Interacting Active Rectangles for Estimation of Intervertebral Disk Orientation

Amer Abufadel[1], Greg Slabaugh[2], Gozde Unal[2], Li Zhang[2], Benjamin Odry[2]
[1]Georgia Institute of Technology, Atlanta GA 30332, USA
amer@ece.gatech.edu
[2]Siemens Corporate Research, Princeton NJ 08540, USA
{greg.slabaugh,gozde.unal,lizhang,benjamin.odry}@siemens.com

Abstract

This paper presents a fast and efficient method to determine intervertebral disk orientation in a magnetic resonance (MR) image of the spine. The algorithm originates from active contour theory and enforces a shape constraint to avoid leaks through weak or non-existent boundaries. The method represents a vertebra as a rectangle, modeled as a semi-affine transformation applied to the unit square. A regional flow integrated along the rectangle's perimeter updates the rectangle's transformation to achieve the segmentation. Further constraints are added so that adjacent rectangles have similar orientation and scale. The orientation of a disk is then inferred from its adjacent vertebrae. Experiments show that the method is fast and effective in detecting the correct intervertebral disk orientation, which is used for transverse image planning.

1 Introduction

MR spine imaging has been widely used for noninvasive detection of different abnormalities and diseases in the spinal column, vertebrae, and intervertebral disks. This paper focuses on setting up transverse image acquisition for diagnosis of intervertebral disk pathologies. In typical MR spine imaging cases, a patient is initially scanned to obtain a set of T2-weighted sagittal images or coronal localizer images. If an abnormality of an intervertebral disk is found, a transverse scan is then performed. The orientation of the transverse images is planned parallel to the major axis of the disk and the center of the transverse images is located on where the disk joins the spinal cord. A saturation band is placed to suppress strong MR signals from abdominal vessels and should not overlap with the spinal column (see Figure 1). Currently transverse imaging planning is done manually. The process, however, is time-consuming and subject to intra- or inter-operator variation. Therefore there is a salient need for automation in transverse imaging planning.

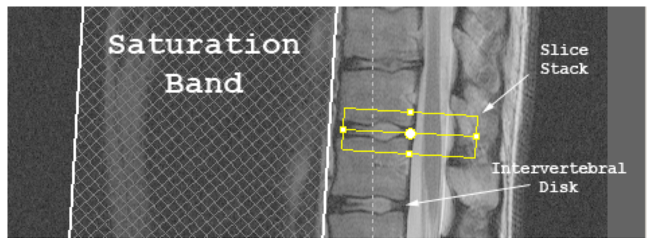

Figure 1. Sagittal view of the vertebral column. The orientation of the intervertebral disk is used to set up the slice stack.

This requires accurate and consistent detection of intervertebral disk orientation and an approximate segmentation of vertebrae. This paper presents a semi-automatic computer-based technique to detect intervertebral disk orientation accurately and to approximate vertebrae by rectangles.

Ideally, the first step in detecting the orientation of an intervertebral disk is to detect the boundaries, or segmentation, of the disk itself. However, this is difficult if there is an abnormality or there are weak or missing boundaries. However, the boundary of the intervertebral disk is closely aligned with the boundaries of the rigid vertebrae it separates. Therefore, we can infer the intervertebral disk orientation by finding the bounding edges of its adjacent vertebrae. Since each vertebrae can be geometrically approximated by a rectangle, we incorporate this a priori shape constraint into our approach to increase the robustness of the solution.

1.1 Related work

Perhaps the most related class of methods are those that perform vertebrae segmentation. A popular imaging modality for vertebrae segmentation and analysis is low-dose X-ray; for example, dual energy X-ray absoptiometry (DXA) [6] and digital videofluoroscopic (DVF) [8] images have been considered. Magnetic resonance images can acquire relatively clear images of the spine without the radiation risk, and is the modality of choice for studying intervertebral disk pathologies. Several authors [1, 2] present segmentation approaches with experiments using this modality.

Given the difficulty of vertebrae segmentation problem, it is desirable to further constrain the solution space. Rather than represent each vertebra as an arbitrary contour, researchers have employed shape templates [3, 5], Fourier descriptors [8], as well as active shape models for individual vertebrae [1] or the entire spinal column [6] built from training data. In this paper, we approximate each vertebra as a rectangle, computed as a semi-affine transformation applied to the unit square. Indeed, for the estimation of intervertebral disk orientation, exact vertebrae segmentation is not necessary since we are interested in the direction of the vertebral edges that are aligned with the disk. Unlike standard active contour methods, the speed function of the contour is integrated along the perimeter of the rectangle, resulting in a rectangle evolution that is more robust to local variations in the speed function and initial placement. This enhances consistency in the results, an important feature for clinical use.

1.2 Our contribution

The method presented in this paper is motivated by the work of Yezzi et al. in [7], which performs simultaneous registration and segmentation of the same object in multiple images that may be acquired by different imaging modalities. However, in this work, we impose the shape constraint of a rectangle by mapping the unit square into the image using a semi-affine transformation. Rectangles are used to segment adjacent vertebrae on the same image rather than using arbitrary contours to segment the same object in different images. In addition, we present interaction forces designed to penalize larger variations in scale and rotation, under the assumption that adjacent vertebrae have a similar size and orientation. Finally, unlike standard level set implementations, our resulting mathematical model is based on ordinary differential equations (ODEs) instead of partial differential equations (PDEs). This allows us to take larger time steps in our numerical implementation.

2 Method

2.1 Active rectangle representation

Let $I : \Omega \subset \mathcal{R}^2 \rightarrow \mathcal{R}$ denote the image of the unit square, formed as a closed polyline with an outward-oriented normal $\mathbf{N}$, as depicted on the left of Figure 2, and let $\hat{I} : \hat{\Omega} \subset \mathcal{R}^2 \rightarrow \mathcal{R}$ be the target MR image. The unit square $\mathcal{C}$ is mapped from I to $\hat{I}$ as $\hat{\mathcal{C}}$ using a transformation $g : \mathcal{R}^2 \rightarrow \mathcal{R}^2$, i.e., $\hat{\mathcal{C}} = g(\mathcal{C})$. The mapping g consists

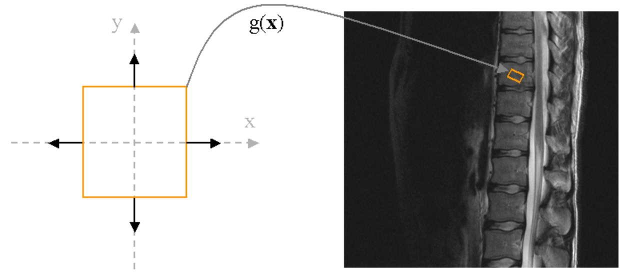

Figure 2. Our atlas shape in image I is the unit square (left), transformed as a rectangle into the image $\hat{I}$ (right) by a semi-affine transformation $g(\mathbf{x})$.

of registration parameters, $g_1 \cdots g_n$, which in this paper are a set of $n = 5$ parameters from a finite-dimensional group represented by a rotation angle θ, non-uniform scale parameters M_x, M_y, and displacement parameters D_x, and D_y. These are used in a semi-affine transformation given as

$$\hat{\mathbf{x}} = g(\mathbf{x}) = RM\mathbf{x} + \mathbf{D}, \quad (1)$$

with rotation matrix $R = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}$, scaling matrix $M = \begin{bmatrix} M_x & 0 \\ 0 & M_y \end{bmatrix}$, and translation vector $\mathbf{D} = [D_x, D_y]^T$, and $\mathbf{x}$ is a point on the unit square. Figure 2 depicts the transformation of the unit square into the MR image.

2.2 Energy function and curve evolution

Segmentation can be achieved by following a gradient descent procedure to minimize a region-based energy functional of the form:

$$E(g) = \int_{\hat{\mathcal{C}}_{in}} \hat{f}_{in}(\hat{\mathbf{x}})d\hat{\mathbf{x}} + \int_{\hat{\mathcal{C}}_{out}} \hat{f}_{out}(\hat{\mathbf{x}})d\hat{\mathbf{x}} \quad (2)$$

where $\hat{f}$ is a function that best represents a certain characteristic of the image such as the mean or variance. We chose the piecewise constant segmentation model of Chan and Vese [4], for which $\hat{f}_{in} = (\hat{I} - \hat{u})^2$ and $\hat{f}_{out} = (\hat{I} - \hat{v})^2$, where $\hat{u}$ and $\hat{v}$ are the mean values inside and outside the segmenting curve respectively. We re-express this functional on the domain Ω as

$$E(g) = \int_{\mathcal{C}_{in}} (|g'|\, \hat{f}_{in} \circ g)(\mathbf{x})d\mathbf{x} + \int_{\mathcal{C}_{out}} (|g'|\, \hat{f}_{out} \circ g)(\mathbf{x})d\mathbf{x} \quad (3)$$

where $|g'|$ is the determinant of the Jacobian of g and $\circ$ denotes functional composition.

Taking the derivative of Equation 3 with respect to the registration parameter g_i gives the following gradient descent minimization,

$$\frac{dg_i}{dt} = \frac{\partial E}{\partial g_i} = \int_{\mathcal{C}} \hat{f}(g(\mathbf{x})) \left\langle \frac{\partial g(\mathbf{x})}{\partial g_i}, mRM^{-1}\mathbf{N} \right\rangle ds, \quad (4)$$

where g_i indicates one element of g, $m = M_x M_y$, $\hat{f} = (\hat{f}_{in} - \hat{f}_{out})$, and $\langle\rangle$ indicates an inner product. Details of this flow can be found in [7]. Intuitively, equation (4) is an ODE whose solution requires us to traverse the contour of the unit square, shown in Figure 2, find its new transformed pose in the image, then update the pose function g until convergence. That is, the segmentation occurs by updating the registration parameters $g_i \cdots g_n$. Unlike [7], there is no contour update $\frac{\partial C}{\partial t}$ since our contour in domain Ω is fixed as the unit square.

To avoid misalignment due to salient features away from the disk, we apply a weighting (empirically set to 4.0) to the edges of the transformed square that are closest to the intervertebral disk. These edges have a similar orientation as the disk itself. For initialization, the algorithm sets the translation to the starting point $\hat{\mathbf{x}}$ in the MR image, the rotation angle to 0 and the scale parameters to 1. An example evolution for a single rectangle appears in Figure 3.

Figure 3. Evolution of a single rectangle. From left to right: 0, 25, and 100 iterations, using time step $\Delta t = 0.5$.

2.3 Interaction forces

While it is possible to independently evolve rectangles in each vertebra adjacent to an intervertebral disk, we can take advantage of the similarity of adjacent vertebrae to further constrain the problem. Under the assumption that adjacent vertebrae have a similar size and orientation, we propose an interaction energy between adjacent rectangles. This energy penalizes large orientation and scale differences, and takes the form $E(g) = f(\nabla g_i)$, where $f(z)$ is a differentiable function that penalizes the variation of the registration parameters of different active rectangles. Differentiation of $E(g)$ with respect to g_i yields the interaction force

$$\frac{dg_i}{dt} = \frac{\partial E}{\partial g_i} = \frac{\partial f}{\partial z}\frac{\partial z}{\partial g_i} \quad (5)$$

We have investigated several forms of the penalty function; however, due to space constraints we only present one function here, namely $f(z) = \frac{1}{2}z^2$, which provides sufficient regularization on the registration parameters. We evolve in the negative gradient direction, yielding the update

$$\frac{dg_i}{dt} = -\alpha \Delta g_i, \quad (6)$$

where Δ is the Laplacian operator and α is a constant used to weight the influence of the interaction force. In all our experiments, we set $\alpha = 0.25$, which has provided sufficient coupling for our data between adjacent active rectangles to jointly perform the segmentation. However, using a lower value of α would decrease the coupling, which could be desirable if the adjacent vertebrae had larger differences in size/orientation.

An example comparing independent vs. coupled segmentation is presented in Figure 4. For the left and middle of the figure, we performed independent evolutions of the two rectangles starting from different initial conditions (seed points), resulting in the active rectangles being attracted to undesirable local minima. On the right we show the coupled segmentation (both sets of initial conditions produced the same result), which achieves a more robust segmentation.

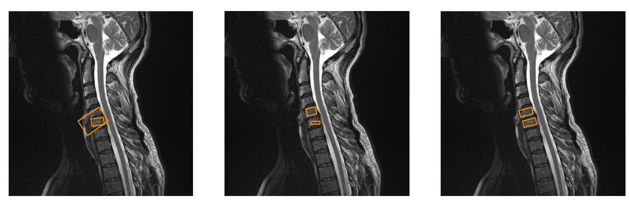

Figure 4. Effect of the interaction force. Left and middle: uncoupled segmentation. Right: coupled segmentation.

3 Results

In this section, we report disk orientation detection results from different parts of the spine. In each case, the user will click on the disk of interest. There is some automatic preprocessing done to get two seed points, one inside each of the upper and lower vertebrae. This is the initialization of the algorithm. Figure 5 shows the initialization and the final detection of a disk in the lumbar region of the spine. In the middle, copies of the unit square are placed at each seed point. Then the segmentation is performed to get the result on the right. Notice how the rectangles align to the edges that are adjacent to the disk. From these results, we compute the orientation of the disk, also shown in the figure. The orientation is found by determining the line equally bisecting the bounding box connecting the detected vertebrae (clinically, manual determination of the orientation is done in a similar fashion). The upper part of Figure 6 shows the result for a sagittal C-Spine image, and the lower part of the figure shows an example for a coronal image. Computing the disk orientation in both the sagittal and coronal views defines a plane that is used for setting up the transverse slice stack. All segmentations complete within a few seconds.

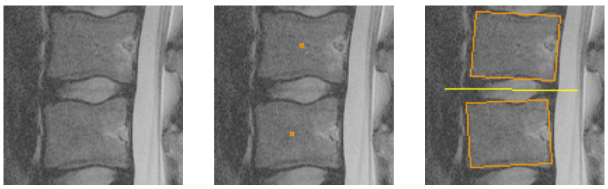

Figure 5. Segmentation approach. Original image (a), seeds overlaid (b), and final segmentation result (c) with disk orientation drawn as a line between vertebrae.

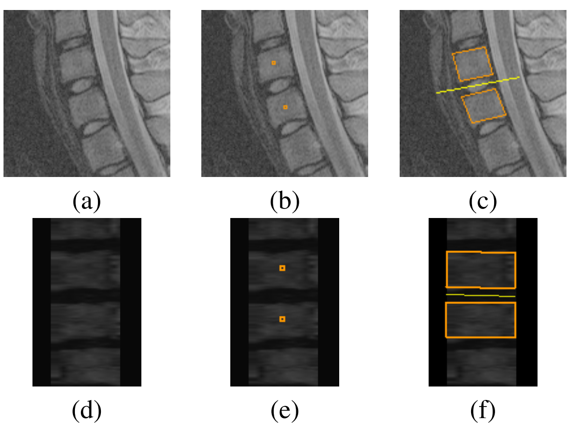

Figure 6. More examples. Saggital C-spine result (a) - (c), and coronal result (d) - (f).

For validation of the proposed method, we used it to determine the orientation of 51 intervertebral disks, coming from 9 different patients. Since ground truth is not available, we compared these orientation results to those estimated by hand, achieved by a user drawing a line over the disk indicating its orientation. The results of these experiments were that on average, the algorithm computed the disk orientation to less than 2.25 degrees of that detected by a human operator.

4 Conclusion and future work

In this paper we presented a simple and efficient method to detect the orientation of intervertebral disks. The method fits a rectangle to each adjacent vertebrae by minimizing an energy functional based on a shape constraint, image data, and coupling between adjacent rectangles. While more comprehensive validation of the algorithm is required, from our experimental results we conclude that the shape constraint combined with the coupled segmentation results in good vertebrae segmentation from which the intervertebral disk orientation can be computed.

Since our method uses gradient descent to minimize an energy functional, it achieves a local minimum of the energy, and can produce different results for different initializations, which is typical for this class of methods. When the vertebrae are imaged so that they have a consistent intensity and their borders have sufficient contrast, our segmentation method typically converges to a reasonable solution. However, for robustness it is certainly possible to include other image statistics (beyond the Chan-Vese model we employ) in our framework. This is left for future work.

The framework presented in this paper is quite general in that any shape representable by a closed polyline is supported. For future work, we are interested considering other segmentation problems with different problem-specific shape constraints, as well as extending the method to polyhedra in 3D space.

References

[1] M. Brejl and M. Sonka. Object Localization and Border Detection Criteria Design in Edge-Based Image Segmentation: Automated Learning from Examples. *IEEE Transactions on Medical Imaging*, 19(10):973–985, 2000.

[2] J. Carballido-Gamio, S. J. Belongie, and S. Majumdar. Normalized Cuts in 3-D for Spinal MRI Segmentation. *IEEE Transactions on Medical Imaging*, 23(1):36–44, 2004.

[3] C. Cardan and R. Allen. Measurement of Spine Motion for Diagnosis of Mechanical Problems. In *J. Simulation Modeling Med.*, volume 1, pages 15–19, 2000.

[4] T. Chan and L. Vese. Active Contours Without Edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.

[5] C. Simonis and R. Allen. Determination Of Instantaneous Centres Of Rotation: Parallel Application In Spine Kinematics. In *Proc. IEEE Conference on Engineering in Medicine and Biology Society*, volume 1, pages 39–40, 1992.

[6] P. P. Smyth, C. J. Taylor, and J. E. Adams. Automatic measurement of Vertebral Shape Using Active Shape Models. *Image and Vision Computing*, 15:575–581, 1997.

[7] A. Yezzi, L. Zöllei, and T. Kapur. A Variational Framework for Integrating Segmentation and Registration Through Active Contours. *Medical Image Analysis*, 7:171–185, 2003.

[8] Y. Zheng, M. S. Nixon, and R. Allen. Automated Segmentation of Lumbar Vertebrae in Digital Videofluoroscopic Images. *IEEE Trans. on Med. Imaging*, 23(1):45–52, 2004.