Available online at www.sciencedirect.com
ScienceDirect
Computers & Graphics 30 (2006) 892–902
COMPUTERS & GRAPHICS
www.elsevier.com/locate/cag

Filling holes in meshes using a mechanical model to simulate the curvature variation minimization

Jean-Philippe Pernot*, George Moraru, Philippe Véron

LSIS-UMR CNRS 6168, CER ENSAM, 2 cours des Arts et Métiers, 13617 Aix-en-Provence cedex, France

Abstract

The presence of holes in a triangle mesh is classically ascribed to the deficiencies of the point cloud acquired from a physical object to be reverse engineered. This lack of information results from both the scanning process and the object complexity. The consequences are simply not acceptable in many application domains (e.g. visualization, finite element analysis or STL prototyping). This paper addresses the way these holes can be filled in while minimizing the curvature variation between the surrounding and inserted meshes. The curvature variation is simulated by the variation between external forces applied to the nodes of a linear mechanical model coupled to the meshes. The functional to be minimized is quadratic and a set of geometric constraints can be added to further shape the inserted mesh. In addition, a complete cleaning toolbox is proposed to remove degenerated and badly oriented triangles resulting from the scanning process.

Keywords: Reverse engineering; Geometric modelling; Holes in meshes; Triangle mesh deformation; Curvature variation minimization; Linear mechanical model; Shape manipulations

1. Introduction

Reverse engineering is a powerful technique used to create a digital representation of an existing physical object. The reconstruction process starts with the acquisition of a point cloud from the outer surface of the physical object. A triangle mesh can be created from these dense unorganized data points [1–4]. B-Spline/NURBS surfaces may also be created either directly from the point cloud or from its triangulation [5–7]. More recently, the use of subdivision surfaces for surface fitting and surface reconstruction has also been explored [8–10]. Here again, most of the proposed approaches begin with a triangulation of the point cloud. As a consequence, it is crucial to find a triangle mesh that best fits the object outer surface.

Depending on both the complexity of the object to be reverse engineered and the adopted data acquisition system technology (e.g. coordinate measuring machines or laser scanning), some areas of the object outer surface may never be accessible. Fig. 1 shows the result of an acquisition with a laser from a single point of view. Using the point of view (a), the scanner is unable to reach some portions of the object visible from another point of view (b). This induces some deficiencies in the point cloud and a set of holes in the triangle mesh. This is not acceptable. The presence of undesired holes may induce unexpected results when doing rapid prototyping or finite element analysis for example.

In this paper, we propose a set of models, methods and tools to fill in undesired holes in meshes. The filling process acts in several steps illustrated on the basic example of Fig. 2. First, the hole contour is identified (a) and cleaned (b) to remove badly oriented and degenerated triangles due to the scanner noise (Section 4). A topological grid is then inserted (c) to fill in the hole (Section 5). Finally, the inserted mesh is deformed (d) to satisfy various blending criteria with the initial mesh (Section 6). We are notably able to minimize the curvature variation across the hole contour while solving a linear equations system. The curvature variation is simulated by the variation between

---

*Corresponding author.

E-mail addresses: jean-philippe.pernot@aix.ensam.fr (J.-P. Pernot), george.moraru@aix.ensam.fr (G. Moraru), philippe.veron@aix.ensam.fr (P. Véron).

0097-8493/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cag.2006.08.020

external forces applied to the nodes of a linear mechanical model coupled to the meshes. The functional to be minimized is quadratic and a set of geometric constraints can be added to further shape the inserted mesh. These are strong improvements with respect to the existing techniques found in the literature (Section 2). The filling process is illustrated with results produced by our prototype software (Section 7). The limits and possible upgrades are finally discussed in Section 8.

2. Related work

Various techniques have been proposed to fill in undesired holes in meshes. Two main categories can be distinguished: the geometric and non-geometric approaches.

Among the non-geometric approaches, Curless and Levoy [11] use a volumetric representation to detect the mesh areas that have to be filled in. Davis et al. [12] apply a volumetric diffusion process to extend a signed distance function through this volumetric representation until its zero set bridges whatever holes may be present. This iterative approach is particularly well adapted for complex geometrical and topological holes. Unfortunately, it does not ensure that the inserted mesh smoothly vanish on the surrounding mesh. A similar approach has been developed by Nooruddin and Turk [13] for the simplification and the repairing of polygonal meshes. Verdera et al. [14] also represent the surface of interest in implicit form, and fill in the holes with a system of geometric partial differential equations derived from image inpainting algorithms. These equations are based on the geometric characteristics of the known mesh (e.g. the curvatures) and are applied only at the holes and a neighbourhood of them. Being these equations anisotropic and geometry based, they lead to a slightly slower algorithm than the one of Davis et al. [12]. Clarenz et al. [15] use a finite element method to minimize the integral of the squared mean curvature (the so-called Willmore energy) of the filled hole. Their process is iterative and can only ensure a tangency continuity with the surrounding mesh.

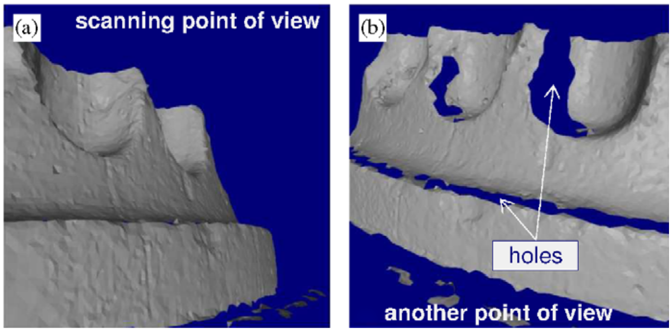

Fig. 1. Example of holes (b) resulting from the acquisition of an object from a single point of view (a).

Considering the geometric approaches, Liepa [16] proposes a filling process quite similar to ours (Fig. 2). The hole is first detected and filled in with a minimum area triangulation of its 3D contour [17]. The triangulation is refined so that the triangle density agrees with the density of the surrounding mesh triangles [18]. The filled hole is finally smooth with a fairing technique based on an umbrella operator [19]. The faired mesh is obtained while solving a system of linear equations. It solely satisfies tangency blending conditions with the initial mesh. Schneider and Kobbelt [20] propose a fairing technique based on solving a non-linear fourth order partial differential equation. G^1 boundary conditions are satisfied but the resolution is iterative. Also the implicit fairing of Desburn et al. [21] requires an iterative resolution process. Other approaches are based on the moving least-squares projection which induces a non-linear minimization process [22,23].

To conclude, the main drawbacks of most of these approaches concern the use of iterative resolution processes to find the shape of the inner mesh and/or the non-satisfaction of curvature blending conditions between the inner and surrounding meshes. In fact, most of these approaches do not use enough geometric information available on the surrounding mesh. This is not true for the context-based surface completion algorithm of Sharf et al. [24] where the inner mesh is defined according to shapes that can be very far from the hole. But their process may lead to very unexpected results. Finally, most of the processes do not enable the prescription of additional constraints inside the inner mesh. These are the limits we try to overcome. We will discuss how successful we are in Section 8.

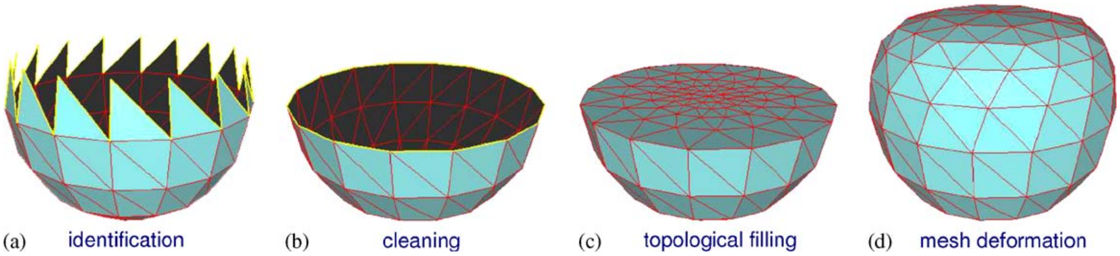

Fig. 2. Overall filling process on a simple example: after the identification and cleaning of the hole contour, the missing area is filled in with a topological grid whose final shape results from a deformation process.