# Reconstruction for Virtual Reality Scenes

Hendrik Annuth and Christian-A. Bohn

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

**Abstract:** A common way of creating a virtual world is scanning a real world environment by a stereo camera set or by a 3D laser scanner. It results in a 3D point cloud which requires a surface reconstruction to generate a triangulated mesh or an implicit surface for following tasks like rendering.

In this paper we focus on challenges concerning the surface reconstruction of arbitrary complex environments in contrast to compact, single objects. These types of reconstructions, that most approaches are not capable of, are vital for Virtual Reality applications. We suggest the *Smart Growing Cells* (SGC) approach satisfying these requirements and demonstrate this by comparison and discussion of classical reconstruction algorithms.

**Keywords:** Virtual Environments, Surface Reconstruction, Growing Cells Structures, Smart Growing Cells, Unsupervised Learning, Artificial Neural Networks

## 1 Introduction

The advantages of virtual towards real world environments are numerous. A virtual environment can be explored independently from time and place. If proper precautions are taken, a virtual environment is timeless and can be studied even if the real world source has been lost. Aspects of a virtual environment can be analyzed and measured easily. Modifications can be done and tested with little effort. These are vital processes in fields like architecture, archaeology and in industrial context. Also in the computer games and movie industry surface reconstruction becomes more and more important. Other applications are crime scene investigation or computer vision for robot navigation.

Reconstruction of an objects surface just from 3D points is surprisingly difficult. Even human need a huge "database" of known objects and related perceived images stored in a complex neural network which is trained over years to achieve this task — and although it works we do not have any idea of how this is being accomplished.

To technically retrieve a three dimensional environment or just one object, a stereo camera set or a 3D laser scanner is used resulting in a 3D point cloud. With this point cloud as a representation of a surface questions concerning lighting, volume, insight or outsight or collisions can not sufficiently be answered. The cloud of zero-dimensional points does not contain any information about the underlying surface as such. Additional assumptions must be drawn from general properties of general surfaces, and since the number of possible

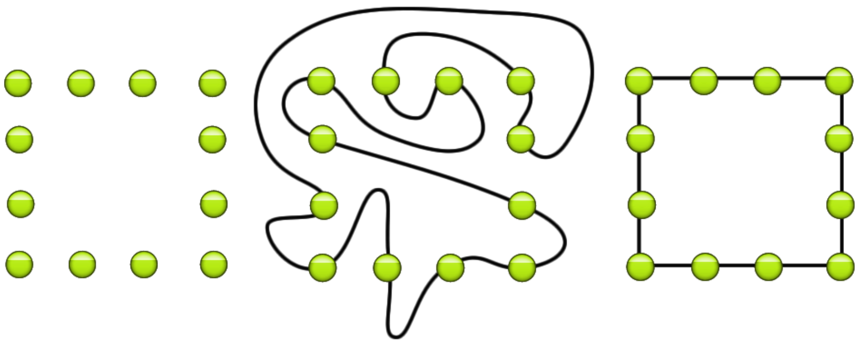

Figure 1: On the left, points originating from a connected but unknown line. In the middle, one assumption for a connected line which is valid but which is not likely if we assume an underlying "usual surface". The geometry on the right seems to be a better approximation since common rules for points on surfaces are applied to create it.

surfaces is not limited, the challenge of surface reconstruction can also be described by finding the most likely surface (see Fig. 1) which matches the point set at its best.

In this paper we concentrate on reconstruction approaches which are suitable for scanning large environments in contrast to methods which are only capable of registering single separate, closed objects. We suggest the SGC meshing approach [AB10b] for this task since from our point of view it is the most effective way to robustly retrieve virtual from real environments. This paper focuses on the abilities of arbitrary surface reconstruction algorithms for the reconstruction of real environments and presents a detailed analysis of the capabilities of the SGC algorithm specifically towards this task and points out the meaning of certain qualities in that context.

In the following we give an overview of classical methods and discuss the SGC meshing approach for real-time or Virtual Reality applications.

## 2 The Challenge of Meshing Large Scenes

Scanning scenes which are not limited to few convex objects like many reconstruction approaches require expose several challenges.

- Multiple scans are needed and the distance of the surface to the scanner varies. Both facts cause anisotropic point densities of the resulting point cloud. Holes can usually not be avoided. I.e., the roof of a house is often missing in the point cloud.

- Outliers arising from objects which are moving through the scene during the scanning process must be recognized as such.

- Vegetation moved by the wind does not have a well defined surface. Lawn for example just produces a thick layer of noise in a scan.

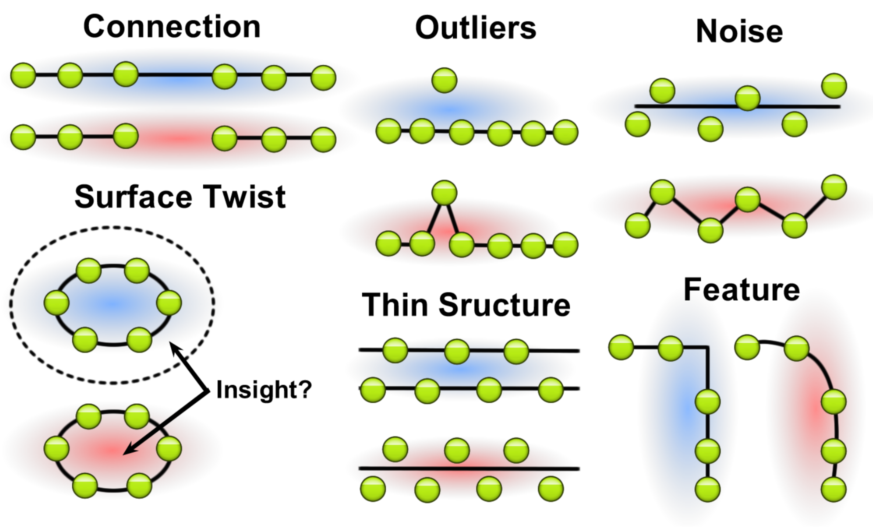


Figure 2: Examples of ambiguous surface reconstruktion cases. These six pictures show a point set twice, which has reconstructed with two different surface topologies.

- In contrast to compact closed objects, large arbitrary areas do not define a reasonable surface orientation. The relations "inside" and "outside" can not clearly be determined which is required in many reconstruction algorithms
- Generally, huge point sets are needed to retrieve large areas since there is more information than usual from the underlying object to be processed. Reconstruction algorithms which are capable of handling billions of 3D points are nearly not known.
- Objects are not necessarily connected, instead there often exist many separate objects to be reconstructed which the most available approaches are not able to accomplish.

See Fig. 2 for some examples how surfaces can be misinterpreted due to sampling artefacts.

## 2.1 Previous Approaches Incapable of Environmental Surface Reconstruction

A huge class of algorithms build a Delaunay tetrahedralization in a first step like [Boi84, EM92, KSO04] or a Voronoi diagram like [ACK01, DG06] of the point cloud. These algorithms have several disadvantages. First, they need to consider all points at once, second they produce an interpolation of the input points such that the mesh can not be scaled in size, third, they are very sensitive to anisotropic point densities and noise and, fourth, most of them are based on the hypothesis that the point cloud comprise a volume.

Some of these drawbacks can be attenuated, but nevertheless this family of algorithms is not adequate for environment reconstruction. One of the most sophisticated implicit surface reconstruction algorithm is based on the Poisson equation [KBH06]. It shows very robust behavior when dealing with noise and outliers. Since it optimizes a function that defines an insight and outsight of the surface, it is strictly limited to reconstruction of closed surfaces and this makes it inadequate for scanning environmental objects.

Another class of algorithms uses radial basis functions (RBF) to represent a surface. In these approaches RBF are placed consecutively inside the point cloud until some closeness criterion is satisfied [Mur91, CBC+01]. The resulting function is an implicit surface. These algorithm have advantages when processing noise and are able to close holes in a surface [TZCO09]. Nevertheless, they produce only closed surfaces.

Another group of algorithms is based on a so called "balloon model". These alter an initial mesh until it fills out the space comprised by the point cloud [SLS+06, CM95]. These approaches are favorable in processing noise and are capable of automatically sealing holes in the point cloud — but still, open meshes cannot be compiled by them.

### 2.2 Previous Approaches Capable of Environment Reconstruction

A well-known reconstruction algorithm is the one by [HDD+92]. Here every point is represented by a tangent plane. The tangent plane normal is calculated as the eigenvector of the covariance matrix of a defined neighborhood of $k$ points which all together define an implicit surface. The plane pieces which are computed from the local $k$-neighborhood are independent from each other, i.e., there is no global information of the surface available, and such, strong variations between neighboring faces are possible. The quality of the resulting mesh can be enhanced by mesh optimization [HDD+93] but nevertheless the problem from above cannot be solved since even orthogonal neighboring faces are accepted as correct topology.

In [Lev03] the moving least squares (MLS) approach is adapted to surface reconstruction. The surface is constructed as a continuous function that is locally adapted to minimize the least squares distances between the input points. The size of the neighborhood used for the calculation determines the smoothness of the surface. The method can handle noise very efficiently but exposes problems in case of outliers. [FCOS05] does not create a mesh but a piecewise smooth surface and the whole surface is a collection of smooth surface pieces. For rendering purposes a final polygonalization is required.

The problem with these approaches is that they might smooth out details since the processes do not adapt to local point cloud conditions. Heterogeneous noise level settings can not be set adequately for all regions of the environment.

Another problem is that the mentioned continuous functions are calculated on a local level — if the local resolution is inappropriate for a certain structure the reconstructed mesh may contain twists in the surface. A twist means that the surface orientation is not homogeneous, such that differently oriented surface areas are directly next to one another.

## 3 Smart Growing Cells Meshing

*Smart growing cells* (SGC) is a novel unsupervised learning technique from the field of *artificial neural networks* (ANN). Our proposed triangulation algorithm is derived from the SGC. In the following we describe its development in consecutive steps starting from general unsupervised learning.

### A) Unsupervised Learning

Unsupervised learning is very similar to *k-means clustering* [Mac67] which is capable of placing $k$ $n$-dimensional reference vectors in a set of $n$-dimensional input samples such that they are means of those samples which lie in the $n$-dimensional Voronoi volume of the reference vectors. Adaption of reference vectors is accomplished by randomly presenting single $n$-dimensional samples from the input sample set and moving the closest reference vector towards these samples.

### B) Self Organizing Map

The *Self Organizing Map* (SOM) [Koh82] can be seen as general unsupervised learning where reference vectors are additionally connected to each other. The learning rule from above is extended to account for the neighborhood of the closest reference vectors. The SOM iteratively adapts its internal structure — the 2D mesh — to the distribution of samples. Thus, besides clustering and dimensionality reduction it enables mapping of the sample input for instance to a 2D space by creating a map of reference vectors.

### C) Growing Cells Structures

To a certain degree *Growing Cells Structures* (GCS) [Fri93, Fri95] may be seen as SOM which additionally is capable of growing and shrinking according to the problem under consideration which is defined by the sample distribution. This mechanism is based on a so called *resource term* contained in every reference vector and which — in the original approach — is a simple counter. It counts the reference vector being the closest vector to a certain training sample. Finally a large counter value signalizes the requirement for insertion of new reference vectors.

### D) Smart Growing Cells

Individual behavior of the cells of the GCS concept is introduced by the *Smart growing cells*. It therefore can adapt to certain phenomena within the network defined by application-focused rules. Whereas cell behaviour defined in a GCS network is limited to rules that concern Euclidian distance only, the behaviour of a smart growing cells can precisely be modeled without any restrictions concerning the underlying data.

## 3.1 The Smart Growing Cells Approach in Detail

SGC is derived from GCS by adaptations necessary to realize surface reconstruction. These modifications are discussed in the following.

To stabilize the surface movement, Laplace smoothing [Tau95] is applied to the neighborhood of the closest reference vector to avoid degenerated triangles within the surface [IJS03]. The process starts with a simple structure like a tetrahedron — with repeated refinement steps it adapts to the structure which the sample set represents (see Fig. 3).

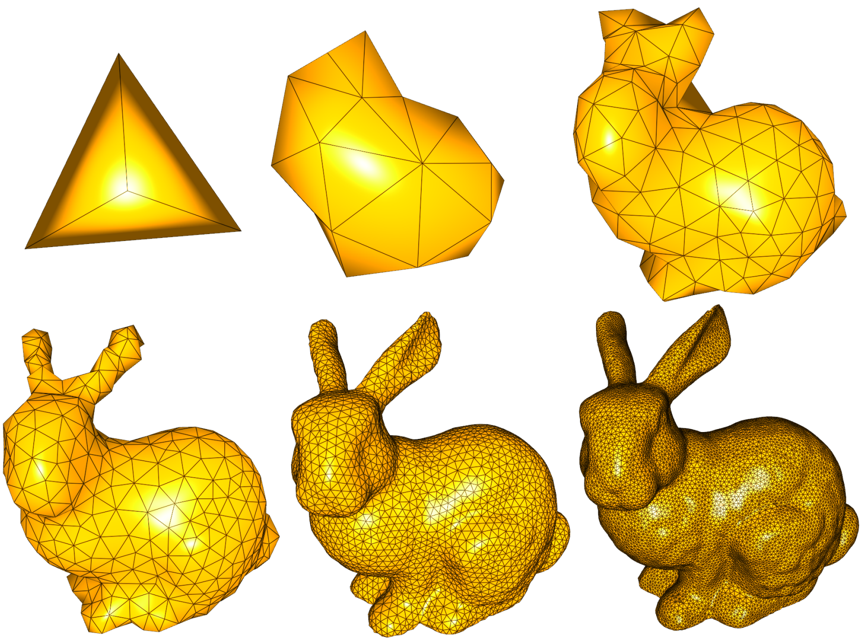

Figure 3: Different training states during reconstruction with the SGC meshing approach.

#### 3.1.1 Vertex Split

The neural vertex carrying the highest resource term is split since it has been selected by many samples and thus, creating more vertices at this location, should represent the underlying sample distribution more adequately. A vertex split operation adds three edges, two faces and a new neural vertex which is placed in the middle of the longest edge at the vertex to be split. The resource term is equally spread between the vertices (see Fig. 4).

#### 3.1.2 Edge Collapse

Neural vertices with resource terms below a certain threshold $r_{min}$ are removed together with three edges and two connected faces in a way that every neuron is an end point of ideally six edges [IJS03].

#### 3.1.3 Smart Growing Cells Features

With these modifications the GCS concept is turned into a surface reconstruction algorithm with one important feature missing — modification of the homeomorphism of the network

Figure 4: On the left, the edge collapse, on the right, the vertex split operation

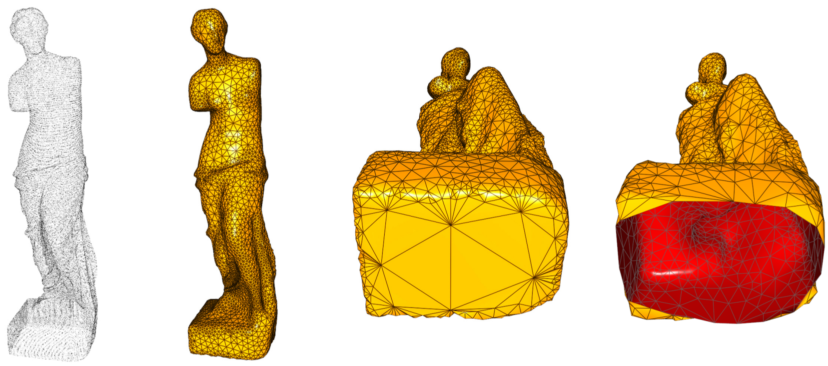

Figure 5: Statue with missing samples at the bottom (see point cloud on the left) and its reconstruction in the second picture. The third picture shows the "degenerated"mesh elements which are detected and destroyed during the SGC meshing approach resulting in a cleanly defined cut out at the empty space (see fourth picture).

structure. Thus, a mechanism for weeding and coalescing of cells is introduced in [AB10a] enabling cutting out unsuitable mesh pieces and rejoining separated elements (see Fig. 5).

### 3.2 Environment Reconstruction with Smart Growing Cells

The SGC meshing approach has a lot of advantages that specially qualify it for the reconstruction of environments.

- The algorithm directly produces a triangulated mesh as a result — no polygonalization is needed for the rendering process.
- With the SGC weeding and coalescing processes any mesh surface (homeomorphism) can by achieved.
- SGC Meshing is based on an iterative refinement process. Any produced surface is based on a preceeding surface topology guiding creation of a new one. The orientation of a surface area is based on all previous refinement levels. Even structures with complex orientations which cannot be derived from the local surface are reconstructed correctly (see Fig. 6). Nevertheless, surface twists on a global level are a remaining problem. A surface twists on a global level means that huge surface areas are oriented differently. This results from mistakes in the early stages of the reconstruction process, that can currently not be solved in the further refinement process.
- In the learning process all vertices are constantly moved to the median point of its sample subspaces. Since all vertices are situated in such a median they generally are distributed harmonically and they are connected forming a Delaunay triangulation. This is favorable for further processing. Methods like texture and bump mapping can be used to reduce the complexity of the rendering process of a scene, while its visual quality stays nearly the same (see Fig. 7). This can be very important in case of large environments and was realized only as a separate postprocess up to now [YLL+09].

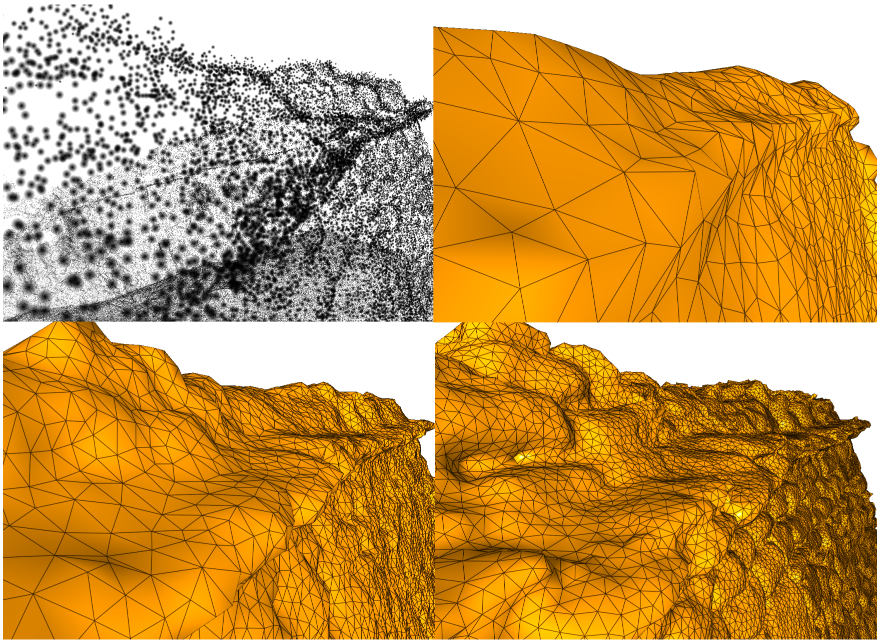

Figure 6: The protruding element of the wall can easily cause twists of the surface, which is avoided through iterative refinement within the SGC process.

- Since vertices represent sample medians, the surfaces in noisy region also are averages of the samples in that area and therefore will be adequately smooth. As shown in section 2, this is a very important property in a reconstruction application.

- Problems of anisotropic point densities are easily diluted since additional mesh elements are created independently of the certain underlying sample distribution but instead depending on the network quality (see Fig. 8).

- The approach is capable of handling arbitrary sized sample sets since it never has to see the sample set on the whole. Instead samples are drawn incrementally from the sample set, and at every state a valid solution in terms of a triangulation mesh is available. This is a vital advantage compared to the most classical reconstruction methods.

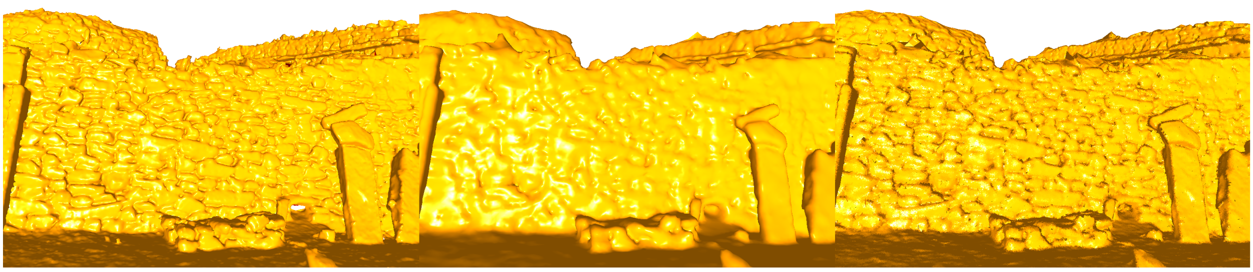

Figure 7: On the left, a high resolution mesh, in the middle, the same mesh with only 1/6 of its original triangles, on the right, again the smaller mesh but with using a bump map drawn from the high-res mesh.

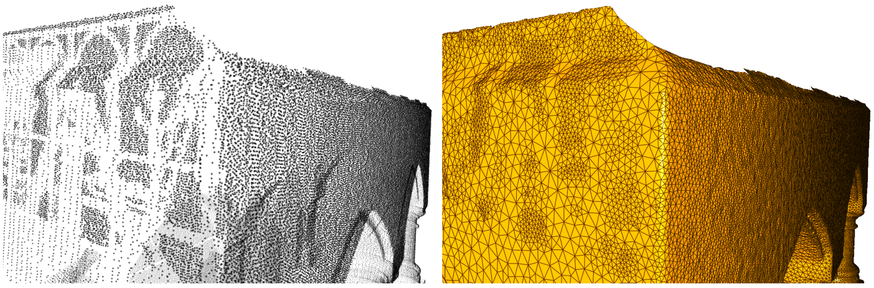

Figure 8: Distinct point densities (exposed on the left) simply lead to varying triangle sizes (see picture on the right) and do not influence accuracy or triangle quality significantly.

- Although training rules are surprisingly simple, they can achieve complex tasks. Thus, the specification of the problem under consideration is very easy — and in many cases the only possible since it frees from defining vast amounts of case discriminations and it delivers a valid algorithm in cases where algorithms simply have not been found yet. For example, see Fig. 9 — here, a variety of different heterogeneous object structures are recognized robustly which can hardly be specified from a global view on an all-embracing algorithm.

- Reconstruction with a neural network is very flexible. The training can be stopped and resumed at any time. If a mesh is not accurate enough the training can simply be extended, even by incrementally taking new samples from the scanned object. An environment can be constructed at the same time the point data is being retrieved. Representations at different training levels can also be utilized in applications where several levels of detail of the scanned object are needed

## 4 Results

Although nowadays there exits a huge amount of technical devices to retrieve a three dimensional representation of an environment, the required subsequent reconstruction process is

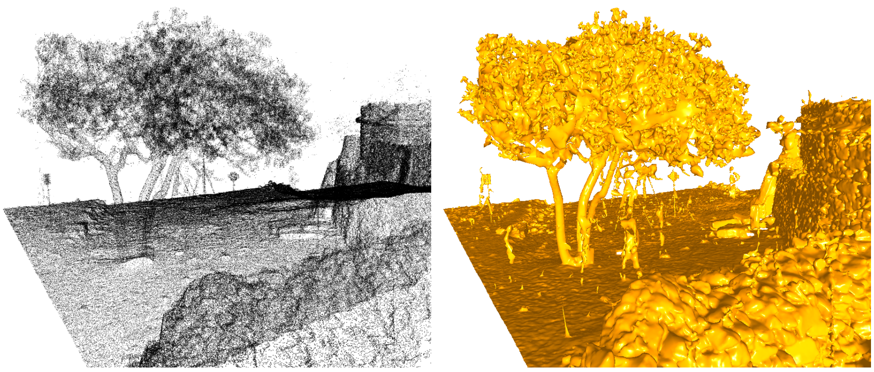

Figure 9: Without any previous knowledge or explicit rules the trees in the scanned environment were nicely been reconstructed.

still very challenging. Only a small margin of existing reconstruction techniques are capable of performing this task automatically. Common problems are anisotropic point densities, holes, outliers and noise. A meshing algorithm needs to find a consistent orientation, handle vast amounts of data, unconnected surface areas, and arbitrary complex homeomorphisms of the underlying topologies.

SGC meshing, in our opinion, is a big step towards an algorithm which solves the most of theses problems and makes retrieval of arbitrary environments robust. Thus, SGC meshing is an adequate tool for real time and Virtual Reality applications.

The approach is able to create open meshes, produces very stable results, is standing up to noise, has nearly no problems with anisotropic point densities, and can handle point clouds consisting of billions of points. The usage is very straightforward, rule sets for matching different reconstruction tasks can easily be created and modified, and the process is very robust concerning complex sample topologies (see Fig. 10). There is no need for any type of post-processing like polygonalization or re-meshing. On the downside, surface twists are the only challenge which has not been solved sufficiently yet. Fortunately they arise very rarely and can mostly be fixed by restarting the algorithm at a different random seed.

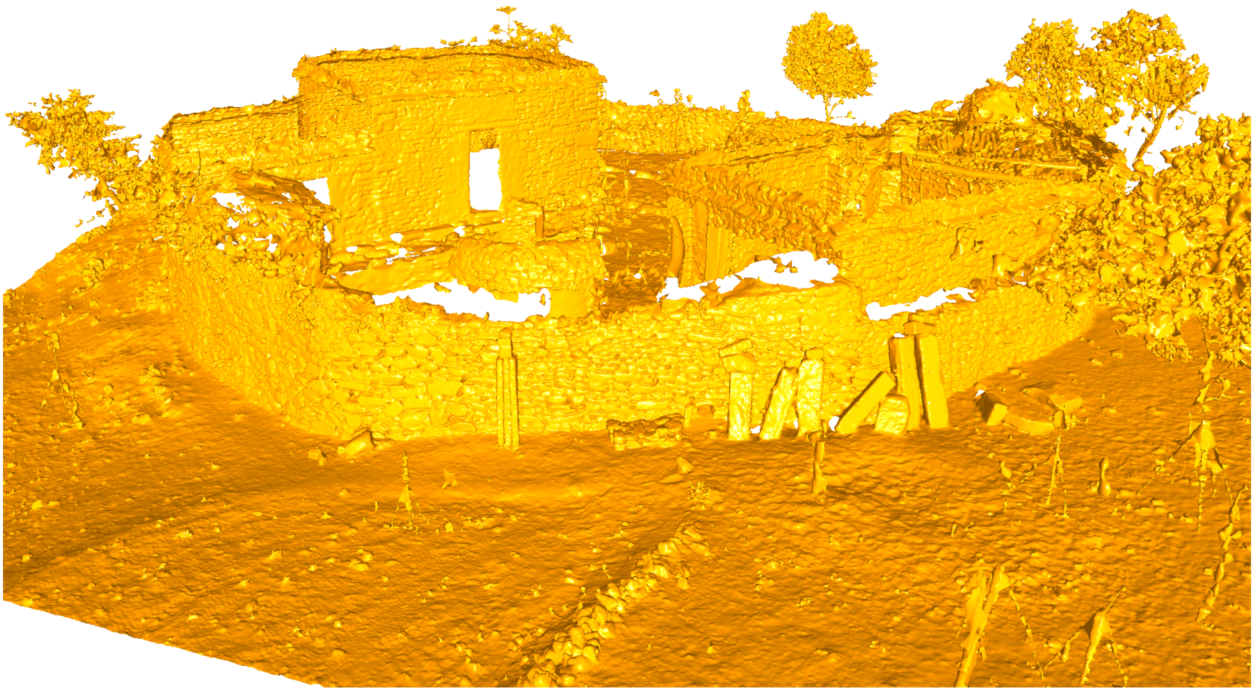

Figure 10: Example for environment reconstruction with the SGC meshing approach. The model is created from a point cloud of 9 millions of points. The resulting mesh contains about the same amount of triangles. Calculation took four hours.

## 5 Future Work

Texturing in the context of environment reconstruction is very important (see Fig. 7) — huge meshes can easily be reconstructed, but nevertheless, they can not efficiently be rendered. Automatic texturing of generated meshes will be a great issue in forthcoming work since it is required for interactive applications with huge data sets.

Another contribute to the reconstruction of environments could be a hybrid approach which combines the strengths of SGC meshing and the MLS method since, in the authors' option, these both seem to be the most promising algorithms for this task.

A great weakness of the SGC meshing approach are global twists of the resulting meshes. The authors are confident to solve this problem in future work.

**Acknowlegment:** The authors like to thank Kai Burjack for his support in rendering the presented images and the *Hafen City University*, Hamburg, in particular Tanja Graeger, for her support of laser scanned data.

## References

[AB10a] Hendrik Annuth and Christian-A. Bohn. Smart growing cells. In *Proceedings of the International Conference on Neural Computation (ICNC2010)*, pages 227–237. Science and Technology Publications, 2010.

[AB10b] Hendrik Annuth and Christian-A. Bohn. Surface reconstruction with smart growing cells. *Studies in Computational Intelligence*, 321:47–66, 2010.

[ACK01] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*, SMA '01, pages 249–266, New York, NY, USA, 2001. ACM.

[Boi84] Jean-Daniel Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graph.*, 3(4):266–286, 1984.

[CBC+01] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 67–76, New York, NY, USA, 2001. ACM.

[CM95] Yang Chen and Gérard Medioni. Description of complex objects from multiple range images using an inflating balloon model. *Comput. Vis. Image Underst.*, 61:325–334, May 1995.

[DG06] Tamal K. Dey and Samrat Goswami. Provable surface reconstruction from noisy samples. *Comput. Geom. Theory Appl.*, 35:124–141, August 2006.

[EM92] Herbert Edelsbrunner and Ernst P. Mücke. Three-dimensional alpha shapes. In *Volume Visualization*, pages 75–82, 1992.

[FCOS05] Shachar Fleishman, Daniel Cohen-Or, and Cláudio T. Silva. Robust moving least-squares fitting with sharp features. In *ACM SIGGRAPH 2005 Papers*, SIGGRAPH '05, pages 544–552, New York, NY, USA, 2005. ACM.

[Fri93] Bernd Fritzke. Growing cell structures - a self-organizing network for unsupervised and supervised learning. *Neural Networks*, 7:1441–1460, 1993.

[Fri95] Bernd Fritzke. A growing neural gas network learns topologies. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems 7*, pages 625–632. MIT Press, Cambridge MA, 1995.

[HDD+92] Hugues Hoppe, Tony DeRose, Tom Duchamp, John Alan McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In James J. Thomas, editor, *SIGGRAPH*, pages 71–78. ACM, 1992.

[HDD+93] Hugues Hoppe, Tony Derose, Tom Duchamp, John J. Mcdonald, and Werner Stuetzle. Mesh optimization. In *Annual Conference on Computer Graphics*, pages 19–26, 1993.

[IJS03] I. P. Ivrissimtzis, 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, 2003. IEEE Computer Society.

[KBH06] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, SGP '06, pages 61–70, Aire-la-Ville, Switzerland, Switzerland, 2006. Eurographics Association.

[Koh82] Teuvo Kohonen. Self-Organized Formation of Topologically Correct Feature Maps. *Biological Cybernetics*, 43:59–69, 1982.

[KSO04] Ravikrishna Kolluri, Jonathan Richard Shewchuk, and James 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.

[Lev03] D. Levin. Mesh-independent surface interpolation. *Geometric Modeling for Scientific Visualization*, pages 37–49, 2003.

[Mac67] J. B. MacQueen. Some methods for classification and analysis of multivariate observations. *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*, pages 281 – 297, 1967.

[Mur91] Shigeru Muraki. Volumetric shape description of range data using "blobby model". In *SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, pages 227–235, New York, NY, USA, 1991. ACM.

[SLS+06] Andrei Sharf, Thomas Lewiner, Ariel Shamir, Leif Kobbelt, and Daniel Cohen-Or. Competing fronts for coarse-to-fine surface reconstruction. In *Eurographics 2006 (Computer Graphics Forum)*, volume 25, pages 389–398, Vienna, october 2006. Eurographics.

[Tau95] Gabriel Taubin. A signal processing approach to fair surface design. In *SIGGRAPH*, pages 351–358, 1995.

[TZCO09] Andrea Tagliasacchi, Hao Zhang, and Daniel Cohen-Or. Curve skeleton extraction from incomplete point cloud. In *SIGGRAPH '09: ACM SIGGRAPH 2009 papers*, pages 1–9, New York, NY, USA, 2009. ACM.

[YLL+09] Dong-Ming Yan, Bruno Lévy, Yang Liu, Feng Sun, and Wenping Wang. Isotropic remeshing with fast and exact computation of restricted voronoi diagram. In *Proceedings of the Symposium on Geometry Processing*, SGP '09, pages 1445–1454, Aire-la-Ville, Switzerland, Switzerland, 2009. Eurographics Association.