Visual Shapes of Silhouette Sets

Jean-Sébastien Franco   Marc Lapierre   Edmond Boyer

GRAVIR–INRIA Rhône-Alpes
655, Avenue de l'Europe, 38334 Saint Ismier, France
{franco,lapierre,eboyer}@inrialpes.fr

Abstract

Shape from silhouette methods are extensively used to model dynamic and non-rigid objects using binary foreground-background images. Since the problem of reconstructing shapes from silhouettes is ambiguous, a number of solutions exist and several approaches only consider the one with a maximal volume, called the visual hull. However, the visual hull is not always a good approximation of shapes, in particular when observing smooth surfaces with few cameras. In this paper, we consider instead a class of solutions to the silhouette reconstruction problem that we call visual shapes. Such a class includes the visual hull, but also better approximations of the observed shapes which can take into account local assumptions such as smoothness, among others. Our contributions with respect to existing works is first to identify silhouette consistent shapes different from the visual hull, and second to give a practical way to estimate such shapes in real time. Experiments on various sets of data including human body silhouettes are shown to illustrate the principle and the interests of visual shapes.

1. Introduction

Recovering shapes from their projected contours in a set of digital images has been a subject of interest for the last three decades in the vision and graphics communities. The main interest of these contours is that they lead to region based modeling approaches which are rapid and do not rely on only local, and sensitive, photometric consistencies between images. They are therefore used to produce models, and especially initial models, in a number of modeling systems in particular dynamic systems which consider moving objects over time. Several methods have been proposed to solve the associated reconstruction problem among which one of the most successful is the *visual hull* [1, 14]. Such an approach consists in computing the maximal volume that projects inside image contours or, in other words, onto *silhouettes*. Straightforward approaches exist to this

purpose [21, 18, 11], some of which are real time [6, 10]. While robust and easy to estimate, the visual hull is not, in general, a good geometric approximation of the observed shape. It can even be rather poor if a reduced number of views are considered. This is due to the fact that the visual hull is merely an extended bounding box, obtained by identifying the region in space where the observed shape can not be with respect to a set of silhouettes. Such a conservative approach does not report on shapes that are consistent with a given set of silhouettes, but on the union of the regions occupied by all such shapes. As a consequence, a number of viewpoints are required to refine this region and ensure that it is reasonably close to the observed object shape.

However, even a few silhouettes provide strong geometric information on shapes under little assumptions. Our intention in this paper is therefore to find better approximations of an object shape given its silhouettes while keeping the ability to model in real time. To this purpose, we introduce the *Visual Shapes* of a set of silhouettes, which are silhouette consistent shapes in the sense that their projected silhouette boundaries, with respect to given viewpoints, match the given silhouette contours. Beside the definition which helps in characterizing silhouette based models, often incorrectly considered as visual hulls in the literature, the main interest of visual shapes is to yield estimations more precise than visual hulls.

While the literature on visual hulls and their computation is vast, less efforts have been devoted to silhouette consistent shapes inside the visual hull. In [7, 22], first solutions were proposed to determine, along viewing lines, single points of contacts with the surface, under local second order assumptions. The associated approaches assume some knowledge on extremal contour connectivities, as well as simple shape topologies, but they allow smooth surfaces to be reconstructed. Our work is founded on the same observation that viewing lines along silhouette contours, and thus the visual hull surface, are tangent to the observed object surface. Following also this observation, approaches [9, 13, 5, 17] exploit the duality that exists between points and planes in 3D space, and estimate the dual

of the surface tangent planes as defined by silhouette contour points. However, these approaches do not account for the fact that surface points lie on known viewing lines, in known intervals, and suffer therefore from various singularities. Also the visual shapes represent a more general concept since a family of plausible shapes, including the visual hull, is defined.

In [15], the topological structure of the visual hull is made explicit in the case of smooth objects. In this work, the mesh describing the extremal contour connectivity on the object surface is called the *rim mesh* and its connection with the visual hull mesh is identified. Unfortunately, this theoretical contribution does not yield a practical method to estimate the rim mesh in general situations, in particular with shapes having complex topologies. Recently, [12] and [20] proposed approaches to estimate the rim mesh on the visual hull surface by adding a photometric consistency constraint. However the rim mesh is not always well defined due to self-occlusions and strong assumptions need to be made on the topology of the observed objects, as stated in [16].

Our strategy is different from the afore-mentioned works. We first define a family of shapes which are consistent with a given set of silhouettes, namely the visual shapes. For one set of silhouettes, the associated shapes differ then by their contact with viewing lines of silhouette contour points: from isolated points, as for extremal contours on the observed shape, to the maximal intervals of the visual hull. Visual shapes are reduced to a single element when an infinite number of viewpoints, outside the shape's convex hull, is considered. In that case visual shapes and the visual hull are equivalent to the original shape, minus its concavities. However, in the general case, additional information is required to identify a single visual shape. Several criteria can be used to that purpose. In this paper, we experiment a very general assumption of local shape smoothness which is true in most real situations. The interest is to provide an approximation of the observed shape which is better than the visual hull, while keeping its robustness advantage over most modeling approaches. Such an approximation is useful not only as a final model but also as the initial input data to several modeling applications including motion capture or model refinement.

The paper is organized as follows. In section 2, geometric entities related to visual shapes are introduced. In section 3, visual shapes are defined and illustrated. In section 4, it is explained how to compute visual shapes, and results with real data are presented, before concluding in 5.

2. Preliminaries

Suppose that a scene, containing an arbitrary number objects, is observed by a set of calibrated pinhole cameras.
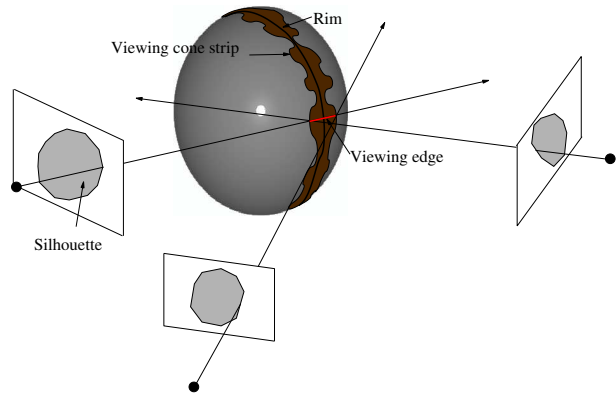

Figure 1. Viewing cone strip of a silhouette.

Suppose also that projections of objects in the images are segmented and identified as foreground. The foreground region of an image i consists then of the union of object projections in that image and, hence, may be composed of several unconnected components with non-zero genus. Each connected component is called a *silhouette*.

Consider the set of viewing rays associated with image points belonging to a single silhouette in one image. The closure of this set defines a cone in space, called *viewing cone*. The viewing cone delimiting surface is tangent to the surface of the corresponding foreground object along a curve called the *rim* (see figure 1). In what follows, we assume that a rim is formally defined as the locus of points on the object surface where viewing lines from one viewpoint are tangent to the surface.

The *visual hull* [1] is then obtained by intersecting viewing cones, possibly with respect to various image visibility domains [11]. It is a generalized polyhedron whose faces are made of cone patches, organized into strips with respect to silhouette contours.

A *viewing cone strip* corresponds then to contributions of a silhouette contour to the boundary surface of the visual hull (see figure 1). By construction, the rim associated with a silhouette contour lies inside the viewing cone strips associated to the silhouette. Observe that for non-smooth objects, the rim can become a strip itself within the viewing cone strip.

Of particular interest for this paper are *viewing edges*, corresponding to contributions of viewing rays to the visual hull surface. For one image point, such a contribution consists of one or several edges along the ray. A viewing cone strip can then be defined as the union of the viewing edges of the points on a silhouette contour. The viewing edges of an image point are easily obtained by finding silhouette contribution intervals along the point's viewing line, and then computing the common intersections of these intervals.

3. Definition

The visual hull is defined as the intersection of the viewing cones. As mentioned in the introduction, our objective is to identify a larger family of shapes associated to a given set of image silhouettes. To this purpose, we will focus on the part of the surface which is observed from silhouette contours, namely rims, and consider that shapes consistent with a set of silhouettes have rims with similar topologies. Hence, the proposed following definition:

Definition. *Let $\mathcal{S}$ be a set of scene silhouettes associated to a set of viewpoints $\mathcal{C}$. Then visual shapes $\mathcal{V}(\mathcal{S}, \mathcal{C})$ of $\mathcal{S}$ and $\mathcal{C}$ are space regions V such that:*

1. *All rim points on the surface of V, belong to viewing cone strips of $\mathcal{S}$.*

2. *All viewing cone strips of $\mathcal{S}$ are tangent to V.*

The two above constraints ensure that, first, visual shapes are consistent with the given silhouettes, and second, that inside any viewing strip there is a rim. Such a definition yields a family of shapes which are consistent with silhouettes and viewpoints, i.e. all the volumes in space for which rims project onto given silhouettes and cover all of them. Intuitively, the visual shapes $\mathcal{V}(\mathcal{S}, \mathcal{C})$ differ by the width, along viewing lines, of their rims, and identifying a single visual shape inside the solution family consists in deciding for the rim width based on *a priori* knowledge. Note that visual shapes include the visual hull as an extremal shape in the family that encloses all the others. We have then the following property:

Property . *Let $\mathcal{S}$ be a set of scene silhouettes associated to a set of viewpoints $\mathcal{C}$, then any viewing edge associated to contours points of $\mathcal{S}$ contains at least one point of any visual shapes $\mathcal{V}(\mathcal{S}, \mathcal{C})$.*

This property means that all visual shapes are tangent to the visual hull surface along viewing cone strips. In particular, we expect better approximations of smooth shapes to be shapes with a single contact point with the visual hull surface along viewing lines. Visual shapes include shapes which satisfy that constraint. This will be used when computing visual shapes as explained in the next section.

Visual shapes could also be seen as dual shapes of the visual hull, by the fact that they are shapes inside the visual hull with tangent contacts. However, the above definition is not restricted to a single shape but identifies a family of silhouette consistent shapes. Also in contrast to duality based approaches [5, 17], visual shapes are well defined shapes which do not suffer from singularities in generic situations. This is due to the fact that visual shape rim points are, by
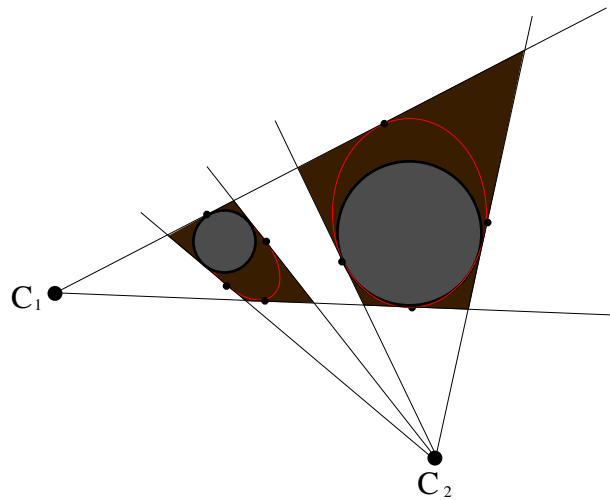

Figure 2. Cross section of a situation where 2 cameras observe 2 spheres. In brown the resulting visual hull. In red, the surface of one of the associated visual shapes with single contact points locally with the visual hull surface. Observe that, by definition, the visual shape is tangent to the visual hull surface, but that the observed objects do not necessarily satisfy that property.

definition, on the visual hull surface which is itself well defined. In duality based approaches, estimated shapes do not necessarily satisfy this containment property since shape point locations are not restricted to viewing edges, or even viewing lines, but to planes. In that sense, visual shapes use all the information provided by silhouettes. The only assumption which is made so far is that observed shapes are tangent to all visual hull faces. Even if this is not always true, as shown in figure 2, it limits the reconstruction solution space in a reasonable way when no additional information are available to decide where the matter is.

By definition, all visual shapes associated with a set of silhouettes share the same topology, that of the visual hull. Note however that the observed objects are not necessarily visual shapes of their silhouettes because of self-occlusions which can hide rim points and unoccupied visible space (see figure 2).

Figure 3 shows examples of visual shapes corresponding to silhouettes of a sphere. Sets with different numbers of silhouettes were used. The figure shows the visual hulls obtained with these sets as well as various visual shapes obtained by: (b) thinning viewing cone strips, (c) choosing a single contact point along viewing edges, and (d) estimating single contact point with local assumptions. Observe that in column (d), well delimited contours always appear on
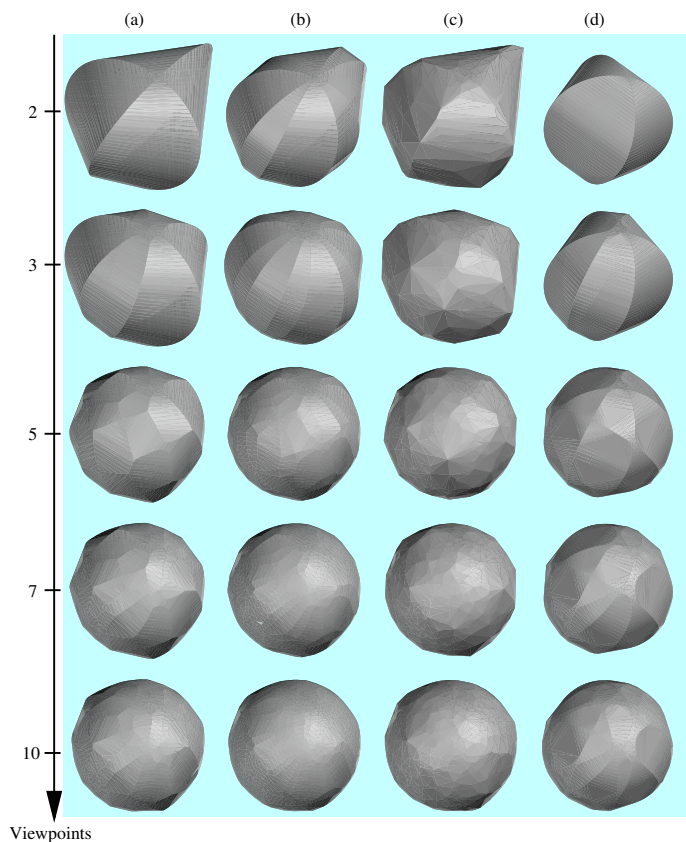
Figure 3. Visual shape examples with silhouette sets of a sphere. The top row correspond to silhouettes from 2 viewpoints, and the rows below show models obtained by progressively adding viewpoints. Columns are: (a) visual hulls; (b) visual shapes obtained by slightly thinning viewing cone strips; (c) visual shapes with one contact point, with the visual hull surface, randomly chosen inside viewing edges; (d) visual shapes with one contact point assuming the surface to be locally of order 2.

visual shapes. They correspond to the observed rims on the sphere. In that case, local assumptions about the observed surface are true, and all the estimated points inside viewing edges belong to the observed sphere. Note also that when increasing the number of views, visual shapes all converge to a single shape. With infinite viewpoints outside the scene convex hull, this *limit* shape becomes the observed shape from which concavities have been removed[1].

---

[1]This is the original definition of the visual hull by Laurentini[14]

4. Computation

In the previous section, we introduced visual shapes of a set of silhouettes. These are shapes with the same topological rims with respect to the considered viewpoints. As a consequence, visual shapes of a set of silhouettes all have contributions inside viewing edges of silhouette contour points. Thus, the computation of visual shapes consists first in identifying these contributions inside the viewing edges, and second to estimate the surface connecting these contributions. This is described in the following sections where we assume polygonal silhouette contours, as generally available in real situations.

4.1. Contributions along viewing edges

As mentioned earlier, viewing edges, or visual hull contribution intervals along viewing rays, are easily computed by intersecting ray projections with image silhouettes (see [4] for how to compute them efficiently). In figure 3, column (b) shows visual shapes obtained by thinning these viewing edges. This is a first solution, however this does not improve the estimation in a significant way with respect to the visual hull. As shown in column (c)-(d) of figure 3, a better estimation is related to the fact that viewing rays along silhouette contours only graze the surface at isolated points. This is true for smooth surfaces, but not only: even if the surface is locally planar, viewing rays will still be tangent at isolated points, except in the specific case where the viewing point belongs to the surface plane.

In the following, we thus assume a single contact point inside viewing edges. To identify the location of the contact point, different assumptions can be made. In [12] and [20], image photo-consistency assumptions are made to determine rim points inside visual hull faces. However photo-consistency applies to true surface points, and in numerous situations where self-occlusions occur there is no such point inside viewing edges, as explained before and shown for instance in figure 2. A shape estimated this way would still be a visual shape by definition, but with an unpredictable local behavior. Another possibility is to assume that the surface is locally of order 2, thus with a predictable local behavior. It is more or less the assumption made in duality based approaches [13, 5, 17] where the surface is assumed to be locally a quadric, or where finite differences are used to estimate derivatives. Our approach differs by the fact that we constrain the points we estimate to be inside well defined intervals along viewing rays, namely viewing edges. In contrast, duality based approaches estimate points dual to planes, and, importantly, can not guarantee that these points belong to the visual hull.

Another advantage of viewing edges is that they naturally define a local neighborhood through epipolar corre-
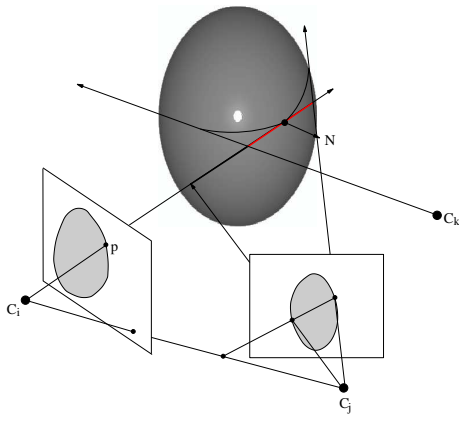
Figure 4. Viewing edges of point p in image i are delimited by viewing rays of epipolar correspondents of p. The neighborhood defined around P by these two correspondents is used to estimate local surface properties in the viewing direction. Note that when p moves along the silhouette contour, the viewing points C_j and C_k change when p reaches rim intersections on the surface.
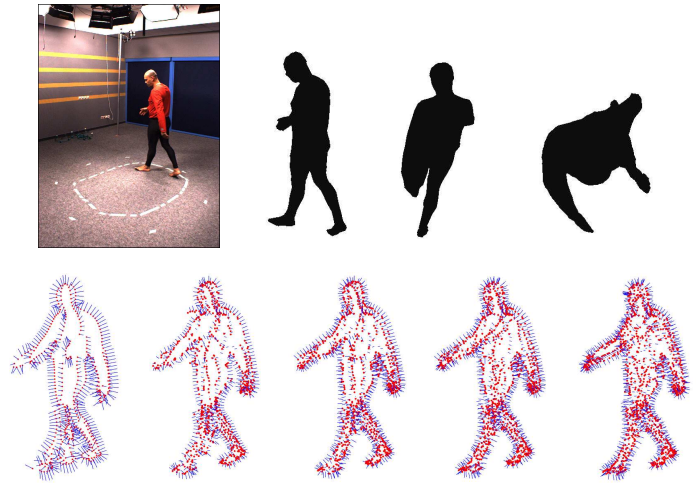

Figure 5. Visual shape points and normals under the assumption that the surface is locally of order 2. On top: one of the image used and 3 of the 6 silhouettes available. Bottom: estimated points (red) and normals (blue) with 2 (left) up to 6 (right) viewpoints.

spondences (see figure 4). Their boundary, i.e. the interval boundary points along viewing rays, identifies the epipolar correspondents, over all input silhouettes, such that the interval, where a surface point can lie, is minimal and not infinite in general[2]. Local neighborhoods defined in this manner are optimal for local estimation of surface properties. Using instead the epipolar parametrization between silhouettes, as in [7, 22, 3] and more recently in [17], does not ensure such a property since correspondences between silhouettes are imposed: points on silhouette at time t are matched with points on silhouettes at time $t \pm \epsilon$, and other silhouettes are not considered. Intervals along viewing rays defined by such correspondences can be infinite even when the visual hull is finite, hence making local surface estimations very difficult.

Each viewing edge defines a neighborhood composed of two epipolar correspondents. Thus for each viewing edge, we have three viewing rays which are locally tangent to the visual shape: 2 viewing lines from the epipolar correspondents and the viewing line supporting the viewing edge. From these three tangents, it is easy to estimate the position of the visual shape point inside the viewing edge, under the assumption that the surface is locally of order 2. To this purpose, we use the algorithm presented in [3]. This algorithm exploits the fact that the three viewing rays define locally two curves on the visual shape surface which present the

---

[2]Viewing edge intervals are finite as long as the visual hull is finite.

same normal curvature at the contact point, and a linear solution for the surface point position inside a viewing edge exists.

Examples

Figure 5 illustrates the above estimation with silhouettes obtained in real conditions. Visual shape points, and their normals to the surface, are shown. Surface normals were classically computed as the cross products of viewing directions and tangents to the silhouette contours in the images. Note in this figure that even with two viewpoints, useful visual information can still be computed from silhouettes. The information computed this way, even if partial, can be useful for various applications. We have in particular successfully used such information, i.e. point locations and surface normals, as input data to a model based motion capture system [19].

4.2. From viewing edge contributions to shapes

In the previous section, we explained how to estimate viewing ray contributions to visual shapes. Several approaches were mentioned, from viewing edge thinning, to single contact point estimations. All these approaches allow visual shape points to be estimated, as well as their normals to the visual shape surface. However, a crucial issue is how to find the visual shape surface interpolating these points.

In the case of the visual hull, the associated mesh is completely defined from silhouette contours. It corresponds to a polyhedral mesh with a constant valence equal to 3[11] and its computation can be achieved from image primitives. For other visual shapes, no a priori information is available apart from their organization into strips onto the surface. However, this information only yields surface patches which are difficult to connect so as to form a valid shape. In [17], a solution is proposed which consists first in re-sampling rims according to parallel slicing planes, and second to solve the simpler problem of surface reconstruction from polygonal contours, for which standard tools exist. While robust, this solution can not guarantee precision since re-sampling introduces errors, nor can it guarantee that the estimated shape has a topology consistent with the observations, since the surface estimation is achieved without any consideration to the image information.

To compute a shape which interpolates visual shape points while being consistent with silhouettes, we use a fairly efficient solution based on the Delaunay tetrahedrization. This has been explored in the case of visual hulls[4] and we extend the idea to general visual shapes. The proposed method computes the Delaunay tetrahedrization of the visual shape points, then carves tetrahedrons of the resulting set, which project outside any image silhouette. Visual shapes are then the union of the tetrahedrons consistent with all the input silhouettes. While simple, the approach still raises a few issues to be discussed:

1. Often tetrahedrons do not project entirely inside or outside a silhouette. To decide whether a tetrahedron is inside or outside a silhouette, we sample several points inside the tetrahedron and verify their projection status with respect to the silhouette. The ratio of points inside and outside the silhouette is then considered for the decision. Another possibility would also be to subdivide the tetrahedron into sub-tetrahedrons and to carve the subdivision.

2. Carving must be achieved with some care if a manifold surface is expected. In some local configurations, tetrahedrons should not be carved to preserve local surface connectivity. These configurations have been identified in [2].

3. The Delaunay tetrahedrization does not necessarily reflect known connections inside viewing cone strips. This is not a critical issue in most cases but sometimes yields annoying visual artifacts in the computed model. To overcome this, a first solution consists in adding vertices to the silhouette polygonal contours, increasing therefore the probability that contour connections appear in the triangulation. While satisfying in most situations, this solution does not give any guar-
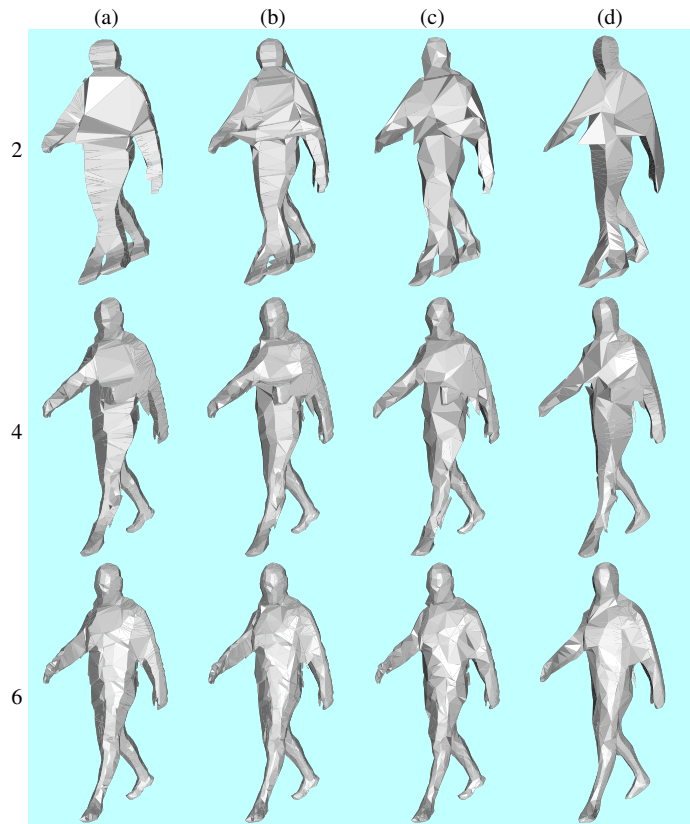

Figure 6. Visual shapes of a body shape with, from top to bottom 2, 4 **and** 6 **viewpoints, and from left to right: (a) visual hulls; (b) thinned viewing cone strips; (c) random single contact point inside viewing edges; (d)single contact point with local smooth assumptions as described in section 4.1.**

anty. One could therefore prefer using a conformal Delaunay tetrahedrization [8], which can ensure that the computed complex includes any predefined rim edges, with however a much higher computational complexity.

Examples

Figure 6 illustrates the method with the same input data than in figure 5. Visual shapes were computed in a way similar to figure 3. In the top row, the visual shapes present a different topology than the human body, because too few viewpoints are used. It shows that visual shapes cover all the visible space, which is a reasonable behavior when no additional information about shape location is available. Note also that since the observed body model has a mostly smooth surface, the visual shapes with a local second order surface model,

in column (d), are the most realistic estimations.

4.3. Texturing Visual Shapes

One of the strong advantage of visual shape models is that they project exactly onto silhouettes in the image, allowing therefore the photometric information inside silhouettes to be entirely mapped onto the model. However some difficulties remain in particular how to map textures on the model ? To this purpose, we developed an original approach. The idea is to consider each camera as a light source and to render the model using a shading model. The contribution of a view to the model textures is then encoded in the illumination values of the light source with the same location than the original view (see figure 7). These contribution values are then combined with texture values to obtain a final image. Depending on the shading model, purely diffuse or with specular like effects, texture mapping will be view-dependent or not. Though simple, this approach appears to be very efficient to texture models in real time (see the video submitted).
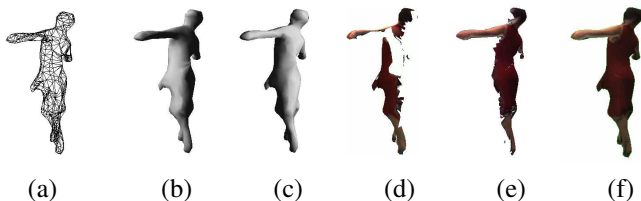

(a)   (b)   (c)   (d)   (e)   (f)

Figure 7. Texturing visual shapes: (a) the mesh, (b)-(c) the mesh rendered from 2 camera viewpoints, (d)-(e) the corresponding view contributions to the textured model, (f) the combined contribution of the the 2 camera images.

Figure 8 shows textured models of visual shapes similar to those used in the previous examples. When textured, all visual shapes share, by construction, the same appearance from viewpoints close to those of the acquisition process. Differences appear, and increase, when moving away from these viewpoints. Then, the visual hull reveals its bounding box like aspect, as viewing cone intersection curves become visible. This effect is made more obvious when considering dynamic visual shapes over time sequences. In that case, differences between chosen visual shapes are more visible, and our visual system naturally considers shapes with more likely local properties, e.g. the right model in figure 8, as more realistic.

5. Conclusion

We have introduced the visual shapes, which are a class of silhouette consistent shapes. The concept is useful to characterize shapes that project onto a set of silhouettes, and which are not necessarily the well known visual hull. This is especially useful when observing shapes with known properties, e.g. smoothness, since local assumptions can easily be used to identify and construct the most appropriate visual shape among a set of solutions. We have proposed an approach to compute points of the visual shape's surface, which are then used to compute this surface. The approach is robust and has been validated over various data sets, showing the interest of the method, in particular when modeling smooth surfaces such as human bodies. Issues we are currently considering include consistency of visual shapes over time sequences, and how to adequately account for photometric information.

References

[1] B.G. Baumgart. A polyhedron representation for computer vision. In *AFIPS National Computer Conference*, pages 589–596, 1975.

[2] J.-D. Boissonnat. Geometric structures for three-dimensional shape reconstruction. *ACM Transactions on Graphics*, 3(4):266–286, 1984.

[3] E. Boyer and M.-O. Berger. 3D Surface Reconstruction Using Occluding Contours. *International Journal of Computer Vision*, 22(3):219–233, 1997.

[4] E. Boyer and J.-S. Franco. A Hybrid Approach for Computing Visual Hulls of Complex Objects. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Madison, (USA)*, volume I, pages 695–701, 2003.

[5] M. Brand, K. Kang, , and D.B. Cooper. Algebraic Solution for the Visual Hull. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Washington, (USA)*, 2004.

[6] G. Cheung, T. Kanade, J.Y. Bouguet, and M. Holler. A real time system for robust 3d voxel reconstruction of human motions. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Hilton Head Island, (USA)*, volume 2, pages 714 – 720, 2000.

[7] R. Cipolla and A. Blake. Surface Shape from the Deformation of Apparent Contours. *International Journal of Computer Vision*, 9:83–112, 1992.

[8] David Cohen-Steiner, Eric Colin de Verdière, and Mariette Yvinec. Conforming Delaunay triangulations in 3d. In *Proc. 18th Annu. ACM Sympos. Comput. Geom.*, 2002.

[9] G. Cross and A. Zisserman. Quadric Reconstruction from Dual-Space Geometry. In *Proceedings of the 6th International Conference on Computer Vision, Bombay, (India)*, 1998.

[10] J.-S. Franco, C. Menier, E. Boyer, and B. Raffin. A Distributed Approach for Real-Time 3D Modeling. In *IEEE CVPR Workshop on Real-Time 3D Sensors and their Applications, Washington DC*, July 2004.

[11] J.S. Franco and E. Boyer. Exact Polyhedral Visual Hulls. In *Proceedings of the 14th British Machine Vision Conference, Norwich, (UK)*, 2003.

[12] Y. Furukawa and J. Ponce. Carved visual hulls for high-accuracy image-based modeling. In *technical sketch at SIGGRAPH 2005*.

[13] K. Kang, J.P. Tarel, R. Fishman, and D. Cooper. A Linear Dual-Space Approach to 3D Surface Reconstruction from Occluding Contours. In *Proceedings of the 8th International Conference on Computer Vision, Vancouver, (Canada)*, 2001.

[14] A. Laurentini. The Visual Hull Concept for Silhouette-Based Image Understanding. *IEEE Transactions on PAMI*, 16(2):150–162, February 1994.

[15] S. Lazebnik, E. Boyer, and J. Ponce. On How to Compute Exact Visual Hulls of Object Bounded by Smooth Surfaces. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Kauai, (USA)*, volume I, pages 156–161, 2001.

[16] Svetlana Lazebnik. Projective Visual Hulls. Master's thesis, University of Illinois at Urbana-Champaign, 2002.

[17] C. Liang and K-Y.K Wong. Complex 3D Shape Recovery Using a Dual-Space Approach. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, San Diego, (USA)*, 2005.

[18] W. Matusik, C. Buehler, R. Raskar, S. Gortler, and L. McMillan. Image Based Visual Hulls. In *ACM Computer Graphics (Proceedings Siggraph)*, pages 369–374, 2000.

[19] M. Niskanen, E. Boyer, and R. Horaud. Articulated motion capture from 3-d points and normals. In Torr Clocksin, Fitzgibbon, editor, *British Machine Vision Conference*, volume 1, pages 439–448, Oxford, UK, September 2005. BMVA, British Machine Vision Association.

[20] S. Sinha and M. Pollefeys. Multi-view Reconstruction using Photo-consistency and Exact Silhouette Constraints: A Maximum-Flow Formulation. In *Proceedings of the 10th International Conference on Computer Vision, Beijing, (China)*, 2005.

[21] R. Szeliski. Rapid Octree Construction from Image Sequences. *Computer Vision, Graphics and Image Processing*, 58(1):23–32, 1993.

[22] R. Vaillant and O. Faugeras. Using Extremal Boundaries for 3-D Object Modeling. *IEEE Transactions on PAMI*, 14(2):157–173, February 1992.
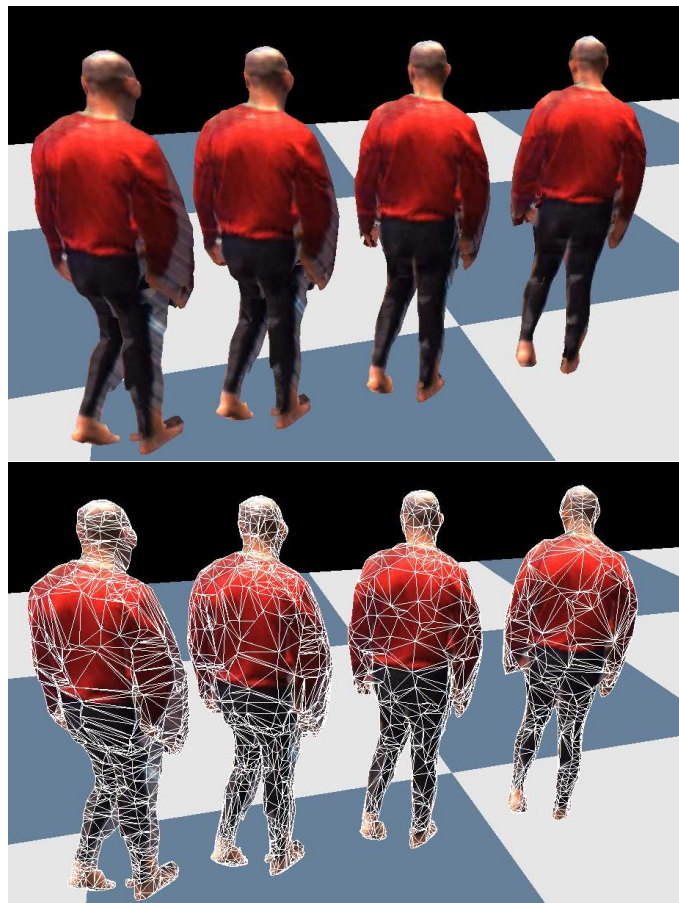
Figure 8. Visual shapes rendered with textures. From left to right: the visual hull, a model with thinned viewing cone strips, a model with random single contact points inside viewing edges and a model with single contact points satisfying local smooth assumptions.