DIMACS Series in Discrete Mathematics
and Theoretical Computer Science

Analysis of Magnetic Resonance Angiography Images Using Skeletonization and Distance Transforms

Ingela Nyström and Örjan Smedby

Abstract. In magnetic resonance angiography (MRA), large 3D data sets of the vascular system are acquired. This paper deals with two techniques to represent relevant information from such data sets in a more compact format. After segmentation of the vessels, the curve skeleton is extracted by an algorithm based on the distance transformation. Using the D^6 distance to the background, the algorithm first reduces the original object to a surface skeleton and then to a curve skeleton, after which "pruning" can be performed to remove irrelevant small branches. Applying this procedure to MRA data from the pelvic arteries resulted in a good description of the tree structure of the vessels with a much smaller number of voxels. To detect stenoses, 2D projections such as maximum intensity projection (MIP) are usually employed, but these often fail to demonstrate a stenosis if the projection angle is not suitably chosen. A new method combines the distance-labelled curve skeleton with the reverse distance transformation to visualize stenoses independent of viewing direction. Experiments with synthetic data indicate that stenoses invisible in an ordinary projection may be seen with this technique.

1. Introduction

Both magnetic resonance angiography (MRA) [**AET93, Pri94**] and computed tomography angiography (CTA) [**NMR+92**] are methods that can produce (3D) volume images of the vasculature in the human body. MRA, which has been used here, is a magnetic resonance imaging (MRI) method to assess vascular disease that has spread widely into the clinical routine.

Volume imaging usually implies large amounts of data, and it is, in general, a non-trivial problem how to present these volumes on 2D devices such as screens and photographic film. When interpreting and analysing MRA images, results are usually presented as maximum intensity projections (MIP), multi-planar reconstructions (MPR), or volume renderings [**DHW+97**]. The choice of viewing direction for these (and other) projections is of great importance for the detection of stenoses.

2000 *Mathematics Subject Classification.* Primary 68T10, 92C50; Secondary 57N99, 52C99.
Key words and phrases. pattern recognition, shape representation, 3D skeleton, reverse distance transform, MR angiography, arteries, stenosis detection.

Despite the large data volumes, the relevant information in an MRA or CTA data set is rather limited. Usually, the diagnostic physician is interested in the vessel anatomy, i.e., in the tree structure of the arteries, and in variations in vessel calibre. If one can identify the vessels with their branching points and the stenoses or narrow parts, this is usually an adequate description of the vessel system. In this paper, we present a method to use a distance-labelled *3D curve skeleton* to extract this important information from a volume image of the vascular system. We will also introduce a new presentation method aiming at visualizing variations in vessel width independent of viewing direction.

Skeletonization, also referred to as thinning, is a process where objects are reduced to structures of lower dimension while preserving topological properties as well as the general shape of the object. Two main classes of skeletonization methods for 2D images have been developed: *iterative thinning* and *distance transform analysis*. In 3D, the published methods have almost exclusively been based on iterative thinning, e.g., [**Ber95**] and [**SCM97**]. Some published methods produce surface skeletons, some of them reduce the objects to curve skeletons. Thin elongated objects, such as blood vessels, are well suited for reduction to curve skeletons. Earlier work with the aim to obtain curve skeletons for applications can be found, e.g., in [**MS96**], in [**GSV96**], where the skeleton concept from [**LKC94**] is adopted, and in [**PK99**]. The methods were applied to visual analysis of computed tomography (CT) lung data, interactive virtual colonoscopy rendered from spiral CT scans, and ventricle extraction from MRI brain data, respectively.

The method used in this paper for reducing objects in binary volume images to curve skeletons is described in detail in [**BNSdB99**]. It is a skeletonization based on iterative thinning, but guided by a distance transform. After skeletonization, the blood vessels consist of thread-like structures, which are one voxel thick and located centrally in the vessel lumen, while containing the distance to the closest background voxel from every voxel. These distances are utilized in our presentation method. For example, a stenosis reducing the anteroposterior diameter of a craniocaudal vessel, which is normally seen only in the lateral projection, will then become visible also in the frontal projection.

This work is a continuation of the work presented in [**Nys98**]. Here, we show a useful application for our earlier work. MRA is briefly introduced in Section 2. Distance transforms are briefly described in Section 3.1, including a method for reversing the skeleton, while the skeletonization method is outlined in Section 3.2. Resulting skeletons will be shown for synthetic and MRA images in Section 4. In Section 5 we show how the curve skeletons may be utilized in a presentation method.

2. Magnetic resonance angiography

Magnetic resonance angiography is a technique, or rather a set of techniques, to non-invasively image the flowing blood in vessels. The physical principles can vary, and there are three methods in common clinical use today:

inflow angiography: also known as time of flight (TOF), which uses flow effects on the signal amplitude,

phase contrast angiography: which uses flow effects on the phase of the signal, and

Gadolinium angiography: which does not rely on flow effects, but on the effect on MR signal by a paramagnetic contrast material, Gadolinium (Gd), which is given as an intravenous injection in order to enhance the arteries.

In all three cases, the most common situation involves acquiring a 3D data set, and then presenting the data on a screen, either as a single 2D slice (which may have a different orientation from the originally acquired slices) or using rendering algorithms such as the most commonly used maximum intensity projection (MIP). An important issue with all projection methods is the choice of viewing direction.

In our study, Gd-MRA images were taken from the pelvic arteries in a patient with atherosclerosis. Figure 1 (top) shows one slice in the data set, showing only parts of the vessel system. Figure 1 (bottom), presents the volume as a MIP, clearly showing the vessel system, but with no depth information in it.

3. Image processing methods

Our method employs two central concepts from image processing: *skeletonization* and (ordinary and reverse) *distance transformation*. It should be noted that here we assume that binary images have been obtained, e.g., by segmentation into object and background.

3.1. Distance transformation. Each object voxel can be labelled with (an approximation of) the *distance* to the closest background voxel. In general, the best metric to choose when computing the distance depends on the application. The three simplest metrics for volume images are the "path-generated" distances D^6, D^{18}, and D^{26}, where the distance is defined as the number of steps in the minimal path between voxels with 6-, 18-, and 26-connectedness, respectively [**Bor96**]. In this paper we will use only D^6, i.e., the 3D equivalent of the city-block distance, where steps only to the six face-neighbours of a voxel are allowed. The D^6 distance transform of the object consists of values representing the number of steps to the closest background voxel. These values are ordered in layers that you may imagine can be peeled off one value at a time. All voxels with a face-neighbour in the background will then have the same distance value. To compute the distances in an image a *distance transformation* (DT) is used, here DT^6. A distance transformation relies on the same principle in the 3D case as in the 2D case [**Bor84**]. The DT^6 is computed by a forward and a backward scan through the image, propagating increasing distances. Each voxel is assigned the minimum of the value of the current voxel and the values of the already visited face-neighbours added by one.

The distance label of any voxel in the distance transform can be interpreted as the radius of a digital ball — a solid sphere — centred on that voxel. Balls look different for different metrics, e.g., the D^6 ball is an octahedron. The definition of the distance transform implies that each ball is as large as possible still included in the object.

The *reverse* distance transformation starts from a number of seed voxels with distance values [**RP66**]. Each distance value corresponds to the radius of a ball. The result is an object with the shape of the union of the balls associated to the distance values. The reverse DT can also (as the ordinary DT) be computed by two raster scans, but this time decreasing distances are propagated. Each voxel is assigned the maximum of the value of the current voxel and the values of the

FIGURE 1. Two different presentations of an MRA image. Top: sample 2D slice. Bottom: MIP.

already visited face-neighbours subtracted by one. Pseudo-code for the forward scan of the D^6 reverse DT (the backward scan is similar) [**NB95**]:

```
/*------------------FORWARD SCAN-------------------------------*/
LOOPXYZ(image) {
   if (x<xLO || x>xHI || y<yLO || y>yHI || z<zLO || z>zHI)
     I(x, y, z) = 0; /* image border omitted */
   else
     I(x, y, z) = MAX(I(x,y,z-1)-1,
                      I(x,y-1,z)-1,
                      I(x-1,y,z)-1,
                      I(x, y, z))
}
/*------------------END--FORWARD SCAN--------------------------*/
```

An object voxel is called a *centre of a maximal ball* (CMB), if the associated ball is not completely covered by any other single ball in the object. The union of the maximal balls, which might be computed by applying the reverse DT using the set of CMBs as seed values, coincides with the whole object. For DT^6, the CMBs are equivalent to the local maxima. (The situation is more complex for other DTs.)

3.2. Skeletonization. The skeletonization method we have used starts with a binary volume image. A number of skeletal properties are fulfilled for our skeletons:

- They have the same topology as the object.
- They preserve the general shape of the object.
- They are centred in the object.
- They are a thin subset of the object.
- They allow reconstruction of the object.

Preserving the topology in this case means that the number of components, the number of tunnels, and the number of cavities all remain the same. In 3D, one must distinguish between two different types of "holes", i.e., tunnels, like what goes through a doughnut, and cavities, like what exists in a hollow ball. Also, in contrast with the 2D case, it is not sufficient to count only the components of the object, but it is necessary to count the components of the background as well. On digital topology, see also [**KR89, MBA93**].

The reconstruction property is true for the surface skeletons, and sometimes also true for the curve skeletons (see below).

The skeletonization method is described in detail in [**BNSdB99**]. It is based on a small number of simple neighbourhood operations; it is fairly time efficient. The skeletonization consists of five parts, where Part I and Part II reduce the original solid 3D object to a surface skeleton consisting of surfaces which are only one voxel thick. In Part III and Part IV, we go further and reduce this surface skeleton to a set of curves which are only one pixel wide. Part V, that may be added, is a "pruning" process, which is performed to remove irrelevant branches.

The method we have used differs in some respects from [**BNSdB99**]. By computing the DT^6 of the object first of all, instead of iteratively identifying the border voxels when computing the surface skeleton, we reduce computation time. Border voxels can be directly defined on the DT^6 by the distance layers, which equals the current iteration. Each successive border layer is 6-connected to the previous layer, exactly as the layers are ordered in the DT^6. Internal voxels are voxels with

higher distance values. For details, see [**BNSdB99**]. The time reduction lies in the iterative process, where now only three passes, instead of four, are needed in each iteration:

Pass 1: Inspect voxels with the current iteration number. Mark those that are *multiple*, i.e., CMBs and voxels necessary for preserving connectivity. See [**BNSdB99**] for details.

Pass 2: Again, inspect voxels with the current iteration number. Mark those that are *tunnel* voxels, which, if removed, would create a tunnel.

Pass 3: Remove all unmarked border voxels.

The process terminates when all border voxels are multiple, and hence no more voxels can be removed.

The rules for detecting multiple voxels prevent disconnecting objects of all connectivities, but they do not prevent creation of tunnels in thin complex surfaces. Therefore, Pass 2 is necessary (see [**BNSdB99**]). The original object can be reconstructed from the resulting surface skeleton by growing balls of radius equal to the DT^6 values, e.g., by efficiently applying the reverse DT^6, as all the CMBs are included. This reconstructibility is lost in the following skeletonization parts.

Part I, which iteratively removes simple (or non-multiple) border voxels, typically results in something which is almost a surface skeleton but not quite, as it may be up to two voxels thick where the original object had a width of an even number of voxels. We therefore have a separate Part II in the algorithm to reduce the thickness from two to one.

Similarly, in the curve skeletonization, Part III removes voxels from the outer edges of the surface until the skeleton is at most two voxels wide:

Pass 1: Identify *outer* voxels of the surface, i.e., voxels whose removal would not create tunnels.

Pass 2: Mark *protrusion* and *break-point* voxels. See [**BNSdB99**] for details.

Pass 3: Sequentially remove all unmarked outer voxels, unless that would disconnect the object.

In case the original surface is an even number of voxels in width, the curve skeleton becomes two-voxel wide. It is reduced to unit thickness by iteratively applying a thinning process in Part IV.

Finally, when the skeleton consists of thin curves, we may post-process the resulting skeleton by removing short branches that are assumed to be caused by noise in the original image. This "pruning" (Part V) is described in Section 4.3. Zig-zags in the curve may also be straightened.

4. Resulting curve skeletons

The resulting curve skeleton will always be 26-connected, if the original object was so [**BNSdB99**]. Below, we give two examples: one synthetic image and one MRA image. The figures for our examples are rendered 2D projections of the 3D objects, i.e., segmented data.

4.1. Synthetic image. The results of skeletonization of a pyramid (base 85×85 and height 122 voxels) can be seen in Figure 2. The curve skeleton consists of the meeting points of all surfaces in the surface skeleton, which results in branches starting in each corner going inwards.

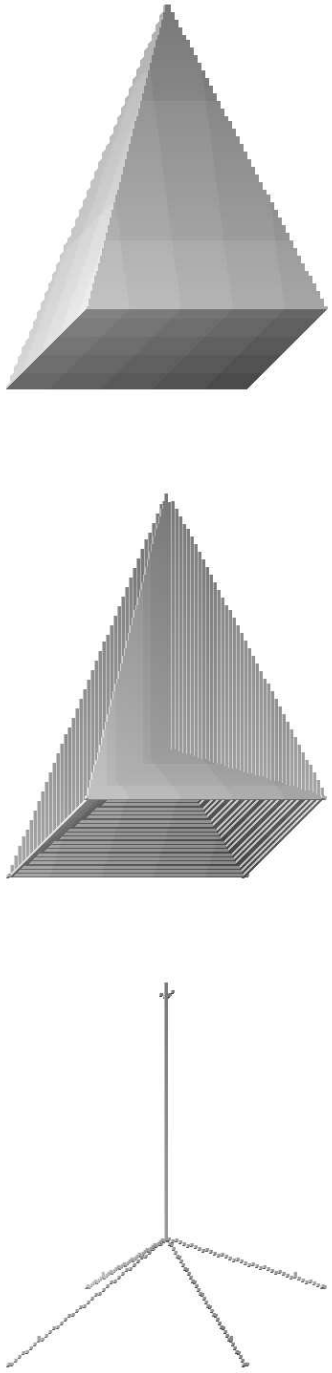

FIGURE 2. A pyramid (top), its surface skeleton (middle), and its curve skeleton (bottom).

4.2. Magnetic resonance angiography image. A $512 \times 512 \times 50$ image of the arteries of the pelvis was obtained with Gadolinium MRA technique in a 1.5 T magnet (Philips Gyroscan ACS-NT). The voxel size was $0.759 \times 0.759 \times 1.799$ mm.

The following image processing methods were performed in the analysis of the blood vessels:

1. Interpolating the image to cubic voxels.
2. Segmentation of the blood vessels from the background:
   (a) simple grey-level thresholding
   (b) connected component labelling
   (c) interactive removal of noise
   (d) morphological smoothing operations
   (e) cavity filling
3. Surface skeletonization.
4. Curve skeletonization.
5. Skeleton pruning.
6. Quantitative and qualitative skeleton analysis.

Computations are simpler if distances between layers are the same for all directions x, y, and z. Therefore, the anisotropy of the image is eliminated by performing a linear interpolation of the image to *cuberille* voxels. There exist other (not as simple) solutions to this problem. In this particular case the interpolated image is of size $512 \times 512 \times 117$ (approximately 30 Mbyte). This will make, among other things, the distance calculations and connectivity aspects a lot easier.

In most cases, segmentation is a difficult task. If a correct segmentation of the blood vessels from the background is not obtained, the results may be completely misleading. In our example, though, this problem was fairly easily solved. Due to the Gd-MRA technique, the contrast between the bright arteries and the darker background was high enough to allow a simple grey-level thresholding. If the segmentation would be more difficult the concept in, e.g., [**US96**] could be used.

However, the bright noise in the image is also thresholded, which is undesirable. To eliminate noise, we applied a 3D connected component labelling, in which a volume threshold discards small objects. Hence, the main (largest) component was selected, i.e., the blood vessels. Finally, a morphological closing operation (three iterations of size $3 \times 3 \times 3$) was performed to reduce artificial irregularities in the vessel border.

Next, the skeletonization algorithm was applied. See Figure 3 (top-left). The surface skeleton and the corresponding curve skeleton are shown for the MRA image in Figure 3, at top-right and bottom-left, respectively.

4.3. Pruning. When we apply the skeletonization to volume images from the real world and not only to synthetic ones, we soon find that the method is quite sensitive to noise in the input data. Typically, the curve skeleton will be rather "hairy", i.e., it will have a large number of short branches. See Figure 3 (bottom-left). Each of these branches corresponds to "unsmoothness" in the original object, either resulting from significant variations or from noise. We remove these in a *pruning* part (analogous to what is done to an apple tree). However, one must be aware of the fact that this is a rather arbitrary decision. It may be that this irregularity of the skeleton reflects an irregularity of the vessel wall, which might be due to atherosclerosis. In other words, we risk to destroy clinically relevant information

FIGURE 3. Rendered 2D projections of arteries of the pelvis (top-left), the surface skeleton (top-right), the curve skeleton (bottom-left), and the pruned curve skeleton (bottom-right).

by making the skeleton too smooth in the "beautifying" process. Consequently, there is a problem in deciding which branches to remove and which to preserve.

The simple pruning method we have used utilizes the distance values of the curve skeleton. The principle is to cut away all short branches by iteratively removing voxels with distance values less than a certain threshold, unless a removal would disconnect the skeleton. See Figure 3 (bottom-right), where threshold value = 5 has been used for the distance labels. A significant improvement is obtained even with this simple pruning. All branches are shortened by the same number of voxels, so that all the short branches disappear. For the long main branches this little shortening is usually of no importance.

After pruning we have an acceptable representation of the arteries and branches in the pelvis. It may be of interest to study the number of voxels in each object. The number of voxels in the pruned curve skeleton is less than 2.5 % of that of the original object (150 000), but still it gives a fairly good impression of the general shape of the arterial tree.

5. Presentation method

A problem with projection images for visualizing 3D angiographic data is that non-circular stenoses are often seen only if the projection is made in a certain direction. For example, a stenosis reducing mainly the anteroposterior diameter of a vessel is difficult to perceive if the projection is made in the frontal plane, simulating anteroposterior rays through the volume. To overcome this limitation, MRA images are often presented as rotation sequences of 2D projections. This approach, however, often implies practical problems in the clinical setting where video display is not always available. One of our goals is to concentrate information about stenoses into one single projection, independent of the viewing direction. Our tool to obtain this is the distance-labelled curve skeleton described above. The skeletal voxels are labelled with their distance to the original background. Each distance value then represents the radius of a ball as large as possible still being a subset of the object. For the surface skeleton, the union of all balls reconstruct the original object. For the curve skeleton, the distance values correspond to the minimum radius of the object at that point, which we will utilize in our presentation method. By surrounding each voxel in the curve skeleton with a ball of radius equal to the distance value, e.g., by applying the reverse DT, the object is displayed with minimum diameter at each point.

The presentation method is first illustrated by a synthetic example. The binary volume image ($256 \times 256 \times 256$ voxels) contains a cylinder of radius 10 voxels and height 240 voxels. Three artificial "stenoses" have been introduced (called upper, middle, and lower). In Figure 4 (left) the object is shown as rendered 2D projections for two different viewing directions. The upper stenosis is visible for the first projection, but not for the second. The middle stenosis should be visible in any projection. The lower stenosis can be seen in the second projection, but is hardly visible in the first (in a MIP it would not be visible). The curve skeleton can be seen in the middle of Figure 4, describing the shape of the original object well. As long as the original surface is a smooth cylinder, the curve skeleton is a simple curve. Where the shape changes skeletal protrusions are introduced. In this application we do not remove any of these branches.

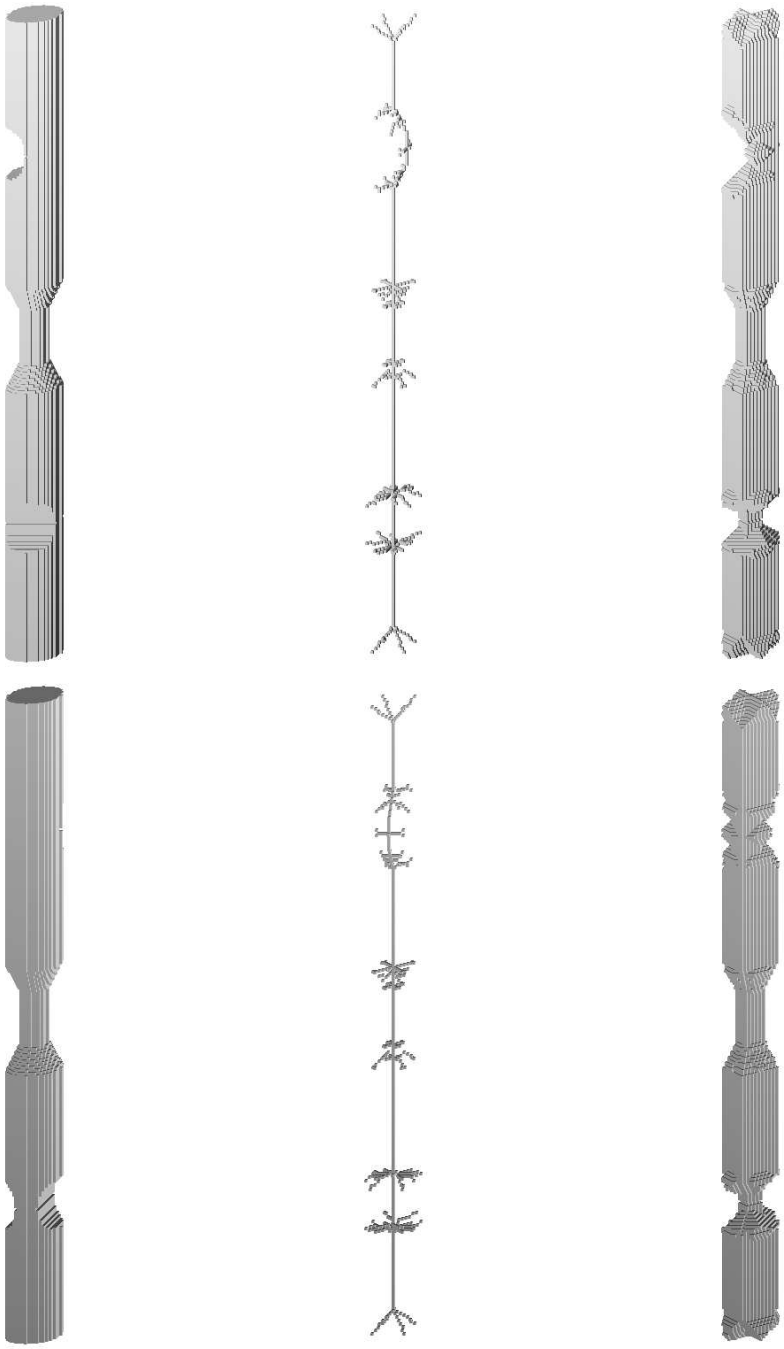

FIGURE 4. A synthetic image with three artificial "stenoses". Left: rendered 2D projections of the object for two different viewing directions. Middle: the corresponding curve skeleton. Right: reconstruction from the curve skeleton, showing that the choice of viewing direction is no longer crucial for detection of the "stenoses."

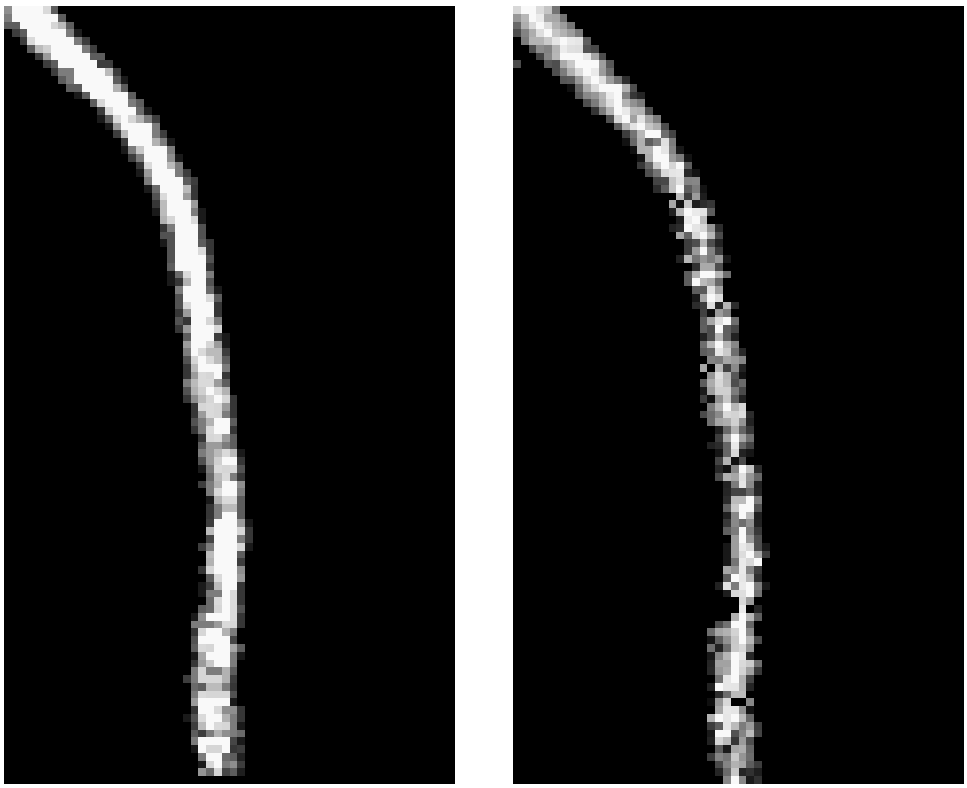

FIGURE 5. Rendered 2D projections of the segmented left external iliac artery before (left) and after (right) curve skeletonization and reverse distance transformation. The lower part of the right image is a possible location for a stenosis, which is not visible in the image to the left.

Finally, the reconstruction is shown to the right in Figure 4. If the reconstruction was performed on the surface skeleton, we would obtain the original object. But as the reverse DT is applied to the curve skeleton, this is not the case. One may notice that the influence of the underlying metric, D^6, is evident in the result; a number of octahedra are added together giving the "diamond shape" of the cross-section. More importantly, all three stenoses are now visible in both projections; the reconstruction uses the minimum distance value to the original background at each voxel.

An example of the application of the proposed presentation method on real MRA data is found in Figure 5. Here, the segmented left external iliac artery is shown before and after the process described above. In this case a segment with reduced minimum radius is found in the lower part of the image only in the processed version, indicating that this region should be studied in more detail in the original data set.

6. Discussion

It seems that the distance-labelled curve skeleton (after pruning) is a meaningful representation of 3D vessel geometry. In addition to the overall tree anatomy, they contain information about the calibre of the vessels, which is the most interesting feature, at least for the diagnosis of atherosclerosis.

We have also described a method to display "stenoses" in 3D images independent of the viewing direction by combining distance transform based skeletonization with the reverse distance transformation. For example, a stenosis reducing the anteroposterior diameter of a craniocaudal vessel, which is normally seen only in the lateral projection, will then become visible also in the frontal projection. It should be emphasized that these rather diagrammatic sketches are not meant to replace traditional rendering methods, but rather to focus the attention of the viewer on certain regions which should be scrutinized with other methods, e.g., viewing the original slices.

One of the advantages of having a distance transform based skeletonization method, is that reversibility of the object might be possible (depending on the method). Reversibility is important when characterizing objects, otherwise much of the shape information is lost. The reversibility property has hardly been considered in the 3D literature. See [**SNB99**]. Other examples of methods based on iterative thinning in distance transforms are [**ST95**] and [**Pud98**]. However, reversibility is not discussed in these references. Obviously, the dependence of the vessel orientation relative to the coordinate system imposed by the D^6 metric is undesirable. An important task for our future research is therefore to develop a skeletonization method based on a less rotation dependent distance transform, more closely approximating Euclidean distance. We are also working to find more intelligent pruning methods which do not shorten main branches.

A serious limitation of this study is that segmented (thresholded) data have been used, since the algorithm has so far been formulated only for binary images. Even though Gd-MRA images are comparatively easy to segment, the difficulty in finding a suitable threshold value must not be underestimated. We plan to combine the presentation approach with the grey-scale connectivity concept developed in a different project for MRA volume visualization [**SSL99**]. Furthermore, any segmented volume image will inevitably contain less information than the grey-scale original. Hence, it is clearly motivated to attempt to transfer the concepts to the grey-scale images, possibly employing some kind of "fuzzy skeleton" built on the concepts from fuzzy digital topology [**Ros79**]. Clearly, this would involve a substantial part of theoretical development.

Acknowledgment

Algorithmic support on skeletonization and distance transformation was given by Prof. Gunilla Borgefors, Dr. Gabriella Sanniti di Baja, and Stina Svensson, which is gratefully acknowledged.

References

[AET93] Charles M. Anderson, Robert R. Edelman, and Patrick A. Turski, *Clinical magnetic resonance angiography*, Lippincott-Raven Publishers, New York, 1993.

[Ber95] Gilles Bertrand, *A parallel thinning algorithm for medial surfaces*, Pattern Recognition Letters **16** (1995), 979–986.

[BNSdB99] Gunilla Borgefors, Ingela Nyström, and Gabriella Sanniti di Baja, *Computing skeletons in three dimensions*, Pattern Recognition **32** (1999), no. 7, 1225–1236.

[Bor84] Gunilla Borgefors, *Distance transformations in arbitrary dimensions*, Computer Vision, Graphics, and Image Processing **27** (1984), 321–345.

[Bor96] Gunilla Borgefors, *On digital distance transforms in three dimensions*, Computer Vision and Image Understanding **64** (1996), no. 3, 368–376.

[DHW+97] C. P. Davis, T. F. Hany, S. Wildermuth, et al., *Postprocessing techniques for gadolinium-enhanced three-dimensional MR angiography*, RadioGraphics **17** (1997), no. 5, 1061–1077.

[GSV96] Yaorong Ge, David R. Stelts, and David J. Vining, *3D skeleton for virtual colonoscopy*, Proceedings of 4^{th} VBC'96: Visualization in Biomedical Computing (Karl Heinz Höhne and Ron Kikinis, eds.), Springer-Verlag, Berlin Heidelberg, 1996, pp. 449–454.

[KR89] T. Yung Kong and Azriel Rosenfeld, *Digital topology: Introduction and survey*, Computer Vision, Graphics, and Image Processing **48** (1989), 357–393.

[LKC94] Ta-Chih Lee, Rangasami L. Kashyap, and Chong-Nam Chu, *Building skeleton models via 3-D medial surface/axis thinning algorithms*, CVGIP: Graphical Models and Image Processing **56** (1994), no. 6, 462–478.

[MBA93] Grégoire Malandain, Gilles Bertrand, and Nicholas Ayache, *Topological segmentation of discrete surfaces*, International Journal of Computer Vision **10** (1993), no. 2, 183–197.

[MS96] Cherng Min Ma and Milan Sonka, *A fully parallel 3D thinning algorithm and its applications*, Computer Vision and Image Understanding **64** (1996), no. 3, 420–433.

[NB95] Ingela Nyström and Gunilla Borgefors, *Synthesising objects and scenes using the reverse distance transformation in 2D and 3D*, Proceedings of ICIAP'95 - Sanremo: Image Analysis and Processing (Carlo Braccini, Leila DeFloriani, and Gianni Vernazza, eds.), Springer-Verlag, Berlin Heidelberg, 1995, pp. 441–446.

[NMR+92] Sandy Napel, Michael P. Marks, Rubin, et al., *CT angiography with spiral CT and maximum intensity projection*, Radiology **185** (1992), 607–610.

[Nys98] Ingela Nyström, *Skeletonization applied to magnetic resonance angiography images*, Medical Imaging 1998: Image Processing (Kenneth M. Hanson, ed.), Proc. SPIE 3338, 1998, pp. 693–701.

[PK99] Kálmán Palágyi and Attila Kuba, *A parallel 3D 12-subiteration thinning algorithm*, Graphical Models and Image Processing **61** (1999), 199–221.

[Pri94] Martin R. Prince, *Gadolinium-enhanced MR aortography*, Radiology **191** (1994), 155–164.

[Pud98] Chris Pudney, *Distance-ordered homotopic thinning: A skeletonization algorithm for 3D digital images*, Computer Vision and Image Understanding **72** (1998), no. 3, 404–413.

[Ros79] Azriel Rosenfeld, *Fuzzy digital topology*, Information and Control **40** (1979), 76–87.

[RP66] Azriel Rosenfeld and John L. Pfaltz, *Sequential operations in digital picture processing*, Journal of the Association for Computing Machinery **13** (1966), no. 4, 471–494.

[SCM97] Punam Kumar Saha, B. B. Chaudhuri, and Dwijesh Dutta Majumder, *A new shape preserving parallel thinning algorithm for 3D digital images*, Pattern Recognition **30** (1997), no. 12, 1939–1955.

[SNB99] Stina Svensson, Ingela Nyström, and Gunilla Borgefors, *Fully reversible skeletonization for volume images based on anchor-points from the D^{26} distance transform*, Proceedings of 11^{th} Scandinavian Conference on Image Analysis, Kangerlussuaq, Greenland, June 1999, pp. 601–608.

[SSL99] Ö Smedby, Stina Svensson, and Tomas Löfstrand, *Greyscale connectivity concept for visualizing MRA and CTA volumes*, Medical Imaging 1999: Image Display (Seong Ki Mun and Yongmin Kim, eds.), Proc. SPIE 3658, 1999, pp. 212–219.

[ST95] Toyofumi Saito and Jun-Ichiro Toriwaki, *A sequential thinning algorithm for three dimensional digital pictures using the Euclidean distance transformation*, Proceedings of 9^{th} Scandinavian Conference on Image Analysis, Uppsala, Sweden, 1995, pp. 507–516.

[US96] Jayaram K. Udupa and Supun Samarasekera, *Fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation*, Graphical Models and Image Processing **58** (1996), no. 3, 246–261.

Centre for Image Analysis, Uppsala University, Uppsala, Sweden
E-mail address: Ingela.Nystrom@cb.uu.se
URL: http://www.cb.uu.se/~ingela

Division of Medical Radiology, University Hospital, Linköping, Sweden
E-mail address: Orjan.Smedby@imv.liu.se