

GLOBALLY OPTIMAL SPINAL CORD SEGMENTATION USING A MINIMAL PATH IN HIGH DIMENSIONS

*Jeremy Kawahara**

Chris McIntosh[†]*

Roger Tam[‡]

*Ghassan Hamarneh**

*Medical Image Analysis Lab., School of Computing Science, Simon Fraser University

[†]Radiation Medicine Program, Princess Margaret Cancer Centre, University Health Network

[‡]Department of Radiology & MS/MRI Research Group, University of British Columbia

ABSTRACT

Spinal cord segmentation is an important step to empirically quantify spinal cord atrophy that can occur in neurological diseases such as multiple sclerosis (MS). In this work, we propose a novel method to find the globally optimal segmentation of the spinal cord using a high dimensional minimal path search. The spinal cord cross-sectional shapes are represented using principal component analysis (in the probability simplex) which captures most of spinal cord’s axial cross-sectional variation and partial volume effects. We propose modifications to the A* minimal path search algorithm that drastically reduce the required memory and run-time to make our high dimensional minimal path optimization computationally feasible. Finally, we validate our results over five vertebrae levels of both healthy and MS clinical MR volumes (20 volumes total) and show improvements on volume agreement with expert segmentations and less user interaction when compared to current state-of-the-art methods.

Index Terms— Segmentation; spinal cord analysis; minimal paths; principal component analysis

1. INTRODUCTION

Multiple Sclerosis (MS) studies have shown that spinal cord atrophy contributes strongly to a patient’s physical disability (e.g. walking speed). Accurately segmenting the spinal cord increases our ability to quantify the physical atrophy of the cord which may allow us to better monitor both the progression of the disease and the effectiveness of treatments [1]. However, spinal cord segmentation is a challenging problem as the signal changes over the length of the cord, images are often noisy due to patient movement, the boundary of the cord is not well defined in regions where it contracts the wall of the spinal canal, and the cord’s shape can change as the disease progresses [2]. As well, a typical cord scan, even at high resolution, has many partial volume voxels relative to the total number representing the cord. While manual segmentations

by experts are generally considered to be accurate, they are more susceptible to operator variability and can take a considerable amount of operator time per scan [2]. This makes semi-automated methods desirable. There are a number of related works on the problem of spinal cord segmentation [2–5].

Schmit and Cole [3] performed segmentation using 3D seeded region growing and observed the narrowing of the spinal cord after injury in MRI scans. Horsfield et al. [4] used an active model of the cord surface where a user marked the center line of the cord on representative slices. McIntosh and Hamarneh [5] used locally optimal 3D deformable organisms to segment the cord. This was extended by McIntosh et al. [2] where the spinal cord segmentation optimization problem was split into two sequential steps: medial path finding (based on two 2D livewires, which can require manual corrections) followed by a medial-guided spinal crawler with cross-sectional shape fitting. Our proposed high dimensional path optimization treats these two steps in a single optimization and guarantees a global solution (from only two seed points).

Globally minimal paths to segment tubular objects were introduced by Li and Yezzi [6] who applied a 4D search to find the spatial coordinates and radius of 3D tubular objects between two user entered seed points. Poon et al. [7] used a similar minimal path approach to interactively segment 2D vessels by moving to a 3D space and searching over the radii.

Our proposed method finds the minimal path in high dimensions between two user entered seed points. This minimal path represents the globally optimal segmentation of the spinal cord. This work is most similar to [6] and [7]. Our contribution is the extension of these methods to higher dimensional space where we demonstrate a novel method to segment the spinal cord using an A* minimal path optimizer. Directly extending these methods to higher (than 4) dimensions can become infeasible due to the high computational needs. Our approach ensures that the globally optimal path can be found by automatically ignoring many suboptimal paths using a heuristic and lowers our space requirements by not requiring an explicit graph representation prior to searching. As well, we describe a probabilistic shape representation (based on principal component analysis, PCA) with implicit regular-

JK, RT, and GH were partially supported by NSERC and Biogen Idec Canada. CM was supported by the Canadian Breast Cancer Foundation and the Canadian Cancer Society Research Institute.

ization capable of capturing partial volume effects that can be extended to arbitrarily high dimensions.

PCA has been used as an efficient method to represent high dimensional shapes by representing complex objects using only a few principal components. Cremers et al. [8] encode an arbitrary shape where each pixel is given a probability it is within the shape and apply PCA to this model. To ensure that PCA does not leave the range of valid probability values (i.e. the unit simplex) Changizi and Hamarneh [9] propose the Isometric Log-Ratio (ILR) transformation where they perform PCA in the ILR space to ensure the probability vectors stay in the valid simplex space. Andrews et al. [10] perform PCA in this ILR space and incorporate this shape prior directly into an energy function for image segmentation. In our work we adopt the ILR-based PCA in the high dimensional minimal path search.

2. METHODS

In this section we describe our PCA based shape representation, followed by a description of our energy function, and finally we give an overview of our minimal path optimizer.

2.1. Probabilistic PCA based Shape Representation

We learn a model of the shape of the axial cross section of the spinal cord using PCA. This PCA based model allows us to learn a mean shape and deform it in ways that respect the main modes of variation found within the training data. By manipulating the eigenweights of the three principal components (PC), we found our model can theoretically describe a mean Jaccard similarity coefficient of 0.874 and a mean area similarity of 97.18% (Section 3 for similarity measures) of the variation found within the training slices. With our PCA based shape model, we represent a single axial slice of the cord using the following six dimensions: the x, y, z spatial coordinates describing the centroid of the cross section, and three PC weight values for describing its shape. We restrict the PC weights to only vary by ± 3 standard deviations.

To capture the partial volume effects, our shape is embedded into an image such that each pixel has a value associated with it between 0 and 1. This value represents the probability that a random point inside a single pixel contains the spinal cord. Hence a value of 0 would signify the pixel does not contain the spinal cord and a value of 1 represents a pixel that is fully part of the spinal cord.

Since we adopted a probabilistic representation, we need to consider that standard PCA is not constrained to the geometry of the unit simplex (space of non-negative vectors with unity sum) and can produce shapes that are invalid (outside the simplex) [9]. To stay within the simplex, we map our data to the ILR space before applying PCA and then map the resulting shape back to our probabilistic shape space [10]. To convert to ILR space, the following ILR transformation is de-

rived for the two label case,

$$q = \text{ILR}(\mathbf{p}, \mathbf{e}) = \ln \frac{p_1}{g(\mathbf{p})} \ln \frac{e_1}{g(\mathbf{e})} + \ln \frac{p_2}{g(\mathbf{p})} \ln \frac{e_2}{g(\mathbf{e})} \quad (1)$$

where, $\mathbf{p} = [p_1, p_2]$ is a vector of our probabilistic values; $\mathbf{e} = [0.804, 0.196]$ is the Aitchison-orthonormal basis in the simplex; and $g(\mathbf{x}) = \sqrt{x_1 x_2}$ is the geometric mean. To convert back to our probability space, we calculate,

$$\mathbf{p} = \text{ILR}^{-1}(q, \mathbf{e}) = (e_1^q + e_2^q)^{-1} [e_1^q, e_2^q]. \quad (2)$$

Given our learned shape space, we can apply varying PC weights to it to deform our shape as can be seen in Fig. 1.



Fig. 1. Sample spinal cord cross-sectional shapes generated by randomly sampling the PC weights. Note the probabilistic pixels and the non-ellipsoidal shapes.

2.2. Energy Function

Our energy function is designed to model two characteristics of the spinal cord: 1) the cord is a lighter object usually surrounded by an outer dark ring of cerebrospinal fluid, E_{∇} , and 2) the intensities found in the spinal cord should be similar, E_{σ} . We define E our energy function as

$$E = \alpha E_{\nabla} + \beta E_{\sigma} \quad (3)$$

where α, β are weights for each energy term (both set to 1 for our experiments). Note that each of these terms are data terms as the regularization is built directly into our shape representation and optimizer; we constrain the PC weights (i.e. shape) to only vary by one standard deviation across slices, and enforce the spatial regularization through graph connectivity.

In E_{∇} , our goal is to ensure the gradient vectors on the object's boundary point towards the center of the cord:

$$E_{\nabla} = \frac{\int_{\Omega} ((f(\nabla P(x, y)) \circ \nabla I(x, y)) \cdot R(x, y)) dx dy}{\int_{\Omega} |f(\nabla P(x, y))| dx dy} \quad (4)$$

where $(x, y) \in \mathbb{R}^2$ is a point in our image domain Ω ; P is the probability of each pixel belonging to the cord (Fig. 1); I represents our image; R is composed of vectors pointing towards the centroid of the shape found in P ; f is a Heaviside step function to only allow for strong edges and to treat all edges equally; and \circ and \cdot denote element wise multiplication and the dot product, respectively.

Our second energy term, E_{σ} , assigns low energy to shapes that contain pixel intensities whose standard deviation $\sigma_p = \text{std}(I(x, y); (p(x, y) > \epsilon))$ agrees with standard deviations, σ , of spinal cord intensities collected from our training data. We

define $\bar{\sigma}$ as the mean learned standard deviation of intensities (i.e. $\text{mean}(\sigma)$); $\text{std}(\sigma)$ as the standard deviation of the learned standard deviation of intensities, σ . We set $\sigma_{lower} = \bar{\sigma} - 3 * \text{std}(\sigma)$, and $\sigma_{upper} = \bar{\sigma} + 3 * \text{std}(\sigma)$. E_σ penalizes shapes whose σ_p exceeds three standard deviations from the learned $\bar{\sigma}$, calculated as:

$$E_\sigma(\sigma_p, \sigma) = \begin{cases} (\sigma_{lower} - \sigma_p), & \text{if } (\sigma_p < \sigma_{lower}) \\ (p_\sigma - \sigma_{upper}), & \text{if } (p_\sigma > \sigma_{upper}) \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

2.3. Minimal Path Optimizer

We optimize our energy function using an A* minimal path search. A standard approach would be to use Dijkstra’s algorithm and apply it to high dimensions (initially we used the N-D Dijkstra’s algorithm provided by [11]). However, this approach is problematic since it (and most traditional minimal path methods) requires as input to the minimal path algorithm a graph composed of nodes and edges. The memory required to explicitly encode all the edges of our graph explodes due to the high connectivity of 6D graphs. For our 6D volume, a fully connected (26-connected in 3D space) *single* node would require $3^6 = 729$ edges. We found that even with modifications to the existing code and underlying data structures of [11], we were unable to lower our memory requirements to make our method computationally feasible (e.g. running our method on a single cropped volume of $25 \times 20 \times 80$ could require over 70 GBs). As well, a further drawback is the considerable run-time needed to create such a graph.

To reduce the graph creation run-time as well as the memory requirements of the A* algorithm, we observed that the structure of a volume allows us to infer, on-the-fly during run-time, the connectivity of a voxel. Thus rather than explicitly storing all the edges as input, we modify the minimal path algorithm to index into our volume and determine the appropriate neighbours at run-time. This removes the need to encode the edges prior to running our minimal path method. As we now only need to keep a subset of the edge values in memory (those in the in the minimal path “open set” queue), our memory requirements drop drastically which allows us to search over larger volumes (e.g. the 70 GB dropped to 4 GB).

We “prune” certain edges of our graph to create a regularization prior. We use the prior knowledge that the spinal cord in the MRI volume extends superior to inferior i.e. along the z dimension. We thus remove any edge that does not transition in the z dimension towards the end (inferior) seed point. As well, we discretize the possible PC weights to use 7 possible multiples of the standard deviation values $\{0, \pm 1, \pm 2, \pm 3\}$.

We implement an efficient low-memory A* search to find the globally optimal path between two user entered seed points where the returned minimal path contains the 6D coordinates that represent the globally optimal segmentation. We chose A* over Dijkstra’s algorithm since A* maintains

globally optimality while allowing us to define heuristics that can speed up the run-time (run-time is between 1 to 5 hours).

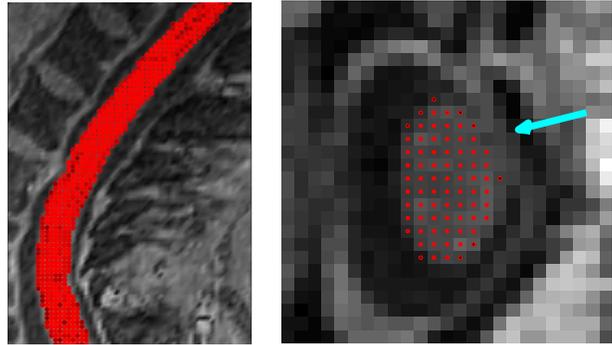


Fig. 2. (Left) Our method’s segmentation shown in a sagittal MR slice. (Right) Axial slice with misleading image data (arrow). Darker areas represent lower segmentation probability of the spinal cord.

3. RESULTS

We validate our method over 20 MRI scans composed of 10 healthy and 10 MS patients scanned with a 3.0T scanner and 1.5T scanner respectively (scans were from different studies) with a voxel size of $0.9760 \times 0.9760 \times 1$ mm. Each scan was segmented by an expert and, given its use in previous studies [2], we consider this a reliable expert segmentation. We validate our method over a challenging area of the cord (C3-C7) of 80 slices.

To seed our method, we extract the approximate center of the expert segmentation for the first and last expertly segmented slice to get the first three spatial dimensions. We use the mean shape of our spinal cord PCA model (we build a separate MS and non-MS PCA model with the testing volume omitted) as the three PC weight dimensions (i.e. zero weights). While this does not give an optimal segmentation for the first and the last slice, our method quickly converged to a reasonable segmentation after one or two slices. A user could get around this limitation by placing the seed a few slices before and after the desired locations. We use our start and end seed point to roughly crop our volume around the cord to help eliminate voxels from our search space that cannot be part of the spinal cord.

Given the small structure of the spinal cord it is important to consider the contributions of partial volume effects (PVE) on the segmentations to ensure accuracy (Fig. 2). The Jaccard similarity index is a commonly used measure to determine the similarity between shapes and is defined as $J(A, G) = \frac{|A \cap G|}{|A \cup G|}$. We modify Jaccard to consider PVE and decrease the score as the PVE for each pixel increases in difference between each shape. We calculate this as,

$$J_{PVE}(A, G) = \frac{\sum_x \sum_y \min(A(x, y), G(x, y))}{\sum_x \sum_y \max(A(x, y), G(x, y))} \quad (6)$$

Method	Mean	Median	Min	Max	Std.Dev
J_{MAP}	0.789	0.800	0.473	0.949	0.076
J_{PVE}	0.784	0.797	0.451	0.906	0.066
$Area_S$ %	94.33	95.15	88.72	99.88	3.30
J_{MAP}	0.855	0.869	0.394	0.989	0.083
J_{PVE}	0.832	0.848	0.379	0.915	0.069
$Area_S$ %	96.88	97.77	90.32	99.81	2.94

Table 1. Similarity in shape overlap (J_{MAP} = binary and J_{PVE} = probabilistic Jaccard) and computed area ($Area_S$) between automated and expert segmentations for multiple sclerosis patients (upper three rows), and healthy (non-MS) patients (lower three rows).

where A and G are the automated and expert (ground truth) shapes embedded into the same sized image. The “min” function represents a probabilistic intersection and the “max” function represents a probabilistic union. We include the crisp version of Jaccard (round PVE) in our results denoted as J_{MAP} .

To determine whether the automated method is a useful surrogate of the manual method for detecting global cord atrophy, we also compare the computed volumes between the automated and expert segmentations regardless of their overlap. As the cross-sectional area (CSA) has been shown to be a good indicator of the progression of MS, a high agreement between an automated and expert segmentation would be a good indication of the clinical utility of a particular method. To measure this, we use a similarity ratio, $Area_S$, as described in [2]. The $Area_S$ is given as $1 - \min(|(A/G) - 1|, 1)$ which returns a 0% accuracy if the automated volume overestimates the expert volume by more than 200%.

In Table 1 we see how our automated segmentation compares with the expert segmentation. As might be expected, our method performs better overall on the healthy patients than the MS patients due to the increased irregularity of the spinal cord as MS progresses and the change from a 3T to a 1.5T scanner. Our results compare favourably with other similar methods and requires less user input. In [2], which requires some user correction of the Live-wire path, they report a mean volume similarity $Area_S$ of 94.31% for MS and 91.60% for healthy patients. Noticeably, for healthy patients they report a min similarity of 77.33%, compared to our reported min similarity of 90.32% which suggests our method is more capable of capturing the cord’s true volume.

4. CONCLUSIONS AND DISCUSSIONS

We proposed a method to segment the spinal cord by finding a minimal path in 6D, whose globally optimal parameters define our segmentation. Although we used a 6D representation, this method could be extended to higher dimensions but at a cost in both memory requirements and computation time. We address some ways to reduce this in the form of our opti-

mized A* search that take advantage of the volume structure and edge pruning to drastically reduce the needed memory. Future work would look at further space and run-time optimizations and correlating the automated segmentation to the progression of MS.

5. REFERENCES

- [1] Rocca et al., “A multicenter assessment of cervical cord atrophy among MS clinical phenotypes,” *Neurology*, vol. 76, no. 24, pp. 2096–2102, 2011.
- [2] C McIntosh, G Hamarneh, M Toom, and R Tam, “Spinal cord segmentation for volume estimation in healthy and multiple sclerosis subjects using crawlers and minimal paths,” *IEEE HISB*, pp. 25–31, 2011.
- [3] B Schmit and M Cole, “Quantification of morphological changes in the spinal cord in chronic spinal cord injury using magnetic resonance imaging,” *Conf Proc IEEE Eng Med Biol Soc*, pp. 4425–8, 2004.
- [4] Horsfield et al., “Rapid semi-automatic segmentation of the spinal cord from magnetic resonance images: Application in multiple sclerosis,” *NeuroImage*, vol. 50, no. 2, pp. 446–455, 2010.
- [5] C McIntosh and G Hamarneh, “Spinal crawlers: Deformable organisms for spinal cord segmentation and analysis,” *MICCAI*, pp. 808–815, 2006.
- [6] H Li and A Yezzi, “Vessels as 4-D curves: Global minimal 4-D paths to extract 3-D tubular surfaces and centerlines,” *IEEE TMI*, vol. 26, no. 9, pp. 1213–1223, 2007.
- [7] M Poon, G Hamarneh, and R Abugharbieh, “Live-vessel: extending livewire for simultaneous extraction of optimal medial and boundary paths in vascular images,” *MICCAI*, pp. 444–451, 2007.
- [8] D Cremers, F Schmidt, and F Barthel, “Shape priors in variational image segmentation: Convexity, lipschitz continuity and globally optimal solutions,” *CVPR*, pp. 1–6, 2008.
- [9] N Changizi and G Hamarneh, “Probabilistic multi-shape representation using an isometric log-ratio mapping,” *MICCAI*, pp. 563–570, 2010.
- [10] S Andrews, C McIntosh, and G Hamarneh, “Convex multi-region probabilistic segmentation with shape prior in the isometric log-ratio transformation space,” *ICCV*, pp. 2096–2103, 2011.
- [11] L Weizman, M Freiman, and L Joskowicz, “Implementation of weighted dijkstras shortest-path algorithm for n-D images,” *The Insight Journal*, 2009.