Journal of Computational and Applied Mathematics 289 (2015) 225–240

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

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

Anisotropic and feature sensitive triangular remeshing using normal lifting

Vincent Nivoliers [a,b,*], Bruno Lévy [c], Christophe Geuzaine [b]

[a] Université Claude Bernard Lyon1, LIRIS UMR 5205, R3AM Group, France
[b] Université de Liège, Montefiore Institute, ACE Group, Belgium
[c] Loria/Inria Nancy Grand Est, Alice Project Team, France

ARTICLE INFO

Article history:
Received 31 August 2014
Received in revised form 22 January 2015

MSC:
65L50

Keywords:
Geometry
Mesh generation
Restricted Voronoï diagram

ABSTRACT

This work describes an automatic method to anisotropically remesh an input bad quality mesh while preserving sharp features. We extend the method of Lévy and Bonneel (2012), based on the lifting of the input mesh in a 6D space (position and normal), and the optimization of a restricted Voronoï diagram in that space. The main advantage of this method is that it does not require any parameterization of the input geometry: the remeshing is performed globally, and triangles can overlap several input charts. We improve this work by modifying the objective function minimized in the optimization process, in order to take into account sharp features. This new formulation is a generalization of the work of Lévy and Liu (2010), which does not require any explicit tagging of the sharp features. We provide efficient formulas to compute the gradient of our objective function, thus allowing us to use a quasi-Newton solver (Liu and Nocedal, 1989) to perform the minimization.

1. Introduction

Mesh generation is a crucial step in order to perform numerical simulations on CAD objects. Our work focuses on *surface* mesh generation, sometimes also denoted as *boundary recovery*. Many techniques have already been developed for this task. Our method takes as input a 3D object provided as a bad quality mesh (like many stereolithography, or STL, meshes), and generates a new mesh with nice triangles in terms of aspect ratio. For CAD objects defined in other formats, fully automatic algorithms exist to generate such a mesh, since these are internally used for the visualization and rendering of the object. In terms of quality, most STL meshes are not suitable for numerical simulations, since their triangles may have very bad shapes. In addition, gaps and T-junctions may occur, especially when adjacent spline patches in the CAD geometry were not discretized accordingly. Our method is an extension of the work of Lévy and Bonneel [1], based on the optimization of restricted Voronoï diagrams. It can be considered as a generalization of Lloyd relaxation [2], to generate a regular sampling of an input mesh. Lévy and Bonneel [1] designed a method capable of automatically generating anisotropic elements in curved regions of the mesh. Our contribution with respect to their algorithm adds several features:

- automatic detection and preservation of sharp features;
- automatic isotropic scaling of the elements in curved regions;

* Corresponding author at: Université Claude Bernard Lyon1, LIRIS UMR 5205, R3AM Group, France.
E-mail addresses: vincent.nivoliers@univ-lyon1.fr (V. Nivoliers), bruno.levy@inria.fr (B. Lévy), cgeuzaine@ulg.ac.be (C. Geuzaine).

http://dx.doi.org/10.1016/j.cam.2015.01.041

- handling of gaps and T-junctions in the input mesh with a robust normal computation;
- new efficient formulation of the gradient of the objective function.

As compared to other state of the art methods, we emphasize that we do not use any preprocessing on the input mesh, and do not require any parameterization of the object. Our method can therefore be applied in cases when other fast and reliable meshing algorithms are not able to handle the input. We however do not prove guarantees on the quality of the elements upon termination.

A full summary of mesh generation techniques is outside the scope of this paper, and several overviews already exist on the field [3,4]. We here review several techniques related to our approach.

1.1. Advancing front algorithms

Advancing front algorithms are very fast and reliable to generate surface and volume meshes. For surface remeshing, the method requires a parameterization of the surface in one or several user defined patches. The boundary of the patches is first sampled, and from this boundary, nicely shaped triangles are generated one by one towards the interior of the patch, given a prescribed edge length. When the meshed region starts to overlap itself, a final procedure is triggered to sew the front curve and close the mesh. Recent advances consider an additional metric field in the parametric space to define how the edge lengths of the generated triangles are computed [5,6], and the *gradation* control, avoiding abrupt transitions between large and small elements [7].

These methods are very efficient, but require the input object to be split into parametric patches. This is sometimes not a limitation since CAD objects are often defined as a set of parametric spline patches. However the mesh is generated per patch, which can be problematic when many patches are used to describe the shape of the object in complicated regions. Marcum [8] propose a volumetric method to merge multiple patches with a unique parameterization while Attene et al. [9] define an advancing front method robust to the quality of the parameterization. Marchandise et al. [10] especially target the remeshing of STL meshes, and study various automatic parameterization techniques. While automatic parameterization methods exist, the topology of the input mesh needs to be properly defined, and gaps cannot be handled. In addition, the parameterization is stored and thus sampled on the vertices of a background mesh. While this is not a problem for dense STL meshes obtained from 3D scanning devices or extracted from voxel data, the STL triangulation automatically generated from CAD data is usually too bad for automatic parameterization.

1.2. Mesh adaptation

An other family of remeshing methods is based on the progressive modification of a bad input mesh. These methods are therefore directly applicable to STL meshes, but require a preprocessing of the input mesh to be able to handle gaps and connectivity problems. The input mesh is generally modified via edge splitting [11], collapsing or swapping, and vertex relocation [12]. Once all the triangles respect some desired quality criterion, a final smoothing step is applied, relocating the vertices. Rassineux et al. [13] study various configurations of edge splitting to simultaneously adapt the surface triangular mesh and the volume tetrahedral mesh. With respect to these methods, our approach replaces the introduction of new mesh vertices using edge splits by a random uniform sampling of the mesh, and our objective function minimization behaves similarly to the final vertex smoothing step.

1.3. Delaunay refinement and optimization

Delaunay refinement techniques are closer to our approach, since the connectivity of the final mesh is obtained via a Delaunay triangulation and is therefore automatically handled through the vertex insertion phase. Borouchaki et al. [14] combine the Delaunay triangulation with an advancing front method for vertex insertion given a certain metric. The connectivity of the mesh is therefore automatically updated during the algorithm. Dey and Ray [15] follow the approach of mesh adaptation by inserting vertices one by one in an existing mesh, updating the connectivity using a restricted Delaunay triangulation. The authors prove guarantees on the quality of the output mesh. Both methods however still require a correct input connectivity to be applied to be able to either parameterize the object, or recover its topology properly.

Chen et al. [16] use the concept of *optimal Delaunay triangulation* to generate volume meshes, which is somehow dual to our approach. An objective function is defined over the restricted Delaunay triangulation of a set of samples, measuring the quality of the output elements. The optimization of the objective function leads to nicely shaped elements. Although the input requirements are as weak as ours for this method, the continuity of the objective function is only C^0, and the optimization procedure requires the use of simulated annealing, whereas our objective function is almost C^2 and can be optimized with a quasi-Newton solver. The objective function however is directly defined on the shape of the output triangles, and can be directly related to the final element quality.

2. Background

We here detail the basic tools we rely on. Our method is based on the minimization of an objective function defined on a restricted Voronoï diagram. From a random initial set of samples, this minimization produces a regular sampling of the