

Incremental Surface Extraction from Sparse Structure-from-Motion Point Clouds

Christof Hoppe¹

hoppe@icg.tugraz.at

Manfred Klopschitz²

manfred.klopschitz@siemens.com

Michael Donoser¹

michael.donosier@tugraz.at

Horst Bischof¹

bischof@icg.tugraz.at

¹ Institute for Computer Graphics and Vision,
Graz University of Technology,
Graz, Austria

² Imaging and Computer Vision
Research Group Video Analytics
Corporate Technology
Siemens AG Austria, Graz

In this paper we propose a new method to incrementally extract a surface from a consecutively growing Structure-from-Motion (SfM) point cloud in real-time. Our method is based on a Delaunay triangulation (DT) on the 3D points. The core idea is to robustly label all tetrahedra into free- and occupied space using a random field formulation and to extract the surface as the interface between differently labeled tetrahedra.

For this reason, we propose a new energy function that achieves the same accuracy as state-of-the-art methods but reduces the computational effort significantly. Furthermore, our new formulation allows us to extract the surface in an incremental manner, i. e. whenever the point cloud is updated we adapt our energy function. Instead of minimizing the updated energy with a standard graph cut, we employ the dynamic graph cut of Kohli et al. [1] which enables efficient minimization of a series of similar random fields by re-using the previous solution. In such a way we are able to extract the surface from an increasingly growing point cloud nearly independent of the overall scene size.

Energy Function for Surface Extraction Our method formulates surface extraction as a binary labeling problem, with the goal of assigning each tetrahedron either a *free* or *occupied* label. For this reason, we model the probabilities that a tetrahedron is free- or occupied space analyzing the set of rays that connect all 3D points to image features. Following the idea of the truncated signed distance function (TSDF), which is known from voxel-based surface reconstructions, a tetrahedron in front of a 3D point X has a high probability to be free space, whereas a tetrahedron behind X is presumably occupied space. We further assume that it is very unlikely that neighboring tetrahedra obtain different labels, except for pairs of tetrahedra that have a ray through the face connecting both. Such a labeling problem can be elegantly formulated as a pairwise random field and since our priors are submodular, we can efficiently find a global optimal labeling solution e. g. using graph cuts. In contrast to existing methods like [2], our energy depends only on the visibility information that is directly connected to the four 3D points that span the tetrahedra V_i . Hence a modification of the tetrahedral structure by inserting new points has only limited effect on the energy function. This property enables us to easily adopt the energy function to a modified tetrahedral structure.

Incremental Surface Extraction To enable efficient incremental surface reconstruction, our method has to consecutively integrate new scene information (3D points as well as visibility information) in the energy function and to minimize the modified energy efficiently.

Integrating new visibility information, i. e. adding rays for newly available 3D points, affects only those terms of the energy function that relate to the 3D point, which are typically only a few. Integrating additional 3D points is more complicated, since it affects the Delaunay triangulation, i. e. new tetrahedra are created and others are deleted. For the update of the energy this implies that certain terms are deleted and new terms are added. Since our energy depends only on rays that are directly connected to a certain tetrahedron, the new costs can be calculated efficiently.

The consecutive integration of additional information leads to a series of energies that typically only differ by a few terms. Instead of solving the labeling each time from scratch using a standard graph cut, we employ the dynamic graph cut proposed by Kohli et al. [1]. In our experiments, we show that the complexity of the minimization does not depend on the overall number of terms but on the number of terms that differ between two subsequent energies. Hence, our surface extraction algorithm is nearly independent of the overall scene size.




Figure 1: Runtime analysis for integrating 1000 new 3D points per step. Although the mesh gets larger, computation time remains independent of the overall scene size.

Experiments In the experimental evaluation, we first compare our approach to the state-of-the-art in static surface extraction from sparse point clouds and demonstrate competitive performance while reducing the computational effort by more than 50%. We furthermore demonstrate that our method is able to incrementally extract the surface from a growing sparse point cloud robustly in real-time while being nearly independent of the overall scene size (Figure 1). Our approach used in an incremental manner requires on average 440 ms to integrate 1000 new 3D points into the existing surface which is up to 20 times faster than the state-of-the-art. Figures 2 and 3 illustrate how the mesh evolves over time for two different point clouds. Our promising prove the applicability of our method for efficiently obtaining surface meshes from consecutively growing sparse point clouds, as e. g. might be provided by a keyframe-based SLAM method.




Figure 2: Surface evolution over time. The gray part of the mesh has been extracted from 40 000 3D points. The red triangles are created by integrating 20 000 further points.




Figure 3: (a) Input image for SfM (b) Surface mesh after integration of 1000 3D points (c) Final surface mesh

- [1] P. Kohli and P.H.S. Torr. Dynamic graph cuts for efficient inference in markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 29(12):2079–2088, December 2007.
- [2] P. Labatut, J.P. Pons, and R. Keriven. Efficient multi-view reconstruction of large-scale scenes using interest points, delaunay triangulation and graph cuts. In *International Conference on Computer Vision (ICCV)*, 2007.