Computers & Graphics 51 (2015) 190–201

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

SMI 2015

Surface-reconstructing growing neural gas: A method for online construction of textured triangle meshes

Tom Vierjahn [a,b,c,*], Klaus Hinrichs [a]

[a] Visualization and Computer Graphics Research Group, University of Muenster, Germany
[b] Virtual Reality Group, RWTH Aachen University, Seffenter Weg 23, 52074 Aachen, Germany
[c] JARA – High-Performance Computing, Aachen, Germany

ARTICLE INFO

Article history:
Received 17 March 2015
Received in revised form
11 May 2015
Accepted 12 May 2015
Available online 30 May 2015

Keywords:
Geometric algorithms
Surface fitting
Shading and texture
Neural nets
Machine learning

ABSTRACT

In this paper we propose *surface-reconstructing growing neural gas* (SGNG), a learning based artificial neural network that iteratively constructs a triangle mesh from a set of sample points lying on an object's surface. From these input points SGNG automatically approximates the shape and the topology of the original surface. It furthermore assigns suitable textures to the triangles if images of the surface are available that are registered to the points.

By expressing topological neighborhood via triangles, and by learning visibility from the input data, SGNG constructs a triangle mesh entirely during online learning and does not need any post-processing to close untriangulated holes or to assign suitable textures without occlusion artifacts. Thus, SGNG is well suited for long-running applications that require an iterative pipeline where scanning, reconstruction and visualization are executed in parallel.

Results indicate that SGNG improves upon its predecessors and achieves similar or even better performance in terms of smaller reconstruction errors and better reconstruction quality than existing state-of-the-art reconstruction algorithms. If the input points are updated repeatedly during reconstruction, SGNG performs even faster than existing techniques.

1. Introduction

Recent developments allow camera-equipped remotely piloted aircraft to be used, e.g., for reconnaissance during an emergency or disaster situation or for search and rescue missions: Aerial images are transmitted to the command and control center and are then used for mission planning. Even the relief-units can be equipped with cameras providing additional images. Three dimensional data can be extracted from these images using stereophotogrammetry. From this data a suitable surface reconstruction algorithm can create a textured 3D model of the operation area supporting decision-making and the units in the field [1,2]. For such scenarios, surface reconstruction has to provide quick initial approximations that will be continuously refined whenever new 3D points become available. Thus—especially during prolonged missions—using an incremental online algorithm is inevitable. Similar scenarios can be conceived for cultural heritage or large scale urban reconstruction where a digital image library is continuously updated either by experts or by crowd work. New images are then automatically used to extend the set of 3D points and eventually the reconstructed 3D geometry. Impressive results have been achieved using offline algorithms [3,4]. However, an incremental online algorithm will be better suited.

Many of the state-of-the-art reconstruction algorithms do not work in the desired way. Even if they can adapt to modified or extended input data they have to recreate at least parts of the model. Many require a regular sampling pattern, e.g., like in depth images, whereas unorganized point clouds as obtained by stereophotogrammetry are a way more general type of input.

In contrast to these approaches, online learning based reconstruction algorithms are generally able to adapt to any modifications to the input data while reconstructing the original surface. They can even start reconstructing as soon as the first input points are available, and they refine their results while more input points are generated. However, existing online learning based approaches require a huge number of input points and rely on post-processing steps for finalizing the mesh representing the reconstructed surface for the points seen so far. If continuous previewing is desired while acquiring data, reconstruction has to be interrupted repeatedly for costly post-processing.

In this paper we propose *surface-reconstructing growing neural gas* (SGNG), an online learning based artificial neural network that iteratively constructs a triangle mesh even from unorganized, sparse point clouds representing an object's surface. Our SGNG

* Corresponding author at: Virtual Reality Group, RWTH Aachen University, Seffenter Weg 23, 52074 Aachen, Germany. Tel.: +49 2418029716.
E-mail address: tom.vierjahn@acm.org (T. Vierjahn).

http://dx.doi.org/10.1016/j.cag.2015.05.016
0097-8493/

does not rely on any post-processing steps. Thus, at any time during the construction process a triangle mesh is available, e.g., for visualization in order to direct further data acquisition or mission planning. During execution, SGNG will adapt to any modifications of the input data. Besides approximating the shape and topology of the original surface, SGNG learns how to assign the available images as textures to the triangles of the constructed mesh. By learning visibility from the input data, instead of deriving it from the constructed mesh, SGNG does not rely on knowing occluding triangles, as previous algorithms did, and thus reduces noticeable texturing artifacts to a minimum. SGNG improves the learned texture assignment whenever new images become available. Since SGNG works incrementally, i.e., it operates at any time only locally on a single input point and its neighborhood within the mesh, it is very well suited for parallel and out-of-core implementations. However, such implementations are left for future work.

This paper gives an overview of the related work in Section 2. The predecessor of SGNG is analyzed in Section 3. SGNG surface reconstruction and texture assignment are described in Sections 4 and 5, respectively. The experiments used to evaluate SGNG and the obtained results are presented in Section 6. Finally, Section 7 concludes this paper suggesting further research directions.

2. Related work

Recent surveys give a good overview of numerous techniques for surface reconstruction from points [5,6], or for completing a surface [7]. An overview over Delaunay-based methods can be found in a survey and a monograph [8,9]. However, these methods pose strong requirements on the input points, whereas our SGNG is far less demanding.

A major group of surface reconstruction algorithms fits a distance function to the input data and triangulates its zero-set. Important approaches using signed distance functions and *Marching Cubes* [10] for triangulation include *tangent planes* [11], *volumetric range image processing* (VRIP) [12], and *Poisson surface reconstruction* (PSR) [13] that was recently refined to *screened Poisson surface reconstruction* (SPSR) [14]. Even virtual depth images [15] that provide an iterative refinement were proposed. Another group of algorithms originated from computational geometry: α-*shapes* [16] and *Ball Pivoting* [17] carve the surface from an initial volume. The related *crust* [18] and *PowerCrust* [19] algorithms fill a volume. They come alongside very elaborate theoretical guarantees and were extensively refined, e.g., [20,21]. The above algorithms require the complete input data set to be available when they start, and do thus not provide any preview during scanning. A real-time model acquisition algorithm provides preview while scanning was proposed [22], but the final mesh is still constructed offline using VRIP. *KinectFusion* [23] provides a GPU implementation of their approach. Recently a dictionary learning approach was proposed [24] that iteratively refines an initial reconstruction. However, the number of vertices in the mesh is predefined and remains fixed during learning.

Surface-reconstructing growing neural gas (SGNG) proposed in this paper uses a neural network for surface reconstruction. With a stochastic learning algorithm, such techniques are able to start reconstructing as soon as the first input points are provided. They refine their results while more input points become available. The neural networks mentioned in the following are closely related to each other. Each adds certain features to its predecessors. However, they were not initially designed for surface reconstruction. The first approach [25] that applied a *self-organizing map* [26] for surface reconstruction uses edge swap and multiresolution learning to determine suitable vertex positions and their connectivity in a mesh with predefined topology from the input points. A later approach [27] improves reconstruction quality by modifying the position and connectivity updates of *growing cell structures* [28], a neural network that adjusts the vertex density while learning. However, it relies on extensions to modify a prespecified topology [29]. Recent approaches [30,31] apply *growing neural gas* [32] for surface reconstruction. They reconstruct arbitrary topologies by *competitive Hebbian learning* (CHL) [33]. Although the original topology update rules are extended in those approaches, CHL is still unaware of the constructed surface. In conjunction with the stochastic nature of the algorithm this renders the above techniques unable to construct all triangles during online learning even for very dense point clouds. Therefore, those approaches add a post-processing step to triangulate the remaining holes that gets ineffective if the point cloud is too sparse. SGNG greatly improves upon its predecessor, the *growing self-reconstruction-map* (GSRM) [30], by introducing surface-aware topology learning, thus creating all triangles during online learning.

SGNG automatically supports texture assignment, if points are extracted from images using, for instance, *structure from motion* (SFM) [34]. Thus, SGNG allows for incremental textured preview while acquiring data. Existing techniques, e.g., [4], that use SFM to create reconstructions from large photo collections might benefit from integrating SGNG turning them into iterative online algorithms.

SGNG is tested against a publicly available benchmark [35]. It is furthermore evaluated using point clouds created from photos by a combination of VisualSFM [34] and CMVS/PMVS [36]. It is compared to SPSR [14], a publicly available state-of-the-art reconstruction algorithm. SGNG does not yet use specialized texture leveling, blending or warping [37–39]. Nevertheless, such techniques can be integrated in a straightforward way into the pipeline proposed in this paper. Texture assignment in SGNG resolves occlusion artifacts even in situations where potential occluders are not represented by the input points at all.

A first prototype implementation of some of the concepts presented in this paper has been sketched by the authors [40]. However, only a rudimentary outline of SGNG reconstruction was presented without any of the improvements that were added recently and which are included in this paper. Furthermore, an extensive evaluation was missing from that previous sketch.

3. Growing self-reconstruction map

In order to make this paper self-contained, the original reconstruction algorithm of a *growing self-reconstruction map* (GSRM) [30] that this work builds upon is outlined in this section. Afterwards, GSRM is analyzed to motivate the improvements leading to our *surface-reconstructing growing neural gas* (SGNG).

3.1. The original algorithm

After creating an initial set $\mathcal{V}$ of three random vertices, but no edges and faces, GSRM reconstruction is iterated in a loop: An input point is selected randomly from the point cloud, and the two vertices that are closest and second-closest to it are determined: $\mathbf{v}_b$ and $\mathbf{v}_c$, respectively. Afterwards, $\mathbf{v}_b$ and its directly connected neighbors are moved towards the selected input point. Then, according to *competitive Hebbian learning* (CHL) [33] $\mathbf{v}_b$ and $\mathbf{v}_c$ are connected with an edge if it does not yet exist. Only upon creation, this edge is guaranteed to be part of the edges of the Delaunay triangulation of $\mathcal{V}$ [33]. GSRM extends this step: If the two vertices are not yet connected by an edge, and if they do not share more than two common neighbors, a new edge that connects $\mathbf{v}_b$ and $\mathbf{v}_c$ is created. If they share two common neighbors that are connected by an edge, this edge is flipped in order to connect $\mathbf{v}_b$ and $\mathbf{v}_c$. If the new edge generates one or two loops of three edges, triangles are created. Additionally, GSRM checks the edges emanating from $\mathbf{v}_b$, and deletes those for which $\mathbf{v}_c$ lies in the