# Context Sensitive Labeling of Spinal Structure in MR Images

Chetan Bhole[a], Suryaprakash Kompalli[b] and Vipin Chaudhary[b]

[a]Computer Science Department, University of Rochester, Rochester, NY;
[b]Computer Science and Engineering Department, University at Buffalo, SUNY, Buffalo, NY

## ABSTRACT

We present a new method for automatic detection of the lumbar vertebrae and disk structure from MR images. In clinical settings, radiologists utilize several images of the lumbar structure for diagnosis of lumbar disorders. These images are co-registered by technicians and represent orthogonal features of the lumbar region. We combine information from T1W sagittal, T2W sagittal and T2W axial MR images to automatically label disks and vertebral columns. The method couples geometric and tissue property information available from the three types of images with image analysis approaches to achieve 98.8% accuracy for the disk labeling task on a test set of 67 images containing 335 disks.

**Keywords:** CAD development, Classifier design

## 1. INTRODUCTION

National Ambulatory Medical Care Survey, 2005 shows that around 11 million people suffered from low back pain in 2005 and Ref. 1 state that the health care expenditures of the people with back pain in US summed up to over 90 billion dollars in 1998. Besides sprains and muscle pains, some of the more serious abnormalities associated with lower back pain include disk herniation, anular tear, disk degeneration, spinal stenosis and spondylolysis. Radiologists usually use the full protocol[2] for investigation of these lumbar spinal disorders. The full protocol refers to the combined use of T1W sagittal and T2W sagittal and axial images. Naturally, the first step in diagnosis is the labeling of disks and vertebrae in these images. Our work outlines a methodology to automatically locate the disks and the vertebrae in clinical MR images of the lower back.

The labeling problem is defined as the function $\mathrm{L(I)} = < l_i; (x_i, y_i, z_i) >$, where I is the set of input MR images for each patient, and $l_i \in$ {L1-L2, L2-L3, L3-L4, L4-L5, L5-S1} are the disk labels and $(x_i, y_i, z_i)$ are the co-ordinates of each label in a 3D co-ordinate frame. In our case, I corresponds to the T1W sagittal, T2W sagittal, and axial images. The labeling of the vertebrae is done using the disk labeling and are labeled as L1 to L5 for the lumbar vertebrae and S1 as the sacral vertebra. The functional mapping is made non-trivial by the challenges posed by pathology and image properties of tissues. Several parts of structures in the vertebral column have similar imaging properties; the disks, spinal cord, and vertebrae are not easily distinguished by gray scale values alone (Figure 1). Equipment variations and contrast agents also add to the challenge.

There are a few aricles for labeling and/or segmenting disks or vertebra. Many of them aim at classification of the entire vertebral column while we concentrate only on the lumbar vertebrae and disks. Schmidt et. al[3] use a combination of classification trees and a graphical model and claim it is robust to scale, orientation changes as well as missing parts. Corso et. al[4] use a two level probabilistic model for labeling the lumbar structure. The lower level models the pixel level information while the higher level models object relationships. We use a superset of the images these authors used and additionally include the images from patients with abnormalities. Weiss et. al[5] use a semiautomated iterative technique to label the entire spine. Yao et. al[6] partition the vertebra on the CT images using the watershed algorithm on the axial slices and use intensity profiles similar to what we use. However, our profiles are generated using normals and we use the valleys for generating adaptive thresholds. Chwialkowski et. al[7] also use intensity profiles to estimate the location of the spine and to detect the disks. Peng et. al[8] use a template of the disk for detection of the vertebral region and vertebra boundaries are detected by edge detection. Statistical shape models like active shape and appearance models are used in some papers[9–12] on X-ray or CT images. A few papers[13–15] deal with the segmentation of the spine and vertebra on MRI or X-ray

Send correspondence to vipin@buffalo.edu

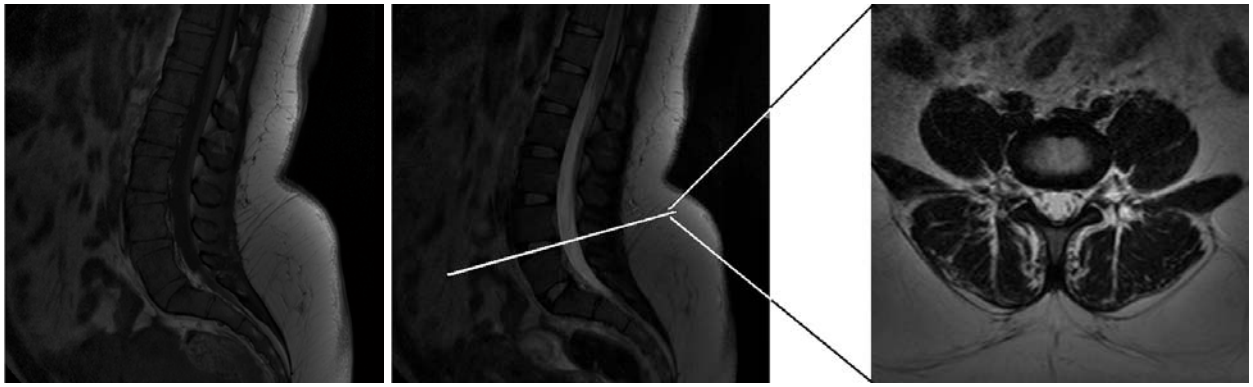

Figure 1: The figure shows sample 512 by 512 pixel T1 weighted and T2 weighted sagittal image slices and a T2 weighted axial slice from left to right. These images are co-registered and the axial slice can be seen in the T2 weighted sagittal image as a bold white line passing through a vertebral disk.

images. Gamio et. al[13] use images of high resolution with inter-slice spacing of 1 mm. Zheng et. al[16] use an optimized hough transform and genetic algorithm while Wong et al.[17] use a combination of edge enhancement, Markov Random Fields (MRFs), support vector machines (SVMs) and snakes for segmentation of the vertebrae on fluoroscopic images. Abufadel et. al[18] discuss a method for detecting the orientation of the disks using active rectangles to fit on the vertebrae. Interactive tools for spine and vertebra segmentation are shown in Ref. 19.

The shape models like active shape models (ASM) and active appearance models (AAM) are susceptible to initialization points. Though they are useful on CT images, they do not perform as well on MRI images. Most of the prior work is on CT or enhanced CT images. We notice that the resolution of clinical images vary from 1.5T to 3T. The spacing between each slice is typically between 4-5.5mm. Many of the above papers deal with very small sample of patient images and use normal patients. Abnormalities cause variations in the vertebral column, making it difficult to generalize the vertebral geometry to standard models. It is more important in a clinical setting to be able to have robust algorithms for labeling and segmenting disks and vertebrae of all types of patients.

Our method accounts for the variations by relying on the three different types of images (Figure 1). T2W sagittal is more sensitive to water and T1W is more sensitive to fat in the body; the corresponding images highlight structures with orthogonal tissue properties. Our method takes advantage of this fundamental difference in the two images. We also note that clinical images have significant variance in quality and resolution, but multiple views are mandatory for diagnosis, and can be used in the labeling procedure. Clinical setting requires different scans to be done in a quick succession; technicians make sure the subjects do not move during the entire scan sequence and hence these images can be considered co-registered. If the subject does move, the scan sequence is repeated again. We extract the spinal cord from these images using preprocessing steps and use it to estimate the boundaries of the vertebrae and disks on the sagittal views. We then use a combination of intensity profiles and axial slice orientations to obtain the centers of the disks.

We discuss our approach in detail in Sec. 2. We then describe the data we used along with the results of our algorithm in Sec. 3. Finally, in Sec. 4 we summarize and provide directions for future work.

## 2. APPROACH

We normalize all our MR images and calculate the difference between all the corresponding T2W and T1W sagittal slices. This difference usually highlights the regions that are high in water content and low in fat content. With this the spinal cord, Nucleus Pulposus, some skin and other tissues get extracted. We can extract the spinal cord by consideration of structure closer to the central vertical axis and use of a thresholding technique. Empirically, a value of 0.15 does well as the threshold value. In future steps, we do take into account the possibility that the spinal cord may be broken either as a natural abnormality of the patient or due to an incorrect threshold value. Hence the algorithm is robust with respect to selection of this threshold value.

We then need to select the most likely sagittal slice on which we will label the center of the disk. We use three candidates for choosing this sagittal slice. Assuming the technician has aligned the scans properly but may not have the true center of the cord as the middle sagittal slice, the best candidate would be amongst the middle sagittal slice or the previous or the next slices to the middle slice. For each of these candidates we examine the spinal cord on that slice. If the number of connected components (pieces) of the spinal cord is one as seen in Fig. 2(a), we declare it as the spinal cord of that slice. However, if there is more than one piece of the spinal cord, for example as in Fig. 2(b) which had severe spinal injury, we perform a join operation. We use pieces larger than 70 pixels in area and closer to the central vertical axis (z-axis, that runs along the height of the patient) as the pieces that would make up the spinal cord and discard all other components. The size constraint imposed above allows us to eliminate the components that might be the Nucleus Pulposus (center of a healthy disk) or other noise pixels. The join operation involves connection of the cord pieces. To do this, we determine the boundary of each of the cord pieces and sort the pieces in order of z-axis values. No z-axis overlap value indicates possibility of the spinal cord of the patient being severely damaged. The shortest distances are calculated from the boundary points of one cord piece to the boundary points of the next cord piece in z-axis order. By this procedure we connect the spinal cord pieces into one. The joint spinal cord is shown in Fig. 2(c). We now select the sagittal slice of our choice amongst our three candidates as that slice that has maximum spinal cord area on each slice. All further processing is done on the selected T2W sagittal slice and we use the joined spinal cord.

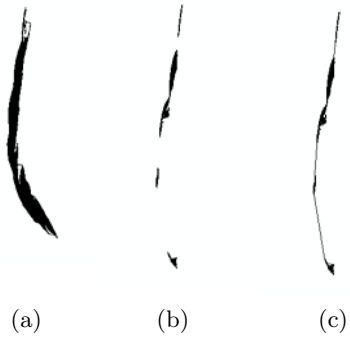

(a) (b) (c)

Figure 2: (a) Normal spinal cord, (b) Severe spinal injury leading to bad segmentation and (c) Interpolation applied to (b).

We continue further processing on the spinal cord extracted from our chosen sagittal slice. If the spinal cord is too short (less than 300 pixels in height does good empirically), then it will usually lead to errors in detecting upper or lower lumbar disks as well as the S1 vertebra. We would like to have the spinal cord extend between 80 and 400 pixels of the height of the image. Hence, we draw tangents on the boundaries at the upper and lower ends (whichever are needed) of the extracted spinal cord and so extrapolate the spinal cord. The tangents give us a coarse direction the spinal cord follows at the ends. Different people have extremely different curvatures of the spinal cord and hence we use the extrapolation approach that we just described.

An intial region of interest (Called ApROI) is selected by considering 140 pixels to the left and 40 pixels to the right of each row of the processed spinal cord as seen in Fig. 3(a). Canny edge detector is used to extract edges from the ApROI. Edges that have less than 15 degree inclination with respect to the spinal cord edge are retained. Since the S1 vertebra is typically more inclined with respect to the spinal cord, edges below (in the lower 40 pixels of the spinal cord) are retained even if they have an angle of 30 degrees. We select longer edges nearer the spinal cord. From observations, there are fewer edges parallel to the spinal cord inside the vertebrae. We connect all these edges to get a first estimate of the left boundary of the vertebral column. The first estimate is rectified using the average distance of the first estimate to the spinal cord. The resultant left boundary is seen in the Fig 3(a). So now we have a more accurate region of interest (AcROI) that consists of the region between the left vertebral column boundary and the spinal cord and this region contains the vertebra and the disks. We then draw normals to the spinal cord (some samples shown in Fig. 3(a)), calculate the average intensity of the pixels on the normals and plot a graph of these values as seen in Fig 3(b). We use the average and standard deviation of the three deepest valleys as the threshold values for extraction of the Anulus Fibrosus (AF) which is the outer darker region of a disk. The extracted AF is shown in the Fig. 3(c). This allows us to take into account the intensity values specific to images.

We make use of the alignment of the slices done by the technician during scanning. The technician makes sure the axial slices pass perpendicular to the vertebral column curvature (very much like our normals) a sample of which is seen in the Fig. 1. The technician arranges them as sets of 4 to 5 slices that pass through a disk and the vertebral region around it. The middle axial slice in each set can thus give an approximate location of the center of the disk along the height of the vertebral column. We map this axial slice on the sagittal view which is a line as seen in Fig. 1. Note that we make use of the fact that the axial and sagittal data is registered.

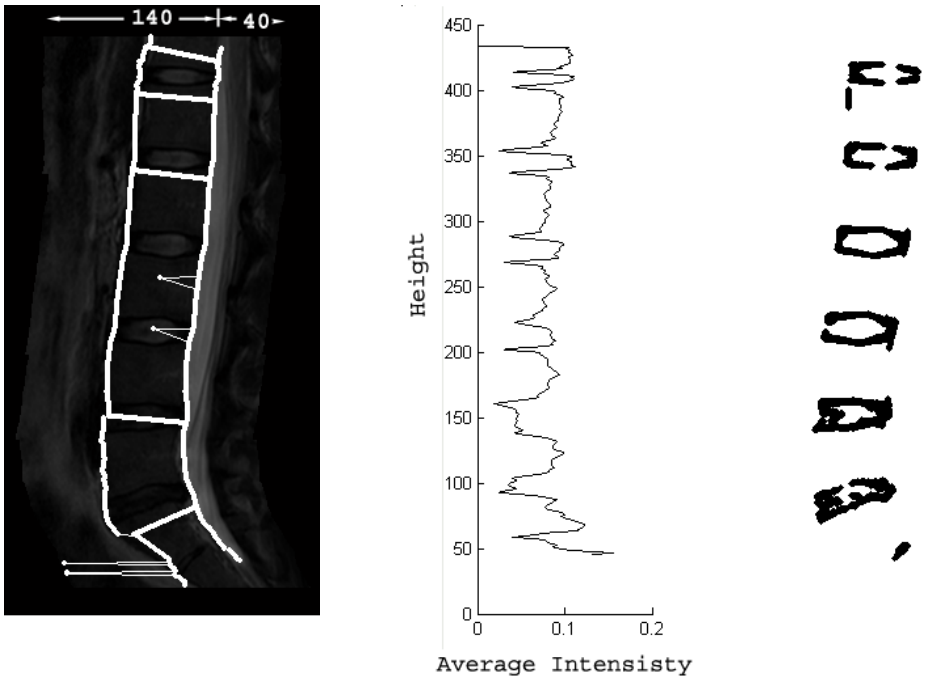


(a) (b) (c)

Figure 3: Fig. (a) shows results of some of the intermediate steps of the algorithm. Firstly, we see the ApROI (explained in text) with the selection of the region 140 and 40 pixels to the left and right of the extracted spinal cord. The bold white lines running vertically show the boundaries of AcROI explained in the text. The bold white horizontal lines are some samples of the normals. Also seen are triangular structures. We expect a higher ratio of lower intensity pixels in the triangles in disk regions compared to vertebra regions. Fig. (b) shows the intensity profile of the average over the pixels in the normals to obtain the threshold value. Fig. (c) shows the extracted Anulus Fibrosus of the disks using thresholding.

We now improve our estimate of the centers of the disks by using the centers from the axial slices and the segmented AF we obtained. We count the number of consecutive pixels belonging to the AF along the center point of each normal. We combine two non-contiguous block of these pixels to a single block if the total pixels (from the start of one block to the end of the other) is less than 5 pixels. We can think of it as a string like BXBXBXB...XB where X stands for pixel counts not belonging to the AF and B stands for pixel counts belonging to AF and we combine BXB to B if the total number of pixels is less than 5 pixels. This means that it is small enough to be a disk and that the X in between cannot be a vertebra. We then perform two more steps to remove spurious disk centers. We consider the black pixels on half of the normals (from the center of the normal to the spinal cord) and remove the spurious centers that don't have black pixels. This is done because the left vertical boundary of AcROI obtained may not be perfect. And as the next step, we consider pixels in a triangle formed by the center of disk and points on the spinal cord boundary. This is shown in Fig. 3 (a). We use the end points of two normals, one 5 normals above the disk center and the other 5 normals below the current normal. If the fraction of pixels with black intensity to other pixels' intensity in the triangle is less than 0.15 remove the corresponding center of the disk. This parameter value is justified empirically.

We can easily estimate the upper and lower disk bounds using the extracted AF we obtained for each disk and centers of disk by traversing along the centers of the normals. We perform special processing on the last disk. If the last candidate disk is less than 50% the length of the average of the other disks, it is marked as the reduced disk formed between the fused sacral vertebrae and remove it from our labeling procedure. We then label the disks and vertebrae. We start labeling bottom up and so label the disks in order as L4-5, L3-4, L2-3

and L1-2. Similarly label the structures in between as the corresponding vertebrae. These are the midpoints of the centers of the consecutive disk centers. We additionally label L1 and S1 by extrapolating the centers and use of normals. Once we have the centers on the T2W sagittal we can easily map the points on the T1W sagittal and T2W axial images.

## 3. EVALUATION

### 3.1 Data description

The clinical MRI data we used is taken from a 3T Philip's Medical Systems MRI scanner which consists of T2 and T1 weighted sagittal views and T2 weighted axial views. The sagittal and axial slices have a thickness of 4.5mm and 4mm respectively and 5mm and 4.4mm of spacing between the slices respectively. There are usually 12 slices of the sagittal views and around 24 to 30 axial view slices. Each image slice is 512 by 512 pixels. The image acquisition takes place as follows. The operator or technician takes a few survey or scout image scans to get an idea of the lumbar region they want to take detailed scans of. They take a quick sequence of T2 and T1 weighted images mentioned above. They make sure that the axial slices pass roughly perpendicular to the vertebral column and close to or through the vertebral disks. They also make sure that the patient does not move during these quick scans otherwise the procedure is repeated again. Because of this, the images we use are co-registered images. Fig. 1 shows T1 and T2 weighted sagittal views and a T2 weighted axial view.

The scans include those of 16 normal patients and of 51 patients that have some kind of abnormalities like disk herniation, anular tear, disk dessication, spinal stenosis and spondylolysis on at least one of the disks. Disk herniation is the bulging or protrusion of the disk into the spinal cord. There are different levels of severities of the protrusion. Disk herniation can be seen easily in the lower most disk (L5-S1) of Fig. 5(d) and Fig. 6(c) and (d). The anular tear refers to a condition of rupture of tissue in the outer regions of the disk. Disk dessication, as seen in Fig. 5(b) and (c) and Fig. 6(c), occurs because of loss of water content of the disk and is common due to aging. Also, Fig. 5(b) shows extreme flattening of the disk L5-S1. Spinal stenosis refers to the compressing of the spinal cord and nerves caused by different reasons. Spondylolysis refers to a defect in the facet joints of the vertebrae. When this is accompanied with misalignment of a vertebra with the rest of the vertebrae, it is termed as spondylolisthesis. There is a clear misalignment between L5 and S1 vertebrae in image of Fig. 5(a). The age of the patients varied from 20 to 80 years with a mean of 44.8 years. Similarly, the weight of the patients varied from 110.2 to 297.5 pounds with a mean of 187.4 pounds. 36 of the patients were female while the rest 31 were male. Experts used Gadolinium as a contrast agent for one of the cases. One case also shows a small movement of the lower lumbar region between the T1 and T2 weighted sagittal views. All the factors above (age, sex, weight and abnormalities) give a sense of variability of the data used including variation in the shapes and sizes of the lumbar vertebrae and disks as well as variations in intensities of the MRI scans.

### 3.2 Results

We count the labeling of a disk as correct if its center lies in the disk. The reason for doing this can be explained with the help of the Fig. 4 (a). We can sample a rectangular box as a region of interest using this center as shown in the figure. This region of interest can be used for further processing to detect abnormalities like herniation and spinal stenosis. Hence it is sufficient to have a reliable point in the disk for further processing. With this measure, we get only 4 centers not located in the disks from a total of 335 disks and thus an accuracy of 98.8%. We also provide results using two other metrics which can be seen in the plot of Fig. 8. The first error metric (d1 measure) is the euclidean distance between the centers obtained by our algorithm and that obtained by calculating the average of ten points marked manually for each disk as shown in Fig. 4(b). The points 1, 3, 6 and 8 are placed at the corners of the disk while 4, 5, 9 and 10 are placed approximately one third distance away from their nearest corners. The second error metric (d2 measure) is the euclidean distance between the center obtained by our algorithm and the closest point in the box formed by points 4,5,9 and 10. The box is shown in Fig. 4(b). We also had 12 points marked manually on the boundaries of the vertebra as seen in Fig. 4(c). The plot of the average error distances for the five disks can be seen in Fig. 8. Some of the labeling results can be seen in Fig. 5 and 6. The white stars and white dots in these figures indicate the centers of the disk and vertebra respectively obtained by our algorithm. Again the white boxes are those as explained earlier using Fig. 4(b). Figure 6(c) and (d) show two cases where the labeling of L5-S1 disk and S1 vertebra is incorrect. We

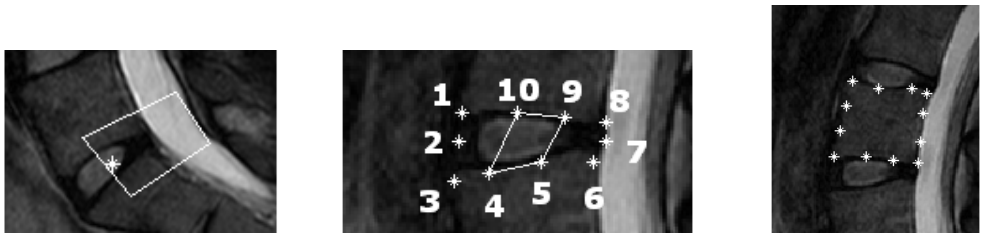


(a) (b) (c)

Figure 4: The fig. (a) shows use of the region of interest for diagnosing abnormalities like disk herniations and spinal problems. Part (b) shows the ten points marked manually to identify the disk boundaries. We define the ideal center of the disk as the center of these points. The white box connecting points 4, 5, 9 and 10 is used in the d2 metric explained in the text. Similarly, part (c) shows the twelve points marked on the boundary of the vertebra.

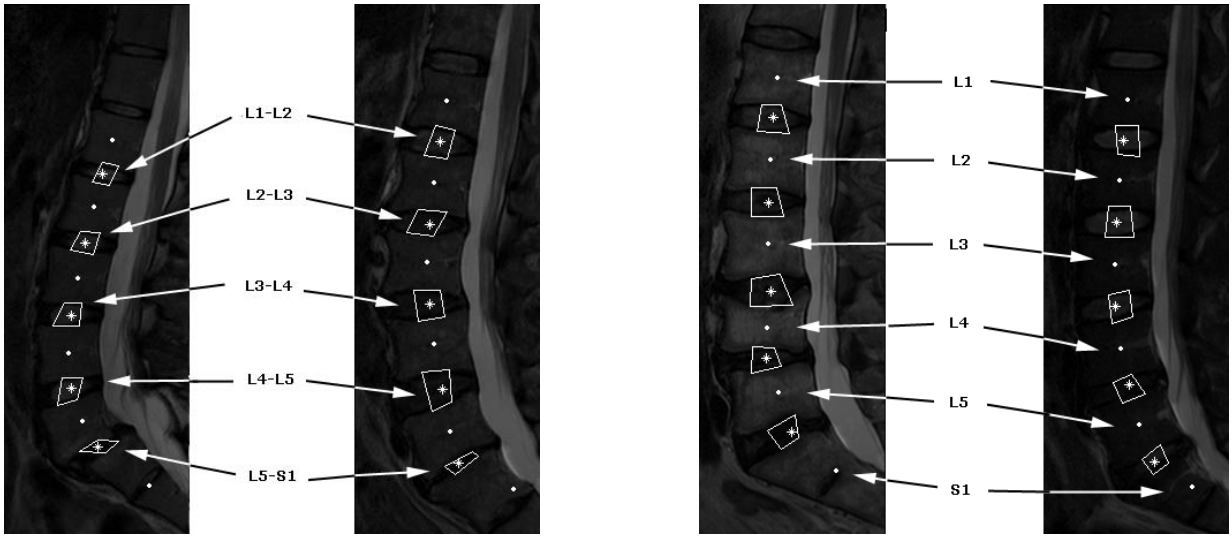


(a) (b) (c) (d)

Figure 5: The figure shows clipped images with labeled disks and vertebrae produced by the algorithm as well as boxes from manually annotated points. White dots are the centers of the vertebrae while white stars indicate centers of the disks obtained from our algorithm. White bordered boxes are generated from points annotated at the edges of the disks and approximately one third away from the ends of the disk both ways. The Figure also shows variability in our data. (a) shows spondylolisthesis, (b) shows disk degeneration and (c) and (d) show disk herniation cases.

used Matlab for creating a prototype of the system. It took less than two minutes on an average for a single patient case study on an Intel core2 duo 2Ghz machine with 2GB RAM for the labeling procedure including the time to load images.

### 3.3 Discussion

We tried initial experiments using an ASM and a variant of an AAM for recognition of the vertebrae and disks. We used the online code provided by Tim Cootes with modifications for our data. Both the ASM and the AAM like model were not able to settle to the correct boundaries of the vertebra and disks. It was even worse for the lower lumbar region. We also tried using the supervised technique using a Markov Random field for segmentation of the images. We used the online code provided by Zoltan Kato with some changes

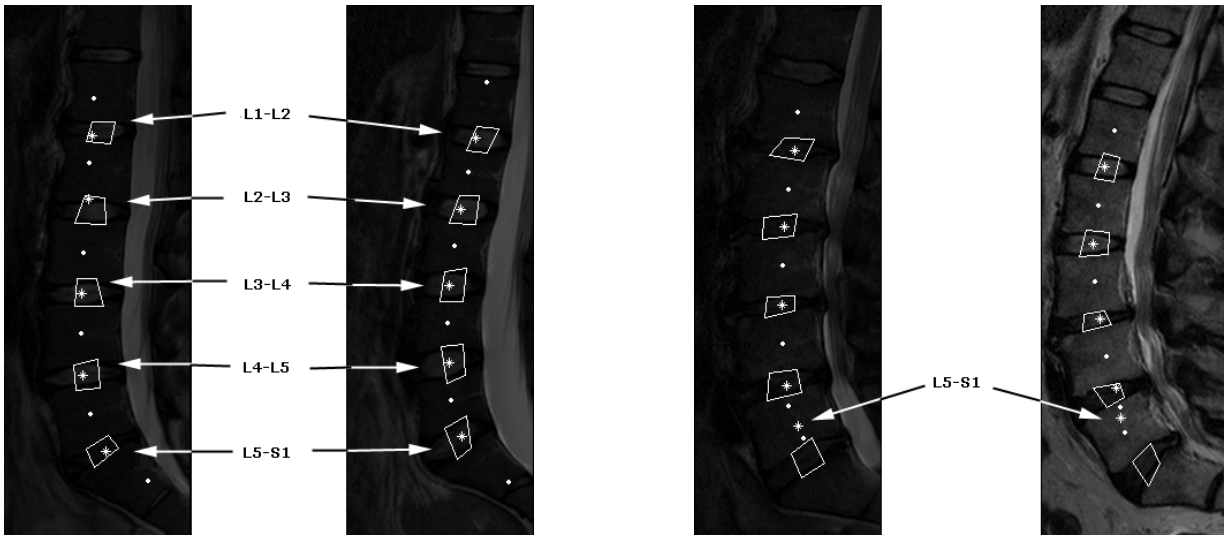


Figure 6: The figure shows few more clipped images with correct labelings (a) and (b) and two cases with erroneous labeling of the L5-S1 disks (c) and (d). As before, white dots are the centers of the vertebrae and white stars are centers of the disks obtained from our algorithm.

to the code. Seed areas were selected manually to perform segmentation with different settings of parameters. However there were large variations for a single class even in a single image which led to bad results. One of the worst variations was where the intensity values differed for vertebrae on the top of the image than the intensity values below. Because of this, it would become necessary to manually provide regions for more vertebrae and disks. Also we would still need a procedure to process and label the segmented parts.

The deviation of the distances from the ideal centers is maximum for the last disk (L5-S1) as seen in Fig. 8. This occurs mainly due to spinal injury being more in that area and the lower lumbar region especially S1 being curved. Figure. 6(c) and (d) show erroneous disk labels. In both these cases, the spinal cord segmented by the algorithm was made of multiple small pieces. This is due to severe spinal injury. In the case of (c), though there was a larger chunk of the spinal tissue at the bottom near the L5-S1 disk, accumulation of errors resulted in detection of a fraction of the disk and led to error in estimating the center. The problem in (d) was much more severe. There were very few and small in size spinal cord pieces near the L3-L5, L4-L5 and L5-S1. The extrapolation of the spinal cord helped to attenuate the errors, but was not able to get L5-S1 correctly. Figure. 7 shows an example of how estimation of the center of the disk can go wrong. The disk is thin and there is a remarkable difference in the z-axis height and the actual thickness (height) of the disk. Inspite of use of normals that run along the width of the disks, the normals for such a small thickness tend to run through the vertebra as well. The average intensity values of points along these normals causes the center to be placed out of the disk though the disk has been correctly located.

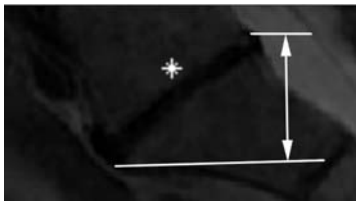

Figure 7: The figure shows the difference in heights of the disk along the z-axis height and actual height of the thin disk leading to errors.

Clearly, our approach relies heavily on the correctness of the MRI technician. Detection of patient movement is critical which would otherwise mean that different views will not be co-registered and our algorithm will produce incorrect results. As discussed previously, there was a case of noticeable movement in the lower lumbar region, but the movement was small and did not result in an error. Some abnormalities include presence of four or six lumbar vertebrae instead of five. Our algorithm does not model this defect. Another important thing to

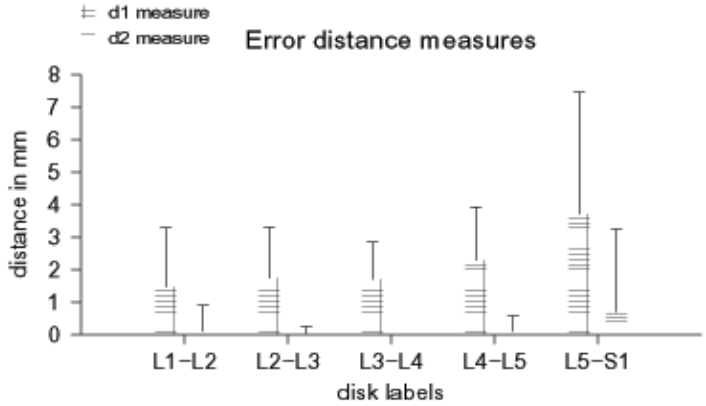


Figure 8: The figure shows the mean error distances for the two metrics defined. d1 is the the euclidean distance between the centers obtained by our algorithm and the average of the manually labeled points. d2 is the euclidean distance between the center obtained by our algorithm and that obtained by the box shown in Fig. 4(b). The upward error bar is one standard deviation in value.

note is that the parameters (usually in terms of pixels or pixel counts) are functions of resolution. And as the description of data suggests, we have not tested our algorithm with the current parameter settings for individuals with extreme physical shapes e.g. very short or very tall people or small children.

## 4. SUMMARY AND FUTURE WORK

We detect the lower lumbar vertebra and disks using clinical data that consists of T2W sagittal, T1W sagittal and T2W axial images. Most techniques reported in literature use only one of the MR images available in clinical data. By combining multiple images, our model attempts to accommodate variations seen in both normal and abnormal pathology. The method can also accommodate variations in slice spacing and thickness. We achieve 98.8% accuracy for the disk labeling task on a test set of 67 images containing 335 disks.

A few things we plan to work on in the near future are making the parameters independent of resolution, use a higher order polynomial for the extrapolation step of the spinal cord, testing the code on more data and speeding up the prototype code. We use a number of parameters, some are adaptive and others are set based on empirical evidence. We would like to look into ways of learning the parameters of the later type as well as integrate more powerful graphical models as used by Refs. 4 and 3 to make the algorithm more robust. We also will be working in directions for diagnosis of different abnormalities based on our current results.

## ACKNOWLEDGMENTS

This research was supported in part by a grant from NYSTAR. The authors would like to thank Dr. Dhillon for kindly providing the MRI image data. They also would like to thank Tim Cootes and Zoltan Kato for their ASM online code and MRF online code, respectively.

## REFERENCES

[1] Luo, X., Pietrobon, R., Sun, S., Liu, G., and Hey, L., "Estimates and patterns of direct health care expenditures among individuals with back pain in the united states," *Spine* **29(1)**, 79–86 (2004).

[2] Mullan, C., Kappes, J. H., and Kelly, B., "Magnetic resonance (mr) imaging of lumbar spine: use of a shortened protocol for initial investigation of degenerative disease," *The Ulster Medical Journal* **74(1)**, 29–32 (2005).

[3] Schmidt, S., Kappes, J., Bergtholdt, M., Pekar, V., Dries, S., Bystrov, D., and Schnörr, C., "Spine detection and labeling using a parts-based graphical model," in [*IPMI*], 122–133 (2007).

[4] Corso, J. J., Alomari, R. S., and Chaudhary, V., "Lumbar disc localization and labeling with a probabilistic model on both pixel and object features," in [*Proceedings of Medical Image Computing and Computer Aided Intervention (MICCAI)*], **LNCS 5241 Part 1**, 202–210 (2008).

[5] Weiss, K., Storrs, J., and Banto, R., "Automated spine survey iterative scan technique," *Radiology* **239(1)**, 255262 (2006).

[6] Yao, J., O'Connor, S., and Summers, R., "Automated spinal column extraction and partitioning," in [*3rd IEEE International Symposium on Biomedical Imaging: Macro to Nano, 2006*], 390–393 (2006).

[7] Chwialkowski, M., Shile, P., Peshock, R., Pfeifer, D., and Parkey, R., "Automated detection and evaluation of lumbar discs in mr images," in [*Proceedings of the Annual International Conference of the IEEE Engineering in Engineering in Medicine and Biology Society, 1989. Images of the Twenty-First Century.*], **2**, 571–572 (1989).

[8] Peng, Z., Zhong, J., Wee, W., and Lee, J., "Automated vertebra detection and segmentation from the whole spine mr images," in [*27th Annual International Conference of the Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005.*], 2527–2530 (2005).

[9] Smyth, P., Taylor, C., and Adams, J., "Vertebral shape: Automatic measurement with active shape models," *Radiology* (1999).

[10] Howe, B., Gururajan, A., Sari-Sarraf, H., and Long, L., "Hierarchical segmentation of cervical and lumbar vertebrae using a customized generalized hough transform and extensions to active appearance models," in [*6th IEEE Southwest Symposium on Image Analysis and Interpretation, 2004*], 182–186 (2004).

[11] Koompairojn, S., Hua, K., and Bhadrakom, C., "Automatic classification system for lumbar spine x-ray images," in [*19th IEEE International Symposium on Computer-Based Medical Systems, 2006. CBMS 2006*], 213–218 (2006).

[12] Vrtovec, T., Tomaevic, D., Likar, B., and Pernu, F., "Statistical shape deformable model of a lumbar vertebra," in [*Proceedings of the 8th Computer Vision Winter Workshop CVWW'03*], (2003).

[13] Carballido-Gamio, J., Belongie, S., and Majumdar, S., "Normalized cuts in 3-d for spinal mri segmentation," *IEEE Transactions on Medical Imaging* **23**, 36–44 (2004).

[14] Hahn, M. and Beth, T., "Balloon based vertebra separation in ct images," in [*Proceedings. 17th IEEE Symposium on Computer-Based Medical Systems, 2004. CBMS 2004*], 310–315 (2004).

[15] Ghebreab, S. and Smeulders, A., "Combining strings and necklaces for interactive three-dimensional segmentation of spinal images using an integral deformable spine model," *IEEE Transactions on Biomedical Engineering* **51**, 1821–1829 (2004).

[16] Zheng, Y., Nixon, M., and Allen, R., "Automatic lumbar vertebrae segmentation in fluoroscopic images via optimised concurrent hough transform," in [*Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2001*], **3**, 2653–2656 (2001).

[17] Wong, S. and Wong, K., "Segmenting lumbar vertebrae in digital video fluoroscopic images through edge enhancement," in [*Control, Automation, Robotics and Vision Conference, 2004. ICARCV 2004 8th*], **1**, 665–670 (2004).

[18] Abufadel, A., Slabaugh, G., Unal, G., L., Z., and Odry, B., "Interacting active rectangles for estimation of intervertebral disk orientation," in [*18th International Conference on Pattern Recognition, 2006. ICPR 2006*], **1**, 1013–1016 (2006).

[19] Kaminsky, J., Klinge, P., Rodt, T., Bokemeyer, M., Luedemann, W., and Samii, M., "Specially adapted interactive tools for an improved 3d-segmentation of the spine," *Computerized Medical Imaging and Graphics* **28**, 119–127 (2004).