Robust and Efficient Triangulation of Anatomical Surfaces from Medical Images

Matthias Meyer, Cristian Lorenz, Vladimir Pekar and Michael Kaus

Philips Research Laboratories, Sector Technical Systems,
Röntgenstraße 24-26, 22315 Hamburg
Email: michael.kaus@philips.com

Abstract. There are numerous techniques for obtaining surface triangulations from 3D image data. Some applications require efficient triangulations, i.e. the source data should be represented as accurately as possible with a minimum number of triangles. This work is based on an existing algorithm, that adapts triangulation density to local surface curvature. To be useful in clinical practice, robustness has been improved and parameterization has been simplified by reducing the number of required arguments from twelve to a single detail control parameter. Triangulation times are 115 sec for 1.5mio source triangles and 3 sec for 60.000 triangles on a 3GHz P4.

1 Introduction

Triangulation of three dimensional digital image data has many applications in engineering and science and a great variety of different approaches exists today [2]. In medical imaging in particular, it is of major interest to obtain triangular surface models from segmented CT or MRI data. The source data consists of large three dimensional voxel arrays ($>10^7$). However, if the surface mesh is to be used for further computations, the amount of detail has to be reduced dramatically. Hence an algorithm is required, that delivers simplified yet faithful triangulations.

In this paper we propose a method that is especially suitable for triangulations of anatomical data, which usually feature surfaces that are mostly smooth and curved. The method deals efficiently with anisotropic properties of the source data, which are typical for customary CT imagery. It is also robust, fast and easy to control.

Previous Work. A common technique to obtain triangulations of 3D imagery is the marching cubes algorithm. Because this algorithm produces very high numbers of triangles, several per surface voxel, the output has to be processed by a decimating algorithm to be useful. For a survey of existing techniques refer to [3].

Another popular approach is to select a set of points on the object's surface, which are in turn used to create the equivalent of a 2D-Delaunay-triangulation on

Fig. 1. Triangulations of hour-glass, lung and thorax.

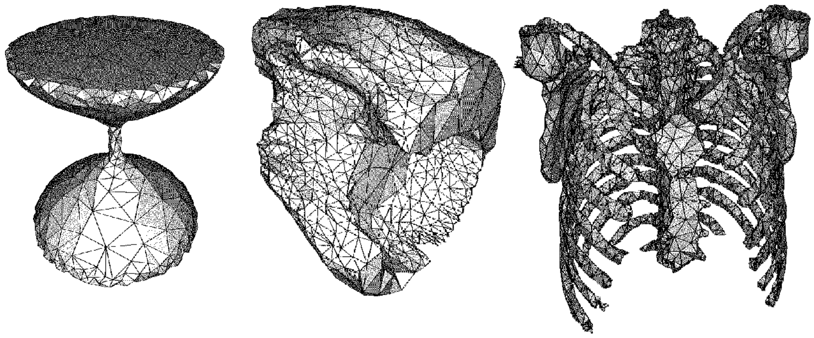

the same surface [7]. Krahnstoever et al. [1] proposed a method that adapts the density of vertex selection to local surface curvature. Their algorithm produces satisfying results in terms of quality and efficiency, but lacked some important properties. First, anisotropic properties of voxel images were ignored, leading to elongated triangles in the output meshes. Second, the original method was controlled by numerous arguments and required a trial and error approach to operate. Third, despite elaborate post-processing repairs, the occurrence of holes in the output mesh was not completly prevented.

Contribution. We propose a new method that is based on [1]. While anisotropic features could be accounted for with small effort, parameterization and robustness proved to be more difficult to improve. Krahnstoever et al. [1] forfeited the option to use adaptive refinement (see [4,5] and section 2.3) for performance reasons. In our approach the internal mesh connectivity was simplified by the removal of edges. Thus the overall performance was critically improved, so that adaptive refinement became feasable at last. Using this technique, the algorithm is now guaranteed to produce meshes that are valid 2-manifolds. Parameterization was reduced from a dozen of parameters to a single target polygon count (TPC).

2 Curvature Adaptive Triangulation

Before we proceed to the improvements in the following subsections, we will shortly introduce the basic method described in [1]. An intermediate representation of the source data is created, the *micro mesh*, which is composed of all voxel surfaces separating object and background voxels. Edges and related connectivity that were a part of the mesh structure in [1] have been eliminated. For each micro mesh face the local face curvature is computed [1,6]. After that, the actual triangulation can start.

A subset of the micro-mesh face centers is selected as vertices for our Delaunay-triangulation. First, all faces are added to a priority-queue according to their local curvature value. Now the face with the highest curvature is extracted and a region of faces of radius R_E (*elimination radius*) around this face will be blocked from selection. This process is repeated until no faces are left for selection and is called *thinning*. The size of the macro mesh triangles is indirectly controlled by R_E. To adapt the triangle sizes to local face curvature a *balance function* maps from curvature to elimination radius.

Now a two dimensional approximation of a Voronoi-graph can be created on the object surface. The selected vertices from the previous step are used as region seed points. These regions are expanded until the entire surface of the object is covered. The dual Delaunay-trianguliation to this Voronoi-graph is then obtained by finding the vertices of the micro-mesh that are connecting three or more different regions.

Anisotropic Data. Several steps of the algorithm - curvature patch, thinning and voronoi graph - involve a variation of region growing. In the original method all regions were expanded by one face at each border simultaneously, which lead to elongated regions for anisotropic source data. By changing the distance measure from counting faces to arithmetic distance, a priority queue was required for region growing. Access times for priority queues are *O(log n),* but in this case $n$is not determined by the size of input data, but by the average region size, which does not scale with the size of input data. Hence the overall complexity for this part of the triangulation remains unchanged.

Parameterization. Parameterization in [1] was composed of 12 arguments. Due to the improved robustness (see 2.3), options concerning robustness became obsolete. For other parameters like minimum and maximum thinning range and curvature patch range, heuristics were found empirically that provided appropiate settings for a given TPC.

The only argument that could not directly be controlled by the TPC was the coefficient of the balance function, that controlled the distribution of detail. However, the number of output triangles was strictly increasing by this coefficient. Thus an appropiate setting could be found by bisecting the coefficient until a face count was achieved that was sufficiently close to the TPC. This can be sped up significantly by not creating actual geometry during the process. Using these techniques, only the TPC argument remains. The balance coefficient can still be set manually to avoid the additional computation time required for bisecting. This is especially usefull if valid settings are already known for a certain class of input data.

Robustness. Using a two-dimensional Delaunay-triangulation on a three-dimensional surface is problematic [4]. The authors of [4,5] solve these problems by using adaptive subdivision. Krahnstoever et al. [1] tried a different approach by applying post-processing repairs. However, it proved to be difficult to repair all occurring artifacts, so we have chosen the adaptive subdivision approach. To do so, we have to identify all triangles that are part of an invalid configuration, i.e. that are connected to more or less than three neighobours. These defective areas are repaired by inserting additional vertices to the initial vertex selection. This process is repeated until no more invalid configurations occur. It can be shown, that this algorithm will always converge to a valid solution.

3 Results

The new algorithm has been tested with a variety of synthetical (cube, sphere, pipe, hourglass, fractal and white noise) and anatomical objects (esp. lung and heart).

The parameterization scheme is simple and requires no knowledge about the algorithm. It delivers accurate detail for all ranges of TPC and a large variety of objects.

The robustness was significantly improved. The 2-manifold condition can always be fulfilled. In some rare artifacts remain, which can be easily and efficiently resolved with the use of standard mesh-fairing algorithms like e.g. centroid-smoothing without significant loss of detail. One artifact is the occurrence of small hooks, that are a by-product of an inappropiate Delaunay-triangulation. Additionally, self-intersections of very close, near parallel surfaces may occur.

The computational performance has been improved. A typical anatomical object (lung, 46x184x57 voxels, 59k micro mesh faces) can be triangulated in 2.9 seconds with a memory footprint of 17 MB (TPC=3000). Our most complex sample (thorax skeleton, 465x285x457, 1.5 mio. micro mesh faces) was processed within 116 seconds with memory footprint of 526 MB (TPC=20000). The method of Krahnstoever et al. required 202 seconds and 437 MB on a Pentium 4 1.6 Ghz. Compared to our results, which were taken on a comparable 3.0 GHz machine, the performance appears to be equivalent. However, our algorithm processes eight partial and eighteen complete triangulations in the time of a single triangulation with the original method. For more benevolent source data the speed gain is even higher, for example the lung requires only five repair passes to complete. In addition, resampling input data to an isotropic grid is no longer necessary. Thus triangulations can be performed on the original data resolution, improving detail and a avoiding the time for resampling (e.g. 30 seconds for the thorax sample).

4 Conclusion

We have presented an algorithm for performing surface triangulations on three-dimensional image data. It is especially suitable to faithfully and efficiently approximating objects featuring curved structures of varying size and is thus useful to process CT or MRI data from the medical imaging domain. Its basic principle is to perform a two-dimensional Delaunay-triangulation on the surface of the input object, whereas the density of vertex selection is adapted to local surface curvature.

Our new algorithm features a higher level of robustness and is controlled by only a single parameter specifying the desired target polygon count. It is capable of accounting for anisotropic features of source data, which are typically found in data obtained by CT or MRI scans. It has been shown that the algorithm can handle very large data sets and even highly pathological topologies. Triangulation of typical data sets can be performed in a few seconds on standard hardware. There is significant gain in speed for anisotropic data sets, since no resampling is required. Memory requirements are linear to the number N of faces of the input surface. Computational complexity is $O(N(log\ N + r^2\ log\ r^2))$, where $r$is the size of the patch used for curvature estimation.

Future Work. There are many applications (e.g. heart modelling) that require the triangulation of multilabel images. Such images feature inner surfaces and T-joints of subsurfaces, which is an exception from the 2-manifold property. We have investigated this particular challenge and conclude that such an extension has to be carefully planned to avoid artifacts at subsurface-joints.

Since our algorithm merely assumes a set of connected faces and there is a distance defined between faces, we have attempted to triangulate arbitrary meshes. Experiments showed promising results, especially after we changed our distance function to Euclidian distance. Potentially, the algorithm can be used for curvature adaptive decimating of arbitrary high-resolution meshes. Finally, it would be interesting to compare our technique to the accuracy and performance featured by other techniques.

References

1. Krahnstoever N, Lorenz C. Computing curvature adaptive surface triangulations of three dimensional image data. The Visual Computer 2003;20(1):17–36.
2. Kumar S. Surface triangulation: a survey. Department of Computer Science, University of North Carolina; 1996.
3. Heckbert P, Garland M. Survey of polygonal surface simplification algorithms. School of Computer Science, Carnegie Mellon University; 1997.
4. Eck M, DeRose T, Duchamp T, et al. Multiresolution analysis of arbitrary meshes. In: Proceedings on Computer Graphics, SIGGRAPH 95; 1995. p. 173–182.
5. Lotjonen J, Reissman PJ, et al. A triangulation method of an arbitrary point set for bio-magnetic problems. IEEE Transactions on Magnetics 1998;34(4):2228–2223.
6. Stokely EM, Wu SY. Surface parameterization and curvature measurement of arbitrary 3-D objects - 5 practical methods. IEEE Transactions on Pattern Analysis and Machine Intelligence 1992;14(8):833–840.
7. Oblonsek C, Guid N. A fast surface-based procedure for object reconstruction from 3D-scattered points. Computer Vision and Image Unterstanding 1998;69(2):185–195.