studies below depend on a minimum progress of the main research track through at least §4.1.5.

4.7.1 Bundlewise Statistics. Bundlewise statistics, analogous to those defined in [18], will be defined for the macrostructure model, including volume, length, and FA distribution. The currently defined measures are based on DTI tractography and therefore do not consider multiple overlapping fiber populations; the macrostructure model, however, will include such regions. Schemes to correct for (or simply exclude) overlapping regions will be necessary for, e.g., the FA distribution computation. **Expected Duration: 1 month**

4.7.2 An Interactive Tool for Macrostructure Editing and Selection. Since it is expected that the model's automatically constructed "bundles" will not correspond to named WM fascicles, an interactive, graphical software tool to divide and select clusters will need to be developed for applications, such as the subsequent two, that require expert selection of particular fascicles. **Expected Duration: 2 months**

4.7.3 Bundlewise Statistics: Evaluation with CADASIL. Working from existing images of healthy controls and CADASIL patients, a domain expert will select corresponding structures from the model instances generated for each of them by either the forward-modeling process of §4.1.5 or the inverse-solving process of §4.3. Chosen statistics will be computed for each structure, and from these the populations will be statistically compared. **Expected Duration: 3 months**

4.7.4 HIV Study with Bundlewise Statistics. The same procedure from §4.7.3 will be repeated with existing images of large populations of healthy controls and HIV patients to identify structures affected by early-stage HIV infection. This study depends on the work up through §4.1.5. **Expected Duration: 5 months**

4.7.5 Fibers at Risk Longitudinal Study. A procedure will be defined to construct a continuous-valued FAR map on the mid-sagittal plane from a macrostructure model instance and a T1-weighted image showing lesions from multiple sclerosis by projecting the lesion image along WM bundles. Maps will be constructed from images that have already been collected of MS patients at various times through the progression of the disease. The FAR maps will be compared to maps of WM damage manually identified by a domain expert and automatically-computed maps of FA at the mid-sagittal plane. The study hypothesizes that FAR maps at earlier times will correlate with WM damage at later times, indicating predictive power of the FAR map. **Expected Duration: 5 months**

4.8 Timeline. Note that all expected durations listed above are estimates, though attempts have been made to include time estimates for preparing publications based on the various contributions.

- **July–December 2009**
  - Begin acquiring high-resolution data (§4.6)
  - Install existing software for Q-Ball tractography (§4.1.1)
  - Develop curve clustering algorithm, write up for ISMRM or a theory workshop (§4.1.2)
  - Implement macrostructure model, mostly from existing software (§4.1.3)
- **January–June 2010**
  - Evaluate macrostructure model, write up for Vis (§4.1.4)
  - Develop macrostructure adjustment (§4.1.5)
  - Develop macrostructure rendering software (§4.1.6)
  - Develop process to generate synthetic images from macrostructure model (§4.2)
  - Begin image-based evaluation of macrostructure model (§4.2.1)
  - Progress report to committee and departmental “proposal” presentation
- **July–December 2010**
  - Finish image-based evaluation of macrostructure model, write up for ISMRM (§4.2.1)
  - Begin chosen clinical applications (§4.7)
  - Develop inverse solver for macrostructure + simple microstructure (§4.3)
  - Evaluate inverse solver for macrostructure + simple microstructure, write up for MICCAI (§4.3.1)
- **January–June 2011**
  - Develop microstructure model (§4.4)
  - Demonstrate microstructure model ambiguity, write up for MICCAI (§4.4.1)
  - Begin developing inverse solver for macrostructure + detailed microstructure (§4.5)
- **July–December 2011**
  - Finish developing inverse solver for macrostructure + detailed microstructure (§4.5)
  - Evaluate inverse solver, write up for NeuroImage or MRM (§4.5.1)
  - Finish writing and defend dissertation

Literature Cited

[1] Andrew L. Alexander, Khader M. Hasan, Mariana Lazar, Jay S. Tsuruda, and Dennis L. Parker. Analysis of partial volume effects in diffusion-tensor MRI. *Magnetic Resonance in Medicine*, 45(5):770–780, 2001.

[2] Daniel C. Alexander. *Visualization and Processing of Tensor Fields*, chapter An introduction to computational diffusion MRI: the diffusion tensor and beyond, pages 83–106. Springer, Berlin, 2005.

[3] Daniel C. Alexander. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine*, 60(2):439–448, 2008.

[4] Yaniv Assaf, Tamar Blumenfeld-Katzir, Yossi Yovel, and Peter J. Basser. Axcaliber: A method for measuring axon diameter distribution from diffusion MRI. *Magnetic Resonance in Medicine*, 59(6):1347–1354, 2008.

[5] Yaniv Assaf, Raisa Z. Freidlin, Gustavo K. Rohde, and Peter J. Basser. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine*, 52(5):965–978, 2004.

[6] Gregory T. Balls and Lawrence R. Frank. A simulation environment for diffusion weighted MR experiments in complex media. *Magnetic Resonance in Medicine*, 62(3):771–778, 2009.

[7] Peter J. Basser, James Mattiello, and Denis LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, Jan 1994.

[8] Peter J. Basser, Sinisa Pajevic, Carlo Pierpaoli, Jeffrey Duda, and Akram Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44(4):625–632, 2000.

[9] Peter J. Basser and Carlo Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative diffusion tensor MRI. *Journal of Magnetic Resonance*, 111:209–219, 1996.

[10] T.E.J. Behrens, H. Johansen Berg, S. Jbabdi, M.F.S. Rushworth, and M.W. Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34(1):144–155, 2007.

[11] T.E.J. Behrens, M.W. Wollrich, M. Jenkinson, H. Johansen-Berg, R.G. Nunes, S. Clare, P.M. Matthews, J.M. Brady, and S.M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50(5):1077–1088, 2003.

[12] Jeffrey I. Berman, SungWon Chung, Pratik Mukherjee, Christopher P. Hess, Eric T. Han, and Roland G. Henry. Probabilistic streamline q-ball tractography using the residual bootstrap. *NeuroImage*, 39(1):215–222, 2008.

[13] Anders Brun, Hae-Jeong Park, Hans Knutsson, and Carl-Fredrik Westin. Coloring of DT-MRI fiber traces using laplacian eigenmaps. In *Proc. EUROCAST*, volume 2809 of *Lecture Notes in Computer Science*, pages 564–572, 2003.

[14] Yunmei Chen, Weihong Guo, Qingguo Zeng, Guojun He, B. Vemuri, and Yijun Liu. Recovery of intravoxel structure from HARD DWI. In *Proceedings of IEEE International Symposium on Biomedical Imaging*, pages 1028–1031, April 2004.

[15] Thomas G. Close, Jacques-Donald Tournier, Fernando Calamante, Leigh A. Johnston, Iven Mareels, and Alan Connelly. A software tool to generate simulated white matter structures for the assessment of fibre-tracking algorithms. *NeuroImage*, In Press, 2009.

[16] P. A. Cook, Y. Bai, S. Nedjati-Gilani, K. K. Seunarine, M. G. Hall, G. J. Parker, and D. C. Alexander. Camino: Open-source diffusion-MRI reconstruction and processing. In *Proc. 14th ISMRM*, 2006.

[17] I. Corouge, S. Gouttard, and G. Gerig. Towards a shape model of white matter fiber bundles using diffusion tensor mri. In *Proc. IEEE ISBI*, pages 344–347, April 2004.

[18] Stephen Correia, Stephanie Lee, Thom Voorn, David F. Tate, Robert Paul, Song Zhang, Stephen P. Salloway, Paul F. Malloy, and David H. Laidlaw. Quantitative tractography metrics of white matter integrity in diffusion-tensor MRI. *NeuroImage*, 42:568–581, 2008.

[19] Olivier Coulon, Daniel C. Alexander, and Simon R. Arridge. A regularization scheme for diffusion tensor magnetic resonance images. In *Proc. 17th IPMI*, volume 2082 of *Lecture Notes in Computer Science*, pages 92–105, Berlin, 2001. Springer.

[20] Maxime Descoteaux, Elaine Angelino, Shaun Fitzgibbons, and Rachid Deriche. Regularized, fast, and robust analytical Q-ball imaging. *Magnetic Resonance in Medicine*, 58(3):497–510, 2007.

[21] Maxime Descoteaux, Rachid Deriche, and Alfred Anwander. Deterministic and probabilistic Q-ball tractography: from diffusion to sharp fiber distribution. Technical Report 6273, INRIA Sophia Antipolis, Valbonne Sophia Antipolis, France, August 2007.

[22] Z. Ding, J.C. Gore, and A.W. Anderson. Case study: reconstruction, visualization and quantication of neuronal fiber pathways. In *Proc. IEEE Vis.*, pages 453–456, 2001.

[23] P. Fillard, C. Poupon, and J.-F. Mangin. Free-spin tracking: A novel global tractography algorithm. In *Proc. 17th ISMRM*, 2009.

[24] Jon E. Grant, Stephen Correia, and Thea Brennan-Krohn. White matter integrity in kleptomania: A pilot study. *Psychiatry Research: Neuroimaging*, 147(2–3):233–237, Oct 2006.

[25] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C.J. Honey, V.J. Wedeen, and O. Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7), 2008.

[26] Christopher P. Hess, Pratik Mukherjee, Eric T. Han, Duan Xu, and Daniel B. Vigneron. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. *Magnetic Resonance in Medicine*, 56(1):104–117, 2006.

[27] Tim Hosey, Guy Williams, and Richard Ansorge. Inference of multiple fiber orientations in high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 54(6):1480–1489, 2005.

[28] M. Jenkinson, P.R. Bannister, J.M. Brady, and S.M. Smith. Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage*, 17(2):825–841, 2002.

[29] Derek K. Jones, Marco Catani, Carlo Pierpaoli, Suzanne J.C. Reeves, Sukhwinder S. Shergill, Michael O'Sullivan, Pasha Golesworthy, Phillip McGuire, Mark A. Horsfield, Andrew Simmons, Steven C.R. Williams, and Robert J. Howard. Age effects on diffusion tensor magnetic resonance imaging tractography measures of frontal cortex connections in schizophrenia. *Human Brain Mapping*, 27(3):230–238, 2006.

[30] D.K. Jones, X.A. Chitnis, D. Job, P.L. Khong, L.T. Leung, S. Marenco, S.M. Smith, and M.R. Symms. What happens when nine different groups analyze the same DT-MRI data set using voxel-based methods? In *Proc. 15th ISMRM*, page 74, 2007.

[31] Jason A. Kaufman, Eric T. Ahrens, David H. Laidlaw, Song Zhang, and John M. Allman. Anatomical analysis of an aye-aye brain (daubentonia madagascariensis, primates: Prosimii) combining histology, structural magnetic resonance imaging, and diffusion-tensor imaging. *Anatomical Record*, 287A(1):1026–1037, November 2005.

[32] Minue J. Kim, In Kyoon Lyoo, Seog Ju Kim, Minyoung Sim, Namkug Kim, Namhee Choi, Do-Un Jeong, Julie Covell, and Perry F. Renshaw. Disrupted white matter tract integrity of anterior cingulate in trauma survivors. *Neuroreport*, 16(10):1049–1053, Jul 2005.

[33] Marek Kubicki, Robert McCarley, Carl-Fredrik Westin, Hae-Jeong Park, Stephan Maier, Ron Kikinis, Ferenc A. Jolesz, and Martha E. Shenton. A review of diffusion tensor imaging studies in schizophrenia. *Journal of Psychiatric Research*, 41(1-2):15–30, 2007.

[34] Mariana Lazar, David M. Weinstein, Jay S. Tsuruda, Khader M. Hasan, Konstantinos Arfanakis, M. Elizabeth Meyerand, Benham Badie, Howard A. Rowley, Victor Haughton, Aaron Field, and Andrew L. Alexander. White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18(4):306–321, 2003.

[35] A. Leemans, J. Sijbers, M. Verhoye, A. Van der Linden, and D. Van Dyck. Mathematical framework for simulating diffusion tensor MR neural fiber bundles. *Magnetic Resonance in Medicine*, 53(4):944–953, 2005.

[36] Kelvin O. Lim, Maj Hedehus, Michael Moseley, Alexander de Crespigny, Edith V. Sullivan, and Adolf Pfefferbaum. Compromised white matter tract integrity in schizophrenia inferred from diffusion tensor imaging. *Archives of General Psychiatry*, 56(4):367–374, Apr 1999.

[37] J.-F. Mangin, C. Poupon, Y. Cointepas, D. Rivière, D. Papadopoulos-Orfanos, C. A. Clark, J. Régis, and D. Le Bihan. A framework based on spin glass models for the inference of anatomical connectivity from diffusion-weighted MR data — a technical review. *NMR in Biomedicine*, 15(7–8):481–491, 2002.

[38] Jennifer A. McNab, Saâd Jbabdi, Sean C.L. Deoni, Gwenaëlle Douaud, Timothy E.J. Behrens, and Karla L. Miller. High resolution diffusion-weighted imaging in fixed human brain using diffusion-weighted steady state free precession. *NeuroImage*, 46:775–785, 2009.

[39] Jadrian Miles, Ronald Cohen, and David H. Laidlaw. Tradeoffs in supersampling of dti metrics. In *Proceedings of ISMRM*, 2009.

[40] Bart Moberts, Anna Vilanova, and Jarke J. van Wijk. Evaluation of fiber clustering methods for diffusion tensor imaging. In *Proc. IEEE Visualization*, pages 65–72, 2005.

[41] Susumu Mori and Peter C. M. van Zijl. Fiber tracking: Principles and strategies – a technical review. *NMR in Biomedicine*, 15(7–8):468–480, 2002.

[42] Susumu Mori and Jiangyang Zhang. Principles of diffusion tensor imaging and its applications to basic neuroscience research. *Neuron*, 51:527–539, 2006.

[43] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Proc. 15th NIPS*, pages 849–856. MIT Press, 2001.

[44] Markus Nilsson, Jimmy Lätt, Emil Nordh, Ronnie Wirestam, Freddy Ståhlberg, and Sara Brockstedt. On the effects of a varied diffusion time in vivo: is the diffusion in white matter restricted? *Magnetic Resonance Imaging*, 27(2):176–187, 2009.

[45] Kenichi Oishi, Karl Zilles, Katrin Amunts, Andreia Faria, Hangyi Jiang, Xin Li, Kazi Akhter, Kegang Hua, Roger Woods, Arthur W. Toga, G. Bruce Pike, Pedro Rosa-Neto, Alan Evans, Jiangyang Zhang, Hao Huang, Michael I. Miller, Peter C.M. van Zijl, John Mazziotta, and Susumu Mori. Human brain white matter atlas: Identification and assignment of common anatomical structures in superficial white matter. *NeuroImage*, 43(3):447–457, 2008.

[46] Ofer Pasternak, Nir Sochen, and Yaniv Assaf. *Visualization and Processing of Tensor Fields*, chapter Variational Regularization of Multiple Diffusion Tensor Fields, pages 165–176. Springer, Berlin, 2005.

[47] M. Perrin, C. Poupon, Y. Cointepas, B. Rieul, N. Golestani, C. Pallier, D. Rivière, A. Constantinesco, D. Le Bihan, and J.-F. Mangin. Fiber tracking in q-ball fields using regularized particle trajectories. In *Proc. 19th IPMI*, volume 3565 of *Lecture Notes in Computer Science*, pages 52–63, Berlin, 2005. Springer.

[48] Carlo Pierpaoli and Peter J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36:893–906, Dec 1996.

[49] C. Poupon, C. A. Clark, V. Frouin, J. Rgis, I. Bloch, D. Le Bihan, and J. F. Mangin. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage*, 12(2):184–195, 2000.

[50] C. Poupon, F. Poupon, L. Allirol, and J.-F. Mangin. A database dedicated to anatomo-functional study of human brain connectivity. In *Proc. 12th HBM*, 2006.

[51] Arish A. Qazi, Alireza Radmanesh, Lauren O'Donnell, Gordon Kindlmann, Sharon Peled, Stephen Whalen, Carl-Fredrik Westin, and Alexandra J. Gobly. Resolving crossings in the corticospinal tract by two-tensor streamline tractography: Method and clinical assessment using fMRI. *NeuroImage*, 2009. in press.

[52] J. Ritchie. On the relation between fibre diameter and conduction velocity in myelinated nerve fibres. *Proc. Royal Society of London, Series B, Biological Sciences*, 217(1206):29–35, 1982.

[53] Denis Rivière, Jean-François Mangin, Dimitri Papadopoulos-Orfanos, Jean-Marc Martinez, Vincent Frouin, and Jean Régis. Automatic recognition of cortical sulci of the human brain using a congregation of neural networks. *Medical Image Analysis*, 6(2):77–92, 2002.

[54] P.E. Roland, Stefan Geyer, Katrin Amunts, Thorsten Schormann, Axel Schleicher, Aleksander Malikovic, and Karl Zilles. Cytoarchitectural maps of the human brain in standard anatomical space. *Human Brain Mapping*, 5(4):222–227, 1997.

[55] Peter Savadjiev, Jennifer S.W. Campbell, G. Bruce Pike, and Kaleem Siddiqi. 3D curve inference for diffusion MRI regularization and fibre tractography. *Medical Image Analysis*, 10(5):799–813, 2006.

[56] Jeremy D. Schmahmann and Deepak N. Pandya. *Fiber Pathways of the Brain*. Oxford University Press, New York, 2006.

[57] John Schreiner, Carlos E. Scheidegger, and Cláudio T. Silva. High-quality extraction of isosurfaces from regular and irregular grids. *IEEE Transactions onVisualization and Computer Graphics*, 12(5):1205–1212, 2006.

[58] David W. Shattuck, Stephanie R. Sandor-Leahy, Kirt A. Schaper, David A. Rottenberg, and Richard M. Leahy. Magnetic resonance image tissue classification using a partial volume model. *NeuroImage*, 13(5):856–876, 2001.

[59] Jack H. Simon, Song Zhang, David H. Laidlaw, David E. Miller, Mark Brown, John Corboy, and Jeffrey Bennett. Identification of fibers at risk for degeneration by diffusion tractography in patients at high risk for MS after a clinically isolated syndrome. *Journal of Magnetic Resonance Imaging*, October 2006.

[60] Greg J. Stanisz, Graham A. Wright, R. Mark Henkelman, and Aaron Szafer. An analytical model of restricted diffusion in bovine optic nerve. *Magnetic Resonance in Medicine*, 37(1):103–111, 1997.

[61] J-Donald Tournier, Fernando Calamante, and Alan Connelly. Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage*, 35(4):1459–1472, 2007.

[62] David S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

[63] David S. Tuch, Timothy G. Reese, Mette R. Wiegell, and Van J. Wedeen. Diffusion MRI of complex neural architecture. *Neuron*, 40(5):885–895, 2003.

[64] Ruopeng Wang, Thomas Benner, Alma Greg Sorensen, and Van Jay Wedeen. Diffusion Toolkit: A software package for diffusion imaging data processing and tractography. In *Proc. 15th ISMRM*, 2007.

[65] Zhizhou Wang, Baba C. Vemuri, Yunmei Chen, and Thomas H. Mareci. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Transactions on Medical Imaging*, 23(8):930–939, August 2004.

[66] V.J. Wedeen, R.P. Wang, J.D. Schmahmann, T. Benner, W.Y.I. Tseng, G. Dai, D.N. Pandya, P. Hagmann, H. D'Arceuil, and A.J. de Crespigny. Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *NeuroImage*, 41(4):1267–1277, 2008.

[67] Song Zhang and David H. Laidlaw. Hierarchical clustering of streamtubes. Technical Report CS-02-18, Brown University Computer Science Department, August 2002.

[68] Y. Zhang, M. Brady, and S. Smith. Segmentation of brain mr images through a hidden Markov random field model and the expectation maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1):45–57, 2001.

[69] Wenjin Zhou and David H. Laidlaw. An analytical model of diffusion and exchange of water in white matter from diffusion-MRI and its application in measuring axon radii. In *Proc. 17th ISMRM*, page 263, 2009.

[70] Z. Zuo, X. Wang, N. Chen, J. Luo, Y. Zhuo, D. Shen, K. Li, and R. Xue. Establishing probabilistic chinese human brain templates using HAMMER elastic registration algorithm. In *Proc. 17th ISMRM*, 2009.