Method for Automatically Segmenting the Spinal Cord and Canal from 3D CT Images*

László G. Nyúl[1], Judit Kanyó[2], Eörs Máté[1], Géza Makay[3], Emese Balogh[1], Márta Fidrich[2], and Attila Kuba[1]

[1] Department of Image Processing and Computer Graphics, University of Szeged,
Árpád tér 2, H–6720 Szeged, Hungary
{nyul, kuba, mate}@inf.u-szeged.hu

[2] Clinical Software Engineering, GE Medical Systems,
Lajos u. 48-66, H–1036 Budapest, Hungary
{marta.fidrich, judit.kanyo}@med.ge.com

[3] Institute of Mathematics, University of Szeged,
Aradi vértanuk tere 1, H–6720 Szeged, Hungary
makayg@math.u-szeged.hu

Abstract. We present two approaches for automatically segmenting the spinal cord/canal from native CT images of the thorax region containing the spine. Different strategies are included to handle images where only part of the spinal column is visible. The algorithms require one seed point given on a slice located in the middle region of the spine, and the rest is automatic. The spatial extent of the spinal cord/canal is determined automatically. An extended region-growing technique is suggested for segmenting the spinal canal while active contours are applied if the spinal cord is to be segmented. Both methods work in 2D and use propagated information from neighboring slices. They are also very rapid in execution, that means an efficient, user-friendly workflow. The methods were evaluated by radiologists and were found to be useful (in reducing/eliminating contouring labor and time) and met the accuracy and repeatability requirements for the particular task.

1 Introduction

In case of radiation treatment (RT) planning, CT imaging is generally used because image voxel gray values (Hounsfield Units) are in direct function of radiation absorption and therefore can be used directly in dose calculation. In RT planning, clinicians (radiologists, dosimetrists or radiotherapists) must trace the outline of a few critical structures on a large number of images. The time and labor increases significantly with the number of image slices, and the number and sizes of the organs in the anatomical area of interest. The quality of the contouring and then the produced 3D objects depend on the resolution and contrast of the 2D images, and on the knowledge and judgment of the clinician performing the segmentation. Using automated image segmentation could

* This work was supported by GE Medical Systems.

A. Gagalowicz and W. Philips (Eds.): CAIP 2005, LNCS 3691, pp. 456–463, 2005.

save tremendous time and effort. Also, automated segmentation could increase precision by eliminating subjectivity of the clinician.

One of the key regions that must be protected during the irradiation treatment is the spinal cord/canal. The difficulty of the automatic segmentation is caused partly by vertebrae that have open cross-sections in the image, partly by the fact that below the level of pelvic bones the cord is no longer situated in the spine.

There are several approaches in the literature to the segmentation of the spinal cord from CT images. The segmentation approach of [1] is based on 2D boundary tracking. It requires an initial point to start tracing the edge. The initial point travels to the vertical or horizontal direction until an edge is reached. Then the algorithm starts to examine the surrounding pixels of that edge and check whether they belong to the current edge or not. The algorithm uses a constant threshold selection which is hard to find (due to the partial volume avaraging effect).

Another approach [2] relies on a knowledge-base which consists of an Anatomical Structures Map and a task-oriented architecture, the Plan Solver. The anatomical structures map contains a frame-like knowledge representation of the macro-anatomy in the human thorax. The plan solver is responsible for determining the position, orientation and size of the structures of interest to radiation therapy. The plan solver relies on a number of image processing operators. However, a general decision making system like the Plan Solver, a method using artificial intelligence could be far from being efficient.

2 Methods

The segmentation procedure described in this paper was devised to work on native CT images of the thorax region containing the spine. Different strategies are included to handle images where only part of the spinal column is visible. The method comprises the following main steps:

1. *Initialization:* The purpose of seed point selection is to determine the starting slice and to provide some localization hint for the segmentation algorithm.
2. *Pre-processing:* Determination of the extent of the cord in the spine between the head and the pelvic bones yields a transaxial slice range, the region of interest (ROI).
3. *2D segmentation on the starting slice:* Cross-section of the spinal cord/canal at the starting slice is segmented by either an active contour or a region growing type algorithm.
4. *Segmentation on other slices:* Repeating propagation and constrain of the 2D segmented region onto the subsequent slice and segmentation on that slice by active contour or region growing, slice-to-slice in upward (toward the head) and downward (toward the feet) directions within the extent yields the final 3D result.

2.1 Determining the Extent of the Cord in the Spine

Seed point is given on a slice located in the middle region of the spine. The slice should contain a 'nice' vertebra, i.e., consisting totally of bone and no cartilage at all. Effectively this assures that the contour of the vertebra is sharp. The extent of the spinal cord/canal is determined by finding the 'upper' end (in the neck region) and the 'lower' end (in the pelvic region) if visible in the image volume. Automatic recognition of the extent is reasonable. Different strategies are used for the two ends.

Extent Towards the Head. Although in many images only the lower part, the pelvic region of the spinal cord is visible there are situations when the upper end is present in the image volume. Here the task is basically the separation of the spinal cord from the brain. This can be reliably done using the bony structures within the region as guides. The region contains the shoulder blades and the collarbone, which are relatively large volumes of bone tissue. The neck contains only a few vertebrae (small bone volume). The skull and jaw are again of relatively large volume. The spinal cord is present in the shoulder and neck region but is not present inside the skull; therefore the extent should be limited to the level below the skull. This can be determined by computing the volume of the bone tissue in each transaxial slice (e.g., by simple thresholding) starting at the seed point level and proceeding toward the head. The level where the bone volume starts to increase considerably after the major decrease indicates the skull base.

Extent Towards the Pelvic Bones. One of the main problems in the automatic segmentation of the spinal cord is to recognize the lower (inferior) end of the spinal canal, where segmentation should be terminated to avoid undesired behavior (e.g., leakage). The main idea of the method is to use the shape of the spine to determine the extent of the cord in the pelvic region. First the apexes of the vertebrae are located on each slice by using the image intensities and some anatomical knowledge of the spine curve. The algorithm seeks for the location of a specific curvature pattern in the spine to determine the lower end of the cord. (see Fig. 1)

2.2 Active Contour Based Segmentation with Propagation

The idea of the active contour or 'snake' algorithm is that a closed curve will best separate the object of interest from its surroundings when its placement and shape is such that an energy function, defined over the boundary, is minimized. [3,4] The contour is described parametrically by $\mathbf{v}(s), s \in [0, 1]$. The snake energy is

$$E_{\text{snake}}(\mathbf{v}(s)) = \int_{s=0}^{1} E_{\text{internal}}(\mathbf{v}(s)) + E_{\text{image}}(\mathbf{v}(s)) + E_{\text{constraint}}(\mathbf{v}(s))\, ds,$$

where internal, the *internal energy* imposes curvature (smoothness) constraints, E_{image}, the *image energy* attracts the contour to the desired features (edges)

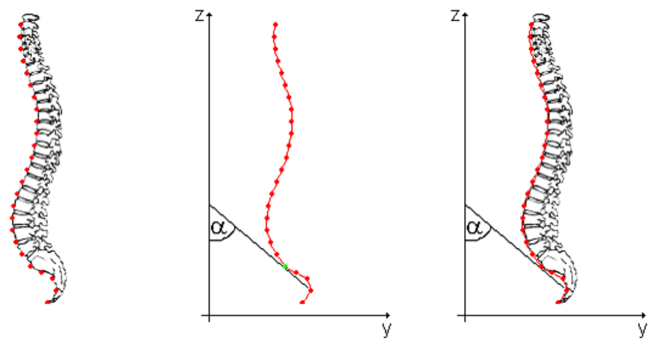

(a) (b) (c)

Fig. 1. Spinal cord with marked apexes of the vertebrae (a) and determination of the stopping angle (b), and the spine image superimposed on the curve function (c)

of the image, and $E_{\text{constraint}}$, the *constraint energy* allows other geometric constraints to be applied.

We chose to solve the discrete problem using a more or less discrete method following some ideas of [5]. Instead of energy minimization, we prefer working with the effects of the energies, i.e., we apply discrete forces on each boundary point and look for equilibrium of these forces.

To obtain the segmented 3D object, first, the user selects a seed point in one of the slices and 2D segmentation is performed on that particular slice. The centroid of the enclosed shape is propagated to adjacent slices and the 2D segmentation is performed on those slices also starting from the propagated seed point.

Forces Used for the Active Contour. At each iteration step, various forces are applied to each contour point p, until points stop moving considerably or maximum number of iteration is reached:

$$\mathbf{F}(p) = \mathbf{F}_{\text{image}}(p) + \mathbf{F}_{\text{constraint}}(p) + \mathbf{F}_{\text{inflation}}(p) + \mathbf{F}_{\text{angle}}(p) + \mathbf{F}_{\text{distance}}(p).$$

Image force is used to attract the contour to edges. We compute the image force as a smoothed gradient from the negative gradient norm of the image. *Constraint force* is used to constrain the size and the shape (circular) of the segmented region. The farther the current point is from the centroid of the enclosed polygon, the larger force is applied to pull it back. *Local inflation force* pushes the contour outwards. When a point does not move considerably (i.e., it is stuck in a local minimum), this force pushes the point outward with a small constant. *Angle force* is used to smooth the contour. We compute the deviation from 180 degrees of the angle at each contour point and penalize large deviations and concavities. *Distance force* aims at evenly distributing contour points. We compute the deviation of the distance of two adjacent points from the average distance of the nearby (or all) adjacent points. This pulls the actual point into its tentative middle position between the neighbors.

Cut and merge operations also aim at evenly distributing contour points and at the same time refining the contour. If the length of an edge is too large, the edge is cut and a new contour point is created. If the length of an edge is too small, it is removed by merging its two endpoints.

Propagation to Other Slices. Since the method works in 2D and it requires a seed point within the slice to start with, we need to provide seed points for the object on each slice. To reduce the user interaction and facilitate automation, user-defined seed point is required only on the initializing slice. For subsequent slices the method automatically generates seed points based on the segmented regions from neighboring slices. The centroid of the enclosed region is propagated to adjacent slices. If it is detected to fall out of the spinal cord (by using image intensity measures), the algorithm will correct for the bending of the spine by extrapolation from the positions of the centroid on the actual slice and on the previous slice.

2.3 Region Growing Based Segmentation with Propagation

This technique uses the standard 'textbook' region-growing algorithm with a stopping criterion that combines local and global differences between gray values of voxels, as well as a technique to add geometrical constraints based on the anatomical knowledge about the organ being segmented.

The region growing method is prone to 'leaking out' if object boundaries are not well defined in the image data. Also, since segmentation is done slice by slice and subsequent iterations are initialized using results on preceding slices, one bad slice could ruin all subsequent results. Our algorithm will mostly prevent these failures.

Segmenting the First Slice. The user is required to select an initial slice where the spinal cord is totally enclosed by a vertebra, thus guaranteeing that the initial segmentation of this slice will not leak out. The object is roughly segmented on the initializing slice, starting from the seed point, and following a conservative thresholding strategy where the thresholds depend on the intensity value of the seed point.

The criterion that is used to stop the region-growing algorithm depends on local and global features, which describe the homogeneity of the segmented region. These are represented by parameters m and M:

$$m = \max_{v \in R} |I(S) - I(v)| \qquad M = \max_{v_1, v_2 \in R} |I(v_1) - I(v_2)|,$$

where R is the segmented region, S is the seed point, and $I(v)$ denotes the gray value at voxel v.

Segmenting the Other Slices. Although this variant of region growing operates in 2D, it can be used to extract a whole 3D volume by slice by slice starting off at the slice containing the initial seed point, and traversing every transaxial slice above and below the initial slice. This is realized by propagating a point to be the new seed point, utilizing the statistical features of the initial slice, and applying several constraints to the new region.

Seed propagation is performed analogously to that used with the active contour method described above. The region growing criterion is set up that a voxel v that is neighbor of a region boundary voxel b is included if

$$\alpha\frac{D}{M} + \beta\frac{d}{m} < T,$$

where $D = |I(v) - I(b)|$ and $d = |I(v) - I(S)|$, α and β are weights and T is a pre-specified threshold.

Before the iteration continues, every segmented region must be evaluated to detect leakage and to embed some a priori information about the organ in question. The spinal canal has a tubular structure and in each slice it appears almost circular, therefore, a circular mask of the approximate size of the spinal canal is applied to the segmented region to reduce false positives.

3 Results

Note that the segmented 'spinal canal' is on average 18–20% larger than the segmented 'spinal cord'. Therefore, the two segmentations cannot be directly compared. For radiotherapy planning, any segmented region falling between the borders of cord and canal can be accepted as accurate. Figure 2 shows a few transaxial cross-sections and a sagittal slice with the segmented spinal cord indicated over the original CT images.

There were 27 image volumes included in our studies and three operators performed manual contouring for producing the 'gold standard'. The measure for accuracy was computed by means of true positive volume fraction (TPVF), false positive volume fraction (FPVF), and false negative volume fraction (FNVF) using the 'gold standard' as the true volume.

Operator A performed the same segmentation task 3 times (A1, A2, A3). For each pair (A1-A2, A1-A3, A2-A3), an overlap measure was computed. The measure for intra-operator reproducibility was computed as the mean and the standard deviation of the computed overlap measure among all performed tasks. Inter-operator reproducibility was computed similarly. Three different operators (A, B, C) performed the same segmentation task. Measures are computed for each pair (A-B, A-C, B-C) and for each task and are pooled for statistics. Table 1 shows the results for accuracy and reproducibility.

Accuracy of the active contour method is approximately the same as that of the manual outlining for the spinal cord. This automatic method also shows higher intra- and inter-operator reproducibility. Accuracy and reproducibility of

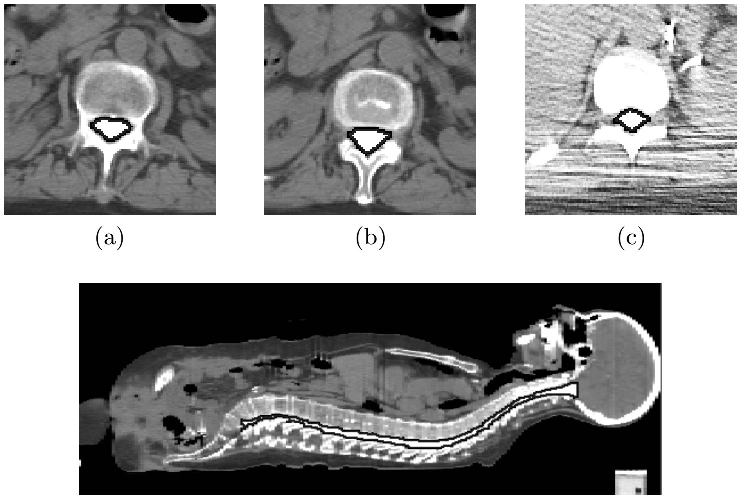

(a) (b) (c) (d)

Fig. 2. Segmentation on the initial slice (a), on an additional slice with open vertebra (b), and sagittal view of the entire segmented spinal canal (d). Segmentation on a slice from a very noisy dataset (c).

Table 1. Accuracy, intra- and inter-operator reproducibility measures for the methods. Mean values and standarad deviation of the measures are expressed in percents.

Region	Method	TPVF	FPVF	FNVF	Intra-operator	Inter-operator
Spinal cord	Manual	97.94 (1.49)	2.06 (1.31)	2.06 (1.49)	93.35 (1.75)	91.90 (2.63)
	Auto (AC)	97.59 (2.59)	2.62 (2.01)	2.41 (2.59)	94.77 (1.94)	94.53 (1.72)
Spinal canal	Manual	97.27 (1.53)	2.73 (1.66)	2.73 (1.53)	Not tested	89.62 (2.54)
	Auto (RG)	95.83 (3.85)	8.79 (6.12)	4.17 (3.85)	Not tested	86.69 (7.20)

the region growing based method is somewhat lower than that of the manual outlining due to leakage in a few cases.

Selecting the initial slice and the seed point can be done in a few seconds per data set. Running time for the semi-automatic segmentations was found to be 25 seconds per study on average for the active contour method and a few seconds per study less for the region growing method. This is much less than the time needed for manual contouring.

4 Discussion

The simple method published in [1] has several drawbacks, which limit its accuracy and usability. It can be very sensitive to where the starting point is placed in a slice image. Also, often images show cross-sections in which the vertebra is not closed, therefore the initial traverse to find a boundary point may fail. Due to the partial volume averaging effect the contour of the spinal canal may not be sharp and clearly identifiable in a slice image at all. It is hard to find proper threshold of general use. Our method requires the selection of a reliable initial section but afterwards it automatically propagates size, shape and intensity constraints.

The size and shape of organs, even the spine, vary a lot and their boundary is not always visible. Thus detection is only possible with some prior information. Instead of using probabilistic approaches (which are known to be not only erroneous but also slow, even in case of a very simplistic model), we incorporate anatomical knowledge into our method. After understanding how radiologists work when they analyze images, we built the radiologists decision making procedure directly into our method.

Since our principal aim was to develop a real-time method that can be used in daily routine work, we optimized on speed instead of generality, unlike the Plan Solver [2]. Also, we use initial human interaction to guarantee the correct starting off of the iterative, propagating procedure, in contrast with [2], where the authors use fully automatic initialization and then try to determine failure of initialization.

5 Conclusions

Both of our algorithms to segment the spinal cord/canal require only one point to start, after that the segmentation is fully automatic and fast. This means a highly efficient, user-friendly workflow (no need to trace a contour on a starting slice, no need for initial model fitting, no need for ROI or slice range selection). The methods were evaluated by radiologists and were found to be useful (in reducing/eliminating contouring labor and time) and met the accuracy and repeatability requirements for the particular task. Since the active contour version performed better of the two in terms of reproducibility and accuracy, that one is currently incorporated into the product.

References

1. G. Karangelis and S. Zimeras. A 3D segmentation method of the spinal cord applied on CT data. *Computer Graphics Topics*, 14(1):28–29, 2002.
2. N. Archip, P. J. Erard, M. Egmont-Petersen, J. M. Haefliger, and J. F. Germond. A knowledge-based approach to automatic detection of the spinal cord in CT images. *IEEE Trans. Med. Imaging*, 21(12):1504–1516, 2002.
3. L. D. Cohen. On active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, 1991.
4. L. D. Cohen and I. Cohen. Finite element methods for contour models and balloons for 2D and 3D images. *IEEE Trans. PAMI*, 15(11):1131–1147, 1993.
5. J. V. Miller, D. E. Breen, W. E. Lorensen, R. M. O'Bara, and M. J. Wozny. Geometrically deformed models: A method for extracting closed geometric models from volume data. In *Computer Graphics (SIGGRAPH'91 Proc.)*, volume 25, pages 217–226, 1991.