Skeletonization applied to magnetic resonance angiography images

Ingela Nyström

Centre for Image Analysis, Uppsala University
Lägerhyddvägen 17, SE-752 37 Uppsala, Sweden

ABSTRACT

When interpreting and analysing magnetic resonance angiography images, the 3D overall tree structure and the thickness of the blood vessels are of interest. This shape information may be easier to obtain from the skeleton of the blood vessels. Skeletonization of digital volume objects denotes either reduction to a 2D structure consisting of 3D surfaces and curves, or reduction to a 1D structure consisting of 3D curves only. Thin elongated objects, such as blood vessels, are well suited for reduction to curve skeletons. Our results indicate that the tree structure of the vascular system is well represented by the skeleton. Positions for possible artery stenoses may be identified by locating local minima in curve skeletons, where the skeletal voxels are labelled with the distance to the original background.

Keywords: pattern recognition, shape representation, skeleton, digital topology, MR angiography, artery stenosis, diagnosis application

1. INTRODUCTION

Magnetic resonance (MR) angiography is a non-invasive imaging method whereby the MR scanner produces images of the flowing blood in blood vessels in the human body.[1,2] When interpreting and analysing these images, their 3D overall tree structure and the thickness of the blood vessels are of particular interest. This information may be easier to obtain from a processed version of the 3D image in which the *skeleton* of the blood vessels is obtained.

Skeletonization, sometimes referred to as thinning, is a process where objects are reduced to structures of lower dimension while preserving the topology and shape. Skeletonization reduces objects in 2D images to a set of curves. A survey on various 2D thinning and skeletonization methodologies can be found in Ref. 3. In 3D, the result of skeletonization is either a set of 3D surfaces and curves,[4–6] or, if reversibility is not required, a set of 3D curves only.[7–10] Reversibility is important when processing and characterizing objects, otherwise much of the shape information is lost, but reversibility is less important when the objects are band-like with known, rather constant width, such as blood vessels.

Earlier work with the aim to obtain curve skeletons for applications can be found in Ref. 11, where objects are reduced directly to 3D curves using four classes, each with a number of deleting templates, and in Ref. 12, where the skeleton concept from Ref. 8 is adopted; the Euler characteristics are stored in look-up tables and used to identify simple points in a six directional sub-iterations algorithm. The algorithms are demonstrated on visual analysis of computed tomography (CT) lung data, and on interactive virtual colonoscopy rendered from spiral CT scans, respectively. In the literature not many realistically sized and/or real images are used for curve skeletonization, which makes it difficult to assess the practical utility of the method.

The method used for reducing objects in binary volume images to skeletons in this paper is described in detail in Refs. 13-14. The first skeletonization step reduces the object to a surface skeleton (Sect. 2.1) from which the original object can be recovered. The second step reduces the surface skeleton further to a curve skeleton (Sect. 2.2). To obtain "cleaner" skeletons an optional pruning step can follow the skeletonization (Sect. 2.3). In addition to the MR angiography data, the method is illustrated by a synthetic image. The image processing steps for the MR angiography application are described in Sect. 3.2. CT angiography images could have been used as well.[15] The method presented here can find applications in other fields, where volume images are obtained from different (tomographic) devices.

Other author information -
E-mail: ingela@cb.uu.se
WWW: http://www.cb.uu.se/

Part of the SPIE Conference on Image Processing • San Diego, California • February 1998
SPIE Vol. 3338 • 0277-786X/98/$10.00

2. SKELETONIZATION METHOD

In a binary volume image the object is reduced to a skeleton in two major steps. The topology (i. e. number of components, tunnels, and cavities), and shape of the object are preserved. The skeletal voxels are labelled with their distance to the original background, which might be useful information in quantitative analysis of the shape. Most often pruning and "beautifying" of the curve skeletons is necessary to remove, or at least shorten, skeletal branches caused by the noise (and rotation) dependency of the skeletonization method.

2.1. Surface Skeletonization

Each voxel has three types of neighbours among its 26 closest neighbours; 6 face-, 12 edge-, and 8 point-neighbours, that share a face, an edge, and a point with the voxel, respectively.

An object component is defined as a 26-connected set of voxels. As a consequence of the 26-connectedness selected for the object, 6-connectedness must be used for the background.

A *border* voxel is an object voxel having at least one face-neighbour in the background. Object voxels, which are not border voxels, are *internal* voxels.

The surface skeletonization is modified and improved with respect to Ref. 16. The surface skeletons are computed in two phases. During the first phase (I), removable voxels are identified and iteratively removed until an at most two voxel thick surface of skeletal voxels is identified:

I-1. Identify border voxels (assign with the iteration number). Also, identify voxels with an edge-neighbour in the background (assign with the iteration number+1).

I-2. Among border voxels, mark *multiple* voxels.

I-3. Among border voxels, mark as tunnel voxels those whose removal would create a tunnel.

I-4. Remove all unmarked border voxels.

A border voxel of the current iteration is defined as multiple if any of the following conditions is satisfied:

A1: The configuration to the left in Fig. 1 does *not* occur in any of the three principal planes.

A2: The configuration in the middle of Fig. 1, or any of its rotations, occurs.

A3: The configuration to the right in Fig. 1, or any of its rotations, occurs.

Condition A1 guarantees that protrusion voxels including the tip (end-point) are not removed. All three conditions prevent disconnecting the object. Condition A3 is unique to the 3D case and is necessary as the object is defined as 26-connected.

Hence, the configuration to the left in Fig. 1 is a necessary, but not sufficient, condition for removal of a voxel. The other two configurations must not occur, and a tunnel through the object must not be created, then a voxel can be removed.

At this stage, the skeletal voxels are labelled with the iteration number they were assigned when identified as border voxels. This coincides with the labels they would have on the D^6 distance transform of the object, i. e. the volume equivalent of the city-block distance transform in 2D.[17]

During the second phase (II), this set is reduced to unit-wide surfaces (and curves) in a thinning process that has to be split into *six* (the number of face-neighbours of a voxel) directional processes, each of them applied once. To prevent shortening of the object, the set is thinned only in directions where it is *exactly* two voxels thick (Fig. 2), unless that would disconnect the object.

The original object can still be recovered, except for a few border voxels of the original object, using the *reverse* D^6 distance transformation on the distance labels.[18]

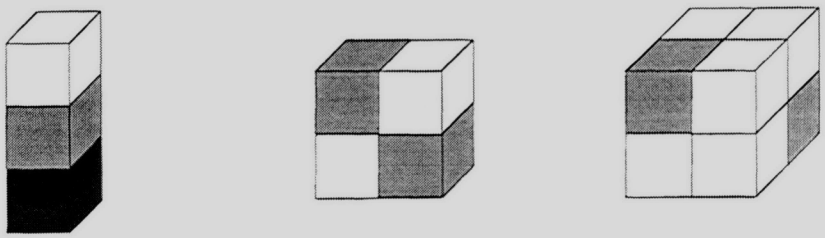

Figure 1. The voxel configurations in conditions A1–A3. Internal voxels are black, border voxels are grey, and background voxels are white. Rotations of the configurations result in six cases for A1 (left), twelve cases for A2 (middle), and eight cases for A3 (right).

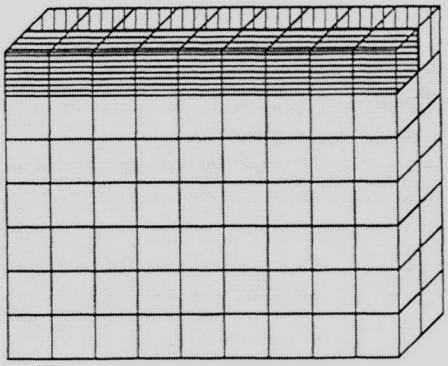

Figure 2. A voxel set, two voxels thick in the Front-Back direction. The Top-voxels (hatched) are not removed in the Top-Down process, but the horizontally hatched Top-voxels are also Front-voxels and will be removed during the Front-Back process.

2.2. Curve Skeletonization

The surface skeleton can be further reduced to a curve skeleton in two iterative phases. After this reduction, the original object can no longer be recovered, but its topology and shape are preserved. The skeletal voxels are labelled with their distance to the original background, which might be useful information in quantitative analysis of the shape.

During the first phase (III), voxels are iteratively removed from the surfaces:

III-1. Identify *outer* voxels of the surface (whose removal would not create tunnels).

III-2. Mark outer voxels as removable, if they are neither a *protrusion* nor a *break-point* voxel. See Ref. 13 for details.

III-3. Sequentially remove marked voxels, unless that would disconnect the object.

In case the original surface is an even number of voxels broad, the skeletal set becomes two voxels wide. During the second phase (IV), it is reduced to unit thickness by iteratively applying a final thinning process, until no more voxels can be removed:

IV-1. Classify *tip-of-protrusion* and break-point voxels.

IV-2. Among non-classified voxels, mark outer voxels.

IV-3. Sequentially remove the outer voxels, unless that would disconnect the object.

The result is a 26-connected unit-wide curve skeleton. The voxels of our curve skeletons are labelled with the D^6 distance to the original background. This is useful information for quantitative analysis of the shape, even though the original object no longer can be recovered.

2.3. Skeleton Pruning

Our curve skeletons need to be pruned and "beautified", in all but very simple cases. Real images will result in "hairy" curve skeletons. All essential parts of the skeleton are present, so the task is to remove branches resulting from weak protrusions and noise.

The skeletons are labelled with the distance to the original background. The following very simple and "brute-force" pruning method uses this information. The shortest skeletal branches, labelled with small distances, probably result from either noise or weak protrusions of the object. They can be pruned in iterated scans. Every skeletal voxel with a distance label less than a threshold is removed, unless its removal would disconnect the skeleton. Imagine a skeletal branch along the scan direction. During the first scan, the skeletal voxels are removed one by one without disconnecting the skeleton, as long as their labels are smaller than the threshold value. But imagine a skeletal branch directed away from the scan direction. Even though the labels are less than the threshold value, voxels can not be removed until the outermost one is reached, as removal would disconnect the skeleton. For the skeleton to be equally pruned from every direction, the number of iterations should be the same as the threshold value.

Even this simple pruning step results in a significant improvement. One drawback, though, is that the main branches are shortened as many voxels as the small branches are. For thin elongated objects this is a relatively minor problem, but it should be kept in mind.

3. SKELETONIZATION APPLIED

The proposed skeletonization method is illustrated by a synthetic example before we show its performance on real MR angiography data. The figures for these examples are rendered 2D projections of the thin 3D skeletons, which is difficult to interpret. A better way to visualise this kind of images is a rotation sequence of projections. The reader can be assured that all skeletons *are* connected and without tunnels, even though the figures might be misleading in this respect.

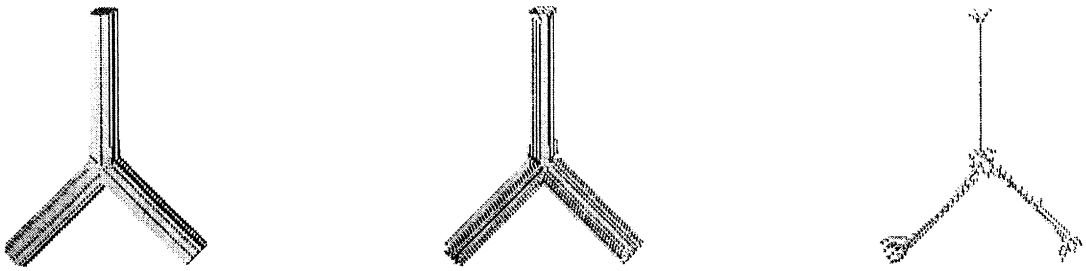

Figure 3. Three fused Euclidean cylinders of radius 5 voxels and height 60 voxels (left), its surface skeleton of labelled voxels (middle), and curve skeleton (right).

3.1. Synthetic Example

One example from the "blocks world" is synthesised in an $128 \times 128 \times 128$ image (2 Mbyte), a binary volume image with three fused solid cylinders. The original object of size 12872 voxels is shown in Fig. 3 to the left. The surface skeleton consisting of 24.4 % of the original voxels can be seen in the middle. The original object can be recovered, except for some (20) original border voxels, as the skeleton includes (almost) all information present in the original image. The curve skeleton can be seen to the right, describing the shape of the original object well with only 4.6 % of the original voxels.

3.2. Magnetic Resonance Angiography Example

MR angiography is increasingly performed as a non-invasive method of evaluating patients with suspected vascular disease.[1] In this study we have used images obtained after intravenous injection of paramagnetic contrast material, i. e. gadolinium chelates, for arterial enhancement.[2]

A $512 \times 512 \times 50$ image of the arteries of the pelvis was obtained with a superconducting 1.5 T magnet (Philips Gyroscan ACS-NT). The voxel size is $0.759 \times 0.759 \times 1.799$ mm.

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

1. Resampling the image to cubic voxels.
2. Segmentation of the blood vessels from the background:
   (a) 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 images should be resampled to *cuberille* voxels. In this particular case the resampled image is of size $512 \times 512 \times 117$ (approximately 30 Mbyte).

In most cases segmentation is a difficult task. It is *very* important to obtain a correct segmentation of the blood vessels from the background. The gadolinium-enhanced images are fairly easy to segment, though. The contrast between the light arteries and the darker background is high enough to allow a simple grey-level thresholding.

The light noise in the image is also thresholded, which is undesirable. To remove the noise, a 3D connected component labelling in which a volume threshold discards small objects could be performed.[19] Hence, the main (largest) component is selected, i. e. the blood vessels. An alternative solution, could be to interactively select a voxel that surely belongs to a blood vessel. From this start voxel, voxels are added to the component recursively if they are 26-connected to a voxel already belonging to the component, an *auto-connection*.

At this stage, some noise-components could be connected to the blood vessels. Interactive editing to erase the largest ones would result in a better segmentation.

By combining three dilations followed by three erosions (the number is image dependent), i. e. the morphological operation *closing*, the border of the object is smoothed. The blood vessels resulting after this processing is shown in Fig. 4, top-left.

Most probably, there are no cavities in the blood vessels, unless an artifact in the image acquistion or a failure of the segmentation. If there are cavities, they need to be filled before the skeletonization is performed, though. One way to fill cavities, is to perform a 3D connected component labelling of the background, skipping the component on the border of the image, and logically OR-combine the rest with the blood vessels.

Thereafter, the blood vessels are prepared for skeletonization. The surface skeleton (Sect. 2.1), and the corresponding curve skeleton (Sect. 2.2) are shown for the MR angiography image in Fig. 4, top-right and bottom-left, respectively.

The "hairy" curve skeleton needs to be pruned. In this pilot study the pruning method described in Sect. 2.3 was used with threshold value = 5 for the distance labels. A significant improvement was obtained even with this simple pruning, as can be seen in Fig. 4, bottom-right.

Even though it is not possible to recover the original object from the curve skeleton its distance labels are useful. Local minima in the skeleton indicate narrowing of the original object, which may identify possible artery stenoses.

4. COMPUTATION COMPLEXITY

The implementations of the described methods are written in C and integrated in IMP (*IMage Processing*),[20] which is a general image analysis software for five-dimensional (x, y, z, t, b) image data developed at the Centre for Image Analysis. The five dimensions are the spatial dimensions (x, y, z), the time dimension t, and the bandwidth (spectral) dimension b. The software runs on any standard UNIX workstation using X-Windows and Motif as user interfaces.

The skeletonization takes on average a few minutes for different objects in images of 2 Mbyte size. This CPU time is on an ordinary (DEC Alpha) workstation.

For the 30 Mbyte MR angiography image the computation times for the various skeletonization steps are listed in Table 1 together with the number of voxels remaining after each step. Despite the huge amount of data to scan many times, the whole process takes less than 15 minutes.

5. DISCUSSION

When describing the anatomy of blood vessels in MR angiography images, the curve skeleton of the blood vessels promises to be useful. To ensure good skeletonization results, it is important with a perfect segmentation method, though. This is possible for gadolinium-enhanced MR angiography images of high contrast. Computed tomography (CT) angiography images could be used for the same purposes.

The topology and shape of the original object are preserved after the skeletonization, based on a small number of simple local neighbourhood operations, which make it fairly time efficient. We observe, though, that it is not invariant under object rotation, as the digital distance metric implicit in the process is the D^6 metric, which is quite

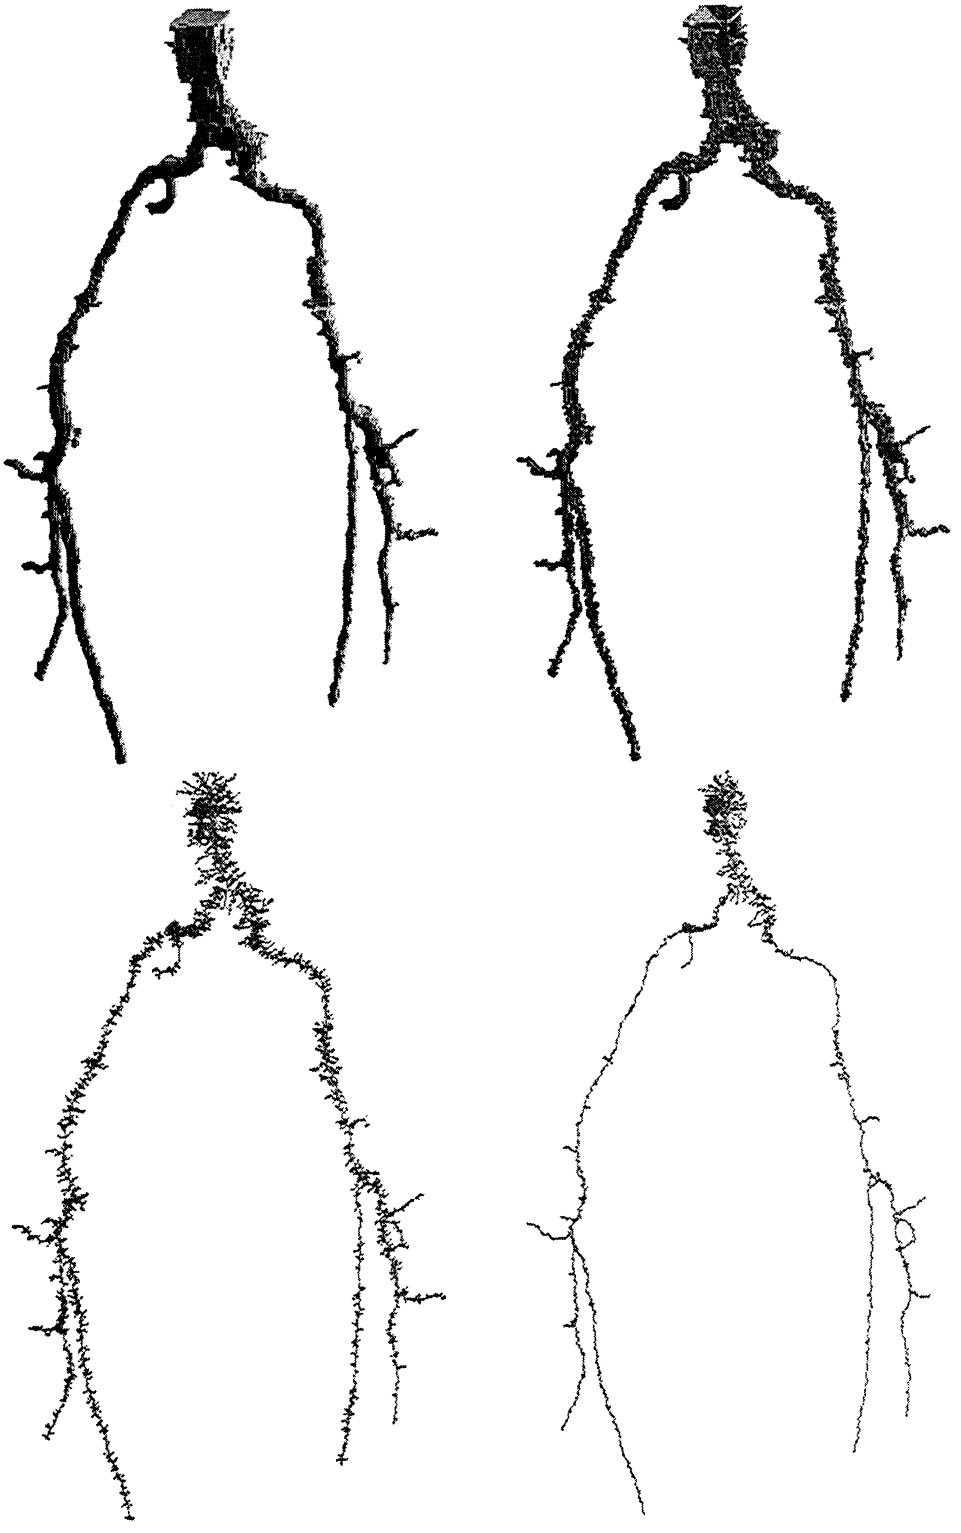

Figure 4. Rendered 2D projections of arteries of the pelvis (top-left), the surface skeleton (top-right), the curve skeleton (bottom-left), and the pruned curve skeleton (bottom-right).

Table 1. The CPU time for computing skeletons of the MR angiography image in Fig. 4 on a DEC Alpha workstation. The original number of object voxels are approximately 150 000.

Method	CPU time	Voxels	%
Surface skeleton	222 seconds	35131	24.3
Curve skeleton	621 seconds	8399	5.8
Pruned skeleton	29 seconds	3493	2.4
Total time	**< 15 minutes**		

rotation dependent, as it is an inaccurate approximation of the Euclidean metric. A solution to this problem is to rotate the object to a normalised orientation in a preprocessing step.[21] Similar objects will get the same orientation, and hence the rotation variant property can (almost) be disregarded. Another weakness is that the skeleton is sensitive to noise, as every little protrusion will generate a skeletal branch. Preprocessing in the form of smoothing morphological (closing or opening) operations will alleviate the noise problem. Future work will be directed towards achieving more rotation, and also noise, independent skeletons.

The curve skeleton needs to be pruned and "beautified" in a more shape preserving way than described in this paper. The task is to remove peripheral branches and to straighten zig-zag parts of the skeleton. To do this, some of the methods working in the 2D case,[22] could be generalised, but probably extra care has to be taken for the 3D case.

The skeletal voxels are labelled with distance information from the original background, which is quantitatively useful. In the results from this pilot study it seems that the tree structure of the skeleton is a good representation of the anatomy of the blood vessels. This will be validated in larger studies. Hopefully, positions for possible artery stenoses may be identified by locating local distance minima in the labelled skeleton.

ACKNOWLEDGMENTS

The author wishes to thank Dr. Örjan Smedby at the Dept. of Diagnostic Radiology, Uppsala University, for supplying the MR angiography image on which the skeletonization was tested, and also for fruitful discussions on the topic. Algorithmic support were given by Prof. Ewert Bengtsson, Prof. Gunilla Borgefors, Dr. Bo Nordin, and Dr. Gabriella Sanniti di Baja, which is gratefully acknowledged. The financial support of the Swedish Research Council for Engineering Science (TFR), grant number 95-182, is also acknowledged.

REFERENCES

1. C. M. Anderson, R. R. Edelman, and P. A. Turski, *Clinical Magnetic Resonance Angiography*, Lippincott-Raven Publishers, New York, 1993.
2. M. R. Prince, "Gadolinium-enhanced MR aortography," *Radiology* **191**, pp. 155–164, 1994.
3. L. Lam, S.-W. Lee, and C. Y. Suen, "Thinning methodologies - A comprehensive survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**, pp. 869–885, Sept. 1992.
4. G. Bertrand, "A parallel thinning algorithm for medial surfaces," *Pattern Recognition Letters* **16**, pp. 979–986, 1995.
5. S. Miguet and V. Marion-Poty, "A new 2-D and 3-D thinning algorithm based on successive border generations," in *Proceedings of* 4^{th} *Conference on Discrete Geometry in Computer Imagery, Grenoble, France*, pp. 195–206, 1994.
6. P. K. Saha and B. B. Chaudhuri, "Detection of 3-D simple points for topology preserving transformations with application to thinning," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **16**, pp. 1028–1032, Oct. 1994.
7. G. Bertrand and Z. Aktouf, "A three-dimensional thinning algorithm using subfields," in *Vision Geometry III*, R. A. Melter and A. Y. Wu, eds., pp. 113–124, Proc. SPIE 2356, 1994.
8. T.-C. Lee, R. L. Kashyap, and C.-N. Chu, "Building skeleton models via 3-D medial surface/axis thinning algorithms," *CVGIP: Graphical Models and Image Processing* **56**(6), pp. 462–478, 1994.

9. S. Lobregt, P. W. Verbeek, and F. C. A. Groen, "Three-dimensional skeletonization: Principle and algorithm," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-2**, pp. 75–77, Jan. 1980.
10. T. Saito and J.-I. Toriwaki, "A sequential thinning algorithm for three dimensional digital pictures using the Euclidean distance transformation," in *Proceedings of* 9^{th} *Scandinavian Conference on Image Analysis, Uppsala, Sweden*, pp. 507–516, 1995.
11. C. M. Ma and M. Sonka, "A fully parallel 3D thinning algorithm and its applications," *Computer Vision and Image Understanding* **64**, pp. 420–433, Nov. 1996.
12. Y. Ge, D. R. Stelts, and D. J. Vining, "3D skeleton for virtual colonoscopy," in *Proceedings of* 4^{th} *VBC'96: Visualization in Biomedical Computing*, K. H. Höhne and R. Kikinis, eds., pp. 449–454, Springer-Verlag, Berlin Heidelberg, 1996.
13. G. Borgefors, I. Nyström, and G. Sanniti di Baja, "Computing skeletons in three dimensions." Submitted for publication.
14. I. Nyström, *On Quantitative Shape Analysis of Digital Volume Images.* PhD dissertation No. 288, Uppsala University, Uppsala, Sweden, May 1997. Available from the author.
15. S. Napel, M. P. Marks, G. D. Rubin, M. D. Dake, C. H. McDonnell, S. M. Song, D. R. Enzmann, and R. J. Jeffrey, "CT angiography with spiral CT and maximum intensity projection," *Radiology* **185**, pp. 607–610, 1992.
16. G. Borgefors, I. Nyström, and G. Sanniti di Baja, "Surface skeletonization of volume objects," in *Proceedings of SSPR'96: Advances in Structural and Syntactical Pattern Recognition*, P. Perner, P. Wang, and A. Rosenfeld, eds., pp. 251–259, Springer-Verlag, Berlin Heidelberg, 1996.
17. G. Borgefors, "On digital distance transforms in three dimensions," *Computer Vision and Image Understanding* **64**, pp. 368–376, Nov. 1996.
18. I. Nyström and G. Borgefors, "Synthesising objects and scenes using the reverse distance transformation in 2D and 3D," in *Proceedings of ICIAP'95: Image Analysis and Processing*, C. Braccini, L. DeFloriani, and G. Vernazza, eds., pp. 441–446, Springer-Verlag, Berlin Heidelberg, 1995.
19. L. Thurfjell, E. Bengtsson, and B. Nordin, "A new three-dimensional connected components labeling algorithm with simultaneous object feature extraction capability," *CVGIP: Graphical Models and Image Processing* **54**, pp. 357–364, July 1992.
20. B. Nordin, "IPAD, version 2.0 & IMP - an IPAD application," internal report No. 6, Centre for Image Analysis, Uppsala, Sweden, 1997. Available from the author.
21. I. Nyström, E. Bengtsson, B. Nordin, and G. Borgefors, "Quantitative analysis of volume images - Electron microscopic tomography of HIV," in *Medical Imaging 1994: Image Processing*, M. H. Loew, ed., pp. 296–303, Proc. SPIE 2167, 1994.
22. G. Sanniti di Baja, "Well-shaped, stable, and reversible skeletons from the (3,4)-distance transform," *Journal of Visual Communication and Image Representation* **5**, pp. 107–115, Mar. 1994.