

Segmentation of Vertebral Bodies in MR Images

Dženan Zukić^{1†}, Aleš Vlasák², Thomas Dukat³, Jan Egger³, Daniel Hořínek⁴, Christopher Nimsky³, and Andreas Kolb¹

¹Computer Graphics and Multimedia Systems Group, University of Siegen, Germany

²Faculty Hospital Motol, Prague, Czech Republic

³University Hospital Gießen and Marburg, Marburg, Germany

⁴Philipps University, Marburg, Germany

Abstract

Segmentation of vertebral bodies is useful for diagnosis of certain spine pathologies, such as scoliosis, spondylolisthesis and vertebral fractures.

In this paper, we present a fast and semi-automatic approach for spine segmentation in routine clinical MR images. Segmenting a single vertebra is based on multiple-feature boundary classification and mesh inflation, and starts with a simple point-in-vertebra initialization. The inflation retains a star-shape geometry to prevent self-intersections and uses a constrained subdivision hierarchy to control smoothness. Analyzing the shape of the first vertebra, the main spine direction is deduced and the locations of neighboring vertebral bodies are estimated for further segmentation.

The method was tested on 11 routine lumbar datasets with 92 reference vertebrae resulting in a detection rate of 93%. The average Dice Similarity Coefficient (DSC) against manual reference segmentations was 78%, which is on par with state of the art. The main advantages of our method are high speed and a low amount of user interaction.

Categories and Subject Descriptors (according to ACM CCS): I.4.6 [Image processing and computer vision]: Segmentation—Pixel classification

1. Introduction and Prior Work

Lower back pain for adults is rather common and its prevalence is rising [FHA*09]. The cancer risk from radiation exposure in computed tomography (CT) imaging [RGMb10] makes magnetic resonance imaging (MRI) preferable in the clinical routine. However, routine MRI poses several challenges for segmentation, such as a low and strongly anisotropic resolution. Unlike CT's Hounsfield unit (HU), MRI does not have standardized units of measurement. MRI also has a non-homogeneous intensity across the image, e.g. the central region has higher intensity and better contrast. All these facts are detrimental to the automation of segmentation procedures for MRI.

Pathologies such as scoliosis (curvature in anatomical left-right direction), vertebral fracture (crushed vertebra) and

spondylolisthesis (misaligned vertebra) can be diagnosed from vertebral shapes, positions and orientations, so their segmentations are needed. Spinous processes are poorly seen on MR images of low inter-slice resolution, which are the routine in clinical practice. Also, transverse processes are usually not seen at all because they are outside of the imaged area. Therefore, we focus on the vertebral bodies instead of the whole vertebrae.

Much research has already been done on spine segmentation. There is a large number of 2D methods due to its relative simplicity and low computational requirements [MCP*09, SSQW07, HCLN09, PZWhL05, CGBM04, EKD*12]. As 2D approaches process individual slices they can miss important information, such as a curvature or a positional shift in the anatomical left-right direction, so 3D approaches are preferable. But most 3D approaches focus on CT datasets only [GS04, KWL*08, WKL*01]. Some are concerned with detection only, not segmentation [KZS*11].

† Contact author, zukic@fb12.uni-siegen.de

We are aware of only three prior fully 3D segmentation methods applied to MRI: Hoar and Martel [HM02], Davatzikos et al. [DLSH02], and Štern et al. [ŠLPV11].

Hoar and Martel [HM02] have developed a segmentation algorithm that combines thresholded region growing with morphological filtering and masking using set shapes. Their algorithm is initialized by one or two ellipses, and can be split into three steps: segmentation of the vertebral bodies, segmentation of the posterior structures, and manual corrections. They used datasets with isotropic voxel size ($1 \times 1 \times 1 \text{ mm}^3$), and they combined two images of a double echo FISP acquisition sequence. Overall, it is a method suited to assist spine surgery, using images quite different from routinely acquired ones. They tested their method on 30 vertebrae. Surface registration error for vertebral bodies was $1.25 \text{ mm} \pm 0.28 \text{ mm}$. They also calculated average percentage of “good” points to be 79.4%, with 3 rejected registrations out of 30 experiments.

Davatzikos et al. [DLSH02] were mainly interested in the registration of different spine images to a manually segmented template image. An isosurface was extracted from this template image resulting in 837 vertices (vertebral bodies L1 to S1 and a corresponding portion of the spinal canal). This deformable model was trained on 13 additional images. In order to determine transformation, which registers a test image with the template image, the surface model is initially placed in the test image (overlapping the true position of the spine segment in the test image) and hierarchically deformed to conform to the edges of the test image. The evaluation was done using the leave-one-out method on routine images ($0.93 \times 0.93 \times 3 \text{ mm}^3$) of healthy volunteers. The average overlap was $81.5\% \pm 3.6\%$.

Štern et al. [ŠLPV11] perform the segmentation by optimizing 29 parameters of a 3D deterministic model of the vertebral body. They maximized dissimilarity between inside and outside intensities, and steered their method by image gradients. The method is initialized with one point per vertebra and accompanying size, depending on anatomical position, i.e. upper thoracic, lower thoracic or lumbar. They evaluated their method on 75 vertebral bodies from nine T2-weighted images. Three of their images were of a routine type ($0.4 \times 0.4 \times 3 \text{ mm}^3$), the others were isotropic ($1 \times 1 \times 1 \text{ mm}^3$). The mean radial Euclidean distance between the segmentation surface and ground truth points was $1.85 \text{ mm} \pm 0.47 \text{ mm}$.

All three previous methods suffer from long execution times: [HM02]–5–10 minutes per dataset, [ŠLPV11]–1–15 minutes per vertebra, and [DLSH02] do not state the execution time, but from the authors’ other papers [SD00] it can be deduced to be in minutes, not seconds. Both Štern et al. and Hoar and Martel segment vertebrae independently, i.e. each vertebra is initialized separately. Hoar and Martel’s initialization is especially detailed. These are significant drawbacks for routine diagnostic usage. Davatzikos et

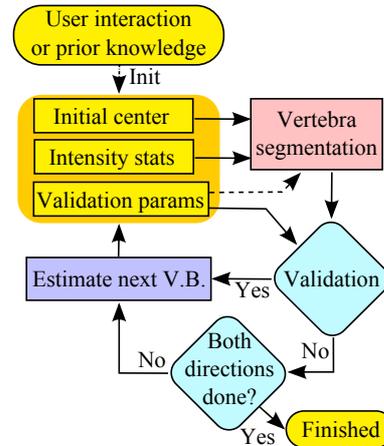


Figure 1: High level diagram of the segmentation system.

al. [DLSH02] use only one initialization (position + orientation) per dataset, and work with routine images thus their work is similar to our work.

We are working towards a fast and automated diagnosis of spine disorders for routine MRI data. In this paper we present a novel method for the segmentation of vertebral bodies in routine MRI datasets with highly anisotropic voxels, e.g. $0.6 \times 0.6 \times 4.4 \text{ mm}^3$. In contrast to commonly used top-down approaches [YOS06, GS04, KWL*08], our method relies on the segmentation of vertebral bodies to build up the global spine shape. This bottom-up approach is more efficient because it does not employ multiple optimization steps. Our inflation-based segmentation incorporates a novel constrained subdivision surface approach for efficient smoothness control and a multiple features boundary-estimator.

Our system addresses the stated drawbacks of prior systems and allows for comparably robust segmentation of pathological spine and vertebra shapes. It also does vertebra detection. The user initializes only the segmentation of the first vertebral body and the segmentation parameters are estimated from previous segmentations. It takes about 30 seconds to segment a whole dataset, and our system runs on datasets of routine quality. The quality of results in terms of precision are comparable to the state of the art.

We also compared our method to a general segmentation method power watersheds [CGNT09] of Couprie et al. Power watersheds are a combination of watersheds, graph cuts, and random walker.

2. Spine Segmentation System

We implemented two different types of initialization: center picking and freehand outline. The initialization should be performed on a slice which is approximately central to the spine. From either initialization we extract center coordinates and take a sample of intensities (in the case of center picking, from a small 2D neighborhood). Then, we remove

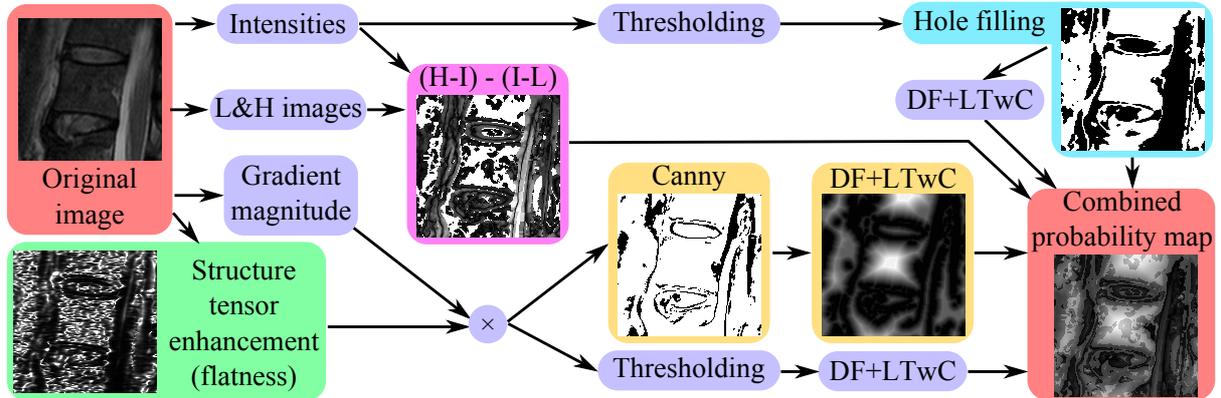


Figure 2: Features explained on T12 vertebra from F02 dataset in Fig. 7. In all images except original: black=edge, white=non-edges (\approx vertebral body interior).

outliers and calculate minimum, maximum, mean and standard deviation. The outline initialization additionally gives an approximate size, otherwise the average human vertebra size is taken [MSM*08]. Basically, it is an improved version of the one used in [ZEB*10] and [ZEB*11]. We only use center picking in this paper, for easier comparison to power watersheds.

The main segmentation loop consists of predicting the position of the adjacent vertebra's center, where the next segmentation is initiated. The required intensity statistics are taken from the previously segmented vertebra. Thus, the method adopts the gradual intensity change within the dataset. See Fig. 1.

Orientation of the vertebral body is determined by fitting an average vertebral body shape to it using the iterative closest point (ICP) algorithm. The initial position for ICP is determined by translating the average shape center to the segmented vertebral body center. Initial orientation of the average vertebral body shape is taken from the previous vertebra's fitted orientation (except the initial vertebra, where anatomical upright orientation is used). Initial scaling of the average shape is determined from the radius of the inflated mesh. The initialization brings the average shape close to the segmented one, which makes ICP robust. ICP is then allowed to optimize position, rotation and scaling to find the best fit. We extract its main axis (head-tail axis) from the orientation of the fitted shape.

The adjacent vertebra center estimate is placed along the main axis at a distance proportional to the radius of the first segmented vertebral body or, for all the others, given by the distance between the last two segmented vertebral bodies.

The main condition for validating the segmentation is that the difference between the radius of the current vertebral body is within 30% of the radius of the previous vertebral body. It also includes several other sanity checks, involving volume, volume/surface ratio and center-center distance of the newly and the previously segmented vertebral bodies.

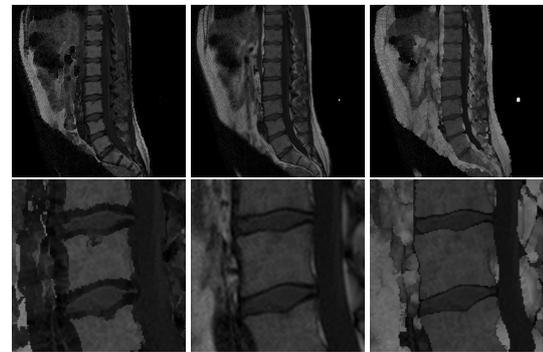


Figure 3: Left: low image. Middle: original slice. Right: high image. Bottom row: closeups. Cross-sections from the only healthy volunteer (dataset DzZ_T1 from Tab. 1).

3. Boundary Classification

The vertebral body boundaries are estimated using multiple features, which are classified into probabilities of the voxel v being at a boundary, and combined to a final probability $p(v)$ using the mean rule (see Fig. 2). Using multiple classifiers and then combining them is known to improve results and robustness [KHDM98]. We apply three edge based and two intensity based features.

The edge features are based on LH (low-high) values [ŠBSG06], Canny edges, and thresholded gradient magnitudes. The difference between low and high values, L, H indicates proximity to a boundary. The boundary probability is deduced from these values and the current voxel v intensity I by $p_{LH}(v) = ((H - I) - (I - L)) / I_{max}$, where I_{max} is the maximum intensity (see Fig. 2 and Fig. 3).

Gradient magnitudes and Canny edges are multiplicatively enhanced using the structure tensor as described by Fernández and Li [FL03] to improve the detection of 2-manifold edges and suppress one-dimensional features. For both features, the respective probability is derived by apply-

ing a distance field (DF) and a linear transform with clipping (LTwC), see Fig. 2.

The two intensity based features directly work on the thresholded MRI intensities. The thresholds are taken from the intensity statistics, i.e. the minimum and maximum values after outlier removal. We apply hole filling using morphological closing. We use the binarized intensities as one feature and the distance field constructed on this binary image as a second feature, thus incorporating both sharp edges and a smooth edge approach for mesh inflation.

After combining probabilities, true edges usually end up having boundary probability around 90%.

4. Segmenting Individual Vertebral Bodies

Each vertebral body is segmented using an iterative inflation algorithm. The algorithm starts with a small triangular surface mesh at the approximate center of the vertebral body. This mesh is enlarged using balloon inflation forces [Coh91], constrained by smoothness, and steered towards a star-shape geometry.

Smoothness is enforced by using a constrained subdivision surface scheme. Star-shape is achieved by allowing the inflation only along center-vertex direction. This deters self-intersections of the inflated surface.

The following steps are performed iteratively:

1. Move the vertices (outwards) towards the boundary with no regard to subdivision rules.
2. Normalize the subdivision hierarchy by moving vertices so that they comply with subdivision rules.
3. Recalculate the polyhedron center from the polyhedron shape to account for different inflation speed in different directions.
4. Stop when convergence is detected.

4.1. Inflating the Mesh

We do the inflation by examining each vertex separately. Depending on the boundary probability, we either inflate or deflate the vertex along the radial (center-vertex) direction. The step size is equal to the minimum voxel spacing, i.e. 0.5–1.1 mm for the examined datasets.

Initially, a vertex v is in the interior and will inflate as long as the boundary probability $p(v) < 0.5$. This way noise inside the vertebral body is ignored. When $p(v) \geq 0.5$ the sign of the probability derivative $\frac{\delta p}{\delta \vec{r}}$ in inflation direction \vec{r} decides on whether to inflate or to deflate, thereby moving the vertex towards maximum boundary probability.

Convergence is achieved when the average center-surface distance (“radius”) stops increasing, i.e. the radius from the current iteration is not larger than the radii from the previous two iterations. An additional check is whether the radius is

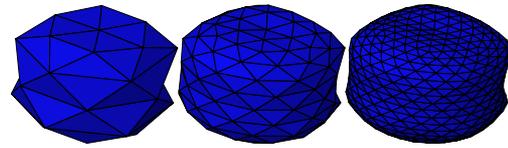


Figure 4: Triangular mesh model which we used. Left: base mesh. Middle: level 1 subdivided mesh (base mesh + additional vertices). Right: level 2 subdivided mesh.

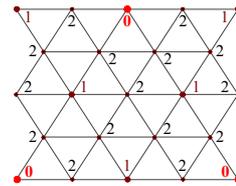


Figure 5: Surface patch showing vertex hierarchy with two refined levels (0–base level).

50% greater than the last vertebra. This saves computation time in case of mis-segmentation.

If the initial center estimate for the segmentation is very close to a boundary, the surface will inflate much more in the opposite direction. This will result in a highly uneven distribution of vertices over the surface, which is detected using the standard deviation of the edge lengths ($\sigma > 1$ mm). In this case, the segmentation is restarted at the current center. In our experiments, one to two restarts with a new center have been sufficient to resolve the problem. This mostly happens with the second vertebra, because intervertebral distance at that time is only an estimate.

4.2. Constrained Subdivision Hierarchy

The butterfly algorithm is the simplest interpolating subdivision scheme working on triangle meshes. Interpolating means that the vertex’s position, once calculated, will be a part of all finer mesh levels and consequently is subject to inflation and deflation. The modified butterfly was presented by Zorin et al. [ZSS96], and it avoids problems with irregular vertices (vertices with valence $\neq 6$).

Our base mesh consists of a closed triangular polyhedron with 32 vertices and 60 triangles (see Figs. 4 and 5). We create new mesh levels through the subdivision rules until the average edge length l is comparable to the voxel size s : $s \leq l < 2s$, where $s = \sqrt[3]{s_x s_y s_z}$ is the geometric mean of the voxel spacings.

The inflation step repositions the vertices with no regard to subdivision rules. Thus, we need to enforce the subdivision hierarchy for all the dependent levels afterwards. We use a global least squares optimization, minimizing the vertex position correction. Only the base mesh vertices are independent, i.e. the base mesh is the control mesh.

Table 1: Information about the datasets used for the quantitative evaluation. S_{xy} – voxel spacing in X and Y directions (millimeters), S_z – spacing along Z axis. xyR – resolution of image along X and Y axes, zR – Z resolution. AF – anisotropy factor $\frac{S_z}{S_{xy}}$. *Path* – pathologies (*St*–stenosis, *SD*–slipped disc, *Sco*–scoliosis, *VF*–vertebra fracture, *SL*–spondylolisthesis).

Dataset	S_{xy}	S_z	xyR	zR	AF	Path.
Ble	0.63	4.4	512	14	7.04	St,SD
C002	1.12	3.3	448	31	2.96	Sco
DzZ_T1	0.68	4.4	512	12	6.44	None
F02	0.5	3.85	768	18	7.70	VF,SL
F03	1.19	3.3	320	25	2.77	VF
Geh	0.63	4.4	512	10	7.04	SD
LC	0.73	4.4	384	14	6.03	SD
S01	0.47	3.85	640	16	8.19	SL
S02	0.47	3.85	640	16	8.19	SL
Sch	0.63	4.4	512	16	7.04	St,SL
St1	0.5	3.85	704	20	7.70	St
Average	0.68	4.0	541	17	6.46	
StdDev	0.25	0.43	135	6	1.90	

If there are many vertices in the base mesh, overshooting effects start to appear, which is frequent in interpolating schemes. This can be avoided by using a low or moderate number of vertices in the base mesh, or by applying slight smoothing after each iteration, thus fighting noise and small ambiguities in the data.

5. Results

Testing was done on a machine with Intel Core i7-920 2.67GHz processor. Preprocessing of typical images (512x512x16) takes about 15 seconds. The individual vertebra segmentation time was about one to three seconds. With about eight vertebrae visible in most images, it puts the total processing time to about 30 seconds per dataset. That is significantly better than the other spine MRI 3D methods, and almost the same as the power watersheds method which is inferior to ours in terms of quality.

We tested our method on ten pathological datasets from various hospitals and one dataset from a healthy volunteer. The reference segmentations were done by neurosurgeons, who manually traced the vertebral body edges in the sagittal plane. The images were mostly T2-weighted, mostly from middle-aged and older patients. The important property of these datasets is the high anisotropy of voxel size between 3x and 8x (see Tab. 1 and Fig. 7). Coronal and axial views show the anisotropy.

The datasets with more severe pathologies (e.g. crushed vertebra) have lower DSC than less noticeable pathologies (e.g. stenosis). Pathologies also reduce the detection rate.

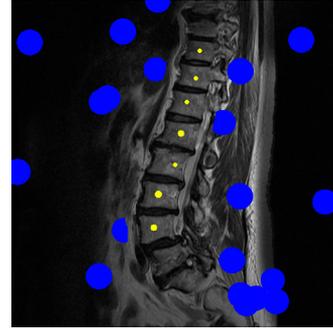


Figure 6: One cross-section of seeds overlaid on image (dataset Ble). Blue are background seeds, yellow are seeds of vertebral bodies. It is noticeable on V.B. seeds that they have different Z -positions (they appear to be of different sizes on this cross-section).

5.1. Test Setup

Besides comparison to methods developed specifically for MRI of spine, we also subjected all our datasets to general segmentation method, power watersheds of Couprie et al. [CGNT09]. The source code for this method is publicly available, so we could run it on the same datasets as our own method. Execution time of power watersheds on a typical dataset (512x512x16) is about 20 seconds.

The graph cuts method (and its extension power watersheds) is well suited for interactive segmentation, a mode where additional seeds are placed until a satisfactory segmentation is achieved. It is not suitable for semi-automatic segmentation, because it requires a lot of seeds and can lead to unplausible results which requires additional seed placement. Initializing this method in 3D manually is quite painstaking so we opted for automatic seed creation, derived from ground truth data, for the purpose of a thorough comparison.

One seed was placed into each vertebral body on a random position near the center, and twice as many seeds into the background (Fig. 6). Since background seeds were large (20 voxel radius), they were clipped by a safety region of interest around vertebral bodies. This safety region was created by morphologically dilating ground truth vertebral body masks with a 20 voxel radius spherical structuring element (Fig. 6). There was one such initialization for each vertebra in the dataset, and we used one of these vertebra seeds centers to initialize our method so it gets initialized once at each of the vertebral bodies of all the datasets.

As there were 92 reference segmented vertebrae, we had 92 different initializations for both power watersheds and our method (“all VB” row in Tab. 2). Power watersheds crashed on datasets F02 and St1 (with ≈ 10 million voxels) due to hitting 2GB user memory limit (Win32), hence the lower number of reference vertebrae in Tab. 2. Since our method should be initialized at a vertebral body near the cen-

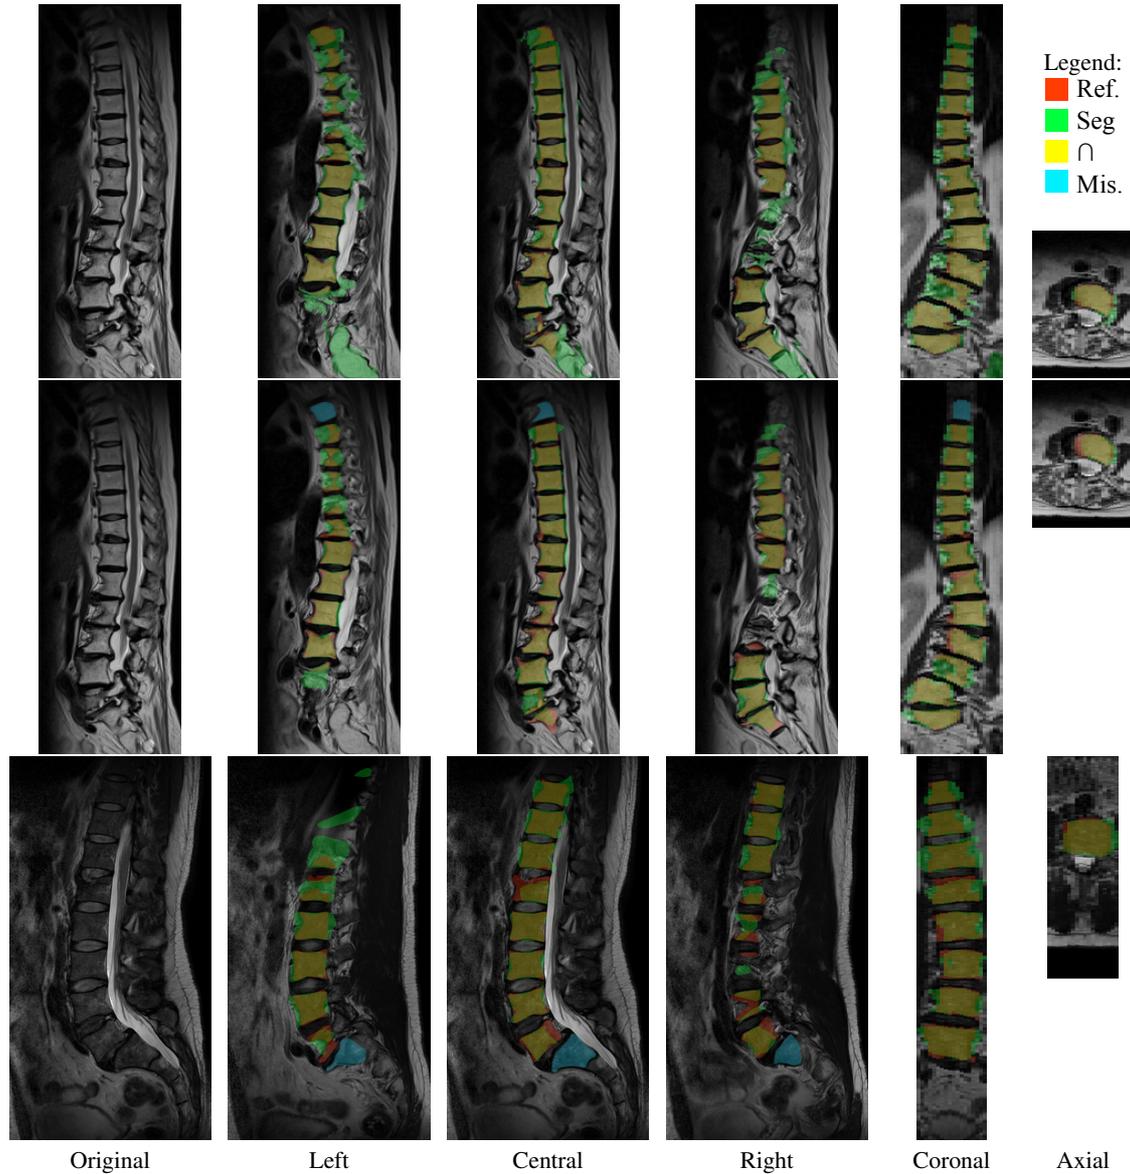


Figure 7: Overlay of reference segmentation and the one produced by power watersheds (top row) and our method (bottom two rows). Images were cropped to save space (unimportant parts were cut off). Real image sizes were $448 \times 448 \times 31$ with $1.1 \times 1.1 \times 3.3 \text{ mm}^3$ voxels (top and middle, C002) and $768 \times 768 \times 18$ with $0.5 \times 0.5 \times 3.85 \text{ mm}^3$ voxels (bottom image, F02). DSC were 65%, 75% and 75%.

ter of the image, and preferably on a segment which is not pathological (so first 2 vertebral bodies are segmented correctly), we also ran the test without initializing edge vertebrae (2 uppermost and 1 lowermost vertebra were excluded, “mid VB” row). Our method provides the best results with manual initialization by an experienced user (“manual initialization” row).

The overview of results is given in Tab. 2. In Tab. 3, per vertebra statistics are given for power watersheds method

along the results of our method for automatic initialization at all vertebrae.

We measured surface distance errors using the Metro mesh comparison tool [CRS96]. Distance error is measured by iterating through vertices and sample points on polygons of the segmentation surface and measuring the distance to the closest point of the reference surface.

The mean distance of segmentation from the reference surface is $2.15 \text{ mm} \pm 1.02 \text{ mm}$. To get a relative measure, we

Table 2: Overview of segmentation results for our method and power watersheds. DSC is the average Dice Similarity Coefficient. Dist is the mean distance of segmented surface to reference surface in millimeters. DR is detection rate (how many vertebral bodies were discovered and segmented). Detection was not done for PW, only segmentation. NR is number of reference vertebrae used. Auto means automatic initialization from ground truth data.

Method, Init.	DSC	Dist	DR	NR
Power Watersheds	68%±17%	7.87	-	659
Our, auto all VB	74%±16%	2.21	50%	804
Our, auto mid VB	76%±13%	2.08	62%	528
Our, manual init.	78%±8%	2.15	93%	92

Table 3: Spine segmentation system results with automated randomized initializations. V.C.–vertebra count in all 11 datasets. pwD and ourD–mean surface distance in mm.

Vert.	V.C.	pwDSC	pwD	ourDSC	ourD
T6	1	61.63%	3.28	19.44%	14.03
T7	1	61.21%	6.39	62.98%	3.38
T8	1	76.98%	2.46	74.34%	1.77
T9	2	63.96%	7.89	73.77%	2.21
T10	4	62.25%	10.32	73.45%	2.55
T11	7	67.29%	7.31	69.51%	2.69
T12	11	69.38%	9.31	71.28%	2.52
L1	11	72.80%	7.10	75.71%	1.98
L2	11	77.45%	4.07	78.11%	1.77
L3	11	74.44%	5.38	78.41%	1.88
L4	11	70.14%	7.83	77.32%	1.72
L5	11	70.81%	7.50	75.11%	2.03
S1	10	44.59%	15.92	64.51%	3.05
Avg.		68.11%	7.87	74.05%	2.21
StdDev		16.57%	9.21	16.09%	1.37

would divide this by the voxel size. In our case of anisotropic voxels, we divide by the edge of isotropic voxel which is equivalent by volume, $A_{iso} = \sqrt[3]{s_x s_y s_z}$. Average A_{iso} for these 11 datasets is $1.21\text{mm} \pm 0.25\text{mm}$, and relative distance is 1.73 ± 0.99 .

5.2. Discussion

The average DSC of 78% is close to Davatzikos et al. with 81%, but they used healthy individuals. Healthy volunteer datasets are easier to segment and achieve better DSC and distance scores. In the three different tests from Tab. 2, the DSC results for the healthy volunteer were 77%, 80% and 86% respectively, or higher than average by 3, 4 and 8 percentage points.

Our method is clearly better than that of Štern et al. at $1.85\text{mm} \pm 0.47\text{mm}$, because they used isotropic voxels of 1mm^3 and this should be compared to our relative distance error of 1.73 ± 0.99 . Hoar and Martel’s surgery-oriented method with thorough initialization and manual correction

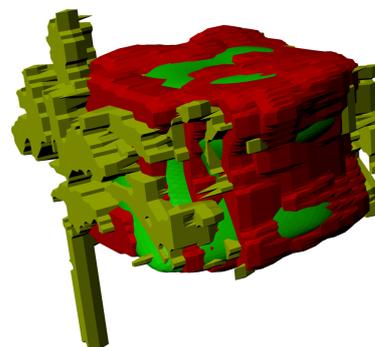


Figure 8: Surface overlay of our segmentation (green) and power watersheds (yellow). Red is the reference segmentation. Reference and PW surfaces were derived from binary masks using marching cubes. L2 vertebral body from DzZ_T1 dataset. DSC: PW–85%, our–88%.

of segmentation at $1.25\text{mm} \pm 0.28\text{mm}$ mean distance error remains more precise than ours. Contributing reason for higher variance in our result is a high anisotropy factor (6.47).

Our method is vastly superior to power watersheds for the purpose of vertebral body segmentation in MRI, as it has similar execution time, higher DSC and significantly lower distance error. Power watersheds have lower DSC than our method in spite of fairly rich initialization (Fig. 6). Power watersheds frequently produce some protruding spikes (Fig. 8). These spikes are thin, so they do not hurt DSC measure so much, but surface distance errors are significant (7.87mm).

The majority of discrepancies between manual and automatic segmentations stem from lateral slices (Fig. 7), where it is harder to algorithmically discern a boundary due to significant partial volume effects. Power watersheds mostly fail on lateral edges of vertebral bodies too (Fig. 7 and Fig. 8).

The quality of the segmentation could have been further improved if we did not do detection based on the segmentation, but rather strongly guided the segmentation towards the expected size. This approach, however, significantly decreases the reliability of any subsequent diagnosis, which we have already started working on but do not present in this paper.

6. Conclusion

We presented a bottom-up approach for spine segmentation in routine MRI. Our segmentation is fast and robust with respect to the low and anisotropic resolution of routine MRI datasets and to pathological spine and vertebra shapes. Our vertebral body segmentation method is inflation-based and incorporates a novel constrained subdivision surface approach for smoothness control and a multiple features boundary estimator. The spine segmentation requires only

minimal user interaction and takes about half a minute to segment the whole dataset.

Our method was tested on a larger set of vertebrae than prior work. Hoard and Martel tested their method on 30 vertebrae, Štern et al. on 75 vertebral bodies in MR modality, and Davatzikos et al. on 84 vertebral bodies from healthy volunteers. We tested our method on 92 vertebral bodies (83 from pathological datasets and 9 from a healthy volunteer).

Acknowledgments

We are grateful to Miriam Bauer from University Hospital Marburg for helping with some of the manual segmentations.

Dr. Dr. Jan Egger has received support from NIH grant 8P41EB015898-08.

References

- [CGBM04] CARBALLIDO-GAMIO J., BELONGIE S., MAJUMDAR S.: Normalized cuts in 3-D for spinal MRI segmentation. *IEEE Trans. Medical Imaging* 23, 1 (2004), 36–44. 1
- [CGNT09] COUPRIE C., GRADY L. J., NAJMAN L., TALBOT H.: Power watersheds: A new image segmentation framework extending graph cuts, random walker and optimal spanning forest. In *ICCV'09* (2009), pp. 731–738. <http://powerwatershed.sourceforge.net/>. 2, 5
- [Coh91] COHEN L. D.: On active contour models and balloons. *CVGIP: Image Understanding* 53, 2 (1991), 211–218. 4
- [CRS96] CIGNONI P., ROCCHINI C., SCOPIGNO R.: *Metro: measuring error on simplified surfaces*. Tech. rep., Centre National de la Recherche Scientifique, Paris, France, France, 1996. <http://vcg.sourceforge.net/index.php/Metro>. 6
- [DLSH02] DAVATZIKOS C., LIU D., SHEN D., HERSKOVITS E. H.: Spatial normalization of spine MR images for statistical correlation of lesions with clinical symptoms. *Radiology* 224, 3 (2002), 919–926. 2
- [EKD*12] EGGER J., KAPUR T., DUKATZ T., KOŁODZIEJ M., ZUKIĆ DŽ., FREISLEBEN B., NIMSKY C.: Square-cut: A segmentation algorithm on the basis of a rectangle shape. *PLoS ONE* 7, 2 (02 2012), e31064. 1
- [FHA*09] FREBURGER J., HOLMES G., AGANS R., JACKMAN A., DARTER J., WALLACE A., CASTEL L., KALSBECK W., CAREY T.: The rising prevalence of chronic low back pain. *Arch Intern Med* 169, 3 (2009), 251–258. 1
- [FL03] FERNÁNDEZ J.-J., LI S.: An improved algorithm for anisotropic nonlinear diffusion for denoising cryo-tomograms. *J. Structural Biology* 144, 1-2 (2003), 152–161. 3
- [GS04] GHEBREAB S., SMEULDERS A.: Combining strings and necklaces for interactive three-dimensional segmentation of spinal images using an integral deformable spine model. *IEEE Trans. Biomedical Engineering* 51, 10 (2004), 1821–1829. 1, 2
- [HCLN09] HUANG S.-H., CHU Y.-H., LAI S.-H., NOVAK C.: Learning-based vertebra detection and iterative normalized-cut segmentation for spinal MRI. *IEEE Trans. Medical Imaging* 28, 10 (2009), 1595–1605. 1
- [HM02] HOARD C. L., MARTEL A. L.: Segmentation of MR images for computer-assisted surgery of the lumbar spine. *Physics in Medicine and Biology* 47, 19 (2002), 3503. 2
- [KHDM98] KITTLER J., HATEF M., DUIN R., MATAS J.: On combining classifiers. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 20, 3 (mar 1998), 226–239. 3
- [KWL*08] KLINDER T., WOLZ R., LORENZ C., FRANZ A., OSTERMANN J.: Spine segmentation using articulated shape models. In *Proc. MICCAI*, vol. 5241 of *LNCS*. Springer, 2008, pp. 227–234. 1, 2
- [KZS*11] KELM B. M., ZHOU S. K., SUEHLING M., ZHENG Y., WELS M., COMANICIU D.: Detection of 3D spinal geometry using iterated marginal space learning. In *Proc. MICCAI-MCV*, vol. 6533 of *LNCS*. Springer, 2011, pp. 96–105. 1
- [MCP*09] MICHOPPOULOU S., COSTARIDOU L., PANAGIOTOPOULOS E., SPELLER R., PANAYIOTAKIS G., TODD-POKROPEK A.: Atlas-based segmentation of degenerated lumbar intervertebral discs from MR images of the spine. *IEEE Trans. Biomedical Engineering* 56, 9 (2009), 2225–2231. 1
- [MSM*08] MASHARAWI Y., SALAME K., MIROVSKY Y., PELLEGG S., DAR G., STEINBERG N., HERSHKOVITZ I.: Vertebral body shape variation in the thoracic and lumbar spine: Characterization of its asymmetry and wedging. *Clin. Anatomy* 21, 1 (2008), 46–54. 3
- [PZWhL05] PENG Z., ZHONG J., WEE W., HUEI LEE J.: Automated vertebra detection and segmentation from the whole spine MR images. In *Inter. Conf. Eng. in Medicine and Biology Society* (2005), pp. 2527–2530. 1
- [RGMB10] RICHARDS P. J., GEORGE J., METELKO M., BROWN M.: Spine computed tomography doses and cancer induction. *Spine* 35, 4 (2010), 430–433. 1
- [ŠBSG06] ŠEREDA P., BARTROLI A., SERLIE I., GERRITSEN F.: Visualization of boundaries in volumetric data sets using LH histograms. *IEEE Trans. Visualization and Computer Graphics* 12, 2 (2006), 208–218. 3
- [SD00] SHEN D., DAVATZIKOS C.: Adaptive-focus statistical shape model for segmentation of 3D MR structures. *Medical Image Computing and Computer 1935* (2000), 206–215. 2
- [ŠLPV11] ŠTERN D., LIKAR B., PERNUŠ F., VRTOVEC T.: Parametric modelling and segmentation of vertebral bodies in 3D CT and MR spine images. *Physics in Medicine and Biology* 56, 23 (2011), 7505. 2
- [SSQW07] SHI R., SUN D., QIU Z., WEISS K.: An efficient method for segmentation of MRI spine images. In *IEEE/ICME Int. Conf. Complex Medical Engineering* (2007), pp. 713–717. 1
- [WKL*01] WEESE J., KAUS M., LORENZ C., LOBREGT S., TRUYEN R., PEKAR V.: Shape constrained deformable models for 3D medical image segmentation. In *Inform. Proc. in Medical Imaging*, vol. 2082 of *LNCS*. Springer, 2001, pp. 380–387. 1
- [YOS06] YAO J., O'CONNOR S., SUMMERS R.: Automated spinal column extraction and partitioning. In *IEEE Int. Symp. Biomedical Imaging* (2006), pp. 390–393. 2
- [ZEB*10] ZUKIĆ DŽ., EGGER J., BAUER M. H. A., KUHN D., CARL B., FREISLEBEN B., KOLB A., NIMSKY C.: Glioblastoma multiforme segmentation in MRI data with a balloon inflation approach. In *Russian-Bavarian Conf. on Bio-Medical Engineering* (2010), pp. 40–44. 3
- [ZEB*11] ZUKIĆ DŽ., EGGER J., BAUER M. H. A., KUHN D., CARL B., FREISLEBEN B., KOLB A., NIMSKY C.: Preoperative volume determination for pituitary adenoma. In *Proc. SPIE Medical Imaging* (2011), vol. 7963. 3
- [ZSS96] ZORIN D., SCHRÖDER P., SWELDENS W.: Interpolating subdivision for meshes with arbitrary topology. In *Proc. ACM SIGGRAPH* (1996), pp. 189–192. 4