Characterization of a spinal cord diffusion tensor imaging pipeline with pathological spine data

Vignesh Sivan, University of Waterloo
Michael Hardisty, Sunnybrook Research Institute
Alexander Wong, University of Waterloo
Stewart McLachlin, University of Waterloo
Email: vsivan@uwaterloo.ca

Abstract

Fiber tractography from magnetic resonance (MR) diffusion tensor imaging (DTI) enables the visualization of white matter bundles. In the presence of pathology, these bundles can be distorted and disconnected, which can reveal clinically significant information about the nature of the underlying pathology. This work studies DTI in the spine in the presence of pathology. A spine DTI pipeline that was developed in an earlier study is evaluated against the pathological data. We study the challenges of adapting the pipeline to pathological spine data, where MRI artifacts and significant distortion in cord shape and contrast from pathology make automated cord segmentation and registration extremely challenging. Moreover, we identify challenges with processing highly anisotropic MRI volumes and the implications this has on DTI processing. Heuristics are developed to handle these issues and are incorporated into the pipeline. Finally, visualizations of the tractography streamlines are generated and the impact of pathology on the streamline trajectories is briefly discussed, awaiting clinical validation.

1 Introduction

Fiber tractography has been studied extensively in the brain, where it has been particularly influential in illuminating the connectivity of different regions in the brain [1]. While there are fewer studies in the spine, spinal cord fiber tractography visualizations have been shown to have clinical significance when pathology is present [2]. Of interest to the work presented in this paper is the use case of pre-surgical planning, where tractography can illuminate regions of interest prior to surgery. The work in this paper builds upon earlier work where an automated spinal cord DTI pipeline was developed and used to show the correspondence of fiber tractography streamlines with the underlying anatomy [3]. This earlier work used a dataset of healthy subjects and the MR imaging parameters were controlled and consistent across all subjects.

The primary contribution of the current study is to extend the aforementioned automated spinal cord DTI pipeline to a dataset of clinical cases with pathology. Complementary goals are to: characterize the dataset, and understand what imaging parameters of the pathological MRI volumes make DTI processing more robust. The short-term goal of this work is to inform MR acquisition parameters for the collection of spinal cord DTI data for future work. Finally, the broader goal is to use spinal cord tractography visualizations to improve patient care, enabling better diagnosis, assessment, and treatment planning.

2 Dataset

The dataset consists of retrospectively collected spine MR-DWI acquisitions from 13 subjects. The images were acquired though collaboration with Synaptive Medical. The image volumes were obtained from a variety of vendors with varying scanning parameters informed by clinical indication. The pathologies vary between subjects, with all subjects showing neurological symptoms. Table 1 summarizes some salient imaging parameters while 2 describes the voxel spacing of the images. Shown in Figure 1 is a sagittal slice from the dataset, where there is evidence of pathology from the 4th - 7th cervical vertebrae (C4-C7), which is evident from the morphological changes, change in tissue contrast, and the lack of a clear definition of cerebrospinal fluid (which appears white) in that region.

Table 1: MRI acquisition properties of the dataset.

Number Of Studies	Magnetic Field Strength	DMRI Directions	Contrast
6	3T	35	T1
3	1.5 T	26	T1
1	1.5 T	25	T1
1	1.5 T	31	T1
2	1.5 T	21	T2

Table 2: Image voxel spacing of the dataset.

Number Of Studies	Anatomical Voxel Spacing (mm)	DMRI Voxel Spacing (mm)
8	1x1x1	2x2x2
2	0.35x0.35x4	0.85x0.85x4.5
2	4x0.42x0.42	1x1x3
1	1x1x1	2x2x3

3 Methods

3.1 Study Design

This paper had two goals. The first was to characterize the performance of the pipeline implemented in an earlier work on spinal cord MRI with pathology. This step will be referred to as 'pipeline testing'. Next, we investigated different approaches to improving the performance of the pipeline on this new dataset. This phase will be refered to as 'pipeline improvements'.

3.2 Pipeline Testing

The steps of the aforementioned automated pipeline is briefly described here. Given a structural MRI volume (T1, T2) and a diffusion weighted MRI (DWI) volume, the pipeline first segments the spinal cord region in both the structural and diffusion volumes. The segmentations are used to inform a deformable registration from the diffusion space to the structural space. Next, 2nd order diffusion tensors are fit to the DWI to create DTI volumes which contain 3D diffusivity tensors at each voxel location. The diffusivities form is shown in equation 1. The elements of the tensor are computed by solving a system of linear equations obtained from the attenuation at least of six diffusion gradients. The pipeline uses weighted least squares to solve this system of equations. Finally, tractography is used to estimate fibre bundles using an unscented kalman filter approach described in [4].

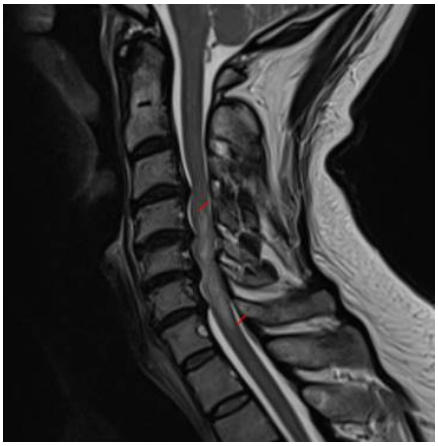

Fig. 1: Spinal cord MRI with pathology at C4-C7.

$$D = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} \quad (1)$$

The pipeline leverages open source software, specifically, the Spinal Cord Toolbox (SCT), which has tools for automatic spinal cord segmentation and also has a spinal cord atlas [5]. SCT also provides a thin wrapper for ANTs, which implements symmetric normalization registration. 3D Slicer's tractography tools are used to perform tensor transformations, streamline estimation and visualization.

The performance of the pipeline was tested against imaging from all subjects. The pipeline was initially run with parameters determined through optimization of performance on healthy volunteer subjects.

3.3 Pipeline Extension

Three steps (segmentation, registration, and tract generation) of the spine MRI DTI tractography pipeline were varied to investigate potential improvement of the performance on the dataset containing clinical cases with pathology. Segmentation: The segmentation of the spinal cord region was varied by considering 3 different algorithms, a 2D deep convolutional neural network (Deepseg2D) a 3D deep convolutional neural network (Deepseg3D) and the previously used propagated segmentation algorithm (Propseg). Registration: The registration steps were adjusted by cropping input volumes to only the spinal cord regions, thereby focusing registration on the relevant features. Further diffusion and structural imaging were directly registered, and the inputs were downsampled. Tract Generation: The stopping criteria used to terminate streamlines was relaxed to allow more weakly connected regions to have streamlines generated. This could potentially be advantageous in regions with pathology because of disruptions in the diffusion signal.

3.3.1 Segmentation

SCT offers two different flavors of spinal cord segmentation. An earlier iteration of the segmentation algorithm, Propseg, propagates a mesh along the image and uses energy minimization to find the ideal structure [6]. A more recent algorithm, called Deepseg, uses a U-Net to segment the cord region [7]. There are two variants of the Deepseg approach: Deepseg2d uses a 2D convolutional neural network and Deepseg3d uses 3D convolutions. The first iteration of the pipeline predates Deepseg. So a natural opportunity to extend the pipeline is to experiment with the newer automatic segmentation algorithm.

To test the performance of segmentation algorithms, we also generate ground truth segmentations for each of the structural scans and test the performance of SCT's Propseg, Deepseg2d and Deepseg3d algorithms. The dice coefficient is computed for each of these generated volumes against the ground truth and the segmentation performance is also qualitatively evaluated.

3.3.2 Registration

Registration is a challenging problem especially when there is a lot of distortion, which we expect to find in the pathological dataset. We, therefore, plan to test a variety of approaches designed to improve registration.

One approach to improve registration is to crop the region outside the cord including the brain. This approach is suggested by the developers of SCT. Fundamentally, the justification for doing this is that it provides a way to bias the registration algorithm towards focusing on the cord region.

Another approach is to register the structural MRI and the diffusion MRI directly with one another. In the earlier pipeline, this was avoided; the diffusion MRI was first registered to SCT's spinal cord atlas, with the atlas then registered to structural MRI scan. The motivation for doing this was to avoid biasing the tractography streamlines with information about the anatomy, as the goal was to measure the correspondence with the anatomy. As this was not the goal of the current work, we test the effect of doing a direct registration between the structural and diffusion MRI.

Finally, owing to the very large anisotropy in the pathological dataset, we experiment with downsampling the structural MRI so that it has the same resolution as the diffusion MRI. We hypothesize that this might improve registration performance, as this will force the images to have similar spatial resolution.

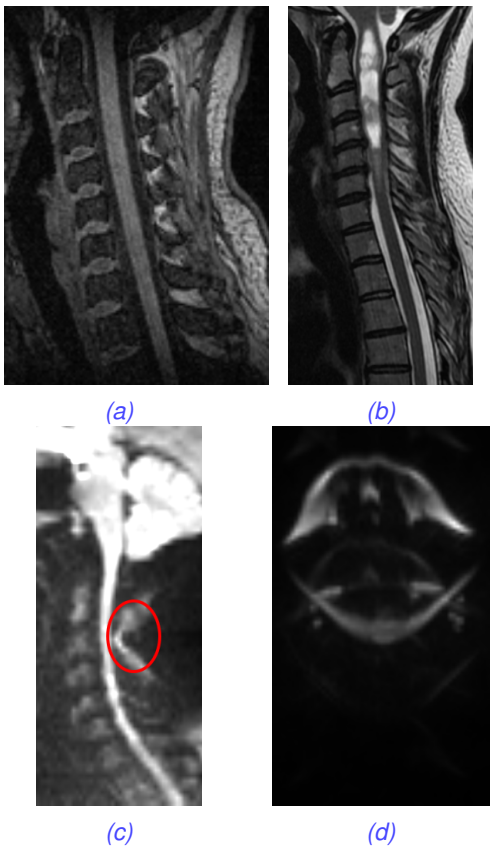

Fig. 2: Study a) is a representative examples of studies that work fine with the pipeline in its initial state. Studies b-d do not register well with the initial pipeline parameters due a) large spinal cord pathology c) have some anatomical feature that makes registration difficult (circled), d) has an MRI artifact

3.3.3 Tractography

We also experiment with parameters of the unscented Kalman filter tractography algorithm. In a previous work, we fixed the tractography stopping threshold at 0.25, which describes the minimum fractional anisotropy (FA) of the diffusion in a voxel we expect for it to be considered part of a fiber bundle. This threshold is set to ensure that spurious tracts are not found. Owing to the high anisotropy and also owing to pathology, we expect the FA values to be lower. We, therefore, experiment with removing this stopping threshold to identify regions that might be weakly connected.

4 Results

4.1 Pipeline Testing

Running the unmodified pipeline on the pathological dataset did not result in the successful production of MRI-DTI tractography in any studies. 7 of the 13 studies, failed at the registration step. The imaging that was successfully registered has isotropic voxels and had only modest morphological changes cause by the pathology. An example of a MRI slice that performs well is shown next to a volume that fails to register with initial pipeline parameters. After analyzing the results, we identified the reasons why the registration step fails. The reasons for not registering can be grouped into two categories: 1) MR artifact, 2) pathology or anatomical features. Table 3 summarizes the distribution of studies failing to register. It should be noted that if both factors are present then the factor we judge to be the greatest contributor is recorded.

We discuss the modifications needed to make the pipeline work for this pathological spine data in the rest of dataset in the next section.

Table 3: Results when running the unmodified pipeline

Reason not registering	Number of studies
MR Artifact	2
Pathology	5

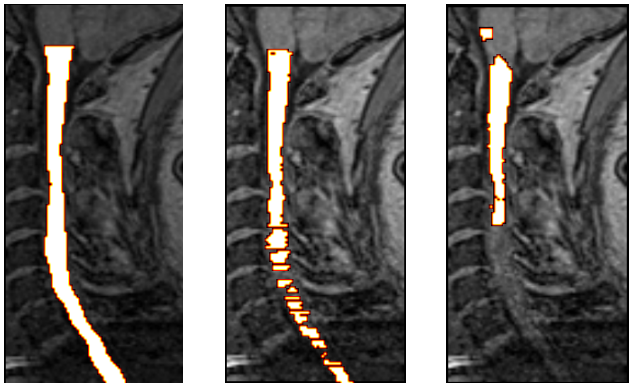

(a) Propseg *(b)* Deepseg2d *(c)* Deepseg3d

Fig. 3: Spinal cord segmentation using a) Propseg, b) Deepseg 2D and c) Deepseg 3D.

4.2 Pipeline Extension

4.2.1 Segmentation

Shown in Table 4 are the mean Dice coefficients between the ground truth segmentation and the segmentation predicted by the various approaches. While the dice coefficients are generally high, we observed that there was no clear choice for all studies, as the segmentation performance varied across each study. Shown in Figure 3 is an example of a study where Propseg outperforms all other segmentation algorithms. More generally, we found that Deepseg3D was more susceptible to contrast changes brought about by pathology (as can be seen in figure 3), but is also able to recover from discontinuity better than the other approaches.

Table 4: Cord segmentation mean and variance of the SCT segmentation algorithms.

Segmentation algorithm	Mean dice
Propseg	0.90
Deepseg 2D	0.91
Deepseg 3D	0.88

4.2.2 Registration

We test a variety of strategies to make registration perform better. Table 5 provides qualitative results of how each modification had an impact on cord registration. Using the techniques, we were able to get 6 out of the 7 studies that could not register with the previous pipeline registering. The study shown in Figure 2d) was unable to register, due to the extreme MRI artifact present.

Based on these results, we can surmise that direct registration of structural and diffusion MRI helps the registration algorithm. This makes sense as the two images share some mutual information that is not present in the SCT atlas, so this information cannot be used by the registration algorithm to perform matching.

Based on the results, direct registration of structural MRI to diffusion is insufficient. We found that when the field of view is drastically different between the diffusion MRI and the structural MRI it is very difficult to correctly register the two volumes. In these cases, we found that an affine transformation of the structural to the diffusion space, which also down samples the structural MRI, helps the registration algorithm to converge.

Finally, it should be noted that convergence of registration does not imply that the registration was perfect. The presence of artifacts still affects the registration, as can be seen in Figure 4, which shows a registered DMRI with significant registration artifacts.

Table 5: Qualitative evaluation of techniques to improve registration

Technique	Effect on Registration
Crop outside cord region	No effect
Register structural MRI directly to diffusion	Allows registration to converge for 5 studies
Downsample structural MRI to diffusion spacing	Allows registration to converge for 2 studies when paired with direct registration of structural MRI and diffusion MRI

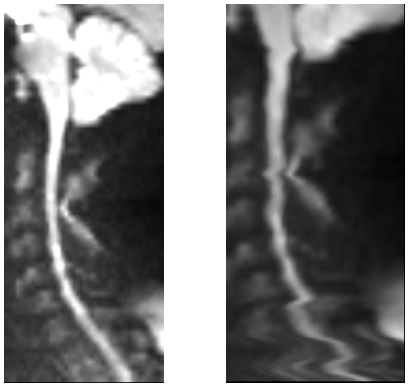

(a) *(b)*

Fig. 4: Spinal cord diffusion MRI a) before and b) after registration. Note the significant registration artifacts.

4.2.3 Tractography Threshold

Once we were able to register the diffusion MRI to the structural volumes, we generated tractography streamlines. We found, perhaps unsurprisingly, that the pathology from different subjects affected the connectivity of white matter bundles. Shown in Figure 5 are the tractography streamlines produced by modifying the stopping threshold. We observe that they produce more spurious tracts and the pipeline lacks functionality to find optimal tract thresholds.

5 Conclusion and Future Work

In this work we tested a spine DTI pipeline against a clinical dataset with pathology and we show that the prior pipeline was insufficient to adapt to this challenging dataset. Registration, in particular, was the step with the worst performance. Registration was likely particularly challenging because of the inherent mobility within the spine. Manually testing strategies to improve performance: segmentation (Deepseg2D, Deepseg3D), registration preprocessing, and tractography parameters demonstrated improvements in the pipeline performance and we developed some heuristics that can be incorporated into the pipeline. However, registration accuracy was still suboptimal. Future work will focus on improving registration for clinical evaluation and eventual translation of this work.

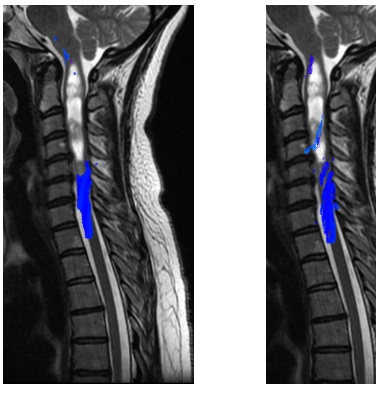

(a) *(b)*

Fig. 5: Spinal cord tractography with a) strict and b) lax FA thresholds.

References

[1] K. H. Maier-Hein, P. F. Neher, J.-C. Houde, M.-A. Côté, E. Garyfallidis, J. Zhong, M. Chamberland, F.-C. Yeh, Y.-C. Lin, Q. Ji, W. E. Reddick, J. O. Glass, D. Q. Chen, Y. Feng, C. Gao, Y. Wu, J. Ma, R. He, Q. Li, C.-F. Westin, S. Deslauriers-Gauthier, J. O. O. González, M. Paquette, S. St-Jean, G. Girard, F. Rheault, J. Sidhu, C. M. W. Tax, F. Guo, H. Y. Mesri, S. Dávid, M. Froeling, A. M. Heemskerk, A. Leemans, A. Boré, B. Pinsard, C. Bedetti, M. Desrosiers, S. Brambati, J. Doyon, A. Sarica, R. Vasta, A. Cerasa, A. Quattrone, J. Yeatman, A. R. Khan, W. Hodges, S. Alexander, D. Romascano, M. Barakovic, A. Auría, O. Esteban, A. Lemkaddem, J.-P. Thiran, H. E. Cetingul, B. L. Odry, B. Mailhe, M. S. Nadar, F. Pizzagalli, G. Prasad, J. E. Villalon-Reina, J. Galvis, P. M. Thompson, F. D. S. Requejo, P. L. Laguna, L. M. Lacerda, R. Barrett, F. Dell'Acqua, M. Catani, L. Petit, E. Caruyer, A. Daducci, T. B. Dyrby, T. Holland-Letz, C. C. Hilgetag, B. Stieltjes, and M. Descoteaux, "The challenge of mapping the human connectome based on diffusion tractography," *Nature Communications*, vol. 8, no. 1, p. 1349, Nov. 2017, bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Computational biology and bioinformatics;Medical research;Nervous system;Neuroscience Subject_term_id: computational-biology-and-bioinformatics;medical-research;nervous-system;neuroscience. [Online]. Available: https://www.nature.com/articles/s41467-017-01285-x

[2] B. M. Ellingson, N. Salamon, D. C. Woodworth, and L. T. Holly, "Degree of subvoxel spinal cord compression measured with super-resolution tract density imaging (TDI) correlates with neurological impairment in cervical spondylotic myelopathy," *Journal of neurosurgery. Spine*, vol. 22, no. 6, pp. 631–638, Jun. 2015. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5565223/

[3] S. McLachlin, J. Leung, V. Sivan, P.-O. Quirion, P. Wilkie, J. Cohen-Adad, C. M. Whyne, and M. R. Hardisty, "Spatial correspondence of spinal cord white matter tracts using diffusion tensor imaging, fibre tractography, and atlas-based segmentation," *Neuroradiology*, vol. 63, no. 3, pp. 373–380, Mar. 2021.

[4] J. G. Malcolm, M. E. Shenton, and Y. Rathi, "Filtered Multitensor Tractography," *IEEE Transactions on Medical Imaging*, vol. 29, no. 9, pp. 1664–1675, Sep. 2010, conference Name: IEEE Transactions on Medical Imaging.

[5] J. Cohen-Adad, B. De Leener, M. Benhamou, D. Cadotte, D. Fleet, A. Cadotte, M. Fehlings, J. Paquette, W. Thong, M. Taso, L. Collins, V. Callot, and V. Fonov, "Spinal Cord Toolbox: an open-source framework for processing spinal cord MRI data," *Proceedings of the 20th Annual Meeting of OHBM, Hamburg, Germany*, Jun. 2014.

[6] J. Cohen-adad, "PropSeg: Automatic spinal cord segmentation method for MR images using propagated deformable models." [Online]. Available: https://www.academia.edu/17482535/PropSeg_Automatic_spinal_cord_segmentation_method_for_MR_images_using_propagated_deformable_models

[7] C. Gros, B. De Leener, A. Badji, J. Maranzano, D. Eden, S. M. Dupont, J. Talbott, R. Zhuoquiong, Y. Liu, T. Granberg, R. Ouellette, Y. Tachibana, M. Hori, K. Kamiya, L. Chougar, L. Stawiarz, J. Hillert, E. Bannier, A. Kerbrat, G. Edan, P. Labauge, V. Callot, J. Pelletier, B. Audoin, H. Rasoanandrianina, J.-C. Brisset, P. Valsasina, M. A. Rocca, M. Filippi, R. Bakshi, S. Tauhid, F. Prados, M. Yiannakas, H. Kearney, O. Ciccarelli, S. Smith, C. A. Treaba, C. Mainero, J. Lefeuvre, D. S. Reich, G. Nair, V. Auclair, D. G. McLaren, A. R. Martin, M. G. Fehlings, S. Vahdat, A. Khatibi, J. Doyon, T. Shepherd, E. Charlson, S. Narayanan, and J. Cohen-Adad, "Automatic segmentation of the spinal cord and intramedullary multiple sclerosis lesions with convolutional neural networks," *NeuroImage*, vol. 184, pp. 901–915, Jan. 2019.