MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Algebraic Solution for the Visual Hull

Matthew Brand, Kongbin Kang, David B. Cooper

TR2004-101   December 2004

Abstract

We introduce an alegbraic dual-space method for reconstructing the visual hull of a three-dimensional object from occluding contours observed in 2D images. The method exploits the differential structure of the manifold rather than parallax geometry, and therefore requires no correspondences. We begin by observing that the set of 2D contour tangents determines a surface in a dual space where each point represents a tangent plane to the original surface. The primal and dual surfaces have symmetric alegebra: A point on one is orthogonal to its dual point and tangent basis on the other. Thus the primal surface can be reconstructed if the local dual tangent basis can be estimated. Typically this is impossible because the dual surface is noisy and riddled with tangetn singularities due to self-crossings. We identify a directionally-indexed loca tangent basis that is well-defined and estimable everywhere on the dual surface. The estimation procedure handles singularities in the dual surface and degeneracies arising from measurement noise. The resulting method has O(N) complexity for N observed contour points and gives asymptotically exact reconstructions of surfaces that are totally observable from occluding contours.

CVPR 2004

IEEE Computer Society Conference on Computer vision and Pattern Recognition (CVPR)

Algebraic solution for the visual hull

Matthew Brand
Mitsubishi Electric Research Labs
201 Broadway, Cambridge, MA 02460

Kongbin Kang David B. Cooper
LEMS, Division of Engineering
Brown University, Providence, RI 02912

Abstract

We introduce an algebraic dual-space method for reconstructing the visual hull of a three-dimensional object from occluding contours observed in 2D images. The method exploits the differential structure of the manifold rather than parallax geometry, and therefore requires no correspondences. We begin by observing that the set of 2D contour tangents determines a surface in a dual space where each point represents a tangent plane to the original surface. The primal and dual surfaces have a symmetric algebra: A point on one is orthogonal to its dual point and tangent basis on the other. Thus the primal surface can be reconstructed if the local dual tangent basis can be estimated. Typically this is impossible because the dual surface is noisy and riddled with tangent singularities due to self-crossings. We identify a directionally-indexed local tangent basis that is well-defined and estimable everywhere on the dual surface. The estimation procedure handles singularities in the dual surface and degeneracies arising from measurement noise. The resulting method has $O(N)$ complexity for N observed contour points and gives asymptotically exact reconstructions of surfaces that are totally observable from occluding contours.

1. Introduction

Can 3D shape be recovered from multiple views without correspondences?

Obtaining reliable correspondences across many images is a notoriously difficult problem that typically contributes the lion's share of the error and compute load in 3D-from-X algorithms. Assuming perfect correspondences, triangulation of the views gives a cloud of 3D points that lie on the true surface, but when meshed, the reconstructed surface is generally an underestimate of the true volume. An alternate route to shape is visual hull: The intersection of a set of osculating projective cones that each "kiss" the object along its visual occlusion contours [1, 2]. The resulting mesh bounds the shape form the outside and is asymptotically exact, i.e., given sufficient views, it will exactly reconstruct a surface that has at least one positive principal curvature everywhere [3]. This has motivated parallel literatures on space carving via ray-tracings through octree or voxel models of space, and on projective approaches to the visual hull. In this paper we characterize the visual hull as a manifold reconstruction problem in differential geometry and find an algebraic solution that can be computed in linear time. The result is a set of exact tangent planes and estimated points of contact to the surface, which is easily converted to a 3D surface mesh whose vertices are exactly photo-consistent with all observed occluding contours.

The method rests on a dual-space formulation that relates a point on a 2D occluding contour to the local tangent plane of the 3D surface. Without correspondences or depth information, it is impossible to fix the exact location of the point on the tangent plane. However, it is possible to deduce its most likely location from continuity principles, given the locations of nearby points on the surface and the local surface curvature. The algebra is nontrivial in world coordinates, but that is moot because we do not know curvatures, depths, or even which points (from other views) are nearby. Fortunately, in the dual space where points represent tangent planes, it is possible to identify nearby points by virtue of similar local tangent structure. Furthermore, all the relevant constraints become linear (differential or algebraic) relations, for example, curvature is the rate of change of the tangent, which is differential in the dual space. The problem becomes one of computing and propagating curvature information along and between contours, which reduces to linear algebra in the dual space.

2. Related work

Our work complements the literature on parametric and volumetric approximations to the visual hull. See [4] and [5] for reviews. Mathematically, there are two frameworks for exact reconstruction that our method is most strongly related to—projective, and dual. Lazebnik et al. [6] point out that observed occluding contours form the edges of a looped graph on the surface. The exact visual hull can be computed given knowledge of the epipolar geometry, the topology of this graph, and the location of its vertices, which requires finding a sparse set of corresponding points observed by pairs and triplets of cameras. Unfortunately, the method does not appear to extend to surfaces of nonzero genus, where the expected correspondences may not exist and the graph may have incomplete loops, and the matching can be very difficult even on simple surfaces. The graph can be repaired [7, 8], though an extensive set of correspondence relations must be computed, requiring $O(N^2)$ algorithms.

Dual space representations have recently been used to good effect in problems where a parametric surface is to be recognized from or fitted to observed occluding contours [9, 10, 11]. Kutulakos [12] estimated, in the dual-space, the 2D visual hull of an object's cross-section. Our work is closest to the contributions of Kang et al. [11]. They recovered small surface patches by binning world-space into small cubes and computing rough matches between occlusion curves viewed by nearby cameras. Because the dual space representation is highly sensitive to noise, Kang et al. fitted a low-degree algebraic surface to the points in each bin in the dual space, then resampled from that surface. There is a subtle interplay between appropriate bin sizes, surface degrees, and sampling densities, all of which need to be specified *a priori*. We will recast the basic framework of the dual space in a differential geometric setting and generalize it to reconstruct *entire* free-form surfaces from occluding contours, without recourse to parametric function-fitting or prior information about the surface.

3. Theory of dual tangent spaces

We shall use typographic styles to denote different kinds of mathematical objects: x is a scalar, $\mathbf{x}$ is a column vector, $\mathbf{X}$ is a matrix, $\mathcal{X}$ is a manifold, $\mathbf{XY}$ is a matrix product, $\mathbf{X}^\top$is a matrix transpose, and $\mathbf{X}^\perp$ is the column nullspace satisfying $\mathbf{X}^\top\mathbf{X}^\perp = \mathbf{0}$.

3.1. Planar curve reconstruction

Consider a differentiable parametric planar curve $\mathbf{x}(t) \doteq [x(t),\, y(t)]^\top \in \mathbb{R}^2$ whose tangent vector at point $\mathbf{x}(t)$ is $\frac{d}{dt}\mathbf{x}(t)$. We shall use a homogeneous coordinate representation $\mathcal{C}(t) \doteq \begin{bmatrix}\mathbf{x}(t)\\ 1\end{bmatrix}$ and refer the curve as $\mathcal{C}$. Using the normal to the curve $\mathbf{n}(t) \doteq [\frac{d}{dt}\mathbf{x}(t)]^\perp$, the equation for the tangent line at $\mathbf{x}(t)$ is

$$[\mathbf{n}(t)^\top,\, -\mathbf{n}(t)^\top\mathbf{v}]\,\mathcal{C}(t) = 0, \tag{1}$$

where $\mathbf{v}$ is a variable 2D point on the tangent line. The set of tangent lines over the entire curve can be represented as

$$\begin{aligned}\mathcal{C}^*(t) &\doteq [n_x(t),\, n_y(t),\, d(t)]^\top \\ &= \tfrac{1}{\|\frac{d}{dt}\mathbf{x}(t)\|}[\mathbf{n}(t)^\top,\, -\mathbf{n}(t)^\top\mathbf{x}(t)]^\top \in \mathbb{S}\times\mathbb{R}^+ \subset \mathbb{R}^3,\end{aligned} \tag{2}$$

where n_x and n_y form a unit normal and d is the shortest distance of the tangent line to the origin. Note that $\mathcal{C}^*$ is a one dimensional manifold in a dual space where points represent tangent lines.

The dual curve $\mathcal{C}^*$ is conveniently calculated as an osculating (kissing) nullspace to the primal curve $\mathcal{C}$ and its tangents $\mathcal{T}_\mathcal{C}(t) \doteq \frac{d}{dt}\mathcal{C}(t)$:

$$\mathcal{C}^* = \frac{1}{\|\frac{d}{dt}\mathbf{x}\|}\begin{vmatrix}\mathbf{e}_1 & \mathbf{e}_2 & \mathbf{e}_3\\ x & y & 1\\ \frac{dx}{dt} & \frac{dy}{dt} & 0\end{vmatrix} = \mathcal{C}\times\mathcal{T}_\mathcal{C} = [\mathcal{C},\, \mathcal{T}_\mathcal{C}]^\perp$$

where parameterization by t is implied and $\mathbf{e}_1 = [1,0,0]^\top$, $\mathbf{e}_2 = [0,1,0]^\top$, $\mathbf{e}_3 = [0,0,1]^\top$.

A key relation between the primal and dual representations is obtained by substituting equation (2) into equation (1) to obtain

$$\mathcal{C}^{*\top}\mathcal{C} = \mathcal{C}^\top\mathcal{C}^* = 0\,. \tag{3}$$

The symmetry of equation (3) suggests that we can swap the role of point and tangent line, such that the original curve in the primal space is deducible from the tangent manifold in the dual space. Kang et al. [13] showed that the tranform from primal to dual is symmetric—applying the dual operator to the dual representation recovers the primal curve. Here we restate their key theorem in terms of a differential operator in order to make it coordinate-free and obtain a simple proof:

Theorem 1 (iterated dual curve) *Given a C^2 (twice differentiable) curve $\mathcal{C}$ described as a family of tangent lines (a dual curve $\mathcal{C}^*$ in a 3 dimension parameter space), the dual to its dual curve, $\mathcal{C}^{**}$, is the primal curve*

$$\mathcal{C}^{**} = [\mathcal{C}^*,\, \mathcal{T}_{\mathcal{C}^*}]^\perp \propto \mathcal{C}\,.$$

$$\begin{aligned} Proof: \quad \mathcal{C}^{**} &= [\mathcal{C}^*, \mathcal{T}_{\mathcal{C}^*}]^\perp \\ &= \mathcal{C}^* \times \tfrac{d}{dt}\mathcal{C}^* \\ &= (\mathcal{C}\times\tfrac{d}{dt}\mathcal{C})\times(\tfrac{d}{dt}\mathcal{C}\times\tfrac{d}{dt}\mathcal{C} + \mathcal{C}\times\tfrac{d^2}{dt^2}\mathcal{C}) \\ &= \det(\mathcal{C}, \tfrac{d}{dt}\mathcal{C}, \tfrac{d^2}{dt^2}\mathcal{C})\cdot\mathcal{C} \\ &\propto \mathcal{C} \qquad \square.\end{aligned}$$

As an example, consider the 2D Archimedean spiral $(t\cos t,\, t\sin t) \in \mathbb{R}^2$. Its dual curve is

$$(n_x,\, n_y,\, -d) = \frac{1}{\sqrt{1+t^2}}(\sin t + t\cos t,\, t\sin t - \cos t,\, t^2),$$

where the scaling gives unit normals ($n_x^2 + n_y^2 = 1$). The dual of the dual is the null space spanned by C^* and T_{C^*}:

$$\begin{aligned}[C^*, T_{C^*}]^\perp &= C^*\times T_{C^*} \\ &= ((t^2+2)t\cos t,\, (t^2+2)t\sin t,\, t^2+2)\,.\end{aligned}$$

Upon rescaling to unit homogeneous coordinates, we recover the original point on the spiral curve, $(t\cos t,\, t\sin t,\, 1)$. Figure 1 illustrates an empirical reconstruction of this curve from a noisy sampling of its tangent lines (using methods of section 4 to suppress the effects of noise).

3.2. 3D surface reconstruction

The iterated dual theorem extends directly to 3D surface reconstruction from tangent planes. Given a surface $\mathbf{x}(u,v)$, the normal is $\mathbf{n}(u,v) \doteq [\frac{\partial}{\partial u}\mathbf{x},\, \frac{\partial}{\partial v}\mathbf{x}]^\perp$. Using homogeneous

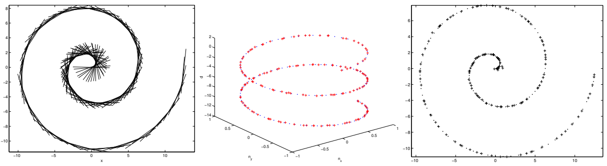

(a) tangents (b) dual manifold (c) recovery

Figure 1: Reconstructing a planar curve from a set of noisy tangent lines. (a) Tangent lines on the primal surface, perturbed with small amounts of noise. (b) The dual manifold $\mathcal{C}^*$is a helix, with '+' representing the noisy estimate and '·' representing the true manifold. (c) Reconstructed points ('+') compared with points sampled directly from the ideal manifold ('·').

coordinates $\mathcal{S}(u,v) \doteq \begin{bmatrix} \mathbf{x}(u,v) \\ 1 \end{bmatrix}$, the tangent plane can be parameterized

$$[\mathbf{n}(u,v)^\top,\ -\mathbf{n}(u,v)^\top \mathbf{x}(u,v)]\mathcal{S} = 0\,. \tag{4}$$

Therefore, the family of tangent planes can be represented as the complete set of 4D vectors:

$$\mathcal{S}^* \doteq \tfrac{1}{||\mathbf{n}||} \begin{bmatrix} \mathbf{n} \\ -\mathbf{x}^\top \mathbf{n} \end{bmatrix} = [\mathcal{S},\ \tfrac{\partial}{\partial u}\mathcal{S},\ \tfrac{\partial}{\partial v}\mathcal{S}]^\perp = [\mathcal{S}, \mathcal{T}_\mathcal{S}]^\perp \tag{5}$$

where $\mathcal{T}_\mathcal{S} \doteq [\frac{\partial}{\partial u}\mathcal{S},\ \frac{\partial}{\partial v}\mathcal{S}]$. Equation (5) tells us that a differentiable 2D manifold in $\mathbb{R}^3$ has a dual surface in the 4D parameter space (n_x, n_y, n_z, d). Strictly speaking the dual space has the topology $\mathbb{S}^2 \times \mathbb{R}^+$ (not to be confused with $\mathbb{P}^2 \times \mathbb{R}\backslash 0$) in which $\mathrm{dist}_\sigma(\mathbf{p},\mathbf{p}')^2 = (d-d')^2 + (\sigma \cos^{-1}(n_x n_x' + n_y n_y' + n_z n_z'))^2$ is a natural error metric that makes distance linear in both angle and displacement from the origin[1]. Our parameterization is an isometric immersion in $\mathbb{R}^4$ that gives the algebraic convenience of a vector space; rescaling to $n_x^2 + n_y^2 + n_z^2 = 1$ projects back onto the correct submanifold. More importantly, Euclidean error in $\mathbb{R}^4$ matches the error metric in $\mathbb{S}^2 \times \mathbb{R}^+$ to second order.

The primal surface $\mathcal{S}$ can be reconstructed from its tangents by computing the dual to the dual surface $\mathcal{S}^*$:

Theorem 2 (iterated dual surface) *Given a C^2 surface S described as a family of tangent planes (a dual surface $\mathcal{S}^*$ in a 4 dimensional parameter space), the dual to the dual surface, $\mathcal{S}^{**}$, is the primal surface*

$$\mathcal{S}^{**} = [\mathcal{S}^*, \mathcal{T}_{\mathcal{S}^*}]^\perp = \mathcal{S}\,.$$

Proof: Same as for iterated dual curve theorem. □

In order to visualize the dual manifold, we will exploit the constraint $n_x^2 + n_y^2 + n_z^2 = 1$ to map a dual space point (n_x, n_y, n_z, d) into spherical coordinates (θ, ϕ, r) with $n_x = \sin\theta\cos\phi$, $n_y = \sin\theta\sin\phi$, $n_z = \cos\theta$, and $d = r$.

[1]Sigma (σ) is a constant that depends on the size of the object.

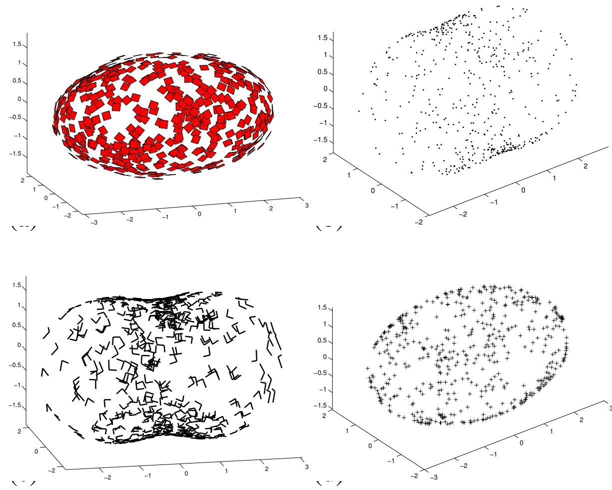

Figure 2: Surface reconstruction of an ellipsoid from its noisy tangent planes. (a) Noisy tangent plane estimates, shown as patches. (b) The dual manifold $\mathcal{C}^*$, visualized in $\mathbb{S}^2 \times \mathbb{R}^+$ coordinates. (c) Estimated tangent bases on the dual surface. (d) The reconstructed ellipsoid.

4. Estimation from image data

The import of the iterated dual theorem is that a point $\mathbf{p} \in \mathbb{R}^3$ on the primal surface $\mathcal{S}$ can be recovered through the calculation

$$[\mathbf{p}^\top, 1]^\top \propto [\mathbf{p}^*, \mathbf{t}_1^*, \mathbf{t}_2^*]^\perp \tag{6}$$

where $\mathbf{p}^* \in \mathbb{R}^4$ is the corresponding point on the dual surface and $\mathbf{t}_1^*, \mathbf{t}_2^*$ are local tangents on the dual surface. The dual point $\mathbf{p}^*$ specifies the local tangent plane on the primal surface while $\mathbf{t}_1^*, \mathbf{t}_2^*$ specify the rate at which that tangent plane is changing, i.e., the local curvature.

In the vision setting $\mathbf{p}^*$ represents the plane that connects the camera center to the tangent line of an observed contour at contour point $\mathbf{p}$. The tangent line, and thus $\mathbf{p}^*$, is determined by differences between points near $\mathbf{p}$ on the countour. The dual tangents $\mathbf{t}_1^*, \mathbf{t}_2^*$ similarly describe differences between nearby points on the dual surface. The differential nature of the tangent suggests that in principle, $\mathbf{t}_1^*, \mathbf{t}_2^*$ should be optimally estimated as an orthogonal basis of a plane in $\mathbb{R}^4$ that is fitted to points in the dual space that happen to be close on the primal surface. Then equation (6) would recover the desired 3D location of the point $\mathbf{p}$, completing the dual-primal 3D reconstruction from contours observed in 2D images. While correct in principle, this *direct* dual-space method rarely works in practice, for two reasons:

First, given discrete data points, the tangents $\mathbf{t}_1^*, \mathbf{t}_2^*$ are essentially second-order differences and as such they are sensitive to measurement error. Indeed, the dual of the dual of noisy data is often garbage.

Second, the dual surface is an topology-varying immersion of the primal surface in $\mathbb{R}^4$: It crosses itself wherever the primal surface has *bitangents*. These are tangent planes that kiss the surface at more than one point, e.g., at two bumps. The dual tangents are undefined or singular on the locii of all such self-crossings. In the data setting, tangents are not even estimable *near* such locii. Thus a direct dual-space reconstruction is infeasible for all but the simplest surfaces (or surface patches).

Although the tangent space is undefined at a singularity, it is well defined along any smooth path *through* a singularity. Thus we can pose the problem of tangent space estimation as one of picking a subset of the observed points in the vicinity of a self-crossing that support a directionally-indexed estimate of the local tangent space.

A neighborhood of points along an observed contour is a natural choice because the contour carries the continuity and local topology of the primal surface into the dual space. In fact, the occluding contour curve is particularly well suited for stable tangent estimation because it is "flatter" in the dual space than other surface contours, in the following sense:

Proposition 1 *Each observed contour spans a 3-dimensional affine subspace of dual space* $\mathbb{R}^4$.

Proof: In an orthographic view, the contour is the intersection of the image plane and an infinite generalized cylinder whose normals all lie in the image plane, and therefore collectively have rank 2. In conjunction with displacements from the origin, these normals specify the tangent planes to the surface. Therefore the contour has rank 3 in the dual space. In a perspective view, every dual point of an image contour defines a plane that passes through the camera center, therefore the camera center satisfies the plane equation (4) and the set of dual points has maximum rank 3. $\square$

We borrow and modify a technique from the manifold modeling literature [14] to estimate $\mathbf{t}_1^*$ along a contour. The contour is viewed as being piecewise approximately geodesic—reasonable in our setting because $\frac{d}{dt}\mathbf{p}^*$ is very small almost everywhere along an occluded contour parameterized by t (and is bounded even where the primal surface has infinite curvature). Geodesics have the property that their projections onto the local tangent plane are straight and identical to their tangents. Thus the tangent estimate is the direction in which the contour samples have the greatest local scatter. Computationally, a Gaussian density $\mathcal{N}_i \doteq \mathcal{N}(\mathbf{m}_i, \mathbf{C}_i)$ is fitted to each neighborhood of points along the contour, and the principal eigenvector of each covariance $\mathbf{C}_i$ is taken to be the local estimate of $\mathbf{t}_1^*$. However this local scheme is vulnerable to noise in the $\mathbb{R}^4$ locations of contour points, which may rotate the principal axis of the covariance away from the true tangent. In the manifold modeling literature, this is handled by a prior that favors tangent spaces that vary smoothly along the curve: The eigenvectors of covariances of two adjacent densities should be similarly oriented, i.e., that the nearby densities have maximal overlap, quantified as cross-entropy. The posterior probability of a set of Gaussians, each parameterized by neighborhood mean $\mathbf{m}_i$ and covariance $\mathbf{C}_i$, is then

$$\prod_i \{\prod_j \mathcal{N}(\mathbf{p}_j^*; \mathbf{m}_i, \mathbf{C}_i)\} e^{-D(\mathcal{N}_i \| \mathcal{N}_{i-1}) - D(\mathcal{N}_i \| \mathcal{N}_{i+1})},$$

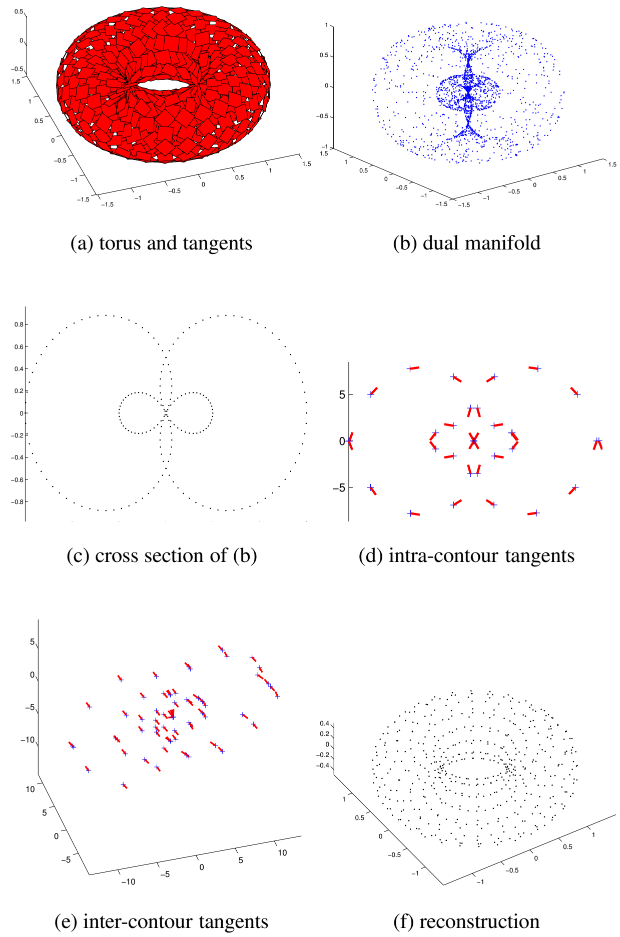

(a) torus and tangents (b) dual manifold

(c) cross section of (b) (d) intra-contour tangents

(e) inter-contour tangents (f) reconstruction

Figure 3: Surface reconstruction of a torus from its tangent planes. (a). A family of noisy tangent planes to a torus is shown as patches. (b). Points on the dual surface $\mathcal{C}^*$, visualized in $\mathbb{S}^2 \times \mathbb{R}^+$, and showing three submanifolds of bitangent singularities. (c) Cross-section of the dual manifold, showing singularities. (d,e) Estimates of dual tangents $\mathbf{t}_1^*$ and $\mathbf{t}_2^*$ along constant-ϕ contours are well-behaved even at the singularities. (f) 3D reconstruction from tangents.

where the term in braces is the data likelihood with j ranging over points in the neighborhood of $\mathbf{p}_i^*$, and the exponentiated cross-entropies are the prior. This makes the covariance estimates globally coupled along the entire contour. The global *maximum a posteriori* (MAP) solution is obtained via least-squares. Let $\mathbf{S}_i$ be the scatter of a neighborhood of n_i local contour points around their mean $\mathbf{m}_i$. Differentiating the log posterior and collecting terms in $\mathbf{C}_i$, we find that

$$(n_i+2)\mathbf{C}_i = \mathbf{S}_i + \sum_{j=\{-1,1\}} (\mathbf{m}_{i-j}-\mathbf{m}_i)(\mathbf{m}_{i-j}-\mathbf{m}_i)^\top + \mathbf{C}_{i-j}$$

This is a banded system of linear equations, solvable in time linear in the total number of contour points. Here we have assumed equal weighting for all points in the neighborhood and the two neighboring Gaussians; see [14] for more sophisticated kernel weighting schemes.

The other spanning vector, $\mathbf{t}_2^*$, is more difficult to estimate. The subspace $[\mathbf{t}_1^*, \mathbf{t}_2^*]$ should maximally span vectors from $\mathbf{p}^*$ to nearby points in nearby occluding contours, but "nearby" is not as easy to determine as it is on a single contour. Here it is useful to observe that nearby points on a smooth primal surface have similar tangent planes and curvatures, and therefore lie close in dual space. Using sorting and recursive splitting methods, nearby points for each point can be computed in $O(N \log N)$ time for N points. In fact, we can do much better, because nearby points in dual space generally lie on contours that come from nearby viewpoints, so a small subset of potential neighbors has be be considered. If we rely on that structure, the neighborhoods can be constructed in $O(N)$ time.

To estimate $\mathbf{t}_2^*$ at point $\mathbf{p}^*$ on a contour, we select the c neighboring contours having closest viewpoints (typically $c = 2$ to limit computations), and compute a weight for each point $\mathbf{p}_j^*$ that declines monotonically with distance in dual space, e.g., $w_j \propto \mathcal{N}(\mathbf{p}_j^*; \mathbf{p}^*, \sigma^2)$. We then seek a tangent plane that maximally spans the weighted tangent directions $\frac{w_j}{\|\mathbf{p}_j^* - \mathbf{p}^*\|}(\mathbf{p}_j^* - \mathbf{p}^*)$. Equivalently, we want each tangent direction to have a minimal component normal to the plane spanned by orthogonal $[\mathbf{t}_1^*, \mathbf{t}_2^*]$. Thus the local tangent estimate $\mathbf{t}_2^*$ should have minimal projection onto the matrix containing the weighted sum of orthogonal projectors

$$\mathbf{Q} \doteq \textstyle\sum_j w_j^2 (\mathbf{p}_j^* - \mathbf{p}^*)^\perp (\mathbf{p}_j^* - \mathbf{p}^*)^{\perp\top} \in \mathbb{R}^{4\times 4},$$

which isolates and sums the normal components. To ensure that we get an orthogonal basis, we project the problem into the nullspace of $\mathbf{t}_1^*$: If $\mathbf{v}_{\min}$ is the minimizing eigenvector of $\mathbf{t}_1^{*\perp\top}\mathbf{Q}\mathbf{t}_1^{*\perp} \in \mathbb{R}^{3\times 3}$, then tangent estimate is

$$\mathbf{t}_2^* = \mathbf{t}_1^{*\perp}\mathbf{v}_{\min} \in \mathbb{R}^4 .$$

It is possible to again assess a Bayesian prior favoring local estimates of $\mathbf{t}_2^*$ that change slowly as we move from contour to contour, but since the inter-contour distance is usually much larger than the distance between points along the contour, the prior is typically very weak, adding little value at the cost of substantially more computation.

4.1. Recovering the primal surface

When returning to the primal space, stronger constraints than equation 6 are available. For photo-consistency, each reconstructed point must lie on the ray that goes through the camera center and the observed image point. The ray is the intersection of two planes and thus can be described as the set of points $\mathbf{p} \in \mathbb{R}^3$ that satisfy $[\mathbf{p}^\top, 1]\mathbf{R} = 0$ for an $\mathbf{R} \in \mathbb{R}^{4\times 2}$ that specifies the contour tangent plane and the contour normal plane. Photo-consistency is then enforced by computing

$$[\mathbf{p}^\top, 1]^\top \propto \mathbf{R}^\perp \mathbf{R}^{\perp\top} [\mathbf{p}^*, \mathbf{t}_1^*, \mathbf{t}_2^*]^\perp .$$

The ray constraint is also useful when the tangents to the dual surface are degenerate. For example, apparent contours of a cylinder are straight lines and therefore $\mathbf{t}_1^* = \mathbf{0}$ along the contour because all the tangents are the same. In this case, $[\mathbf{p}^*, \mathbf{t}_1^*, \mathbf{t}_2^*]^\perp$ is a rank-2 subspace (signifying ambiguity) rather than a vector. The orthogonal projector $\mathbf{R}^\perp\mathbf{R}^{\perp\top}$ annihilates the unwanted degree of freedom.

The result is a set of 3D points that lie on the visual hull, plus tangent plane normals at each of those points. The visual hull mesh is then constructed by locally intersecting the planes, using neighborhood information determined when estimating the tangents. The resulting mesh is an (asymptotically tight) polyhedral envelope of the true smooth visual hull; alternatively, it is the exact hull assuming (nonsmooth) piecewise linear contours [6].

5. Experiments

We demonstrate first with synthetic test cases. Figures 2 and 3 illustrate surface reconstructions from tangent planes noisily sampled from an ellipsoid and a torus, respectively. In the ellipse example, naive tangent estimation suffices, while the torus requires our constrained solution.

5.1. 3D reconstruction from image data

Figure 4 shows a 3D visual hull of a pear recovered from occluding contours in 15 images taken from a fixed camera and a turntable. For purposes of illustration, the figure shows a surface obtained by directly meshing 634 points regularly sampled from the 7200 recovered 3D points. Like the polyhedral hull, this surface of meshed points also asymptotically matches the exact smooth visual hull, but it gives a tighter finite approximation to the true surface wherever the principal curvature is positive. Although 15 views is a rather sparse set of occluding contours, the computed visual hull is good enough to be usable as a model of the actual 3D surface—accurate enough for texture-mapping from multiple views. This is largely due to the tangent estimators presented in section 4—without these methods, a "direct" dual reconstruction is not even recognizable as a surface.

Figure 4: Reconstruction of a pear from 15 views. Left: The mesh superimposed on some original images. Bottom: 3D views of the mesh, with the camera rotating around and then elevating for a view from above. Right: Synthetic views of the pear from viewpoints not spanned by the original camera centers. (Again, the camera orbits and then elevates.) The bottom of the pear is clipped because the occluding contours are incomplete.

6. Discussion

We have used a differential characterization of a surface's tangent manifold to estimate the 3D visual hull of an object from occluding contours. In contrast to prior art, there is no need for any kind of image correspondences, point matching, topological analysis, or discretization. Instead, we estimate missing depth information on contours from curvature information, and obtain this by solving for locally consistent estimates of curvature at points that are found to be nearby on the dual manifold. Effectively, we are assuming that curvature changes slowly and smoothly between observed points on the primal surface. Note that this does not preclude surfaces with edges, since the dual-space tangent estimation procedure correctly handles primal surfaces that are smooth along an edge and on either side of it. The price we pay for the smoothness assumption is that surfaces should be sampled more densely where their curvature changes rapidly. Intuitively, this is exactly what will happen if the cameras are distributed uniformly around the object's view-sphere.

6.1. Acknowledgments

We are grateful to Paul Beardsley and Ramesh Raskar for assistance with the capture rig. K.K. and D.B.C. are supported by NSF grant IIS-0205477.

References

[1] B. Baumgart, *Geometric modeling for computer vision*. PhD thesis, Stanford University, 1974. TR AIM-249.

[2] A. Laurentini, "The visual hull concept for silhouette-based image understanding," *IEEE PAMI*, vol. 16, no. 2, 1994.

[3] K. N. Kutulakos and S. M. Seitz, "A theory of shape by space carving," *International Journal of Computer Vision*, 2000.

[4] R. Cipolla and P. Giblin, *Visual Motion of Curves and Surfaces*. Cambridge University Press, 1999.

[5] W. Matusik, C. Buehler, and L. McMillan, "Polyhedral visual hulls for real-time rendering," in *Proceedings of Eurographics Workshop on Rendering*, 2001.

[6] S. Lazebnik, E. Boyer, and J. Ponce, "On computing exact visual hulls of solids bounded by smooth surfaces," in *Conference on Computer Vision and Pattern Recognition*, 2001.

[7] E. Boyer and J.-S. Franco, "A hybrid approach for computing visual hulls of complex object," in *Conference on Computer Vision and Pattern Recognition*, 2003.

[8] J. Franco and E. Boyer, "Exact polyhedral visual hulls," in *British Machine Vision Conference (BMVC'03)*, 2003.

[9] G. Cross and A. Zisserman, "Quadric surface reconstruction from dual-space geometry," in *Proc. 6th International Conference on Computer Vision*, (Bombay, India), pp. 25–31, January 1998.

[10] D. Renaudie, D. Kriegman, and J. Ponce, "Duals, invariants, and the recognition of smooth objects from their occluding contour," *European Conference on Computer Vision*, 2000.

[11] K. Kang, J.-P. Tarel, R. Fishman, and D. B. Cooper, "A linear dual-space approach to 3d surface reconstruction from occluding contours using algebraic surface," in *International Conference on Computer Vision*, vol. 1, (Vancouver, Canada), pp. 198 – 204, 2001.

[12] K. N. Kutulakos, "Shape from the light field boundary," in *Proc. CVPR97*, 1997.

[13] K. Kang, J.-P. Tarel, and D. B. Cooper, "A unified linear fitting approach for singular and non-singular 3d surface from occluding contours," in *Proc. of International Conference on Computer Vision*, (Beijing, China), 2003.

[14] M. Brand, "Charting a manifold," in *Proc. NIPS-15*, vol. 15 of *Advances in Neural Information Processing Systems*, 2003.