Spinal Cord Segmentation for Volume Estimation in Healthy and Multiple Sclerosis Subjects using Crawlers and Minimal Paths

Chris McIntosh*, Ghassan Hamarneh*
*Medical Image Analysis Lab
Simon Fraser University
Burnaby, BC
{cmcintos,hamarneh}@sfu.ca

Matthew Toom*†
†MS/MRI Research Group
University of British Columbia
Vancouver, Canada
mtoom@alumni.sfu.ca

Roger C. Tam†‡
‡Department of Radiology
University of British Columbia
Vancouver, Canada
roger.tam@ubc.ca

***Abstract*—Spinal cord analysis is an important problem in the study of various neurological diseases. Current segmentation and analysis methods in clinical use are slow and labor-intensive, especially for pathological data. "Spinal Crawlers" are a recently developed technique based on an artificial life framework for medical image analysis that complements classical deformable models (snakes and deformable meshes) with high-level control mechanisms. Our method extends Spinal Crawlers to better function in a clinical setting in which images of variable quality and challenging anatomy are encountered. We augment the Spinal Crawler's local optimality with that of globally optimal paths using the live-wire technique. This fusion of globally optimal paths, with locally optimal filtering allows our method to better adapt to contrast changes compared to other methods and therefore allows a larger section of the cord to be measured. Our improvements are validated on 5 vertebral levels of both healthy and pathological spinal cords from clinical MR data. This is the first study to validate a spinal cord segmentation method over a large region encompassing the length of five vertebrae.**

***Keywords*-spinal cord analysis; medical image analysis; segmentation**

I. Introduction

Atrophy of the central nervous system (brain and spinal cord) is an important contributor to cognitive and physical disability in multiple sclerosis (MS) [1]. The relationship between physical disability and spinal cord atrophy is particularly strong [1], and in vivo studies done with magnetic resonance imaging (MRI) have found significant correlations between a reduction in spinal cord cross-sectional area (CSA) and clinical disability [2]. These results motivate further study of the spinal cord in MS, but there are also certain expected MRI findings for the spinal cord that remain unobserved, especially in the early stages of MS. For example, while a number of longitudinal studies have revealed reductions in CSA in both relapsing-remitting (RR) and progressive MS [2], the large majority of cross-sectional studies comparing RRMS patients to healthy control subjects have failed to show the expected reduction in CSA [2], with only one such study showing a small difference [3]. In addition, while no study has yielded conclusive evidence that current disease-modifying therapies for MS are able to slow the rate of spinal cord degeneration [4], a number of current therapies have shown at least some benefit on whole brain atrophy [5]. Considering the importance of the spinal cord for fine motor skills and locomotion, it is necessary to try to eliminate methodological reasons why spinal cord measures are not revealing the expected findings. Accurate and precise measurement methods are particularly important for the spinal cord because it is a very small structure (about 80mm^2 for a healthy adult male), and small absolute errors can translate into large relative errors.

Spinal cord segmentation is, however, a difficult problem. Though the cord is basically cylindrical in shape, its diameter varies at different vertebral levels. Progressing from superior to inferior there is a natural bending of the cord causing it to abut the spinal canal, often in multiple locations. Finally, a decreased signal resulting from increased distance from the center of the MR receiving coil causes a change in contrast in the inferior regions of the cord that proves difficult for nonadaptive segmentation methods.

The most commonly used methods for measuring the spinal cord on MRIs [6]–[8] are limited in that they require heavy user interaction or manual intervention for accurate delineation of the cord boundary, which makes them susceptible to intra-rater and inter-rater variability, and labor intensive. As a result most studies are focused on only a single vertebral region of the cervical cord. However, there is evidence from a recent post-mortem study that suggests MS affects different parts of the cord differently, and thus it would be beneficial to analyze a larger region [9]. A primary challenge in studying a larger section of the cord is that the cord becomes more difficult to segment at increasingly inferior vertebral levels; with a typical neck coil used for cervical cord imaging, the contrast at C6-C7 is significantly lower than that at C2-C3, which is the most common region measured. Attempts at automating spinal cord measurements [10]–[13] are either less sensitive to real volumetric changes than existing semi-automatic methods [10] or have not been validated with a statistically significant number of scans [11], [12]. In [12] Schmit and Cole investigate whether an observable change in the spinal cord occurs after spinal

cord injury. Segmentation is performed using seeded region growing, concluding that spinal narrowing in MRI is indeed observable. However, no validation of the segmentation results is presented. In [10], a deformable model approach is used for spinal cord segmentation. However, they only compared with a semi-automatic method on a 15mm region at the C2/C3 vertebral level. The method also underestimated the cord very significantly, with the volume measurements being an average of only 74% of that produced by the semi-automatic method. A possible explanation for the above is that fitting a single B-spline surface to the whole cord may cause the result to be overly smooth, motivating the need for a more localized approach. In [14], an active surface model is used to fit the cord between the C2/C5 vertebral levels, but then only validated on 5-slices at C2. There exists a need for a method, and a corresponding study, that focus on a larger vertebral range.

In previous work, McIntosh and Hamarneh demonstrated the merit of locally optimal filtering and deformable model fitting in vessel segmentation and analysis [11]. They then built a customized Hessian-based filter for spinal cord segmentation that intuitively encodes the geometrical properties of the cord, along with an elliptical mesh that explicitly encodes the major and minor axes. The result was "Spinal Crawlers" (SC), a deformable model fused with high-level control mechanisms that iteratively progresses along the medial axis of the spinal cord, segmenting it layer-by-layer. It is thus a localized approach, traversing the spinal cord, making localized optimizations at each point. These local decisions can, however, be misleading, when put in the context of global information. As a consequence, their method struggled with the lower areas of the spinal cord (past vertebra C5), and is thus not directly applicable to studying a larger area of the cord (C3-C7). Optimal path approaches contrast this by optimizing globally over the scope of the entire cord. Optimal path approaches have been previously developed [15], and in some cases have been specifically designed for tubular structures, including a radius, r, in addition to the position variables x and y [16], [17]. However, methods expressly designed for tubular structures are ill-suited to spinal cord segmentation because of the elliptical nature of the cord's cross-section.

In what follows we describe how we extend upon and improve the SC technique to be applicable in a clinical setting via its fusion with a popular optimal path technique, live-wire (LW) [18]. Our results demonstrate superiority of our approach in both accuracy and speed. It compares favorably to the original proposed SC on the task of segmenting a full 5 vertebral region, C3-C7, of 20 real clinical MR datasets. Specifically, it overcomes problems where the fully-automatic SC would crawl off-axis in regions of low contrast or spinal cord abutment. The result is superior boundary delineation, and volume estimation of both healthy and diseased spinal cords in MR. Our main contribution is thus a new method for large-region spinal cord segmentation, that we will show compares favorably to other state-of-the art methods in volume estimation, despite focusing on a larger more difficult region of the spinal cord.

II. Methods

A. Spinal Crawlers:

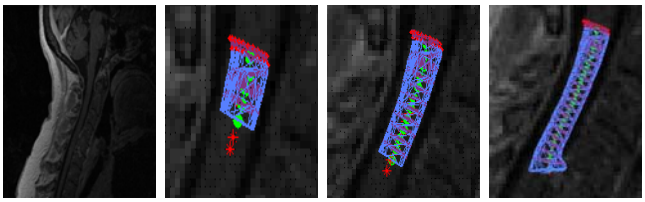

Figure 1. A SC in action. From left to right, a typical saggital plane from an input MR scan, the original SC technique after a few iterations, ten more iterations of the method, and the point at which its local decision leads it off the cord.

As already noted, SC is a localized approach, traversing the spinal cord, making localized optimizations at each point. Though a full description of SC is outside the context of this paper, we provide a high-level summary of the concepts needed to understand our improvements.

The SC begins as a two-layered elliptical mesh embedded in the volume at a user specified co-ordinate, p, and crawling direction. Its primary execution loop makes use of the following steps:

1) Deform to the local boundary of the spinal cord.
2) Measure the local image statistics.
3) Use the estimated spinal cord intensity and size characteristics to perform an optimal filtering operation for that particular location.
4) Use a region growing method to roughly segment the next area of the optimally filtered spinal cord.
5) Estimate from that segmentation the next medial point of the cord, p_{SC}.
6) Create a new elliptical cross-sectional layer at $p_{new} = p_{SC}$ with its normal, $\vec{v}$, along the SC grow direction, i.e. $\vec{v} = \vec{v}_{SC} = p_{SC} - p$.
7) Repeat from step (1) setting $p = p_{new}$.

Fig. 1 shows the repeated execution of this loop over an MR scan.

The strength of SC is their locally optimal filtering, as they avoid assuming a single filter-shape and -size will fit the entire SC. However, the method's focus on locally optimal decisions is also its main drawback. In the presence of misleading structures, the SC will make local decisions, which do not reflect the best global choice given the data it has yet to see. In step (4), the SC roughly segments the nearby cord, and thus examines the adjacent data, but only a few pixels ahead. However, once the optimization is made, it never goes back to the start and re-evaluates the optimization decisions it made earlier on. In contrast, the LW approach

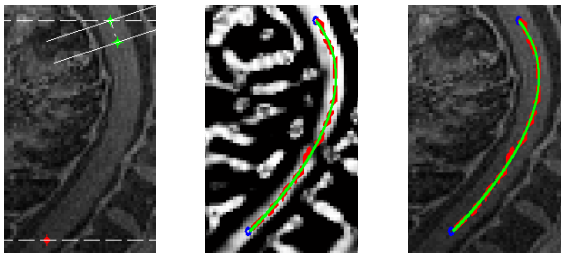

Figure 2. The application of the LW method to a MR scan of the spinal cord. From left to right: the initial start points, in green, and the end point, in red, of the user selecting the C3 and C7 vertebra; the resulting path overlaid on the vesselness measure image; and the same path on the MR scan. Red lines along the path indicate the discrete LW path, while the green curve represents the spline interpolated result.

[18] sacrifices local optimality for global optimality; it ensures the best optimization across the path, but does not locally filter the data in an optimal manner. A hybrid approach is, thus, a natural extension of both methods, where we trust the LW to make the overall decision, but the SC to make the localized ones. In the remainder of this section we will show how we equip SC with a guidance path through a minimal path technique that circumvents the SC myopia problem. We first review the LW minimal path approach, before we detail how we fuse it with SC.

B. Minimal Path Guidance:

In order to improve the SC's path in 3D, we create a user-driven SC guidance path. This guidance path is a 1-D smooth curve in 3D that is designed to be maximally medial to the spinal cord. To this end, we develop a minimal path approach that is a variant of the original LW [18]. In the proposed approach, the cost of the path is inversely proportional to the path's smoothness and to the tubularness values of the image along the path. The input from the user to this minimal path algorithm is simple and attained in mere seconds. In what follows we outline how the process works.

i) The user selects a single planar sagittal slice, easily positioned through the approximate middle of the cord. The cord should roughly lie within a single sagittal plane, as prior to each scan the patient is positioned and strapped in a way such that the neck is held straight. Though the cord's medial-axis may deviate slightly from the selected plane, we found this method of initialization to be very fast, and yet accurate enough to give a global context to the SC that greatly improved its accuracy, see section III.

ii) The user supplies three seed points in the selected plane, with three quick mouse clicks. The area of the spinal cord we study begins at the cervical segment of the cord bounded by the vertebral disc separating the C2 and C3 vertebrae, and ends at the border of the thoracic segment separated by the C7 and T1 vertebrae (Fig. 2). The first two points, green points in Fig. 2, are at the superior end of the target area C2/C3, spaced a few millimeters apart. The first point is where the LW and the SC will start, and the second point is simply used to provide an initial direction for the SC. The third point, red point in Fig. 2, is placed at the end of the study area, C7/T1.

iii) A minimal path algorithm is employed to connect the first and third seeds.

iv) The minimal path, C, is smoothed using a polynomial spline for subpixel accuracy, since the discrete optimization of C forces it to pass through pixel centers only. If desired, the spline points may be adjusted by the user. The end result is a LW curve (in 3D), under the assumption that the sagittal slice on which the curve lies approximately bisects the spinal canal.

In more detail, the path optimization takes place as follows. The mid-sagittal plane pixels are represented as vertices V and the pixels' neighborhood relations (8-connected in 2D) as edges, E, in a graph $G(V, E)$. The globally optimal path is then found based on Dijkstra's algorithm. The discrete path $C = \{v_1, v_2, \cdots, v_N\}$ minimizes the cost function

$$E(C) = \sum_{(v_i, v_{i+1}) \in C} E_b(v_i, v_{i+1}) + \lambda \sum_{v_i \in C} E_u(v_i) \qquad (1)$$

where $E_b(v_i, v_j) = |v_i - v_j|_2$ is the Euclidean distance between v_i and v_j favoring shorter and smoother contours. $E_u(v_i)$ is a measure inversely proportional to the tubularness at the pixel in the sagittal plane corresponding to vertex v_i, and is calculated based on Frangi's vesselness measure [19]. In particular, the vesselness measure at vertex v_i is obtained by:

$$E_u(v_i) = 1 - \max_{s_{min} \leq s \leq s_{max}} \nu(v_i, s) \qquad (2)$$

where $[s_{min}, s_{max}]$ is the range of scales over which the filter operates, and $\nu(v_i, s)$ is the vesselness defined by eqs. (14) and (15) in [19]. In essence, the filter measures ratios between eigenvalues calculated from the Hessian matrix computed at scale s, and pixel v_i. An example response image is shown in Fig. 2. λ controls the trade off between E_b and E_u, and was empirically set to a single value for all data. Though different values might improve the curve slightly, the user can easily correct an incorrect path with a few simple clicks to account for a slightly suboptimal λ (see step iv). The whole process including any necessary corrections took less than one-minute in total for each scan.

C. LW-Guided SC Hybrid Approach:

With the path obtained from LW, we must integrate that information into steps (5) and (6) of the SC. First, the disagreement between the SC's hypothesis, p_{SC}, and the nearest point on the LW guidance path, p_{LW}, is measured under Euclidean distance. A small degree of disagreement

can be attributed to local error in the LW guidance path, resulting from local optimality being sacrificed for global optimality. A large degree of disagreement could only be attributed to erroneous local optimization of the SC which should lead to a decrease in confidence of the SC's hypothesis. Therefore, we use a scoring function to decide the relative weights between the SC's hypothesis and the LW path. The function is built as follows:

5a) Calculate $m = \|p_{SC} - p_{LW}\|$, the Euclidean distance between the LW guidance path and the SC.

5b) Let n be a threshold of distance past which the SC's proposed direction of movement is ignored. This value should be set relative to some fraction of the vertebral width (we empirically set it to 20% of the anteroposterior width of the vertebra).

5c) Set $x = m/n$, and $y = x$ if $x < 1$, and otherwise $y = 1$.

5d) Calculate a scoring function and obtain the final result, $p_{new} = (1-y)p_{SC} + y \times p_{LW}$, and $\vec{v} = (1-y)\vec{v}_{SC} + y \times \vec{v}_{LW}$ where $\vec{v}_{LW}$ is the vector tangent to the LW curve at point p_{LW}.

These steps, (*a-d*), are thus incorporated into the original method as steps (*5.a-5.d*), and our new p_{new}, $\vec{v}$ are used as the new location and normal of the next elliptical layer, respectively.

In addition to the above improvement, when applying the SC to novel data we also observed that, at times, bright neighboring structures can cause problems with the deformation step (step 1), in which the boundary of the SC mesh is supposed to be fit to the boundary of the spinal cord. The primary problem was parts of the mesh being pulled out of place, causing an almost sheering effect. We solve this problem by re-fitting an ellipse to the cross-section of the mesh every few deformations using the least-squares ellipse fitting method detailed in [20].

III. Results and Discussions

To validate our method, we have acquired 20 T1-weighted Magnetization-Prepared Gradient Echo (MP-RAGE) scans with a dedicated spine coil in the sagittal plane. Each volume is 256×256×60 in dimension, with voxel sizes of 0.9760×0.9760 and 1mm thickness. There are 10 healthy subjects acquired with a 3.0T scanner, and 10 MS patients acquired with a 1.5T scanner. We used two scanners with different field strengths and two groups with different disease states to test the robustness of our method under various conditions. We expect the MS cords to be more difficult to segment due to atrophy and patient movement during the scan. For each scan, a ground truth segmentation was computed using a modified version of the semi-automatic method by Tench et al. [7], which uses edge detection and region growing. This method has been used in a large multi-center MS clinical trial [21] to show the relationship between spinal cord atrophy and patient disability in long-term MS patients. It is also highly interactive, allowing the user to work slice-by-slice and adjust the edge-strength thresholds to find values that yield the optimal representation of the cord/CSF boundary. If necessary, the user can also add or delete points to refine the boundary, which is often required if the scan is noisy or if the slice is in a region in which the cord abuts the spinal canal. With all these factors combined, we believe the method is thus appropriate for use as our gold standard when performed by an expert operator. Results with this method took on average 35 minutes of user-interaction for the MS scans and 15 minutes for the healthy scans.

To isolate improvements from LW versus those from the elliptical fitting, we compare the original SC to both our new SC with and without dynamic LW paths. For all results, SC initialization required under a minute of user-interaction per scan; a significant improvement over the semi-automatic method of Tench et al. [7] at 25 minutes-per-scan on average, and the initialization time for the semi-automatic method of Horsfield et al. [14] at 5 minutes-per-scan. Once initialized, our method took 2.5 minutes-per-scan running fully automatically, for a total of 3.5 minutes-per-scan including user-interaction; a full 1.5 minutes faster than state-of-the art methods [14], despite segmenting an extra vertebral level. Our per-slice run-time is also very fast, at 1.9 seconds-per-slice.

We examine both surface accuracy, and volume estimation accuracy for the three tested methods. Our results for surface accuracy of the three configurations of the SC are shown in Table I and Fig. 3. The *Orig. SC* is the original SC method proposed by McIntosh and Hamarneh in [11], *SC 0.0 LW* is our new SC with elliptical refitting but no LW contribution, and *SC Dyn. LW* is the proposed method of weighting the LW contributions by its difference from the SC's hypothesis as described in section II. We note that for 2 MS scans both the *Orig. SC* and *SC 0.0 LW* failed to produce any kind of usable segmentation, and these datasets were left out of Fig. 3 and all of the resulting statistics for all methods. The proposed method, *SC Dyn. LW*, did not fail for any scans, a very significant improvement that already shows its increased applicability to more challenging data. For reference, we will make note of the results that *SC Dyn. LW* obtained for these two data sets as we discuss each accuracy measurement.

Table I presents two related error measurements, both in millimeters. The first is the forward-Hausdorff distance, *H. Dist.*, which is the maximum inter-surface pairwise distance calculated between the automatic segmentation and the closest point in the ground truth segmentation. More formally, $H.Dist = \max_{a \in A} \left(\min_{b \in G} D(a,b) \right)$ where A is the automatic segmentation, G is the ground truth, and $D(a,b)$ is the Euclidean distance between two 3D points a, b. In other words, each point in the automatic segmentation is assumed to correspond to the closest point in the ground truth, and

Table I
SURFACE ERROR COMPARISONS FOR HEALTHY/MS SUBJECTS. SEE TEXT FOR DETAILS. TOP PERFORMANCE FOR EACH COLUMN, SUBJECT TYPE, AND METRIC IS HIGHLIGHTED IN RED.

	Method	Mean	Median	Min	Max	Std. Dev.
H. Dist.	**Orig. SC**	4.02 / 6.43	4.54 / 2.31	1.75 / 1.41	6.60 / 28.82	1.73 / 9.41
	SC 0.0 LW	3.50 / 4.77	3.02 / 2.75	1.66 / 1.44	7.61 / 19.09	1.94 / 5.92
	SC Dyn. LW	1.48 / 1.59	1.40 / 1.39	1.33 / 1.33	1.95 / 2.93	0.20 / 0.51
M. Dist.	**Orig. SC**	1.02 / 2.45	0.98 / 0.71	0.72 / 0.59	1.71 / 13.95	0.28 / 4.66
	SC 0.0 LW	0.90 / 0.97	0.86 / 0.71	0.69 / 0.53	1.22 / 2.94	0.17 / 0.81
	SC Dyn. LW	0.52 / 0.46	0.56 / 0.45	0.33 / 0.37	0.68 / 0.62	0.12 / 0.07

Table II
VOLUME ACCURACY COMPARISONS FOR HEALTHY/MS SUBJECTS. SEE TEXT FOR DETAILS.TOP PERFORMANCE FOR EACH COLUMN AND SUBJECT TYPE IS HIGHLIGHTED IN RED.

Method	Mean	Median	Min	Max	Std. Dev.
Orig. SC	82.27% / 71.49%	85.54% / 86.60%	46.91% / 23.14%	99.34% / 96.62%	18.35% / 30.63%
SC 0.0 LW	89.99% / 79.66%	92.93% / 77.37%	74.35% / 62.56%	97.63% / 98.44%	8.28% / 12.19%
SC Dyn. LW	91.60% / 94.31%	94.83% / 94.24%	77.33% / 87.20%	98.06% / 98.83%	7.64% / 4.16%

then the largest distance between all corresponding pairs is reported. Columns of the table represent different statistics measured over the data sets, presented as Healthy/MS. As already mentioned *Orig. SC*, *SC 0.0 LW* failed for a few data sets, and the results for those sets are excluded from the table so as not to distort their statistics. The proposed *SC Dyn. LW* performed well with this metric, with a mean Hausdorff-distance across all 18 data sets of 1.48mm for healthy scans and 1.59mm for those with MS, more than twice as good as the other methods (4.02mm/6.43mm for *Orig. SC* and 3.50mm/4.77mm for *SC 0.0 LW*). For the two MS scans the comparison methods were unable to segment, *SC Dyn. LW* achieved 1.39mm/1.95mm, better than the average performance of the failing methods.

There is measurable reduction in mean Hausdorff-distance from *Orig. SC* to *SC 0.0 LW*, which can be attributed to the elliptical shape re-fitting step. A much larger improvement is obtained with *SC Dyn. LW*, attributed to the better localization of the spinal cords medial-axis. Specifically, the largest improvement over *Orig. SC* occurs in areas where localized decisions are misleading. Fig. 4 shows one such example (bottom two panels). Without LW to provide a global context of the path, local image noise caused by bright extraspinal structures pulls the SC off course. Hence why both the maximum errors and std. dev. are much lower for the *SC Dyn. LW*.

Table I also presents the mean inter-surface pairwise distance computed against the ground truth segmentation, *M. Dist.*. In comparison to the Hausdorff-distance, which reports the maximum, this metric reports the average distance, $M.Dist = \underset{a \in A}{mean}\left(\min_{b \in G} D\left(a, b\right)\right)$. This gives a better impression of the average inter-surface distance as opposed to the maximum inter-surface distance. For example, the mean mean-distance across all 18 data sets for the proposed *SC Dyn. LW* was 0.52mm for healthy scans and 0.46mm for those with MS, approximately twice as good when compared to 1.02mm/2.45mm for *Orig. SC* and 0.90mm/0.97mm for *SC 0.0 LW*. For the two MS scans the comparison methods were unable to segment, *SC Dyn. LW* achieved 0.40mm/0.49mm, better than the average performance of the failing methods.

Spinal cord volume computations in Table II are given as ratios of $1 - \min\left(\left|(O/G) - 1\right|, 1\right)$ where O is the obtained volume and G is the ground truth volume, respectively. The $\min$ function treats results of more than a 200% overestimate, i.e. $O \geq 2G$, as being equivalent to 0% accuracy. Statistics are calculated across all the 18 data sets, presented as Healthy/MS. So, on average *SC Dyn. LW* achieved 92% accuracy in volume estimation for healthy scans, and 94% for MS scans. Again, significant improvments are seen over *Orig. SC* and *SC 0.0 LW*, which acheived $82\%/72\%$ and $90\%/79\%$, respectively. For the two MS scans the comparison methods were unable to segment, *SC Dyn. LW* achieved 98%/89%, which is again better than the average performance of the failing methods.

These results represent state-of-the-art degrees of accuracy. Namely, Horsfield et al. report a systematic overestimation of 14%, or an accuracy of 86%, when compared to the Losseff method [6], which already overestimates the cross-sectional area by 4.3% to 10.8% depending on the cord diameter, for a mean of 7.6% as shown by Losseff [6]. Our

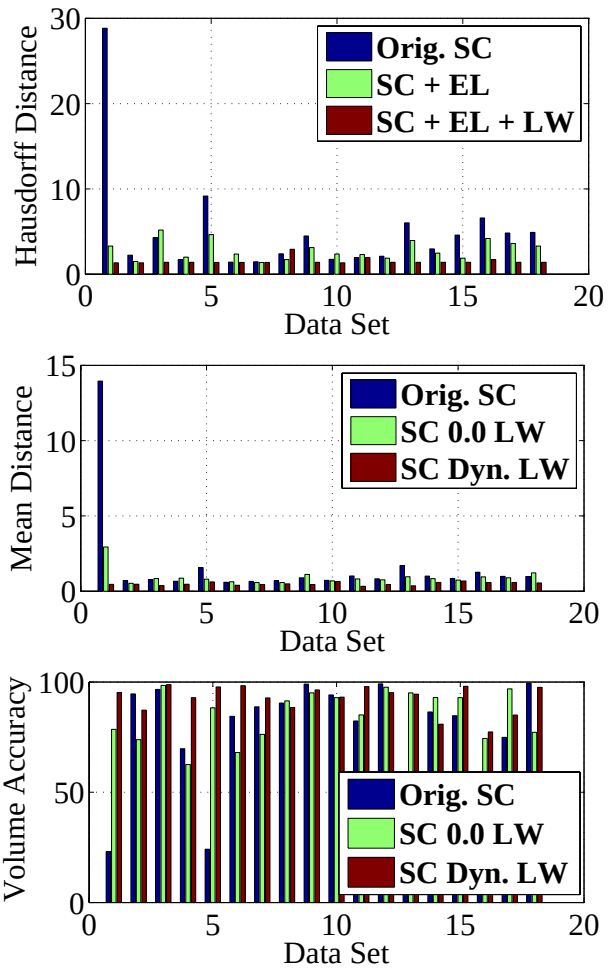

Figure 3. Bar plots comparing results obtained under the three metrics. See text for details.

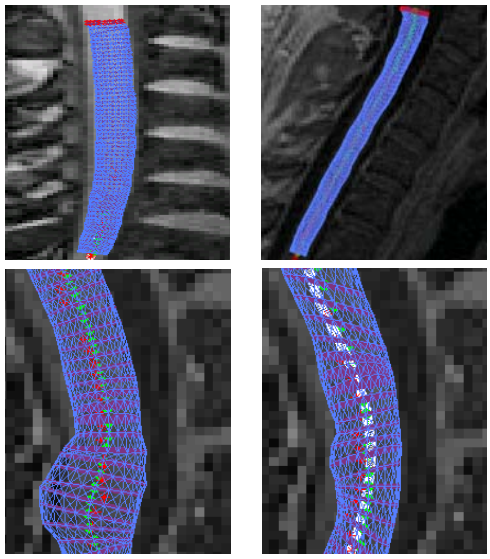

Figure 4. Sample 3D results. From left to right, top to bottom, a healthy scan, a MS scan, a close up showing an MS scan without LW, and a close up of the same scan with LW influence. Notice how the LW influence pulls the SC back on track and reduces bulging of the mesh where the cord abuts the canal.

chosen gold standard, the method developed by Tench et al., has important advantages, most notably for partial volumes, and is thus much more consistent with a lower mean error of 3.15% [2]. We believe the method by Tench et al. is a better comparison because it provides more accurate ground truth. In summary, in combination with the estimated ground-truth errors, our method overestimates volume by approximately $8\% + 3.15\% = 11.15\%$, in comparison to Horsfield et al.'s method at $14\%+ 7.6\% = 21.6\%$. Further, as we noted, there is evidence from a recent post-mortem study that suggests MS affects different parts of the cord differently, and thus it would be beneficial to analyze a larger region [9]. We validate on a much larger (70-80 slices) and more difficult (C3-C7) area of the cord than analyzed by Horsfield et al. who segmented C2-C5, but only validated 5 slices at C2.

Furthermore, our accuracy in volume estimation compared favorably to [10], where they reported an accuracy of 74% in comparison to our 92% for healthy scans and 94% for those with MS, a significant improvment. It should, however, be noted that [10] used a 1.5T scanner, whereas we used a 3.0T for our healthy scans and a 1.5T for MS. They also focused their volume analysis on C1-C3, whereas we study C3-C7, under the premise that different areas may respond to MS differently [9]. Lastly, we note that C1-C3 is typically the easiest region to segment due to the strong signal and low chance of the cord contacting the canal, while C5-C7 is the most difficult. We choose C3-C7 simply because the C2/C3 disc makes a stable superior landmark.

IV. Conclusions

A variety of medical image segmentation algorithms exist, but seldom is any single algorithm able to address a complex image segmentation problem. In this paper, we focused on the challenging problem of spinal cord segmentation and made the following observation that allowed us to propose a novel and robust algorithm. Namely, we observed that different image regions require different algorithm parameter settings, e.g. filtering parameters, and, for those cases, a locally optimal algorithm is ideal. However, local optimality came at a price: missing the global image context. Therefore, we combined an enhanced crawler method (artificial life segmentation framework with optimal local tubularness filters), with LW (minimal path guaranteeing global optimality). We obtained superior results using our hybrid method compared to those without any LW guidance. A noticeable improvement is how our guided-crawler can segment down to C7, whereas competing methods often fail around C5. For example, note how the failure in Fig. 1 is rectified in the second image of Fig. 4. Note also the close-up of LW

improving the SC's path in the last two images in Fig. 4. Our method also proved to very fast; requiring less than a minute of user interaction, followed by 2.5 minutes to obtain the final result. Furthermore, our accuracy in volume estimation compared favorably to state-of-the-art methods [10], [11], [14]. Future work includes demonstrating clinical significance of our method by correlating with disease progression, and conducting a study on an even larger dataset.

V. Acknowledgements

We would like to thank Emilie Mackie for coordinating data acquisition for the healthy volunteers; and Biogen Idec Canada and the Natural Sciences and Engineering Research Council of Canada (NSERC) for providing funding for RT. GH is funded by NSERC, and CM is funded by an NSERC Graduate (Doctoral) Scholarship and a Michael Smith Foundation for Health Research (MSFHR) Senior Graduate Studentship.

References

[1] R. Zivadinov and R. Bakshi, "Central nervous system atrophy and clinical status in multiple sclerosis," *Journal of Neuroimaging*, vol. 14(3 Suppl), pp. 27S–35S, 2004.

[2] R. Mann, C. Constantinescu, and C. Tench, "Upper cervical spinal cord cross-sectional area in relapsing remitting multiple sclerosis: application of a new technique for measuring cross-sectional area on magnetic resonance images," *Journal of Magnetic Resonance Imaging*, vol. 26(1), pp. 61–65, 2007.

[3] F. Song, Y. Huan, H. Yin, Y. Ge, G. Wei, and et al., "Normalized upper cervical spinal cord atrophy in multiple sclerosis," *J Neuroimaging*, vol. 18, pp. 320–327, 2008.

[4] R. Rudick, "Impact of disease-modifying therapies on brain and spinal cord atrophy in multiple sclerosis," *Journal of Neuroimaging*, vol. 14(3 Suppl), pp. 54S–64S, 2004.

[5] D. Miller, D. Soon, K. Fernando, D. MacManus, G. Barker, T. Yousry, and et al., "Mri outcomes in a placebo-controlled trial of natalizumab in relapsing ms," *Neurology*, vol. 68, pp. 1390–1401, 2007.

[6] N. Losseff, S. Webb, J. O'Riordan, R. Page, and et al., "Spinal cord atrophy and disability in multiple sclerosis. a new reproducible and sensitive mri method with potential to monitor disease progression," *Brain*, vol. 119(Pt 3), pp. 701–708, 1996.

[7] C. Tench, P. Morgan, and C. Constantinescu, "Measurement of cervical spinal cord cross-sectional area by mri using edge detection and partial volume correction," *J Magn Reson Imaging*, vol. 21, pp. 197–203, 2005.

[8] J. Carbonell-Caballero, J. Manjn, L. Mart-Bonmat, J. Olalla, B. Casanova, M. de la Iglesia-Vay, and et al., "Accurate quantification methods to evaluate cervical cord atrophy in multiple sclerosis patients," *MAGMA*, vol. 19, pp. 237–246, 2006.

[9] C. Gilmore, G. DeLuca, L. B, T. Owens, and et al., "Spinal cord atrophy in multiple sclerosis caused by white matter volume loss," *Arch Neurol.*, vol. 62, pp. 1859–1862, 2005.

[10] S. Hickman, O. Coulon, G. Parker, G. Barker, and et al., "Application of a b-spline active surface technique to the measurement of cervical cord volume in multiple sclerosis from three-dimensional mr images," *J Magn Reson Imaging*, vol. 18, pp. 368–371, 2003.

[11] C. McIntosh and G. Hamarneh, "Spinal crawlers: deformable organisms for spinal cord segmentation and analysis," *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2006*, pp. 808–815, 2006.

[12] B. D. Schmit and M. K. Cole, "Quantification of morphological changes in the spinal cord in chronic human spinal cord injury using magnetic resonance imaging," *IEEE EMBS*, vol. 26, pp. 4425 – 4428, 2004.

[13] O. Coulon, S.J.Hickman, G. Parker, G. Barker, D. Miller, and S. Arridge1, "Quantification of spinal cord atrophy from magnetic resonance images via a b-spline active surface model," *Magnetic Resonance in Medicine*, vol. 47, pp. 1176–1185, 2002.

[14] M. A. Horsfield, S. Sala, M. Neema, M. Absinta, A. Bakshi, M. P. Sormani, M. A. Rocca, R. Bakshi, and M. Filippi, "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.

[15] L. Cohen and R. Kimmel, "Global minimum for active contour models: a minimal path approach," in *Computer Vision and Pattern Recognition, 1996. Proceedings CVPR '96, 1996 IEEE Computer Society Conference on*, Jun. 1996, pp. 666 –673.

[16] O. Wink, W. J. Niessen, and M. Viergever, "Multiscale vessel tracking," *IEEE Transactions on Medical Imaging*, vol. 23, no. 1, pp. 130–133, 2004.

[17] H. Li and A. Yezzi, "Vessels as 4-d curves: Global minimal 4-d paths to extract 3-d tubular surfaces and centerlines," *Medical Imaging, IEEE Transactions on*, vol. 26, no. 9, pp. 1213 –1223, 2007.

[18] W. A. Barrett and E. N. Mortensen, "Interactive live-wire boundary extraction," *Medical Image Analysis*, vol. 1, no. 4, pp. 331 – 341, 1997.

[19] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, "Multiscale vessel enhancement filtering," *Lecture Notes in Computer Science*, vol. 1496, pp. 130–137, 1998.

[20] M. Fitzgibbon, A. W.and Pilu and R. B. Fisher, "Direct least-squares fitting of ellipses," *IEEE PAMI*, vol. 21, no. 5, pp. 476–480, May 1999.

[21] D. L. D, G. Ebers, A. Traboulsee, R. Tam, D. Goodin, and A. Konieczny, "Interferon beta-1b 16-year long-term follow-up study: Mri outcomes," *Multiple Sclerosis (special supplementary issue of ECTRIMS 2006 oral presentations)*, vol. 12, pp. S188–189, 2006.