

Growing Cells Meshing

Hendrik Annuth Christian-A. Bohn

Wedel University of Applied Sciences
Feldstr. 143, D-22880 Wedel, FR Germany
E-Mail: {annuth,bohn}@fh-wedel.de

Abstract

We propose *Growing Cells Meshing* (GCM) — a reconstruction algorithm which creates triangle meshes from clouds of arbitrary point samples registered on object surfaces. GCM is different to classical approaches in the way that it uses an artificial neural network (ANN) to represent the triangle mesh together with an iterative learning technique. Based on the *Growing Cell Structures* (GCS) approach [Fritzke, 1993] we introduce the *Smart Growing Cells* (SGC) network as extension to fulfill the requirements of surface reconstruction.

Our method profits from the well-know benefits entailed by neural networks, like autonomy, robustness, scalability, the ability of retrieving information from very complex data, and adaptability. On the downside, typical drawbacks like undesirable smoothing of information, inability to exactly model detailed, discontinuous data, or a vast amount of computing resources at big network sizes are overcome for the application of surface reconstruction.

The GCM approach creates high-quality triangulations of billions of points in few minutes. It perfectly covers any amount and distribution of samples, holes, and inconsistent data. It discovers and represents edges, manages clusters of input sample points, and it is capable of dynamically adapting to incremental sample data.

Keywords: Surface Reconstruction, Mesh Optimization, Geometric Modeling, Neural Networks, Unsupervised Learning, Neural Mesh, Growing Cells Structures

1 Introduction

The demand for efficient high quality reconstruction algorithms has grown significantly in the last decade, since the usage of 3D point scans has widely been spread into new application areas. These include geometric modeling to supplement interactive creation of virtual scenes, registering landscapes for navigation devices, tracking of persons or objects in virtual reality applications, medicine, or reverse engineering.

3D points, retrieved by laser scanners or stereo cameras, introduce two vital questions. First, how can one recognize a topology of the originating 2D surfaces just from 0-dimensional 3D samples and without any other information from the sampled objects? Second, for further processing, how is it possible to project this topological information on a data structure like a triangle mesh — meeting given constraints concerning mesh quality and size?

Although this issues have intensely been tackled since the early eighties a general concept that addresses all problems of surface reconstruction has not been determined up to now. Noise contained in the sample data, anisotropic point densities, holes, and discontinuities like edges, and finally, handling vast amounts of sampling data with adequate computing resources are still a big challenge.

Previous Work. The issue of surface reconstruction is a major field in computer graphics. There have been numerous approaches with different algorithmic concepts. In [Hoppe et al., 1992] and [Hoppe, 2008] an implicit surface is created from point clouds which then is triangulated by the marching cubes approach. [Edelsbrunner and Mcke, 1994] and [Kolluri et al., 2004] reduce a delaunay tetrahedralization of a point cloud until the model is carved out. Approaches like [Storvik, 1996] or [Huang et al., 2007] utilize techniques based on the Bayes' theorem.

In the area of artificial neural networks a famous

work is [Kohonen, 1982]. They propose the *Self Organizing Map* (SOM) which iteratively adapts its internal structure — a 2D mesh — to the distribution of a set of samples and enables clustering or dimensionality reduction of the sample data. While a SOM has a fixed topology, the growing cell structures concept [Fritzke, 1993, 1995] allows the network for dynamically fitting its size to the sample data complexity. SOM and GCS are suitable for processing and representing vector data, like point samples on surfaces. [Hoffmann and Vradý, 1998] uses a SOM and [Vradý et al., 1999] and [Yu, 1999] a GCS for the purpose of surface reconstruction. Further improvements in are made by [Ivrissimtzis et al., 2003b] where constant Laplace smoothing [Taubin, 1995] of surfaces is introduced, and in [Ivrissimtzis et al., 2003a] the curvature described by the input sample distribution is taken to control mesh density. In [Ivrissimtzis et al., 2004a] the GCS reconstruction process is further enhanced in order to account for more complex topologies. [Ivrissimtzis et al., 2004b] use several meshes of the same model for a mesh optimization process, and [Yoon et al., 2007] present a concept for combining common deterministic approaches and the advantages of the GCS approach.

In the following, we describe the basis of our approach — the growing cell structures — and then derive our idea of the smart growing cells which matches the specific requirements of reconstruction.

2 Growing Cells Meshing Approach

GCM is based on using the internal structure of a smart growing cells network as triangulation of object surfaces which are described by a set of surface sample points. The SGC is a neural network approach which is derived from the growing cells structures network, and which is presented for the first time in this work.

The reason for using a neural network scheme for reconstruction tasks are its obvious advantages compared to deterministic approaches.

- They robustly handle arbitrary sample set sizes and distributions which is important in case of billions of unstructured points from scanned objects.
- They are capable of reducing noise and ply discontinuities in the input data.
- They are capable of adaption — it is not required to regard all points of the sample set on the whole. Further, incrementally retrieved samples can be used to retrain the network without starting the triangulation process from scratch.
- They guarantee to — theoretically — find the best solution possible. Thus, approximation accuracy and mesh quality are automatically driven to the maximum.
- Training can be designed by simple learning rules which nevertheless solve complex problems. This is important for our application case, since the converged neural network structure and the general learning task only match few requirements of a reconstruction task.

Nevertheless, these advantages partly clash with our needs. On the one hand, discontinuities are often desired (for example, in case of edges or very small structures on object surfaces). On the other hand, smoothing often destroys important aspects of the model under consideration (for example, if holes are patched, if separate parts of the underlying objects melt into one object, or if the object has a very complex, detailed structure). In such cases, neural networks tend to generalize which may be advantageous from the physical point of view, but which mostly lets vanish visually important features which the human is quite sensitized to.

Growing cells meshing, from our point of view, is a breakthrough in this area. Through the use of a neural network, it delivers a general, robust, high-quality reconstruction algorithm which entails several advantages compared to classical approaches. Problems which often arise when using neural networks are solved for the application of surface reconstruction.

In the following, we outline the SGC algorithm. Then, our modifications to meet the requirements of surface triangulation are described.

2.1 Unsupervised Learning and Growing Cells

Unsupervised learning is accomplished by certain types of artificial neural networks which are able to organize its internal structure automatically depending on an arbitrary input sample distribution. After training, a set of *reference vectors* match the input sample distribution — classification is a typical task for this type of networks. In case of this work, reference vectors are interpreted as a set of vertices located on object surfaces.

Adaption of reference vectors in iterative unsupervised learning approaches is called *learning* or *training* and is generally accomplished by randomly presenting single n -dimensional samples from the input sample set to n -dimensional reference vectors and moving them in n -dimensional space, described as follows.

Place k reference vectors $\mathbf{c}_i \in \mathbb{R}^n$, $i \in \{0..k-1\}$ randomly in the n D space of input samples.

repeat

Chose sample $\mathbf{s}_j \in \mathbb{R}^n$ randomly from the input set.

Determine reference vector \mathbf{c}_b (*best matching vector* or *winning unit*) which lies closest to \mathbf{s}_j .

Move \mathbf{c}_b in the direction of \mathbf{s}_j according to a certain strength ϵ_{bm} , like $\mathbf{c}_b^{\text{new}} = \mathbf{c}_b^{\text{old}}(1 - \epsilon_{bm}) + \mathbf{s}_j \cdot \epsilon_{bm}$.

Decrease ϵ_{bm} .

until $\epsilon_{bm} \leq$ a certain threshold ϵ_0 .

Surface reconstruction with pure unsupervised learning would place a set of reference vectors on object surfaces, but does not determine information about the underlying surface topology. This leads to the *Kohonen Self Organizing Map* (SOM) described in the following.

Kohonen Self Organizing Map. The SOM is based on reference vectors which are randomly connected in a regular 2D mesh. The learning

rule is extended to account for the direct neighborhood of a best matching unit as follows.

```

for all  $c_{nb} \in$  neighborhood of  $c_b$  do
  Move  $c_{nb}$  in the direction of  $s_j$  according to a
  certain strength  $\epsilon_{nb}$ , like  $c_{nb}^{new} = c_{nb}^{old}(1 - \epsilon_{nb}) +$ 
   $s_j \cdot \epsilon_{nb}$ .
  Decrease  $\epsilon_{nb}$ .
end for

```

Insertion of the above “neighborhood loop” into the general unsupervised learning algorithm (after moving of c_b) leads to the phenomenon that the reference vertices are moved by accounting for the regular 2D mesh topology of the SOM. For example, training a plane-like arranged sample set leads to an adaption of the SOM grid to this implicit plane — the sample topology is recognized and finally represented by the SOM mesh.

Nevertheless, the mesh size of a SOM is fixed and cannot adjust to the sample structure complexity. The growing cell structures approach overcomes this drawback.

Growing Cell Structures. To a certain degree, GCS may be seen as SOM which additionally are capable of growing and shrinking according to the problem under consideration defined by the sample distribution. This mechanism is based on a so called *resource term* contained in every reference vector which is a simple counter. It counts the reference vector being a best matching unit. A high counter value signals the requirement for insertion of new reference vectors. Generally, “resource” may be defined differently, for example by accumulating an approximation error.

But using a GCS for surface reconstruction still exposes vital problems. There is only one given type of topology available. In other words, starting training with a regular mesh can best approximate the topology of plane-like structures, and a tetrahedron would be adequate for sphere-like objects — only objects which are homeomorphic to the start object can be represented satisfactorily. Further problems arise at discontinuities like sharp edges and holes. These

are commonly difficult to model through neural network type algorithms like mentioned above. This brings us to our approach, the smart growing cells.

2.2 Smart Growing Cells and GCM

Smart Growing Cells overcome the most problems mentioned above by introducing the three basic mechanisms of (a) *aggressive cut out* (ACO), (b) *discontinuity vertices*, and (c) a specific curvature criterion, explained below.

The SGC basic structure is identical to general GCS. There are n -dimensional cells which we now term *neural vertices* connected by *links* through an m -dimensional topology. We let $n = 3$ since neural vertices are directly taken as vertices of the triangulation mesh and $m = 2$ since we aim at 2D surfaces to be reconstructed.

2.2.1 General Training

In the following, we expose the simplified main training loop which is very similar to the classical GCS approach (see section 2.1).

```

repeat
  for  $j = 1$  to  $k_{del}$  do
    for  $i = 1$  to  $k_{ins}$  do
      Choose a sample randomly, find closest neural vertex  $c_b$  and move it with its neighbors towards the sample.
      Increase signal counter of  $c_b$  (the resource term mentioned above), decrease the signal counters of all other vertices.
    end for
    Find the best performing neural vertex — which is the one with the highest signal counter value — and add a new vertex at this position (see Fig. 1).
  end for
  Find the worst performing neural vertices, delete them and affected edges (see Fig. 1).
until a certain limit like approximation error, or number of vertices is reached.

```

For drawing the neighboring vertices, a smoothing process like described in [Ivrissimtzis et al.,

2003b] and [Taubin, 1995] is applied which replaces the classical movement, and which makes the adaptation of the topology more robust.

As initial network, usually a tetrahedron or a plane with random vertices is suitable.

Vertex Split. A vertex split operation adds three edges, two faces and a new neural vertex. The longest edge at the neural vertex with the highest resource term is split and a new vertex is added in the middle. The signal counter value is equally spread between the two vertices (see Fig. 1).

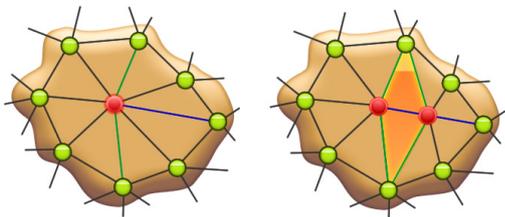


Figure 1: Vertex split operation (from left to right) to increase mesh granularity locally, and edge collapse (from right to left) to reduce mesh locally.

Edge Collapse. Neural vertices with resource terms below a threshold r_{min} are removed together with three edges and two connecting faces (see Fig. 1). Determination of the edge to be removed is driven by connectivity irregularities as proposed in [Ivrissimtzis et al., 2003b].

2.2.2 Aggressive Cut Out

Aggressive Cut Out (ACO) may be the most important new scheme we added to the standard GCS algorithm for achieving sufficient flexibility to match any topology of the training samples (the “homomorphism problem”). Before the edge collapse operation is applied to a vertex, it will be tested if the vertex is contained in a *degenerated mesh region* (definition follows below). If so, the ACO process deletes it and additionally all connected faces.

It has been shown that degeneration of a part of the mesh serves as perfect indicator for a mesh

topology which does not fit the underlying sample structure correctly. For example, consider a region where sample densities equal zero. Although vertices are not directly drawn into it by training adjustment, their neighbors may be moved there through their mesh connections. Due to the resource term, these vertices will be deleted by edge collapse operations, but their links remain and mistakenly represent the existence of a topology. In this case, the structure of the links is degenerated, i.e., it usually shows a surpassing number of edges with *acute-angled*¹ vertices (see Fig. 2).

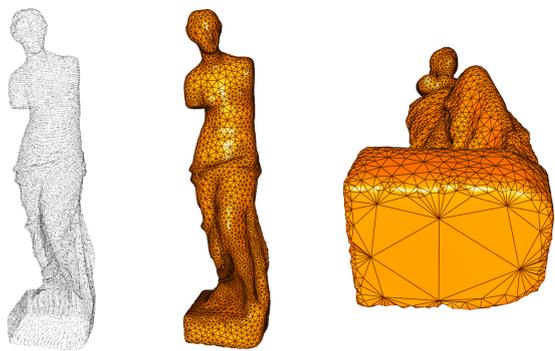


Figure 2: The statue’s bottom is not represented by samples, on the right, the acute-angled triangles expose a degenerated mesh region.

The term “aggressive” is chosen since triggering properties are often matched during training — suspicious neural vertices will be cut out early.

Criterion for Degenerated Mesh Regions. In [Ivrissimtzis et al., 2004a] a large area of a triangle is taken as sign for a degenerated mesh structure, but it has been shown that this criterion warns very late. Also, anisotropic sample densities are mistakenly interpreted as degenerated mesh regions. Our proposal is a combination of *vertex valence*², triangle quality, and quality of neighboring vertices. If all of the following conditions hold, an ACO is started at the vertex to be deleted.

1. Vertex valence rises above a certain threshold $n_{devalence}$.

¹A triangle is termed acute-angled if the ratio of its area and the area which is spanned by a second equilateral triangle built from the longest edge of the first lies below a certain threshold ϵ_{acute} .

²Vertex valence is the number of connected vertices.

2. Vertex is connected to at least $n_{degacute}$ acute-angled triangles.
3. Vertex has more than n_{degnb} neighboring vertices for which condition (1) or (2) hold.

The latter condition says that ACO is only started if at least one or two neighbors have the same inconsistencies in their local mesh structure. This is reasonable since single degenerated vertices do not necessarily expose a problem, but may arise by accident.

Curing Boundaries after an ACO. Obviously after an aggressive extinction of a neural vertex and its surrounding faces, a boundary will be left behind which may consist of unfavorable mesh structure elements. Curing detects those structures along a boundary and patches them discriminating between the following four cases.

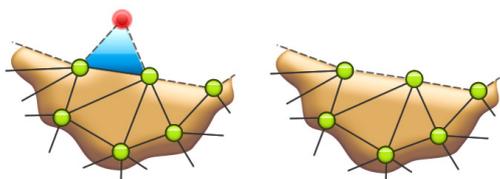


Figure 3: Curing a boundary with a spike.

Spike. A boundary vertex with a valence of 2 (see Fig. 3) is termed spike. Such a vertex is very unlikely to support a correct reconstruction process since it will be adjusted to an acute-angled triangle after few iteration steps. A spike must be deleted completely.

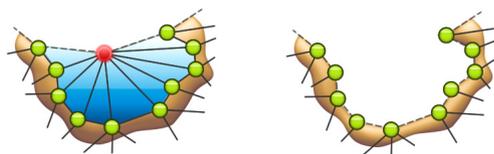


Figure 4: Cut out processes of a nasty vertex.

Nasty Vertex. A nasty vertex is a neural vertex with at least $n_{nastyacute}$ acute-angled triangles and/or triangles with a valence greater than

$n_{nastyval}$ (see Fig. 4). These vertices are suspected to be part of a degenerated mesh region and are deleted.

Needle Eye. A needle eye is a neural vertex that is connected to at least two boundaries (see Fig. 5). At these locations the mesh does not have a valid mesh structure. To delete a needle eye, all groups of connected faces are determined. From these, the group with the most faces is kept and all others are deleted.

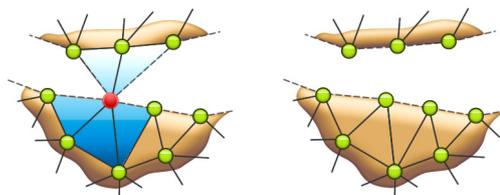


Figure 5: Cut out processes of a needle eye.

Bridge. A bridge is very likely to be part of a degenerated mesh region. If a mesh has a hole that consist of three vertices, then it would soon be closed by a coalescing process (see section 2.2.3). This is not allowed if exactly one of the edges of this hole would additionally be connected to a face (which we term “bridge”, see Fig. 6) since an invalid edge with three faces would arise. The entire bridge structure is deleted and the hole will be closed with a new face.

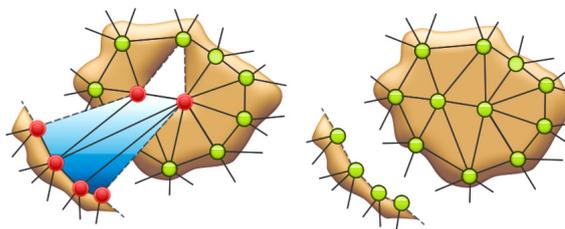


Figure 6: Curing a bridge.

Multiple Boundary Search Through. After the deletion of a neural vertex by the ACO process the curing mechanism will look for unfavorable structures along the boundary. There is more than one boundary to be considered, if

the deletion destroys a coherent set of faces and multiple separate groups of faces arise.

Four cases may appear. First, the usual case with no additional boundaries. Second, when a needle eye is destroyed, the boundaries of all groups of connected faces need to be tested. Third, when surrounding faces of a vertex are interrupted by boundaries. And fourth, when a needle eye is connected to the surrounding faces of a vertex (see Fig. 7). In other words, these cases happen since the faces that are deleted may not necessarily be connected to another face due to another deletion process.

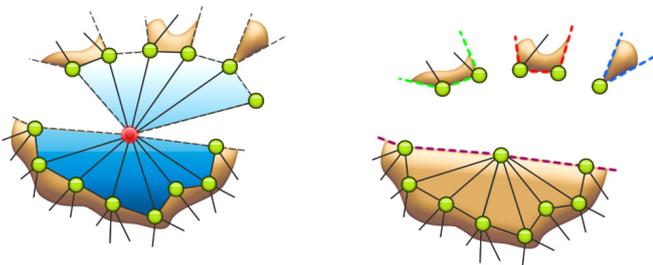


Figure 7: Cut out of a needle eye with a row of faces. Here, each face is not necessarily connected to another face. In contrast, if a needle eye has several groups of connected faces then there are some omissions of faces around the needle eye.

2.2.3 Coalescing

Like the mesh can be split by deletion of vertices or by the ACO process, it must also be possible to merge two mesh boundaries during the training process. For that, a coalescing test is accomplished when a boundary vertex is moved.

Coalescing Test. It determines if two boundaries are likely to be connected to one coherent area. For that, a sphere is created with the following parameters. Given the neighboring boundary vertices v_1 and v_2 of c_b , then we define $c = 1/2(v_1 + v_2)$. A *boundary normal* n_c is calculated as the average of all vectors originating at c and ending at neighbors of c_b , where v_1 and v_2 are not taken into account. The boundary normal can be seen as a direction pointing to the opposite side of the boundary. We define a sphere with the center at $c + n_c r$ with a radius r as the average length of the edges at c_b .

The coalescing condition at two boundaries hold, i.e., merging of the boundaries containing c_b and q on the opposite side happens, if q is contained in the defined sphere, and the scalar product of the boundary normals at c_b and q is negative.

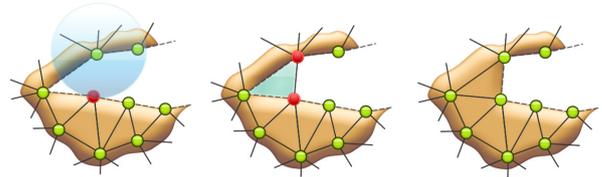


Figure 8: Coalescing process at a mesh corner. On the left, the search process of a coalescing candidate. In the middle, one edge is created, on the right, the only face capable of being added is the corner face.

Coalescing Process. After detecting the neutral vertex q to be connected with c_b , the according faces must be created, starting with one edge from c_b to q . There are two cases which have to be considered.

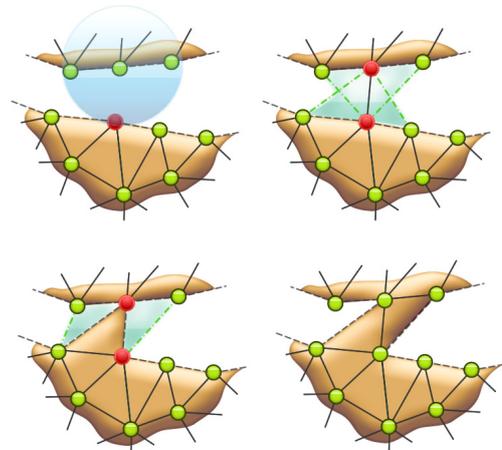


Figure 9: Coalescing of two separate boundaries. In the second picture, the edge is determined, in the third picture, the triangle with the smallest variance of edge lengths is added, in the fourth picture, another triangle must be added to avoid a needle eye.

Corner. A corner of the same boundary arises when c_b and q have one neighboring vertex in common. Here, it is only possible to generate the face which closes the corner (Fig. 8). Otherwise a needle eye would be created which is not a satisfactory mesh structure.

Long Side. Here, two boundaries appear to be separated. After determining the new edge, there are four possibilities for insertion of a new face containing the edge (see second picture in Fig. 9). The triangle with edge lengths which vary fewest is taken in our approach (see third picture in Fig. 9) since it is the triangle with the best features concerning triangle quality. Finally, to avoid a needle eye, a further triangle must be added — again, we take, the face with the highest edge similarity (see fourth picture in Fig. 9). Coalescing sometimes creates unfavorable bridges that need to be deleted through a further ACO process.

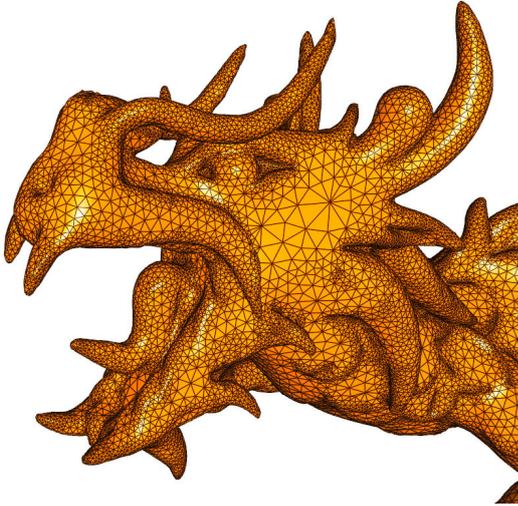


Figure 10: Granularity adaption correlates surface curvature to mesh density, details of the model are accentuated.

2.2.4 Granularity Adaption

Up to now, the SGC are able to approximate an arbitrary sample set by a 2D mesh. What remains is an efficient local adaption of the mesh density in a way that areas with a strong curvature are modeled by a finer mesh resolution (see Fig. 10). This also relieves the influence of the sample density on the mesh granularity making the SGC less vulnerable to sampling artefacts like holes or regions which are not represented by a uniform sample distribution.

Each time a vertex is adapted by a new sample we calculate the estimated normal \mathbf{n}_k at a neural vertex \mathbf{v}_k by the average of the normals at the

surrounding faces. The curvature $c_k \in \mathbb{R}$ at a vertex is determined like

$$c_k = 1 - \frac{1}{|\mathcal{N}_k|} \sum_{\mathbf{n} \in \mathcal{N}_k} \mathbf{n}_k \cdot \mathbf{n} \quad (1)$$

with the set \mathcal{N}_k containing the normals of the neighboring neural vertices of \mathbf{v}_k . Each time a neural vertex is chosen as winner, its curvature value is calculated and a global curvature value \bar{c} is adjusted. Finally, the curvature dependent resource term r_k at \mathbf{v}_k is adjusted through $r_k^{new} = r_k^{old} + \Delta r_k$, and

$$\Delta r_k = \begin{cases} 1, & \text{if } (c_k < \bar{c} + \sigma_{r_k}) \\ [c_k / (\bar{c} + \sigma_{r_k})] (1 - r_{min}) + r_{min} & \text{else,} \end{cases}$$

with the deviation σ_{r_k} of the resource term r_k , and a constant resource r_{min} that guarantees that the mesh does not completely vanish at plane regions with a very low curvature.

2.2.5 Discontinuity Modeling

A sampled model that exposes discontinuities like edges is difficult to be approximated by a neural network mesh. Discontinuities are smoothed out since the network tries to create a surface over them. This might be acceptable in many application areas since the approximation error is fairly small, but the effect is unfavorable in computer graphics since it is clearly visible. And even worse: edges are quite common in real world scenarios.



Figure 11: Discontinuity modeling provides for mesh edges which map to the object edges.

Therefore, we propose handling vertices at regions classified as discontinuities of the underlying scanned object differently — first, they are

only allowed to move in the direction of an object edge to represent it more properly, and second, the smoothing process from section 2.2.1 is not applied to them.

Recognizing these *discontinuity vertices* is accomplished as follows. We determine the curvature values of those neighbors which have a distance of two connections from the vertex v under consideration — the “second ring” of neighbors. Then the average δ_{ring} of the squared differences of consecutive curvature values on the ring is calculated.

If a curvature value clearly deviates from the average curvature value, then a discontinuity at this location is assumed if the average of the neighbors’ (second ring) curvature gradient differs to a certain amount, i.e., if

$$(c_k > 2\sigma_{c_k}) \wedge (\forall c \in \mathcal{C}_k : \delta_{ring} > 4\sigma_{c_k}^2)$$

with \mathcal{C}_k the set of curvature values of the second ring of neighbors.

For approximating the edge normal at a discontinuity vertex v , we take the average of the normals of two of the neighboring vertices, either those with the highest curvature value, or those which are already detected as being a discontinuity vertex. Finally, the normal is mirrored if the edge angle lies above 180° , which is indicated by the average of the surrounding vertex normals; in the first case it points in the direction of v .

3 Results

Complete Algorithm. In Fig. 12 the whole GCM algorithm is listed as pseudocode. To keep it comprehensive, the outermost loop of the algorithm is neglected, vertex split and edge collapse operations are triggered by using counters (compared with the basic algorithm from section 2.2.1).

Parameters. The following parameters have been proven to be reliable for almost all sample sets we took for reconstruction: $\epsilon_{bm} =$

```

Adjust samples regarding curvature.
Calculate average curvature and deviations.
Determine and sign discontinuity vertices.
for all Boundary vertices do
  if  $\exists$  coalescing candidate then
    Melt boundary.
    for all ACO candidates do
      Clean by ACO.
    end for
  end if
end for
if Edge collapse operation triggered then
  Collapse edge.
  for all ACO candidates do
    ACO-cleaning
  end for
end if
if Vertex split operation triggered then
  Split vertex.
end if

```

Figure 12: The complete GCM algorithm. The outermost loop of the algorithm is neglected, vertex split and edge collapse operations are triggered through counters.

0.1 , $\epsilon_{nb} = 0.08$, $r_{min} = 0.3$, $\epsilon_{acute} = 0.5$, $n_{degacute} = 4$, $k_{ins} = 100$, $k_{del} = 5$, $n_{degnb} = 1$, $n_{nastyacute} = 4$, $n_{nastyval} = 3$.

Test Hardware. We used a *Dell*[®] Precision M6400 Notebook with *Intel*[®] Core 2 Extreme Quad Core QX9300 (2.53GHz, 1066MHz, 12MB) processor with 8MB 1066 MHz DDR3 Dual Channel RAM. The algorithm is not parallelized.

Some visual results are exposed on page 12. All pictures are drawn from one SGC mesh and they are rendered using Phong Shading. Most models stem from the *Stanford 3D Scanning Repository*.

Besides visual results, GCM comes up with impressive numbers, listed in Fig. 13. It can be seen that mesh quality, i.e. the percentage of perfect triangles in the mesh lies at 96% at average. This is an outstanding result, nevertheless this is typically expected when using an approach from the field of unsupervised learning, since it guarantees an ideal representation of the underlying training sample distribution.

	Samples	Vertices	Time [min:sec]	Quality [Delaunay]	RMS/ Size
	36K	30K	1:19	95.6%	4.7e-5
	438K	100K	5:33	95.5%	3.3e-5
	544K	260K	18:33	93.1%	1.7e-5
	14,028K	320K	24:27	98.5%	1.3e-5
	5,000K	500K	42:10	95.9%	2.7e-5
	511K	10K	0:21	99.8%	6.6e-5
	38K	5K	0:12	99.0%	15e-5
	346K	5K	0:12	98.3%	0.7e-5

Figure 13: Results with objects from the *Stanford 3D Scanning Repository*. “Quality” means the percentage of triangles which hold the Delaunay criterion. RMS/Size is the root of the squared distances between the original point samples and the triangle mesh, divided by the longest edge of the maximum expansion of the sample set.

Further, the distance (RMS/object size) between samples and mesh surface is negligible low — far below 1% of the object size at average. This is even more pleasant, since usually, the problem at edges generate large error terms. It proves that discontinuity vertices are worth it. Also the computing times needed are very short, i.e., few minutes in almost all case.

All those measurements seem to be far better than those from classical approaches, as long as we could extract them from the regarding papers. The presented algorithm works very robustly. There are nearly no outliers visible in the mesh.

4 Summary and Future Work

We presented a new neural network approach for surface reconstruction purposes — the smart growing cells. The algorithm, on the one hand, profits from the common outstanding facilities of neural networks, on the other hand, well-known problems when using such iterative tech-

niques are avoided successfully. Growing cells meshing is capable of recognizing and representing arbitrary topologies of objects. The network training generates valid meshes at any size and is able to handle 3D point samples at an arbitrary amount. The mesh accounts for discontinuities like sharp edges and holes, smoothes over missing sample data, handles difficult distributions of sample data, and its granularity is efficiently adapted to the curvature of the object shape. The iterative approach enables to dynamically change the training sample set to adapt the network to small changes in a shorter time.

To achieve this, we introduced several mechanisms like, first, aggressive cut out, which can be seen as a brute cleaning mechanism for ill-formed structures. We found out that these structures are also a great hint for changing connectivities to match the underlying surface topology. Second, face normals are regarded and included in the neural network training loop to efficiently adapt mesh granularity and to make the reconstruction process independent from the sample distribution. Third, we proposed a specific learning strategy for vertices which lie in a sample region which represents discontinuities in the underlying, scanned model.

The proof of concept of our approach is enriched by the achieved quality and performance measures. For the tested geometries which each hold specific challenges of reconstruction, we got approximation errors for comparable mesh resolutions that lie far below 1% at average. Mesh quality, measured by the percentage of triangles which comply the Delaunay criterion, lies at 96% at average, and the time needed to compute meshes of several hundreds of thousands of polygons were just few minutes.

Thus, with these results, we are confident that we could show that the growing cells meshing approach is a considerable alternative to common reconstruction methods.

Future Work. There are several critical object structures which can hardly be modeled adequately by the presented method and which will be challenged in future work. These are, for ex-

ample, planes which are thinner than the average sample distance. From the SGC point of view, these surfaces are not clearly distinguishable.

Another issue will be reducing vertex density on flat planes whereas avoiding invalid mesh structures due to the extreme variation in vertex granularity.

Computation times are that small that one could think of a real-time reconstruction approach. Scanning users in VR applications for sending their shape to other places could be a nice feature. Here, an effective parallelization of the algorithm would be required and could be accomplished by executing sample adjustment on several processing units concurrently.

Generally, the SGC approach should be able to bring a lot of new ideas, because of its flexibility and since it is capable of learning arbitrary sample sets. It delivers arbitrary sized meshes as well as meshes with a very low resolution in a very short time — independent from the sample set size.

Acknowledgment. The authors would like to thank Kai Burjack for his support in rendering the images.

References

- H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes, 1994.
- B. Fritzke. Growing cell structures - a self-organizing network for unsupervised and supervised learning. *Neural Networks*, 7:1441–1460, 1993.
- B. Fritzke. A growing neural gas network learns topologies. In G. Tesauero, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems 7*, pages 625–632. MIT Press, Cambridge MA, 1995.
- M. Hoffmann and L. Vrády. Free-form surfaces for scattered data by neural networks. *Journal for Geometry and Graphics*, 2:1–6, 1998.
- H. Hoppe. Poisson surface reconstruction and its applications. In *SPM '08: Proceedings of the 2008 ACM symposium on Solid and physical modeling*, pages 10–10, New York, NY, USA, 2008. ACM.
- H. Hoppe, T. DeRose, T. Duchamp, J. A. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. In J. J. Thomas, editor, *SIGGRAPH*, pages 71–78. ACM, 1992.
- Q.-X. Huang, B. Adams, and M. Wand. Bayesian surface reconstruction via iterative scan alignment to an optimized prototype. In *SGP '07: Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 213–223, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.
- I. Ivrišsimtzis, W.-K. Jeong, and H.-P. Seidel. Neural meshes: Statistical learning methods in surface reconstruction. Technical Report MPI-I-2003-4-007, Max-Planck-Institut für Informatik, Saarbrücken, April 2003a.
- I. Ivrišsimtzis, W.-K. Jeong, S. Lee, Y. Lee, and H.-P. Seidel. Neural meshes: surface reconstruction with a learning algorithm. Research Report MPI-I-2004-4-005, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, October 2004a.
- I. Ivrišsimtzis, Y. Lee, S. Lee, W.-K. Jeong, and H.-P. Seidel. Neural mesh ensembles. In *3DPVT '04: Proceedings of the 3D Data Processing, Visualization, and Transmission, 2nd International Symposium*, pages 308–315, Washington, DC, USA, 2004b. IEEE Computer Society.
- I. P. Ivrišsimtzis, W.-K. Jeong, and H.-P. Seidel. Using growing cell structures for surface reconstruction. In *SMI '03: Proceedings of the Shape Modeling International 2003*, page 78, Washington, DC, USA, 2003b. IEEE Computer Society.
- T. Kohonen. Self-Organized Formation of Topologically Correct Feature Maps. *Biological Cybernetics*, 43:59–69, 1982.
- R. Kolluri, J. R. Shewchuk, and J. F. O'Brien. Spectral surface reconstruction from noisy point clouds. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 11–21, New York, NY, USA, 2004. ACM.
- G. Storvik. Bayesian surface reconstruction from noisy images. In *In Interface 96*, 1996.
- G. Taubin. A signal processing approach to fair surface design. In *SIGGRAPH*, pages 351–358, 1995.
- L. Vrády, M. Hoffmann, and E. Kovcs. Improved free-form modelling of scattered data by dynamic neural networks. *Journal for Geometry and Graphics*, 3:177–183, 1999.
- M. Yoon, Y. Lee, S. Lee, I. Ivrišsimtzis, and H.-P. Seidel. Surface and normal ensembles for surface reconstruction. *Comput. Aided Des.*, 39(5):408–420, 2007. ISSN 0010-4485.
- Y. Yu. Surface reconstruction from unorganized points using self-organizing neural networks. In *In IEEE Visualization 99, Conference Proceedings*, pages 61–64, 1999.



Figure 14: Upper lines: mesh training stages with number of vertices, lower lines, some assorted pictures of reconstructed models.