

Spinal Crawlers: Deformable Organisms for Spinal Cord Segmentation and Analysis

Chris McIntosh and Ghassan Hamarneh

Medical Image Analysis Lab, School of Computing Science Simon Fraser University,
BC, Canada

{cmcintos, hamarneh}@cs.sfu.ca

Abstract. Spinal cord analysis is an important problem relating to the study of various neurological diseases. We present a novel approach to spinal cord segmentation in magnetic resonance images. Our method uses 3D “deformable organisms” (DefOrg) an artificial life framework for medical image analysis that complements classical deformable models (snakes and deformable meshes) with high-level, anatomically-driven control mechanisms. The DefOrg framework allows us to model the organism’s body as a growing generalized tubular spring-mass system with an adaptive and predominantly elliptical cross section, and to equip them with spinal cord specific sensory modules, behavioral routines and decision making strategies. The result is a new breed of robust DefOrgs, “spinal crawlers”, that crawl along spinal cords in 3D images, accurately segmenting boundaries, and providing sophisticated, clinically-relevant structural analysis. We validate our method through the segmentation of spinal cords in clinical data and provide comparisons to other segmentation techniques.

1 Introduction

The spinal cord is a crucial part of the nervous system that resides within the vertebral canal of the spinal column and acts as a relay to convey information between the brain and the rest of the body. However, there are numerous clinical problems related to the spinal cord, including multiple sclerosis (MS), meningitis, neural tube defects, syringomyelia, transverse myelitis, and spinal cord injury (SCI). Of specific interest, MS affects more than a quarter million people in the US, and is suspected of shrinking the cord. A necessary prior for statistically analyzing the spinal cord is segmenting it. For example, [1] performs texture analysis on manually segmented spinal cords to study the effects of MS.

However, spinal cord segmentation is a difficult problem. Though the cord is approximately cylindrical in shape, its diameter varies at different vertebral levels. Furthermore, progressing from superior to inferior the information transfer requirement decreases resulting in decreased white matter and, therefore, decreased contrast. This change in contrast proves difficult for nonadaptive segmentation methods. Consequently, the majority of previous work focuses on segmenting the canal and the vertebra [2, 3, 4] (which is a much simpler problem given the high contrast of bone in CT).

There are, however, a few existing works on spinal cord segmentation. In [5] Schmit and Cole investigate if an observable change in the spinal cord occurs after SCI. Segmentation is performed using seeded region growing, concluding that spinal narrowing in magnetic resonance images (MRI) is indeed observable. However, no validation of the segmentation results is presented, and the assumption that extraspinal structures do not touch the spinal cord is made, which often does not hold (Fig 3). In [6] a deformable model approach is used for spinal cord segmentation. However, the method requires initializations along the entire cord, is only validated on a single phantom data set, and takes approximately 10 hours for one segmentation. In [7] Aylward *et al.* use a hessian-based approach for strictly tubular (vessel) structures and, thusly, do not address the variability in the major and minor axes of the cord. Furthermore, no quantitative validation of the proposed technique on spinal cord data is provided. As we motivate in section 2.1 this is not an optimal approach as the spinal cord can vary from tubular to elliptical in shape, and local extraspinal structures distort derivatives of Gaussians along the minor-axis of the cross-sections.

In previous work we demonstrated the merit of locally optimal filtering in vessel segmentation and analysis [8]. In this work, we address the clinical need for a fast semi-automatic spinal cord segmentation and analysis technique by extending and adapting our previous work for this new problem. We demonstrate the use of hessian-based filters for spinal cord segmentation that intuitively encode the geometrical properties of the cord, along with an elliptical mesh that explicitly encodes the major and minor axes.

The remainder of the paper is structured as follows. We begin with an overview of our earlier work on DefOrgs for the segmentation and analysis of tubular structures. We then describe our method of spinal cord segmentation using optimal hessian scales and our extension to adaptive, elliptical, cross-sectional filtering (section 2). In section 3 we present results comparing our method to other to state-of-the art filtering and segmentation techniques. Finally, we summarize and draw conclusions in section 4.

1.1 Tubular Crawlers

In [8] we presented a novel approach to the segmentation and analysis of vasculature from volumetric medical image data. Our method extended the original DefOrgs introduced by Hamarneh *et al.* in [9] to 3D, modelled their bodies as tubular spring-mass systems and equipped them with a new repertoire of sensory modules, behavioral routines, and decision making strategies specifically designed for vascular segmentation. The result was a new breed of robust DefOrgs, “vessel crawlers”, that crawl along vasculature in 3D images, accurately segmenting vessel boundaries, detecting and exploring bifurcations, and providing sophisticated, clinically-relevant structural analysis. However, our methods were restricted to high-contrast tubular structures on dark backgrounds (vessels) and are, thus, ill-suited for segmenting the spinal cord; a low-contrast elliptical structure.

2 Methods

In this section we provide an overview of the spinal crawler and its use in the segmentation and analysis of spinal cords in 3D MRI. We begin with an overview of the geometrical properties of the spinal cord, then describe the spinal crawler, followed by details of our implementation of the various layers of the DefOrgs framework.

2.1 Geometrical Properties of Spinal Cords

As previously mentioned, the cross-section of the spinal cord is tubular in some regions and elliptical in others (Fig 1). This geometrical property motivates the use of an adaptive filtering procedure that takes advantage of the local geometry. Here we describe such a procedure.

It is well known that the eigenvalues of the Hessian $H(I * G_\sigma)(x)$ computed at scale σ describe the principle curvature at location x [10]. Given an ordering of eigenvalues $|\lambda_1| \leq |\lambda_2| \leq |\lambda_3|$ the ratios between the eigenvalues yield measurements of the underlying geometrical structure. This property motivates Frangi *et al.*'s vesselness measure at a single scale, under which R_A differentiates between plates and lines, R_B measures deviation from blob like structures, and S emphasizes areas of high contrast (for details see [11]). The final filtering result at each voxel is the maximum of a search across scales with a fixed value for c , which is ill-suited to the spinal cord as it changes shape and dims along its medial axis.

In order to adapt Hessian based approaches to spinal cord segmentation, we locally adapt the gradient calculation method and filtering equations. We define an anisotropic Hessian $H(I * G_{\sigma_x, \sigma_y, \sigma_z})(x)$ computed at scales $\sigma_x, \sigma_y, \sigma_z$ where $G_{\sigma_x, \sigma_y, \sigma_z}$ is an anisotropic Gaussian with σ_x, σ_y equal to the major and minor axes of the elliptical structure and σ_z the average of the two; a somewhat arbitrary choice indicating how far the structure is expected to remain straight. However, a direct requirement of this type of filtering is that the major and minor axes of the structure be known. Furthermore, the eigenvalues are based on the difference taken inside and outside the structure and, as such, are influenced by the bright extraspinal structures that lie (anterior to posterior) around the spine (Fig 1). Consequently, we define a 'spinalness' filter as:

$$S_{\sigma_x, \sigma_y, \sigma_z}(s) = \left\{ \begin{array}{l} 0 \text{ if } \lambda_2 > 0 \text{ or } \lambda_3 > 0 \\ \left(e^{\left(-\frac{(R_A - 0.5)^2}{2\alpha^2} \right)} \right) e^{\left(-\frac{R_B^2}{2\beta^2} \right)} \left(1 - e^{-\frac{S^2}{2c^2}} \right) \quad o/w \end{array} \right\} \quad (1)$$

in which R_A , the ratio of the largerest two eigenvalues, is allowed to deviate from its expected value of 0.5. In our framework we derive locally optimal values for σ_x, σ_y and c (section 2.2), thereby greatly enhancing the robustness of the method to image noise by adapting to the local cord shape and intensity level.




Fig. 1. Left pair: Axial slices of a spinal cord showing circular and elliptical cross-sections. Right pair: Simple user interaction on a cross-sectional slice can be used to repair false boundary identification. Right pair also shows the extraspinal structures, with gray level intensity higher than that of the spinal cord, that are often encountered in the minor-axis derivatives (horizontal-direction).

2.2 Spinal Crawlers

Our spinal crawler DefOrg was built under a multilevel Artificial Life modelling approach consisting of four primary levels: cognitive, behavioral, physical, and geometrical (Fig 2) [9]. Specifically, the cognitive layer makes decisions based on our model's current state, anatomical knowledge, and its surrounding environment (the image volume including the spinal cord). Decisions could be made to sense information, to deform based on sensory data, or to repair an erroneous layer. All of these actions are described under the behavioral layer of the organism, and they rely upon both the physical and geometrical layers for implementation. For example, the act of progressing down the spine is described by the *'growing'* behavioral method. The cognitive center gathers sensory input using the *'spine sense'* sensory module, decides the correct location via the *'grow'* decision module, elicits the act of *'growing'*, and then conforms to the spinal cord by *'fitting'*. In turn, each of these methods relies upon the physical and geometrical layers to carry out tasks, such as maintaining model stability through the application of stabilization springs (Fig 2 middle three). Consequently, we have a framework with many independent layers of abstraction, each built upon the implementation of independent models and processes. What follows is an explanation of each layer.




Fig. 2. Left: The layered architecture of DefOrgs showing all of the modules used by our spinal crawler. Middle three: Topology of the spinal crawler showing: Masses (left), radial springs (middle), and stability springs (right) across sequential layers of the organism. Right: A spinal crawler (gray) utilizing a hemispherical “off-board” (detached from body) sensor showing an idealized output (light gray).

Geometrical. The elliptical geometric module is parameterized by both section length (distance between neighboring medial masses) and circumferential resolution (proportional to the number of circumferential boundary masses) (Fig 2). Each layer supports major and minor axes lengths, and thus, is able to conform locally to the spinal cord.

Physical. We use a 3D Newtonian physically-based deformation module [8]. External forces are provided by the volumetric image gradient and a drag force, while internal forces are supplied through Hooke’s law and dampening spring forces. This physical layer enables intuitive user-interaction, under which the user can interact with individual nodes to deform the entire system in a natural manner. This enables a quick and easy method of repairing erroneous segmentations in areas of non-existent boundary information (Fig 1)[12].

Sensory. The primary sensory module is controlled by the spinal crawler’s decisions, wherein optimal parameters are derived through basic sensory modules (local image intensity, standard deviation, etc.). As noted in section 2.1 the Hessian and its associated filters must be computed at locally optimal scales. Our approach estimates the scales from the major and minor axes of the leading (front-most) layer of the spinal crawler after it has deformed to the local gradient information.

Specifically, our ‘**Spine Sense**’ module captures volume intensity information on the surface of a hemisphere centered around the spinal crawler and normal to its front-most layer (Fig 2). A locally optimal connected-components filter is used to ensure only relevant information to the crawler is present on the surface. The filter returns all connected voxels with intensities and ‘*spinalness*’ measurements that lie within \pm multiples of standard deviations of maximum-likelihood-estimates of Gaussian distributions’ means sampled from within the local segmented region of the spinal cord. Intensities from the filter now lie on the surface of the hemisphere, and the next location to grow to is taken as the centermost centroid of all circular regions on the surface. As shown in figure 3, this works quite well at sequentially detecting the local medial axis of the spinal cord.

Cognitive. The spinal crawler can make three key *decisions*. Namely, where to grow, when to repair, and when to stop. Each decision can be based on sensory input, anatomical knowledge, or user interaction. The user is able to override any of these key decision functions at any time during the organisms life cycle, and hence can illicit intuitive real-time control over the segmentation process. This presents a strong advantage over those methods that force the user to wait for a potentially incorrect output before re-running the segmentation process with different parameters (region growers, level sets, neural networks, etc.).

The ‘**Grow**’ decision dynamically decides the next grow location via the *spinal sense* module, and activates the *growing* behavior with appropriate parameters (major and minor axis lengths, location for next layer, etc.). ‘**Repair**’ decides to repair a given layer if its eccentricity exceeds the average eccentricity over

the past few layers. While ‘**Terminate**’ decides when to terminate execution, which can be at the request of the user, when reaching a target point, or when a provided threshold of spinalness is not met.

Behavioral. Each of the spinal crawlers key decisions results in the execution of the appropriate *behavior* using the concluded locally optimal parameters, scales, estimated spinalness mean and variance, etc. ‘**Growing**’ creates new layers and subsequently connects them to the current front-most layer (Fig 2). Once connected the model can be fit to the image data. ‘**Repairing**’ takes advantage of the spinal cord’s symmetric reflection about its major axis. The section of a layer exhibiting the highest eccentricity is replaced by a reflection of the opposite side. ‘**Fitting**’ is accomplished using 3D image gradient driven deformations simulated by the physics layer. Connections to the previous layer provide smoothness, while stiffer circumferential springs provide local stability to noise, and flexible radial springs allow deformation to the spinal boundary.

3 Results

We present quantitative validation of our method on four $256 \times 256 \times 60$ MRIs, each of which has been manually segmented at approximately 60 minutes a volume¹, and has as a voxel size of $0.98 \times 0.98 \times 1mm^3$. As a point of comparison we also segment the data using connected components and ITK-SNAP’s level-set method [13]. The connected components proved ill-suited to the task, even with heavy manual intervention, as the intensity levels of nearby structures were to similar to that of the cord (Fig 1, 3 far right). However, the level-set from SNAP was partially successful with high-degrees of manual intervention. Specifically, we had to partition the volume into pieces and segment each individual piece with slightly different parameters and seed point placements along most of the spine. Furthermore, in regions where boundaries are not clear, we had to re-process the image to create an improved edge image. Essentially, segmentation with SNAP required manually performing the automatic features of our framework. Overall, the whole process took on average 30 minutes for a single image. We expect this time will be even higher for a clinical expert with minimal knowledge of image segmentation algorithms and their parameters. Finally, the output from these techniques is simply a binary mask that itself cannot be intuitively interacted with nor easily analyzed (for example to obtain cross-sectional slices normal to the medial axis of the extracted spinal cord, skeletonization and pruning must be performed).

In contrast, segmentation with our method required on average 10 minutes per image with the user only providing two seed points. As can be shown in table 3 our method outperforms the SNAP method. Furthermore, as noted in section 2.2, our final results can be easily improved as shown in figure 1. In addition to a binary mask, the final segmentation without post-processing presents the

¹ Reported times performed by a medical image analysis software literate person.




Fig. 3. Sample 3D segmentation results (left to right): manual, spinal crawler, SNAP level set, and a region grower with seed point shown in red that has leaked into nearby extraspinal structures

user with the medial axis, and cross-sectional slices of the spine taken along its normal². Typically, 10 to 15 cross-sections of approximately 200 taken along the spinal cord were found in noticeable error, and can easily be corrected within a few minutes.

Table 1. Error comparison between the 2-seed-point, automatic spinal crawler, and the heavily adjusted/tuned SNAP. Errors are measured as minimal inter-surface distance computed against manual segmentations. Statistics are calculated across all inter-surface pairs and the 4 data sets. For example, the mean Hausdorff distance (maximum of inter-surface pairwise distances) across all 4 data sets for the spine crawler was 5.39.

Error	Mean	Median	Min	Max	Std
Spine Crawler Mean distance	0.5662	0.455	0.26	1.095	0.3694
SNAP Mean distance	0.812	0.734	0.415	1.39	0.4252
Spine Crawler Max error	5.39	5.485	4.1	6.5	1.07
SNAP Max error	5.71	4.83	3.58	9.6	2.73

4 Conclusion

We set out to develop a semi-automatic tool for spinal cord segmentation suitable for clinical studies of spinal cord-related diseases [1, 5]. To this end, we have presented an elliptical DefOrg that incorporates the cord’s geometrical properties to delineate its medial axis and boundary. Our method has proven faster and more accurate than other segmentation techniques (connected components, level sets). Furthermore, through its intuitive and interactive control structures, real-time analysis, and post-segmentation interactions, it keeps the doctor *‘in the loop’*; in contrast to numerous black box approaches that provide no interaction, no recourse in case of error, and whose inner-workings are nonintuitive.

² See <http://mial.fas.sfu.ca/researchProject.php?s=157> for a video and larger images.

Acknowledgements

We would like to thank the MS/MRI Research Group at the University of British Columbia and Berlex for providing the spinal cord MR images, Lisa Tang for running tests on various data sets, and funding from the Natural Sciences and Engineering Research Council of Canada.

References

1. Mathias, J., Tofts, P., Losseff, N.: Texture analysis of spinal cord pathology in multiple sclerosis. *Magnetic Resonance in Medicine* **42** (1999) 929:935
2. D'Haese, P., Niermann, K., Cmelak, A., Duay, V., Li, R., Dawant, B.: Atlas-based spine segmentation for radiotherapy planning. *American Society of Therapeutic Radiology and Oncology* (2004)
3. Mancas, M., Gosselin, B., Macq, B.: Segmentation using a region-growing thresholding. *Proceedings of the SPIE* **5672** (2005) 388–398
4. Karangelis, G., Zimeras, S.: An accurate 3d segmentation method of the spinal canal applied to CT data. In: *Bildverarbeitung fur die Medizin*. (2002) 370–373
5. Schmit, B.D., Cole, M.K.: Quantification of morphological changes in the spinal cord in chronic human spinal cord injury using magnetic resonance imaging. *IEEE EMBS* **26** (2004) 4425 – 4428
6. Coulon, O., S.J.Hickman, Parker, G., Barker, G., Miller, D., Arridge1, S.: Quantification of spinal cord atrophy from magnetic resonance images via a b-spline active surface model. *Magnetic Resonance in Medicine* **47** (2002) 1176–1185
7. Aylward, S., Bullit, E., Pizer, S., Eberly, D.: Intensity ridge and widths for tubular object segmentation and description. *IEEE/SIAM workshop Mathematical Methods Biomedical Image Analysis* (1996) 131–138
8. McIntosh, C., Hamarneh, G.: Vessel crawlers: 3d physically-based deformable organisms for vasulature segmentation and analysis. *IEEE Conference on Computer Vision and Pattern Recognition* **1** (2006) 1084–1091
9. Hamarneh, G., McInerney, T., Terzopoulos, D.: Deformable organisms for automatic medical image analysis. In: *MICCAI*. (2001) 66–76
10. Lindeberg, T.: *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Norwell, MA, USA (1994)
11. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. *Lecture Notes in Computer Science* **1496** (1998) 130–137
12. Hamarneh, G., McIntosh, C.: Physics-based deformable organisms for medical image analysis. *SPIE Medical Imaging* **5747** (2005) 326–335
13. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., Gerig, G.: User-guided 3d active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage* (2006)