Quantitative analysis of volume images - electron microscopic tomography of HIV

Ingela Nyström, Ewert Bengtsson, Bo Nordin

Centre for Image Analysis
Uppsala University, Sweden

Gunilla Borgefors

Centre for Image Analysis
Swedish University of Agricultural Sciences, Sweden

ABSTRACT

Three-dimensional (3D) objects should be represented by 3D "images". So far, most of the evaluation of images of 3D objects have been done visually, either by looking at slices through the volumes or by looking at 3D graphic representations of the data. In many applications a more quantitative evaluation would be valuable. Our application is the analysis of volume images of the causative agent of the acquired immune deficiency syndrome (AIDS), namely human immunodeficiency virus (HIV), produced by electron microscopic tomography (EMT). A structural analysis of the virus is of importance. The representation of some of the interesting structural features will depend on the orientation and the position of the object relative to the digitization grid. We describe a method of defining orientation and position of objects based on the moment of inertia of the objects in the volume image. In addition to a direct quantification of the 3D object a quantitative description of the convex deficiency may provide valuable information about the geometrical properties. The convex deficiency is the volume object subtracted from its convex hull. We describe an algorithm for creating an enclosing polyhedron approximating the convex hull of an arbitrarily shaped object.

1 INTRODUCTION

During the last decade the increasing computing power and storing capacity have made the representation of 3D objects possible. As a result, 3D volume images are becoming increasingly common, particularly in many biomedical applications. So far, most of the evaluation of such images has been done visually, either by looking at slices through the volumes or by looking at rendered 3D graphics representations of the data. In some applications more quantitative comparisons between different 3D objects are of interest. The features to extract from the objects may be of structural type, such as size, shape, placement, and density distribution. Some of these features will depend on the orientation and position of the object relative to the digitization grid. But these aspects can not be controlled when the images are obtained. In order to extract these features we therefore need a method of defining position and orientation of the objects. This paper presents such a method, based on the moment of inertia of the objects in the volume image. Also presented is a method of describing the shape of an object. Both methods are applicable to a wide range of biomedical problems and to volume data from many kinds of tomographic devices, e. g. CT, MRI, PET, EM and CM. The application presented in this paper is volume images produced by the HIV Structure Group at the Dept. of Biochemistry, Uppsala University. The algorithms are used to assist in the structural analysis of the causative virus of the AIDS disease, HIV. Especially interesting is a study of the internal morphology of the virus at different stages and a comparison of infectious viruses to non-infectious ones.

2 THE APPLICATION

The HIV has the geometry of a sphere with a diameter of approximately 150 nm. The virus' genetic material, RNA, and enzymes are contained within a protein core, surrounded by a lipid envelope. Assembly of infectious virus particles is initiated upon budding of the virus from the host cell membrane. After virus release the core proteins are assembled along the interior of the viral envelope. Upon cleavage of the proteins they are transported to the centre of the virus, accompanied by the condensation of a core structure. The degree of compactness of the core is correlated to the maturation stage of the virus. In late stages the core has developed into a 100 nm cone shape. The core is connected to the viral envelope by a core-envelope-link (CEL) through the narrow part of the core. The CEL seems to be necessary for stabilizing the core and protecting the genetic material, and could therefore be a suitable target for attack.[1]

The transmission EMT results in a series of two-dimensional (2D) electron micrograph projections from which a 3D reconstruction of an individual virus can be obtained in the following way: Specimens with HIV infected cells are cut into 60–80 nm thick sections to allow analysis of the entire core structure. Approximately 30 micrograph projections are produced in a tilt series taken from $-60°$ to $+60°$. Each micrograph is digitized and thereafter the images are aligned, using the coordinates of applied 10 nm gold particles as reference points. The reconstruction of the density distribution is carried out as a series of 2D reconstructions, each a combination of the radius-weighted Fourier transform of one pixel line of each micrograph. A reconstruction can be seen in Figure 1.

If the section has a thickness of 80 nm a resolution of approximately 5 nm is attainable. Sections thicker than 80 nm will cause a loss in resolution of the volume image, so at present we are unable to expose the whole virus, even though it would be interesting to draw conclusions from the relative information of core and whole virus. On the other hand, it is hard to get better resolution than we have. A section can not be less than 60 nm, as the electron beam then will destroy it. The 3D density map obtained is low-pass filtered to proper signal-to-noise ratio through Fourier filtering. The volume data obtained is normally a cube with $200 \times 200 \times 80$ voxels (more than 3 Mbyte).

3 THE NORMALIZATION OF ORIENTATION

3D objects are described in the 3D density map through some segmentation operation, in the simplest case density thresholding. Once a 3D object has been defined, parameters describing its shape can be extracted. Most parameters vary with the object orientation relative to the digitization grid. This is due to digitization effects and to non-perfect algorithms. In order to compare structures in different volume images the structures must be defined in position and orientation. Described here is a procedure to resample volume data. Each structure is normalized to a coordinate system, where the x-axis is the axis with minimal moment of inertia, the y-axis is the axis orthogonal to the x-axis with the minimal moment, and the z-axis is the axis orthogonal to these two axes.[2] In this way, similar structures will always get the same normalized orientation. Hence, a starting position for several 3D methods for quantitatively evaluating and comparing volume images is achieved. These methods are applied to the comparison of different virus particles in EMT images of HIV.

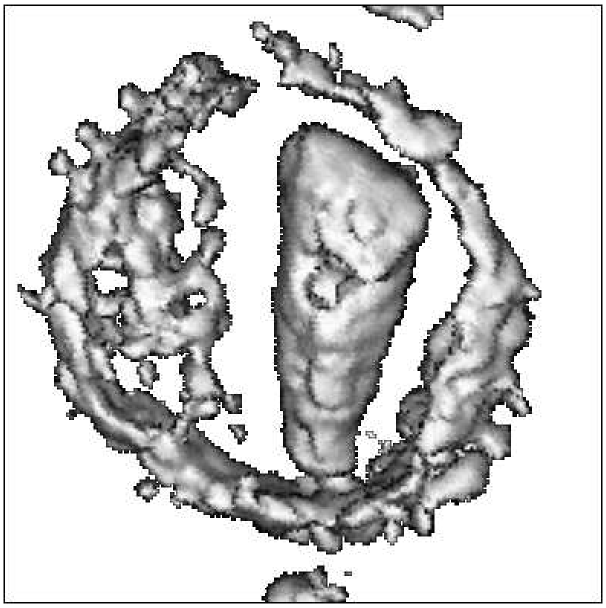

Figure 1: Reconstruction of an infectious HIV particle. The core-envelope-link (CEL) can be seen to connect the core and the envelope.

Our 3D normalizing method can be divided into three phases as outlined below:

A. Calculation of centroid and covariance matrix

In order to create a frame of reference for object normalization, the centroid of the object and moments around the centroid have to be computed.[3] Assume that the object to be analyzed has been segmented from the image and is defined as a single (connected) 3D binary structure. The method will mathematically work also with several disjoint objects, but the resulting normalization will then usually be of little value. With the object defined as a binary structure the mass (gray-level in this case) of the object has no relevance, only the position of each voxel within the object. The moments around the centroid form a covariance matrix, S, according to the following equation:

$$S = \begin{bmatrix} S_{xx} & S_{xy} & S_{xz} \\ S_{yx} & S_{yy} & S_{yz} \\ S_{zx} & S_{zy} & S_{zz} \end{bmatrix}, \tag{1}$$

where

$$\begin{aligned} S_{xx} &= \frac{1}{n}\left(\sum x^2 + \frac{\left(\sum x\right)^2}{n}\right) \text{ and} \\ S_{yz} &= S_{zy} = \frac{1}{n}\left(\sum yz + \frac{\sum y \sum z}{n}\right), \end{aligned}$$

where x, y and z denote the positions for the n voxels belonging to the object in the image. The remaining terms are evaluated using the same schemes.

Alternatively, we can work directly on the gray-scale volume. One way of defining the moments of inertia is then to use a gray-scale weighted mass of each voxel in the original volume, without thresholding. This will make the normalization method less dependent on segmentation problems, but require the object to be the dominant structure in the image. The terms in the covariance matrix (Equation 1) will then change to

$$\begin{aligned} S_{xx} &= \frac{1}{\sum p}\left(\sum px^2 + \frac{(\sum px)^2}{\sum p}\right) \text{ and} \\ S_{yz} &= S_{zy} = \frac{1}{\sum p}\left(\sum pyz + \frac{\sum py \sum pz}{\sum p}\right), \end{aligned}$$

where p denotes the gray-value at position (x, y, z).

A compromise between the two methods can be achieved by weighting the gray-scale values in a non-linear way based on some application or preliminary segmentation data.

B. Principal component analysis (PCA)

PCA[4] is a way of explaining the variance-covariance structure in a data set through a linear combination of the original variables. Geometrically, it is a selection of a new coordinate system obtained by rotating the original system where the original variables are the coordinate axes. For any structure it is always possible to find a set of three orthogonal axes, such that the new axes represent the directions with maximum variability. The principal components are computed from the eigenvalue-eigenvector pairs of the covariance matrix and form a transform matrix with the eigenvectors as columns in the order of corresponding descending eigenvalues. By performing a PCA on our covariance matrix we can thus define the normalization transformation.

C. Rotation and translation

Once a unique object oriented frame of reference, a new coordinate system, is obtained, the 3D image is resampled to put the coordinate axes along the reference dimensions. This requires a 3D resampling algorithm that handles the problems of aliasing and interpolation properly. The following steps have to be applied to the image: First, translate the centroid to the origin. This new system is rotated according to the transform matrix computed above and then translated to the middle point of the image. In practice, however, these steps have to be done in reverse, as the voxel positions after transformation in the discrete world hardly ever coincide with the integer positions in the result image. It is thus necessary to scan through the result image and compute the corresponding, possibly non-integer, position in the original image, $(\hat{x}, \hat{y}, \hat{z})$. In this way each voxel gets a defined gray-value from the original image, according to some interpolation. Applying the inverse of the transform matrix is straightforward as it is the transpose, if the lengths of the eigenvectors are normalized to 1.

The simplest scheme for the interpolation is based on a nearest-neighbour technique, also called zero-order interpolation. The gray-value in the closest integer position to $(\hat{x}, \hat{y}, \hat{z})$ is selected. However, to get results with less aliasing artifacts, a trilinear interpolation technique can be applied. The trilinear interpolation used here is a generalization of a bilinear interpolation technique that uses the gray-values of each of the four nearest integral neighbours. As the name implies, the trilinear interpolation is performed in all three directions. A non-integer position in a volume, $(\hat{x}, \hat{y}, \hat{z})$, has eight integral neighbours, and are thereby surrounded by six planes, two in each direction. The gray-values of the eight neighbours are weighted according to the orthogonal distance to each of the six planes from $(\hat{x}, \hat{y}, \hat{z})$. The weight is inversely proportional to the distance. The interpolated gray-value is assigned to the position in the result image.

Rotating a binary object causes some problems. If we use zero-order interpolation we may loose fine details and have narrow parts break up because of aliasing problems. With trilinear interpolation, the resulting image will not be binary, but has to be thresholded again.

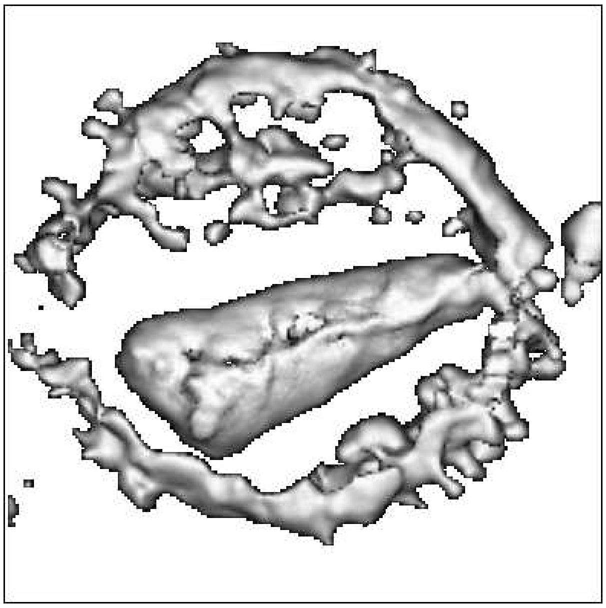

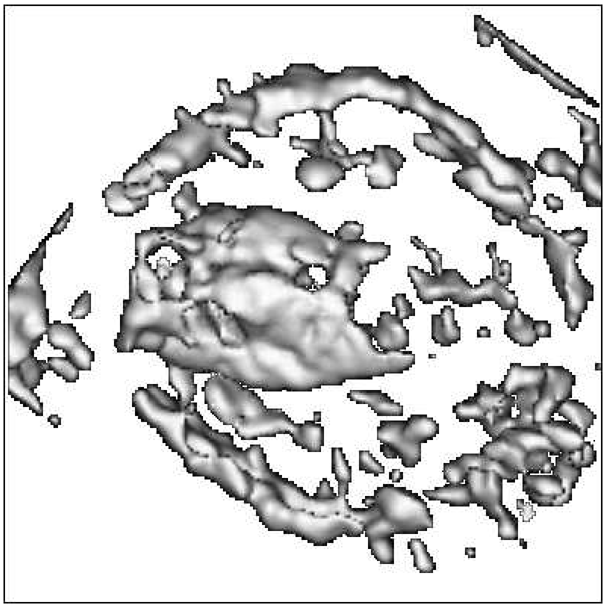

Figure 2: Normalized reconstructions of infectious HIV particles. The virus to the left is a "nice" example. To the right is an example of a virus which is more difficult to threshold and normalize correctly.

After normalizing the objects it is straightforward to visually make a comparison. See Figure 2.

Computationally, we will start from this normalized position and continue by developing methods for describing size, various shape measurements, and binding angle of the CEL. From this position it will also be possible to identify corresponding points in different objects through automatic or interactive correlation methods. These methods will also be useful in other applications as the need for better general methods of quantitatively evaluating and comparing volume images has increased as tomography and other ways of generating volume images are used in more and more fields.

4 FILLING CONCAVITIES

One way of quantifying the structure of an object is to calculate geometrical properties of the convex hull of the object and of its concavity regions. A covering polyhedron, see Figure 3, which is an approximation of the convex hull, can be computed by iterative, local operations. The concavity regions, or convex deficiency, can then be easily computed by subtracting the original object from the convex hull.

The covering polyhedron is computed using an iterative process, where each iteration consists of two passes through the image volume. The algorithm is inspired by a 2D concavity filling algorithm.[5] In the first pass, all voxels in the background that have at least one neighbour in the object are labelled with a vector consisting of the numbers of neighbours in the object in the $x-$, $y-$, and $z-$planes, respectively. In the second pass, each labelled voxel is inspected. If one of the three labels in the vector is at least four, then the voxel is located in a local concavity, and is changed to become an object voxel. In this way, we use the local curvature in a $3 \times 3 \times 3$ neighbourhood. The process is repeated until no further voxels are changed. The resulting covering polyhedron is equal to the smallest enclosing rhombicubeoctahedron (one of the Archimedian solids), which means that new, spurious, external corners are created. This can be visualized by the fact that in each coordinate plane the covering polygon is the closest enclosing octagon. However, for many applications this approximation of the

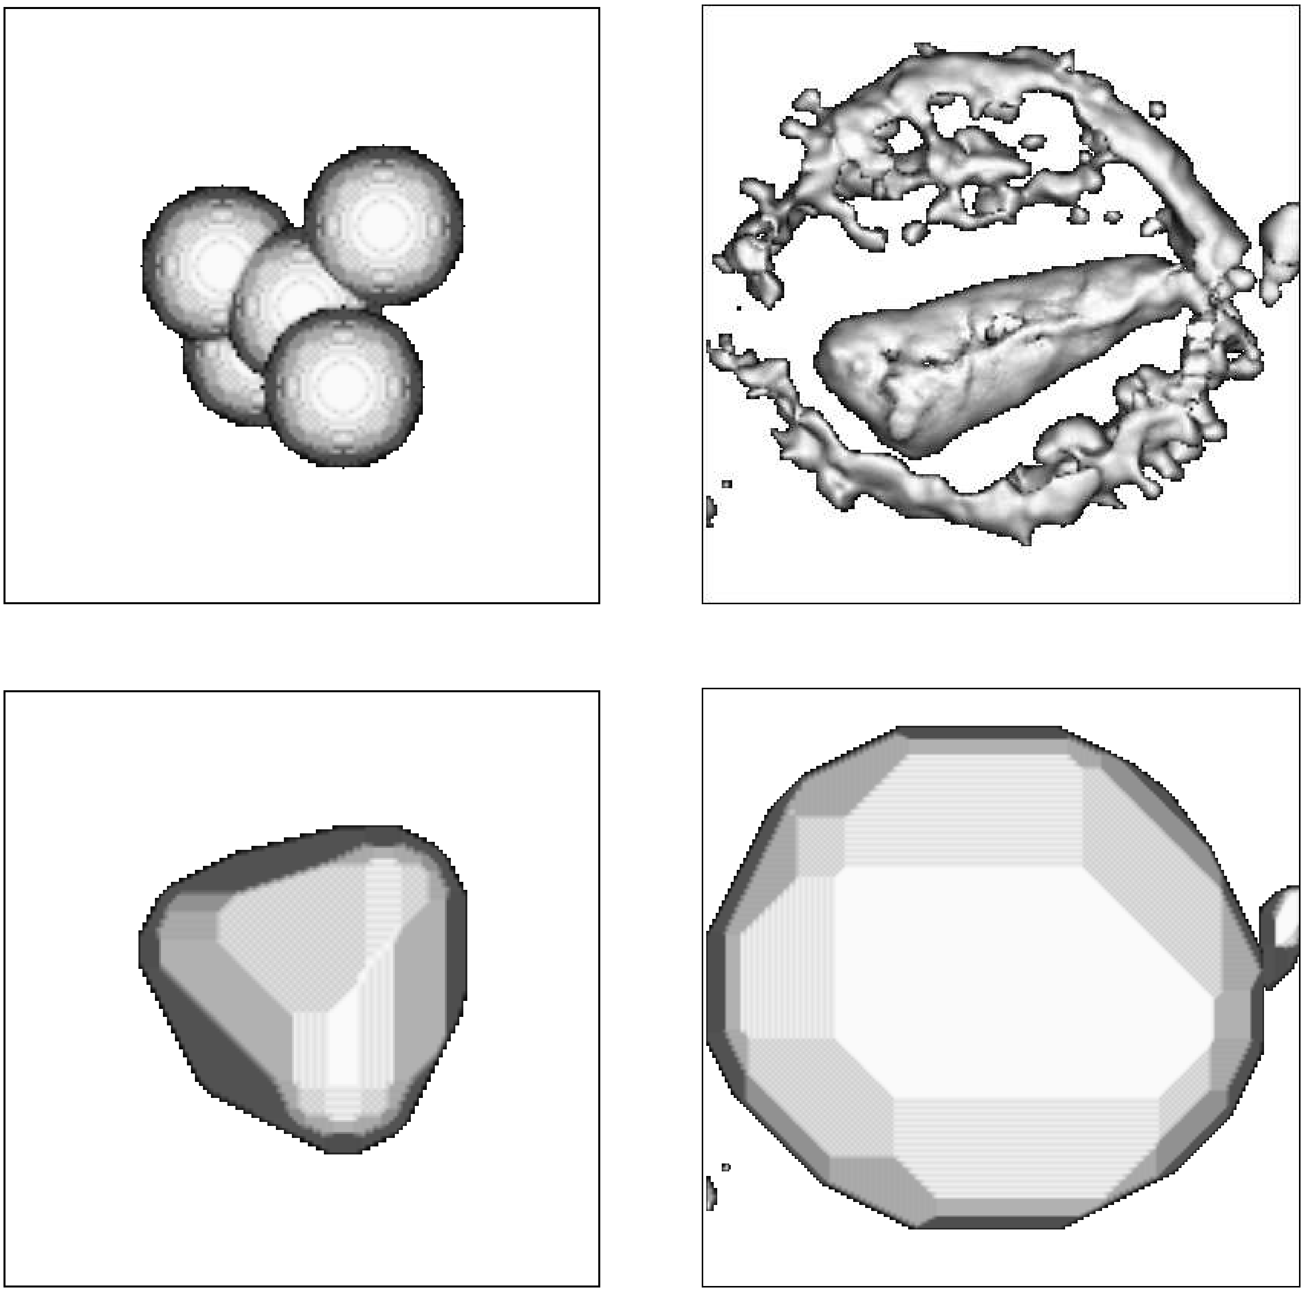

Figure 3: Objects (upper) and their covering polyhedrons (lower). To the left a synthetic image with spheres. To the right a reconstruction of HIV.

convex hull is good enough. In the case described here the object is oriented in a known and constant way, so the moderate rotation dependence of this covering polyhedron should not be serious.

The covering polyhedron can be improved by using the local curvature in a larger neighbourhood. Using the $5 \times 5 \times 5$ neighbourhood is equivalent of using not only the label of each voxel after the first pass, but also inspecting the values of its labelled neighbors. The exact rules that should be used can be found elsewhere.[6] The resulting covering polyhedron is a close approximation of a digital sphere, as the covering polygon in each coordinate plane now is the closest fitting hexadecagon (16 sides). The covering polyhedrons in Figure 3 are computed using this $5 \times 5 \times 5$ neighbourhood algorithm.

Once the covering polyhedron is computed, it can be used to compute a number of shape characteristics of the object. The polyhedron itself can be used to compute the maximal dimension of the object, as well as its principal directions. The concavity regions are even more interesting. The different concavities are identified using a connected component labelling, where each concavity gets a unique label. Then each concavity can be distinguished by a number of properties, e. g., volume, external surface area (number of voxels in the concavity having a neighbour in the background), internal surface area (voxels in the surface of the region that are not in the external surface), and depth. The depth can be computed using a constrained[7] distance transform.[8] The distance is computed from the background into the concavity region, where the object is used as a non-traversable region (constraint). The maximum distance value in each concavity is the depth. Finally, the concavity regions themselves can be considered as objects and their concavities, the meta-concavities, can be identified. In this way, a concavity tree is built, that describes the shape of the object in detail.

5 IMPLEMENTATIONS

The implementations of the described methods are written in C and integrated in IMP (IMage Processing), a general image analysis software for five-dimensional ($x,\ y,\ z,\ t,\ b$) image data developed at our Centre. The five dimensions are the spatial dimensions ($x,\ y,\ z$), the time dimension t, and the bandwidth dimension b, e. g. red (R), green (G), and blue (B) in color images. The software is using the X-Windows standard and Motif for the user interface.

For the visualization of the images we are using $ANALYZE^{TM}$, developed at the Mayo Foundation Clinic, USA.[9]

6 DISCUSSION

In our application we are primarily interested in comparing certain features at different stages of the maturation of the virus particle. We have chosen three stages: clearly immature, half-mature, and mature. The structure of the virus changes as the core develops into a compact cone during the maturation. In order to know whether structural differences are individual variations or actually represent systematic changes during the maturation process, several particles from each stage need to be compared.

By using the described normalizing method, similar structures will always get the same normalized orientation and, hence, a starting position for several 3D methods for quantitatively evaluating and comparing volume images is achieved. One way of quantifying the structure of an object is to use the described method of obtaining the concavity regions.

It is important to get as much information as possible out of each reconstructed image since the procedure of preparing the specimen and doing the EMT puts a limit to how many reconstructions can be made; in addition to being tedious the procedure does not always succeed in giving micrographs suitable for reconstruction. So far we have processed a handful of virus particles. The project is thus still in progress.

7 ACKNOWLEDGMENTS

Thanks to Stefan Höglund and Åsa Öhagen (née Nilsson), members of the HIV Structure Group at the Dept. of Biochemistry, Uppsala University, for providing us with the 3D images. Part of this project was financially supported by the Swedish National Board for Industrial and Technical Development (NUTEK) under contract no. 92-04453P.

8 REFERENCES

[1] S. Höglund et al., "Spatial Visualization of the Maturing HIV-1 Core and Its Linkage to the Envelope," *AIDS Research Human Retroviruses* , 8(1), pp. 1–7, 1992.

[2] I. Nyström, E. Bengtsson, B. Nordin, Å. Nilsson, S. Höglund, "Normalizing Volume Images Based on the Moment of Inertia for an Object," In *Proceedings of SSAB (Swedish Society for Automated Image Analysis) Symposium on Image Analysis, Gothenburg, Sweden*, pp. 127–130, 1993.

[3] K. Bosch, *Mathematik Taschenbuch*, Kap. 4.1.3.2, 6.11.9.8, 3. Aufl., R. Oldenbourg Verlag, München, 1991. In German.

[4] R. A. Johnson, D. W. Wichern, *Applied Multivariate Statistical Analysis*, Ch. 8, Prentice-Hall, Inc., Englewood Cliffs, 2nd Ed., 1988.

[5] G. Borgefors, G. Sanniti di Baja, "Methods for Hierarchical Analysis of Concavities," In *Proceedings of the 11th IAPR International Conference on Pattern Recognition, The Hague, The Netherlands*, Vol. III, pp. 171–175, 1992.

[6] G. Borgefors, I. Nyström, G. Sanniti di Baja. Under preparation.[1]

[7] J. Piper, E. Granum, "Computing Distance Transformations in Convex and Non-Convex Domains," *Pattern Recognition*, 20(6), pp. 599–615, 1987.

[8] G. Borgefors, "Distance Transformations in Arbitrary Dimensions," *Computer Vision, Graphics, and Image Processing*, 27, pp. 321-345, 1984.

[9] R. A. Robb, C. Barillot, "Interactive Display and Analysis of 3-D Medical Images," *IEEE Transactions on Medical Imaging*, 8(3), pp. 217–226, 1989.

[1] G. Borgefors, I. Nyström, G. Sanniti di Baja, "Computing Covering Polyhedra of Non-Convex Objects," In *Proceedings of* 5^{th} *British Machine Vision Conference, York, UK*, pp. 275-284, 1994.