Pattern Recognition Letters 30 (2009) 1198–1207

Contents lists available at ScienceDirect

Pattern Recognition Letters

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

Shape from silhouette using topology-adaptive mesh deformation ☆

Y. Yemez *, Y. Sahillioğlu

Multimedia, Vision and Graphics Laboratory, College of Engineering, Koç University, Sarıyer, Istanbul 34450, Turkey

ARTICLE INFO

Article history:
Received 20 November 2008
Received in revised form 22 May 2009
Available online 21 June 2009

Communicated by M. Couprie

Keywords:
Shape from silhouette
Surface reconstruction
Mesh deformation
Robust topology control

ABSTRACT

We present a computationally efficient and robust shape from silhouette method based on topology-adaptive mesh deformation, which can produce accurate, smooth, and topologically consistent 3D mesh models of complex real objects. The deformation scheme is based on the conventional snake model coupled with local mesh transform operations that control the resolution and uniformity of the deformable mesh. Based on minimum and maximum edge length constraints imposed on the mesh, we describe a fast collision detection method which is crucial for computational efficiency of the reconstruction process. The topology of the deformable mesh, which is initially zero genus, can be modified whenever necessary by merging operations in a controlled and robust manner by exploiting the topology information available in the silhouette images. The performance of the proposed shape from silhouette technique is demonstrated on several real objects.

1. Introduction

The shape from silhouette technique has recently regained interest since its robust output forms a solid initial model for further tasks of computer vision. Early examples of this technique were presented in (Chien and Aggarwal, 1986; Szeliski, 1993) and later much improvement has been established concerning accuracy and efficiency issues (Niem and Wingbermuhle, 1999; Tarini et al., 2002; Boyer and Franco, 2003; Fang et al., 2003; Yemez and Schmitt, 2004; Erol et al., 2005). Silhouette models cannot capture hidden surface concavities, but can further be carved or deformed so as to achieve a more accurate object representation by incorporating stereo or optical triangulation information (Esteban and Schmitt, 2004; Yemez and Wetherilt, 2007). Moreover, the shape from silhouette technique, as a passive reconstruction method, can successfully be used to model time-varying scenes and objects (Mueller et al., 2004; Carranza et al., 2003; Bilir and Yemez, 2008) whereas most of the active scene capture technologies become inapplicable in the dynamic case.

Shape from silhouette techniques existing in the literature commonly make use of an intermediate volumetric representation, e.g., in terms of cubic voxels or tetrahedra, which is then converted to a surface mesh representation by means of a triangulation method such as the marching cubes algorithm (Newman and Yi, 2006) or Delaunay triangulation (Boyer and Franco, 2003). These techniques can be computationally very efficient and accurate, and can generate topologically consistent (i.e., manifold) mesh representations. However, triangulation methods that they use rely only on local surface information and hence they always have to deal with topological ambiguities. Adhering to very high resolution representations, which can be obtained in an efficient manner by adaptive sampling strategies as in (Erol et al., 2005), may alleviate but cannot totally eliminate the topological ambiguity problem. Since these techniques do not have an explicit or global control on topology, they can easily yield representations which are topologically incorrect, i.e., which do not have the same topology as the object to be reconstructed. Moreover, since volumetric visual hull techniques are commonly based on the idea of space carving with no inherent mechanism of smoothing, their reconstructions, especially at high resolutions, are very sensitive to imperfections of the silhouette extraction and camera calibration processes, and may hence severely suffer from geometric and topological distortions.

Deformable mesh models, which in general yield smooth and robust representations, were successfully used before for the problem of 3D shape recovery in different contexts such as shape from range data, shape from stereo and segmentation of volumetric images (Terzopoulos et al., 1991; Park et al., 2001; Lachaud and Taton, 2003; Esteban and Schmitt, 2004; Duan et al., 2004). However, the use of mesh deformation, which we advocate in this paper, has not been considered so far for the shape from silhouette problem. The reason seems to be twofold. First, mesh deformation methods are considered as computationally inefficient as compared to visual hull techniques based on space carving. Second and more

---

☆ This work has been supported by TUBITAK under the project EEEAG-105E143 and by the European Network of Excellence 3DTV (www.3dtv-research.org).

* Corresponding author. Tel.: +90 212 3381585; fax: +90 212 3381548.
E-mail addresses: yyemez@ku.edu.tr (Y. Yemez), ysahillioglu@ku.edu.tr (Y. Sahillioğlu).

0167-8655/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.patrec.2009.05.012

importantly, it is a challenging problem for deformable mesh models to recover the shape of an object with arbitrary genus based on only silhouette information. The primary contribution of this paper is a shape from silhouette technique based on mesh deformation, which addresses these two challenges.

The deformable model that we employ is based on the deformation scheme proposed in (Lachaud and Taton, 2003) for 3D segmentation of complex anatomical structures. In this scheme, an initial mesh model is iteratively deformed towards the target boundary by external and internal forces. The resolution and structural uniformity of the deformable mesh are controlled during surface evolution via edge collapse, split and flip operations. The most problematic issue in this method, as also in most mesh-based deformation schemes, concerns the topology control which enables recovery of the shapes with arbitrary genus. The common practice is to incorporate topology modifying operators (splitting and merging) into the deformation scheme, such as in (Lachaud and Taton, 2003and Duan et al., 2004), whenever a collision is detected between distant parts of the deformable surface. However, this strategy relies on a false assumption which may easily lead to incorrect and/or redundant topology modifications. Collisions may indeed occur due to the limited flexibility of the deformable model as well as due to a need for topology change. The deformable model may not be fast or flexible enough to penetrate into deep concavities at a given resolution and even partially get stuck at some part of the surface at some iteration of the evolution while distant surface parts of the model keep getting close to each other, yielding unexpected collisions. The challenge is to be able to correctly decide whether a detected collision requires modification of the topology or it is simply a misguided self-intersection of the deformable model.

In view of the above discussion, the contributions of this paper can be summarized as follows:

- We adopt the deformable mesh described in (Lachaud and Taton, 2003) to address the shape from silhouette problem. We define an isolevel function based on bilinear interpolation of the silhouette images to be used as the external force of the deformable model along with a fine-tuning strategy that accurately places the model onto the boundary surface. This isolevel function is also used to assess the geometrical accuracy of the deformable mesh to support adaptive resolution. The resulting reconstructions are accurate and always smooth even when the silhouettes are corrupted with noise.
- We propose a robust topology control mechanism. The reconstruction scheme starts with an initial deformable mesh of genus zero. Topological changes are introduced during deformation whenever necessary in a controlled manner based on the topology of silhouette images. Our scheme also allows an optional user interaction mechanism to improve robustness of the topology control when modeling geometrically very complicated objects. Our method does not guarantee topological correctness for all possible types of surfaces but can ensure that the topology is correctly recovered for most of the typical real-world objects with arbitrary genus.
- We describe a fast and effective method for collision detection and handling, based on the minimum and maximum edge length constraints imposed on the deformable mesh, which is crucial for computational efficiency of the deformation scheme.

In Section 2, we describe the surface deformation method that we employ and then in Section 3, we address the shape from silhouette problem. Our topology control mechanism is described in Section 4 along with the overall reconstruction algorithm. Section 5 presents the experimental results and Section 6 finally provides concluding remarks and perspectives for future research.

2. Deformation algorithm

The deformation method that we use is based on the active contour models, or so called "snakes", which were first developed in (Kass et al., 1988). The method basically follows the Lagrangian approach: An initial deformable mesh S_0 representing the bounding sphere is iteratively evolved towards the boundary of the object under the guidance of internal and external forces that try to minimize an overall energy. This surface evolution can be expressed by the following equation:

$$S_k = S_{k-1} + \boldsymbol{F}_{\text{int}}(S_{k-1}) + \boldsymbol{F}_{\text{ext}}(S_{k-1}), \tag{1}$$

where $\boldsymbol{F}_{\text{int}}$ and $\boldsymbol{F}_{\text{ext}}$ denote the internal and external forces. By iterating the above equation, the surface S_k converges to its optimum at the equilibrium condition when all the forces cancel out to 0. The external force component, $\boldsymbol{F}_{\text{ext}}$, is application-specific and commonly set to be in the direction of the surface normal:

$$\boldsymbol{F}_{\text{ext}}(P) = v(P) \cdot \boldsymbol{N}(P), \tag{2}$$

where $\boldsymbol{N}(P)$ is the normal vector and $v(P)$ is the force strength at vertex P of the deformable mesh.

The internal force component, $\boldsymbol{F}_{\text{int}}$, controls the smoothness of the mesh. We employ a combination of tangential Laplacian operator (Wood et al., 2000) and Taubin's fairing technique (Taubin, 1995), which is easy-to-compute, and yields a smooth deformation with no significant geometrical shrinkage and bias in the final surface estimate. The internal force $\boldsymbol{F}_{\text{int}}(P)$ is defined in terms of its components along two orthogonal directions,

$$\boldsymbol{F}_{\text{int}}(P) = \boldsymbol{F}_{\text{T}}(P) + \boldsymbol{F}_{\text{N}}(P), \tag{3}$$

where the components $\boldsymbol{F}_{\text{T}}$ and $\boldsymbol{F}_{\text{N}}$ correspond to smoothing along tangential and normal directions of the surface, respectively. We obtain the tangential component, $\boldsymbol{F}_{\text{T}}$, by tangential Laplacian smoothing:

$$\boldsymbol{F}_{\text{T}}(P) = \mathscr{L}(P) - (\mathscr{L}(P) \cdot \boldsymbol{N})\boldsymbol{N}, \tag{4}$$

where $\mathscr{L}(P)$ denotes the Laplacian operator that moves a vertex to the centroid of its one-ring neighbors. The component $\boldsymbol{F}_{\text{N}}$ is obtained by fairing the surface along its normal direction:

$$\boldsymbol{F}_{\text{N}}(P) = (\mathscr{F}(P) \cdot \boldsymbol{N})\boldsymbol{N}, \tag{5}$$

where $\mathscr{F}(P)$ is the displacement created by the non-shrinking surface fairing algorithm described in (Taubin, 1995).

Following (Lachaud and Taton, 2003), we incorporate three local mesh restructuring operators (Kobbelt et al., 2000), namely edge collapse, edge split and edge flip, into the surface deformation process, to handle degenerate edges and irregular vertices and to impose uniformity on the mesh structure. At each iteration of the above algorithm, split operations are first applied to the deformable mesh to remove edges with length longer than $\varepsilon_{\max}$. The same process is then repeated with collapse operations to remove edges shorter than $\varepsilon_{\min}$. Next to split and collapse operations, which inevitably change the valence distribution of the mesh structure, edge flip operations are applied by swapping the common edge of any two neighboring triangles with the one joining the unshared vertices of the triangles as long as this operation favors the existence of the vertices of valence close to 6. We set $\varepsilon_{\max} = 2\varepsilon_{\min}$ to have uniformly sized triangles with small aspect ratios as proposed in (Kobbelt et al., 2000). Since the edge length ratio is then bounded by $\varepsilon_{\max}/\varepsilon_{\min} = 2$ and the valence distribution preserves its uniformity by flip operations, the deformable mesh keeps its initial high quality in terms of the aspect ratio of the triangles during surface evolution.

For a stable mesh evolution process, it is important to comply with the minimum and maximum edge length constraints. Unfortunately, it is usually impossible to achieve this via successive