Intervertebral Disc Segmentation Using Mathematical Morphology

Edwin Carlinet and Thierry Géraud

EPITA Research and Development Laboratory (LRDE), France

Abstract. We propose a method to segment intervertebral discs (IVDs) from 3D magnetic resonance images, based on simple image processing operators. Most of these operators come from the mathematical morphology framework. Driven by some prior knowledge on IVDs (basic information about their shape and the distance between them), and on their contrast in the different modalities, our method has four stages. 1. Prepare a single 2D image, where IVDs appear more prominently than in the input multi-modality volumes. 2. Localize a 2D region of interest (ROI) per IVD. 3. Perform the segmentation on the 3D volume, constrained by the ROIs, thanks to a morphological tree-based representation of the volume image. 4. Regularize and remove disconnected regions.

Keywords: mathematical morphology · tree of shapes

1 Method Description

1.1 IVDs Gross Estimation

The first step of the method aims at getting a gross estimation of the IVDs in 3D which will be refined later. At this stage, we do not need a precision at pixel level, only the bounding box of the IVDs.

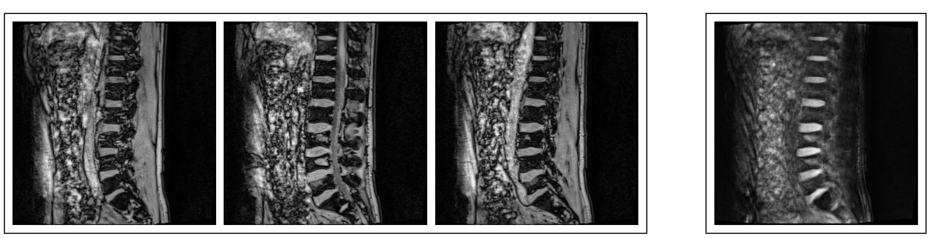

Fig. 1. Left: *opp* slices at $z = 8, 16, 24$. Right: 2D projection after preprocessing.

Image preprocessing. We work with the *opp* volume only. In slices which reveal the IVDs the most, IVDs appear as bright oriented blobs which are at least 7-pixels high. Thus, for each slice, a *top-hat* with a flat vertical structuring element of size 7×1 allows filtering out the background and highlights the IVDs. Then, the slices are summed up (similar to an Average Intensity Projection along the z-azis) to produce a consensus image. The projection serves as a Temporal Noise Reduction to reduce noisy structures that could have passed the tophat filtering in some frames (see fig. 1).

IVD selection. We compute the Tree of Shapes (ToS) on the preprocessed image. The latter enables a hierarchical representation of the inclusion of the hole-filled connected components of the image. The tree is then filtered by some prior-knowledge-based basic criteria:

- bounding box position & dimension of the shape
- position of the center of the shape
- orientation of the shape
- height of the shape
- average gray level of the shape

Only about 20 *maximal* (i.e., *not included in any other shape*) shapes S_i are able to pass these requirements. In these shapes, we then look for the sub-shapes S_i^* the most *compact* (ratio of the surface over the enclosing oriented rectangle surface) to favor more *regular* shapes. We then need to select 7 shapes among these candidates (because 7 IVDs are expected). The candidates are sorted by decreasing average gray value. The first one serves as a reference and is augmented with shapes taken from S_i^* satisfying some relative positioning constraints:

- the y-distance between the shape center and the current bounding box is between 15 and 45 pixels
- the x-distance between the shape center and top/bottom selected shapes is below 15 pixels

An example is shown in fig. 2.

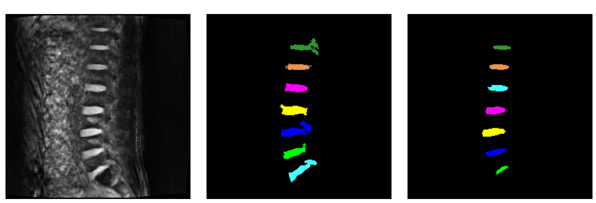

Fig. 2. Markers & IVDs selection. Left: input preprocessed image. Middle: some *maximal* candidate shapes. Right: selected subshapes with a *compactness* criterion.

1.2 3D Volume Segmentation

The previously detected seeds are used to guide the search in the 2D slices. We now work on an image combining the *opp*, *fat* and *wat* modalities as IVDs contours may be spread among these images. Yet, the tophat filtering is used to enhance the contrast of IVDs.

$$\text{input} = \textit{tophat}(\text{opp}) + \textit{tophat}(\text{wat}) - \text{fat}$$

Basically, we apply the same method as in the preprocessing. A ToS is computed on each slice of *input* and we look for the best *regular* shapes passing some geometric criteria and matching the IVDs Regions of Interest computed previously.

1.3 Post-processing: 3D Regularization

Z-axis regularization. In some slices, when no shape can be found for a given IVD, it may lead to a missed detection. If a pixel (x, y) is labeled at $z = k-1$ and $z = k+1$, but not at slice $z = k$, it is likely a miss-detection. As a consequence, the regularization applies:

$$f(z, x, y) = f(z, x, y) \vee (f(z-1, x, y) \wedge f(z+1, x, y))$$

Isolated pixels removal. On the contrary, false-detection may appear (especially at the beginning and the end of the sequence) in the noise. These shapes are generally disconnected in 3D from the *real* IVDs. Thus, as a final step, we perform a 3D connected component labeling and only retain the 7 largest ones.