Synthesising Objects and Scenes Using the Reverse Distance Transformation in 2D and 3D

Ingela Nyström[1] and Gunilla Borgefors[2]

[1] Centre for Image Analysis, Uppsala University,
ingela@cb.uu.se
[2] Centre for Image Analysis, Swedish University of Agricultural Sciences,
gunilla@cb.uu.se

[1,2] Lägerhyddvägen 17
S–752 37 Uppsala
Sweden

Abstract. The reverse distance transformation has proved useful in image synthesis. This paper describes how digital objects are created from a number of seed labels in an image. The shape of the obtained objects depends on the metric used. In 2D the Euclidean and the 3-4 metrics are mentioned, and in 3D the D^6, the D^{26}, and the 3-4-5 metrics are discussed. The proposed method has no need of expensive CAD systems. It is an excellent image synthesising tool when developing image processing algorithms, i. e. shape quantification, visualisation, scene analysis and range imaging, as the obtained objects are well-defined in the image. The method is most advantageous in 3D, as there is an increasing need for volume images, but synthesising objects in 2D can also be useful.

1 Background

The approach described in this paper is an aid in obtaining synthetic volume objects without using expensive and complex CAD systems.

When implementing algorithms for the 3D image processing tool box it is of importance to have a number of mathematically well-defined objects for algorithm evaluation. With a "known" object you know what result to expect and can debug your algorithms accordingly. This also gives the possibility to perform quantitative comparisons between algorithms.

Often one can see generalisations of the world into blocks, spheres, cylinders, cones, etc. Anyone who has been involved in applications will agree with the fact that the world does not consist of any of these geometrical models (especially not cylinders!). It is therefore desirable to use more general objects.

In a biomedical application it can also be the case that a sufficient number of images can not be provided fast enough; some tomography methods are still tedious and expensive.

The objects resulting from our method have smooth surfaces (depending on the metric used), which is consistent with many biomedical applications.

2 Reverse Distance Transformation

This section will give an overview of how to obtain a *reverse* distance transform.

A distance transformation (DT) is an operation that gives every object pixel a label that approximates the distance to the closest background pixel, or, *vice versa*, every background pixel the distance to the closest object pixel. In a *weighted* DT a mask, with weights corresponding to the distance between neighbouring pixels, is used [2]. Each type of neighbour is given a different local distance from the central pixel. When applying a reverse DT the same mask can be used. Similarly to the algorithm for computing the weighted DT, the sequential algorithm to compute the reverse DT requires a forward and a backward scan, during which half of the neighbours of every pixel are taken into account.

It is also possible to make a parallel implementation of the algorithm, similarly to the weighted DT case [3], but we have only implemented the sequential algorithm so far.

The input to the reverse DT is an image with seed labels in different positions. The algorithm performs a propagation with decreasing labels. Every pixel is assigned the maximum label of the current pixel and the neighbouring pixels subtracted by the local distances to the current pixel.

The reverse DT, using different metrics, has been used earlier in the literature when reconstructing the shape of an object from its local maxima (skeleton). Examples are the 3-4 weighted DT [1] and the Euclidean DT [4].

When extending to 3D the principle is the same as in the 2D case, but the algorithm now uses a 3D mask [2]. The weights in the mask are chosen according to the metric used. The D^6 metric is obtained by setting the distance to the 6 face-neighbours to 1 and omitting the other neighbours. The D^{26} metric is obtained by setting the distance to the closest 26 neighbours to 1. The 3-4-5 metric, where the distance to the face-, the edge-, and the point-neighbours is set to 3, 4, and 5, respectively, is a reasonably good integer approximation of the Euclidean metric. Pseudo code for forward pass of the 3-4-5 reverse DT:

```
/*-----------------FORWARD PASS-------------------------------*/
LOOPZ(image) {
 LOOPY(image) {
  LOOPX(image) {
   if (x<xLO || x>xHI || y<yLO || y>yHI || z<zLO || z>zHI)
    image = 0; /* image border omitted */
   else
    image = MAX(I(x-1,y-1,z-1)-5, I(x,y-1,z-1)-4, I(x+1,y-1,z-1)-5,
                I(x-1,y,z-1)-4,   I(x,y,z-1)-3,   I(x+1,y,z-1)-4,
                I(x-1,y+1,z-1)-5, I(x,y+1,z-1)-4, I(x+1,y+1,z-1)-5,
                I(x-1,y-1,z)-4,   I(x,y-1,z)-3,   I(x+1,y-1,z)-4,
                I(x-1,y,z)-3,     I)
  }
 }
}
/*-----------------END--FORWARD PASS--------------------------*/
```

Extending the Euclidean DT to 3D when running on a sequential computer requires a large amount of calculation and data storage; it is quite possible though, if desired. Even a parallel algorithm uses four sub-iterations in the image with large masks [5]. Also, the reverse Euclidean DT is not simply the reversal of the Euclidean DT. Unexpected difficulties occur, that can be overcome, but requires special care in designing the algorithm already in 2D, see [4]. The 3-4-5 DT should be a fair approximation in many cases.

3 Synthesis Algorithm

The two-step algorithm is very straightforward.

First of all, seed labels are placed in the image, randomly or manually. Placing and labeling the seeds randomly are done according to an appropriate probability distribution. The seed label equals the radius of the corresponding "sphere".

In the second step a reverse DT of some metric is applied. The metric used decides what the associated "sphere" looks like. In 3D, the D^6, the D^{26}, and the 3-4-5 reverse DT generate octahedrons, cubes, and 24-sided polyhedrons, respectively.

It might be tricky to place the seeds manually and really obtaining the planned object, as there are so many degrees of freedom. How close the seed labels are positioned decides how "grown together" the spheres are. An isolated seed label with a low value generates a single "sphere". If the seed labels are placed along a line an elongated object is generated, e. g., a block when using the D^{26}.

If the obtained 3D object has internal cavities, not apparent from simply viewing the object, they may bias quantitative measurements. But as the process of placing the seeds may be done in a manual manner, this situation can be avoided by a careful planning of the labelling and positioning of the seeds. To be certain that there are no internal cavities, a 3D connected component labelling [7] can be applied to the object and the background.

4 Examples

The shown examples are 3D objects synthesised in $128 \times 128 \times 128$ images. For the visualization we are using $ANALYZE^{TM}$, developed at the Mayo Foundation Clinic, USA [6].

Fig. 1 shows a cluster of five partially fused "spheres" with different radii obtained by using the D^6, the D^{26}, and the 3-4-5 reverse DT, respectively. By choosing the seed positions and labels manually a (more or less) planned object is created. In the three cases the same seed labels have been used, except for a rescaling of the labels for the D^6 image.

Another way of creating an object, is to place seeds of a special or a random label at random in the image, and then performing the reverse DT. In Fig. 2 randomly chosen seed labels are placed at random positions and the 3-4-5 reverse

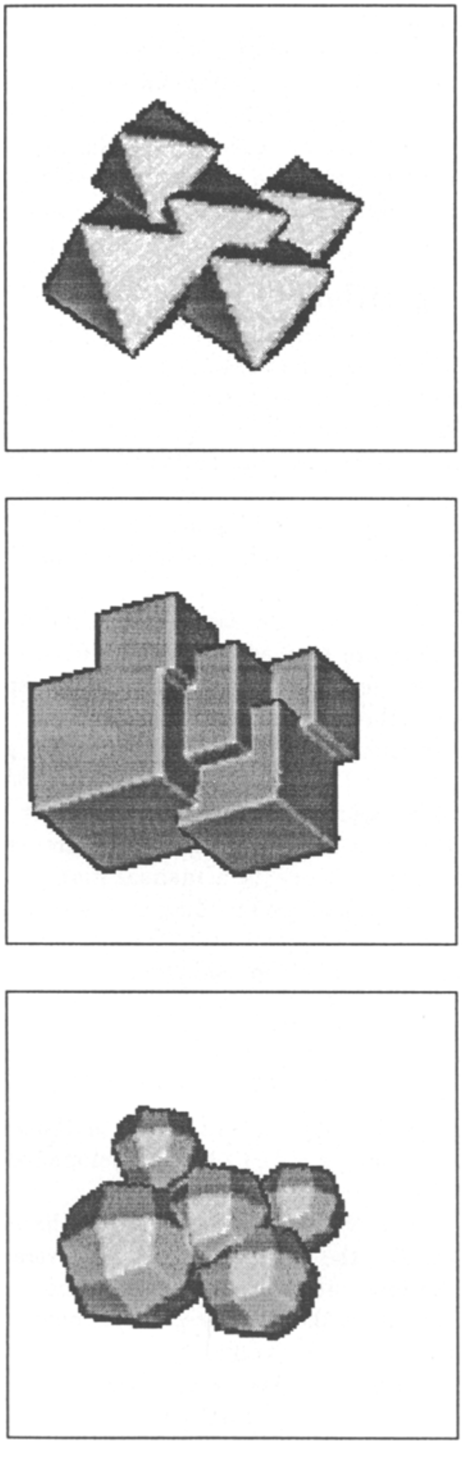

Fig. 1. **A cluster of partially fused "spheres". The metrics are D^6 (top), D^{26} (middle), and 3-4-5 (bottom).**

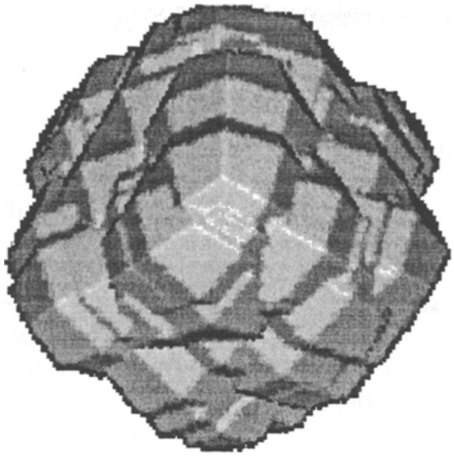

Fig. 2. A random 3D "blob" object.

DT is used. The positions and labels are chosen according to the rectangular probability distribution.

The method can also be used to obtain volume images suitable for scene analysis. A "blocks world" is generated with the D^{26} reverse DT on some seed labels. Volume rendering techniques may then produce 2D projections in desired directions. An example is shown in Fig. 3. Range images may also be obtained from the volume images.

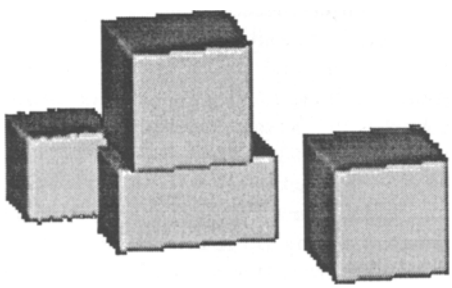

Fig. 3. A blocks world generated with the D^{26} reverse DT on a few seed labels.

5 Conclusion

We have found a way of creating useful objects in 2D or 3D by using a reverse DT in different metrics. After some practice, seed images can be created that generate complex objects and scenes.

At our Centre the method has been useful when testing tools for 3D shape quantification, i. e. approximation of the convex hull and different skeletonisation methods. Various objects have been synthesised.

An additional area where "known" objects may be useful is during development of visualization algorithms. You know what the algorithm should produce and if it does not, then something is wrong. There are certainly more areas where this method can be used, i. e. scene analysis and range imaging.

The implementation of the described method is written in C, so the image synthesis is easily portable to any computer equipment. There is really no need for expensive computer graphics software, unless you want to do some visualisations.

6 Acknowledgments

A special thanks to Dr. Gabriella Sanniti di Baja, who came up with the idea of using "small bunches of grapes" as almost convex test objects, when we were working on a 3D concavity filling algorithm.

References

1. Arcelli, C., Sanniti di Baja, G.: Finding local maxima in a pseudo-Euclidean distance transform. Computer Vision, Graphics, and Image Processing. **48** (1988) 361–367
2. Borgefors, G.: Distance transformations in arbitrary dimensions. Computer Vision, Graphics, and Image Processing. **27** (1984) 321–345
3. Borgefors, G., Hartmann, T., Tanimoto, S. L.: Parallel distance transforms on pyramid machines: theory and implementation. Signal Processing. **21** (1990) 61–86
4. Borgefors, G., Ragnemalm, I., Sanniti di Baja, G.: The Euclidean distance transform: finding the local maxima and reconstructing the shape. Proceedings of 7^{th} Scandinavian Conference on Image Analysis, Aalborg, Denmark. (1991) 974–981
5. Ragnemalm, I.: The Euclidean distance transform in arbitrary dimensions. Pattern Recognition Letters. **14** (1993) 883–888
6. Robb, R. A., Barillot, C.: Interactive display and analysis of 3-D medical images. IEEE Transactions on Medical Imaging **8(3)** (1989) 217–226
7. Thurfjell L., Bengtsson E., Nordin B.: A new three-dimensional connected component labelling and object feature extraction algorithm. Computer Vision, Graphics, and Image Processing **54** (1992) 357–364