Erratum

Correction to "Recovery of SHGCs From a Single Intensity View"

In the February 1996 issue of this Transactions, in the above-titled article authored by Ari D. Gross and Terrance E. Boult, the correct versions of the figures were inadvertently omitted by the printer. The figures and relevant sections of text should read as follows:

Abstract—Generalized Cylinders are a flexible, loosely-defined class of parametric shapes capable of modeling many real-world objects. Straight Homogeneous Generalized Cylinders are an important subclass of Generalized Cylinders, whose cross-sections are scaled versions of a reference curve. Although there has been considerable research into recovering the shape of SHGCs from their contour, this work has almost exclusively involved methods that couple contour and heuristic constraints. A rigorous approach to the problem of recovering solid parametric shape from a single intensity view should involve at least two stages: 1) deriving the contour constraints, and 2) determining if additional image constraints, e.g., intensity, can be used to uniquely determine the 3D object shape. In this paper, the authors follow the approach just described. This methodology is also important for the recovery of object classes like tubes, where contour and heuristic constraints are shown to be insufficient for shape recovery. First, we prove that SHGC contours generated under orthography have exactly two degrees of freedom. Next, we show that the remaining free parameters can be resolved using reflectance-based constraints, without knowledge of the number of light sources, their positions, intensities, the amount of ambient light, or the surface albedo. Finally, the reflectance-based recovery algorithm is demonstrated on both synthetic and real SHGC images.

Index Terms—Computer vision, shape recovery, generalized cylinders, shape from shading, shape from contour.

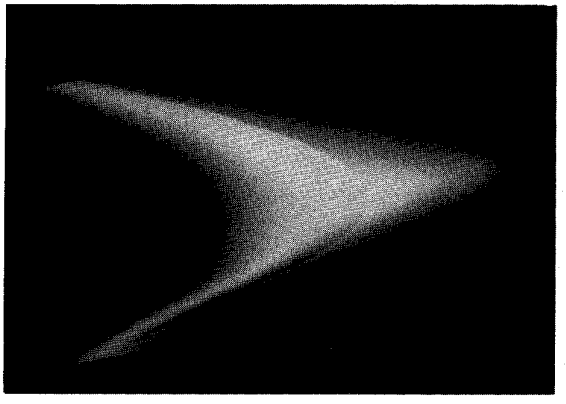

Fig. 2b. A "banana" SHGC constructed by taking the axis outside the cross-section curve.

5 Reflectance-Based Constraints

In this section, a method is developed that uses the SHGC intensity image to recover the parameters that are unconstrained by contour. In adding intensity-based methods to the recovery algorithm, we are at once concerned that the resulting algorithm not require detailed, a priori knowledge of the imaging model, such as the direction and strength of illuminant sources. Except for highly-controlled research environments, such information is generally unavailable. We have tried to avoid this pitfall by keeping our assumptions as general as possible.

Since this section of the paper deals with shape from shading, a reflectance model is required. The reflectance model used in this paper assumes that there are three reflectance components, ambient, Lambertian, and specular that, when added together, determine the image intensities. If we use the Phong reflectance model [8], then the shading rule can be expressed as the ambient component plus the sum of the Lambertian and specular components from each light source. This expression can be written as

$$I = I_a k_a + \sum_j I_j \left[k_d \left(\vec{L}_j \cdot \vec{N} \right) + k_s \left(\vec{R}_j \cdot \vec{V} \right)^n \right], \qquad (23)$$

where N is the surface normal, I_a is the ambient intensity, k_a the ambient albedo, I_j, $\vec{L}_j$, and $\vec{R}_j$ are, respectively, the intensity, direction, and reflected direction of the jth light source, $\vec{V}$ is the viewing direction, k_d and k_s are, respectively, the diffuse and specular surface albedo, and n is a constant proportional to the degree of specularity of the surface. The Phong model is used to generate the synthetic SHGC images in this paper and to develop the shape from shading invariants for recovering the slant and translation parameters. However, since the recovery method is also demonstrated on real images in Section 6, similar shape from shading invariants are probably derivable (at least approximately) for a more physics-based reflectance model, e.g., the generalized Lambertian model presented in Oren and Nayar [25].

5.1 Solving for SHGC Rotation

The parallel extrema of an SHGC (assuming they exist) can be determined directly from the ruled contour image since in [29] it is shown that contour distance extrema with respect to the SHGC image axis correspond to extrema of the sweeping rule function r, i.e., $r' = 0$. Let C be an extremal parallel curve on the SHGC surface, as shown in Fig. 7a. Consider an SHGC surface point P_1 lying on C. From (3) and the extremal constraint $r' = 0$, the surface normal at P_1 is given by

$$\vec{N}_{P_1} = q'\vec{i} - p'\vec{j}. \qquad (24)$$

If N_{P_1} intersects the $\vec{i}$ axis at some angle δ_1, then p' and q' can be written as

$$p' = c_1 \sin\delta_1,\ q' = c_1 \cos\delta_1.$$

Substituting these expressions for p' and q' into (24) yields

$$\vec{N}_{P_1} = c_1 \cos\delta_1 \vec{i} - c_1 \sin\delta_1 \vec{j}.$$

Two quantities that are computable directly from the SHGC ruled contour are the image parallel and image meridian tangents. From (2) and (5), we derive an expression for the 3D parallel tangent at P_1 with respect to the viewer-centered $(\vec{u}, \vec{v}, \vec{w})$ coordinate system given by

$$\frac{\partial \vec{OP}_1}{\partial t} = -\cos\beta\, rp'\vec{u} + \sin\beta\, rp'\vec{v} + rq'\vec{w}.$$

Let Q_1 be the projection of P_1 onto the $(\vec{u}, \vec{w})$ image plane and let $\vec{\tau}_1$ be the image parallel tangent vector at Q_1. Since $\vec{v}$ is the viewing direction, $\vec{\tau}_1$ has the form

0162-8828/96$05.00 ©1996 IEEE

$$\vec{\tau}_1 = -\cos\beta\, rp'\vec{u} + rq'\vec{w}.$$

After substituting for p' and q', we obtain

$$\vec{\tau}_1 = -\cos\beta\, rc_1 \sin\delta_1\, \vec{u} + rc_1 \cos\delta_1\, \vec{w}.$$

The slope of $\vec{\tau}_1$, computable directly from the ruled contour image, is given by

$$m_1 = -\cos\beta \tan\delta_1.$$

Thus, if we knew how to compute the value of δ_1 from the image, the desired slant angle β could also be determined from image-computable m_1. We now proceed to show that the value of δ_1 is also image-computable using the SHGC intensity image.

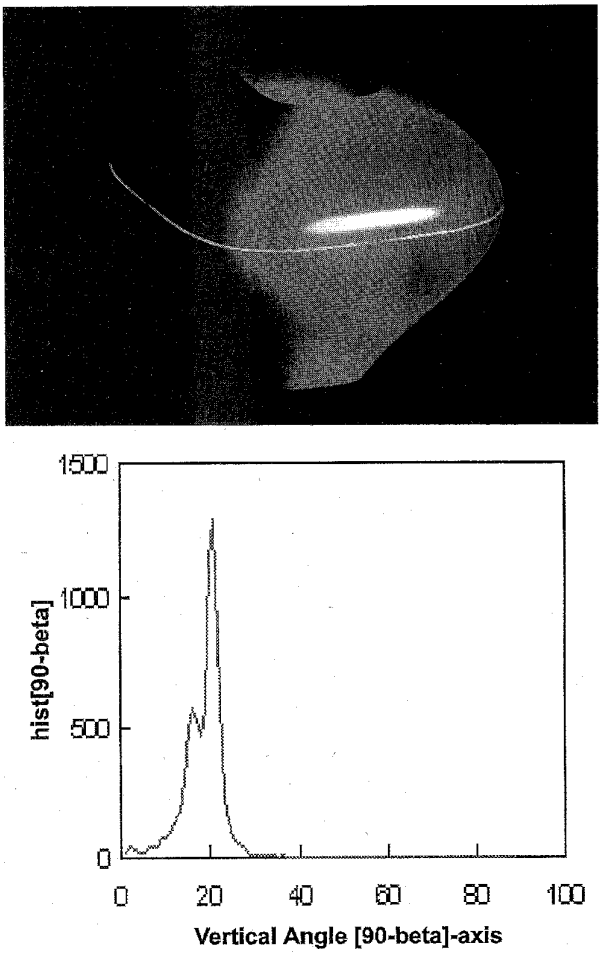

Fig. 7. Recovering the SHGC rotation: an image parallel (top) and the computed histogram (bottom).

First, we need to find an image point Q_1' that we refer to as the *complement* of Q_1 since it satisfies the property

$$m_1' = -m_1.$$

It is clear that a point's complement can be determined directly from the image. If P_1' is the surface point that projects to Q_1', then we will also refer to P_1' as the complement of P_1. Since Q_1' is a visible image point, the surface normal $\vec{N}_{P_1'}$ must satisfy

$$\vec{N}_{P_1'} \cdot \vec{v} \geq 0.$$

From the normal (7) and the extremal constraint that $r' = 0$, we have

$$\sin\beta\, q' \geq 0. \tag{25}$$

Since $m_1' = \cos\beta \tan\delta_1$ and (25) needs to be be satisfied, $\vec{N}_{P_1'}$ must have the form

$$\vec{N}_{P_1'} = c_2 \cos\delta_1\, \vec{i} + c_2 \sin\delta_1\, \vec{j}.$$

Thus, if we can find a pair of complementary image points Q_1 and Q_1', then we know that these points correspond to complementary surface points whose normal vectors intersect the $\vec{i}$ axis at opposite equal angles.

Let $\vec{L}$ be a unit point light source vector given by

$$\vec{L} = l_1\vec{i} + l_2\vec{j} + l_3\vec{k},$$

where I_L is the intensity of the light source and k_d is the diffuse surface albedo. Using Lambert's Law, the diffuse reflectance at P_1 and P_1', respectively, are given by

$$I_{P_1} = I_L k_d (l_1 \cos\delta_1 - l_2 \sin\delta_1),$$
$$I_{P_1'} = I_L k_d (l_1 \cos\delta_1 + l_2 \sin\delta_1).$$

We define a function $\Delta(i, j)$ that, given two image points, computes their intensity difference. Assuming a strictly Lambertian reflectance model, the intensity difference between Q_1 and Q_1' can be expressed as

$$\Delta(Q_1, Q_1') = -2 \sin\delta_1 l_2 I_L k_d.$$

If we still assume a strictly Lambertian reflectance model, but allow for multiple light sources with varying illuminant strength and position, then we obtain

$$\Delta(Q_1, Q_1') = \sum_i -2 \sin\delta_1 l_{2i} I_{Li} k_d,$$

where I_{Li} and l_{2i} are, respectively, the intensity and $\vec{j}$ directional component of the ith unit light source vector. Assuming that the specular reflection at points Q_1 and Q_1' is negligible, then we have the approximation

$$\Delta(Q_1, Q_1') \approx \sum_i -2 \sin\delta_1 l_{2i} I_{Li} k_d.$$

Next, consider a second pair of complementary image points Q_2 and Q_2' that are the projections of the corresponding complementary surface points P_2 and P_2' such that

$$\vec{N}_{P_2} = \cos\delta_2\, \vec{i} - \sin\delta_2\, \vec{j},\ \vec{N}_{P_2'} = \cos\delta_2\, \vec{i} + \sin\delta_2\, \vec{j},$$

where δ_2 is the angle of intersection between their respective surface normals and the $\vec{i}$ axis. The intensity difference is approximated by

$$\Delta(Q_2, Q_2') \approx \sum_i -2 \sin\delta_2 l_{2i} I_{Li} k_d.$$

Taking the image-computable ratio of these two intensity differences, we obtain

$$a = \frac{\Delta(Q_1, Q_1')}{\Delta(Q_2, Q_2')} \approx \frac{\sin\delta_1}{\sin\delta_2}.$$

Let $\vec{\tau}_2$ be the image parallel tangent vector at Q_2 and m_2 its slope. Then the image-computable value of m_2 can be expressed as

$$m_2 = -\cos\beta\tan\delta_2.$$

A second image-computable ratio, obtained by dividing m_1 by m_2, is given by

$$b = \frac{m_1}{m_2} = \frac{\tan\delta_1}{\tan\delta_2}.$$

After some simple algebraic manipulation, we can derive a δ_2 approximation of the form

$$\delta_2 \approx \sin^{-1}\sqrt{\frac{b^2-a^2}{a^2\left(b^2-1\right)}}.$$

Since a and b are both image-computable, an approximation for δ_2 is computable directly from the intensity image. But now that estimates for m_2 and δ_2 have been computed, the value of β can be computed using

$$\beta = \cos^{-1}\left(\frac{-m_2}{\tan\delta_2}\right).$$

This method can be used to compute an approximation for the value of β from two sets of complementary extremal image points without knowledge of the number of light sources, their respective strength and direction, the diffuse or ambient surface albedo, or the intensity of the ambient light.

The algorithm used to recover the value of β is given below:

1) Group points along the parallel extrema into complementary pairs of points Q_i and Q_i'.
2) For every pair of points Q_i and Q_i'.
   For every pair of extremal points Q_j and Q_j', $i \neq j$:
   a) Compute a solution for β.
   b) Increment hist[β].

Each pair of complementary pairs of points votes on a β solution, and the bin corresponding to the peak in the histogram is considered to be the approximate β solution.

Let us illustrate the algorithm with an example. For the SHGC shown in Fig. 7a, the algorithm just described was applied to points along the extrema. The $90-\beta$ histogram is shown in Fig. 7b. The peak of the histogram occurs at 20°, which corresponds to the 20° rotation towards the viewer (from the vertical direction) that was used in generating the SHGC. Additional experiments on both real and synthetic SHGC images are presented in Section 6.

5.2 Solving for SHGC Axis Translation

In this section, we present an intensity-based method for the recovery of the 3D axis translation parameter h. A closed-form intensity-based solution for h was given in [12] and [13], however, unlike the method described in this section, it did not allow for ambient and specular reflective components. We assume at this point in the recovery process that the image axis has been recovered, the contour has been ruled (Section 4), and the slant parameter β has been recovered (Section 5.1). These are prerequisites for recovering translation parameter h. In addition, we assume (as in the previous section) that there is an extrema of the sweeping function.

We would like to derive some intensity-based error-of-fit (hereafter EOF) measure that is minimized for the correct value of h. This EOF measure does not correspond to a distance metric so that some bias may be induced by the measure (see [11]). Further study is required to determine whether such distance-metric EOF measures can be derived in the intensity domain. The algorithm presented in this section iterates through the h solution space and accepts as the solution for h the value that minimizes the intensity-based EOF measure.

Assume for a moment that the translation parameter h is known. Then the underlying 3D SHGC can be fully parametrized (modulo scale). An expression for the surface normal is analogous to the one given in (3). One of the tools we consider helpful in finding photometric invariants is to use an alternative orthogonal basis for the light sources such that the reflectance equation is simplified (i.e., reduced from 3D to 2D). We first observe from (3) that all the points along a given meridian M_i will have surface normals perpendicular to the vector

$$p'\vec{i} + q'\vec{j}.$$

In analyzing the intensities of points along the image of M_i, then, it is natural to parametrize each light source with respect to the orthogonal basis given by

$$\vec{e}_1 = q'\vec{i} - p'\vec{j},\ \vec{e}_2 = p'\vec{i} + q'\vec{j},\ \vec{e}_3 = \vec{k}.$$

Then for any point P_1 on M_i, its surface normal $\vec{N}_{P_1}$ satisfies the constraint

$$\vec{N}_{P_1}\cdot\vec{e}_2 = 0.$$

Thus, for any point light source in the scene, we need only be concerned with its $\vec{e}_1$ and $\vec{e}_3$ components. If the translation parameter is known then, using (19), the normal can be computed for any point on the ruled SHGC image. For the purpose of analyzing the intensities on an image meridian, we consider two image points Q_1 and Q_2 to be complements if their respective surface normals intersect the $\vec{e}_1$ axis at opposite equal angles, such that

$$\vec{N}_1 = c_1\cos\delta\,\vec{e}_1 + c_1\sin\delta\,\vec{e}_3,\ \vec{N}_2 = c_2\cos\delta\,\vec{e}_1 - c_2\sin\delta\,\vec{e}_3.$$

Now consider a unit light source direction $\vec{L}$ given by

$$\vec{L} = l_1\vec{e}_1 + l_2\vec{e}_3.$$

The image intensities at Q_1 and Q_2, where I_L is the strength of the light source and the reflectance is purely Lambertian, are given by

$$I_{Q_1} = I_L\left(l_1\cos\delta + l_2\sin\delta\right),\ I_{Q_2} = I_L\left(l_1\cos\delta - l_2\sin\delta\right).$$

The intensity difference between Q_1 and Q_2 is given by

$$\Delta\left(Q_1, Q_2\right) = -2I_L l_2\sin\delta.$$

If there are multiple light source and ambient lighting, with no specularity, then this expression for the intensity difference between Q_1 and Q_2 generalizes to

$$\Delta\left(Q_1, Q_2\right) = -2\sin\delta\sum_i I_{L_i} l_{2i},$$

since the ambient component is eliminated by taking the intensity difference. But if h is known (the recovery algorithm iterates through h space so that there is always a current value for h) and β is known (the value of β has already been recovered), then the value of δ is also known. Dividing the above expression by $-2\sin\delta$, we obtain

$$\Theta\left(Q_1, Q_2\right) = \sum_i I_{L_i} l_{i,2}. \tag{26}$$

The above expression should have a constant value for any set of complementary meridian points that we choose. Thus, an EOF measure that should be minimized for the correct value of h is obtained by computing the variance of the function given in (26) over all pairs of complementary points on a given image meridian. This

EOF measure is given by

$$EOF_I(h) = \overline{\Theta^2(i,j)} - \overline{\Theta(i,j)}^2.$$

This EOF measure is not particularly robust, and we could probably obtain considerably more robustness using a least median of squares EOF measure. The general algorithm used to recover the translation parameter h is as follows:

1) For each discrete value of h.
   a) Group points of an image meridian into complementary pairs.
   b) For each pair of complementary points i, j compute $\Theta(i, j)$.
   c) Compute $EOF_I(h)$.
2) Iterating through h, find the value of h that minimizes $EOF_I(h)$.

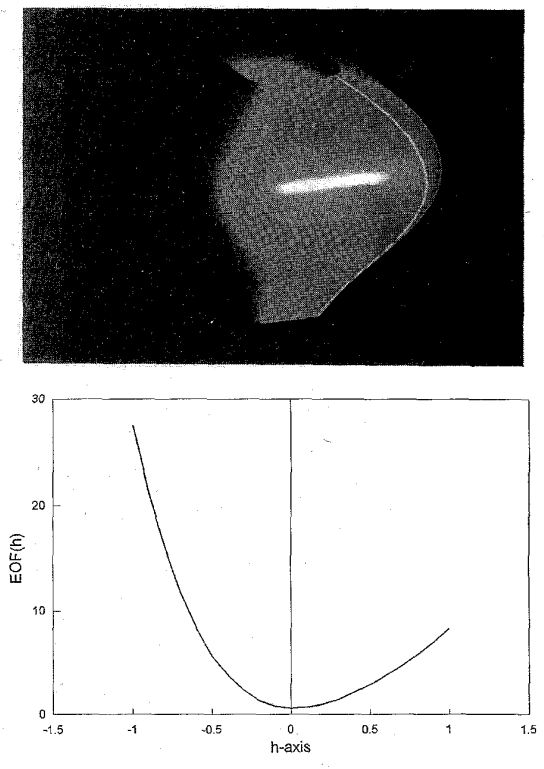

Fig. 8. Recovering the SHGC translation: an image meridian (top) and the graph *EOF*(*h*) obtained by iterating over *h*.

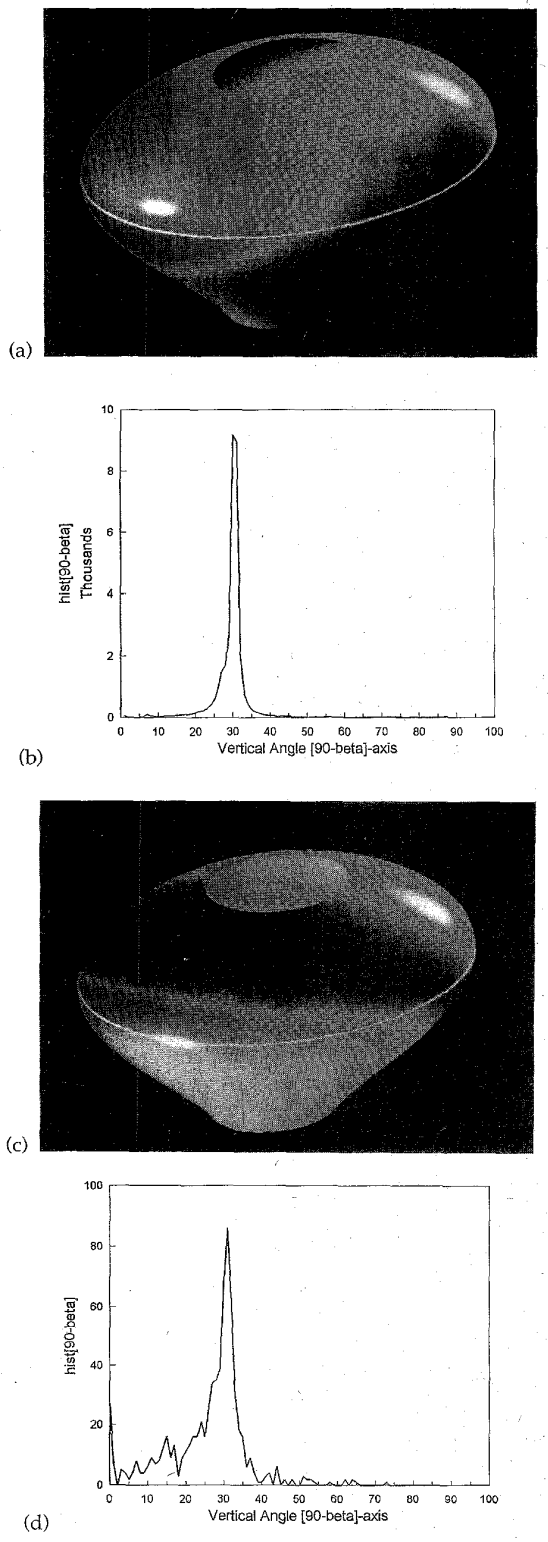

Fig. 9. Recovering the rotation angle for an SHGC with two light sources: (a) an image parallel of an SHGC with two light sources is used in (b) to compute a reliable histogram, while in (c) an image parallel intersecting a specular region results in (d) a less reliable β histogram.

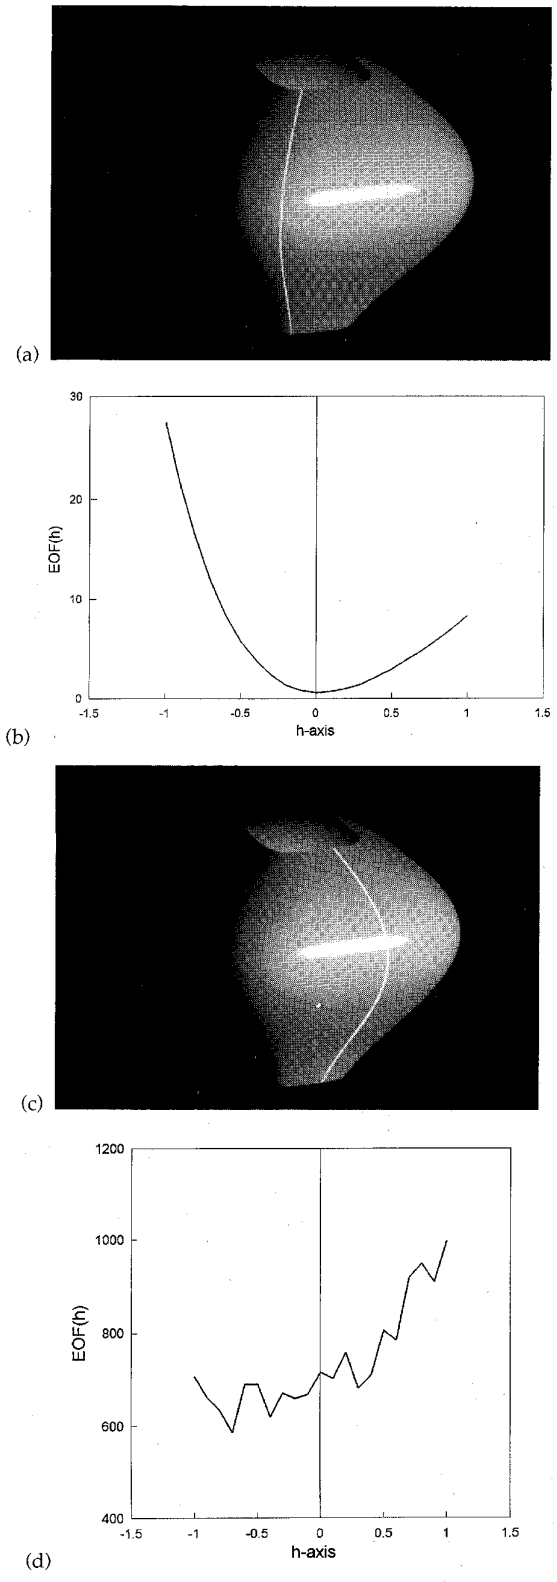

Fig. 10. Recovering the translation parameter is also affected by specular regions: (a) Lambertian meridian and (b) the recovered graph for *EOF*(*h*) compared to (c) a nonLambertian meridian and (d) its *EOF*(*h*) graph.

Consider the SHGC shown in Fig. 8a. This is the same SHGC intensity image shown in Fig. 7a, except that one of the SHGC meridians has been marked. Running the algorithm on the image intensities of the meridian shown, the graph in Fig. 8b was obtained. The graph has a minima at $h = 0$ which is the correct solution for h, corresponding to the fact that in this example the SHGC was generated with the SHGC axis intersecting the centroid of the cross-section curve. In the next section, additional experiments are presented on both synthetic and real SHGC images.

6 Experimental Results

Consider the elliptical SHGC shown in Fig. 9a. It should be noted that in this image there are two light sources. The extremal cross-section curve used in recovering β is also shown. The slant angle β is 90° when the SHGC $\vec{k}$ axis is parallel to the vertical image axis $\vec{u}$. For the purpose of displaying the recovered rotational parameters in this section, we will display the vertical angle $(90 - \beta)°$, which we will refer to as the *vertical angle*. The vertical angle used to generate the SHGC intensity image was 30°. The histogram shown in Fig. 9b attains its maximum value at 30°. In this example, it is worth noting that a symmetry-based shape from contour method would not recover a unique solution since the contour is not bilaterally symmetric (i.e., it is not a surface of revolution) and there are an infinite number of elliptical SHGCs that could have generated the contour (i.e., an ellipse is infinitely skew symmetric).

How robust is the algorithm to specularities? In Fig. 9c, the same elliptical SHGC is shown, except that this time one of the light sources has been moved so that it generates a specular highlight that covers part of the parallel extrema. The histogram is shown in Fig. 9d, where the maximum vertical angle is attained at 31°, one degree off the correct solution used in generating the image.

Next, we consider recovering the axis translation h, where $h = 0$ corresponds to an SHGC axis that intersects the centroid of the cross-section curve. In Fig. 10a, the SHGC from Fig. 8 is shown, this time with a different meridian being analyzed. The graph in Fig. 10b is the result of running the translation recovery algorithm of Section 5.2 on the meridian shown. The minima occurs at $h = 0$, as expected. When the meridian goes through a specular region, however, some degree of error is introduced. For example, in Fig. 10c a different meridian is considered, one that intersects a specular region. The resulting graph in Fig. 10d no longer has a minima at $h = 0$. Of course, since each meridian yields a solution for h, there are generally sufficient correct meridian solutions for h so that there is a clear cluster of solutions at the correct h value. Similarly, for the elliptical SHGC of Fig. 9, three meridians are shown in Fig. 11, together with their respective minima. The two "nonspecular" meridians correctly obtain a minima at $h = 0$, while the "specular" meridian does not. Once the rotation and translation parameters have been recovered, the underlying 3D SHGC is now full-determined, modulo scale. Two of the recovered SHGCs are shown in Fig. 12.

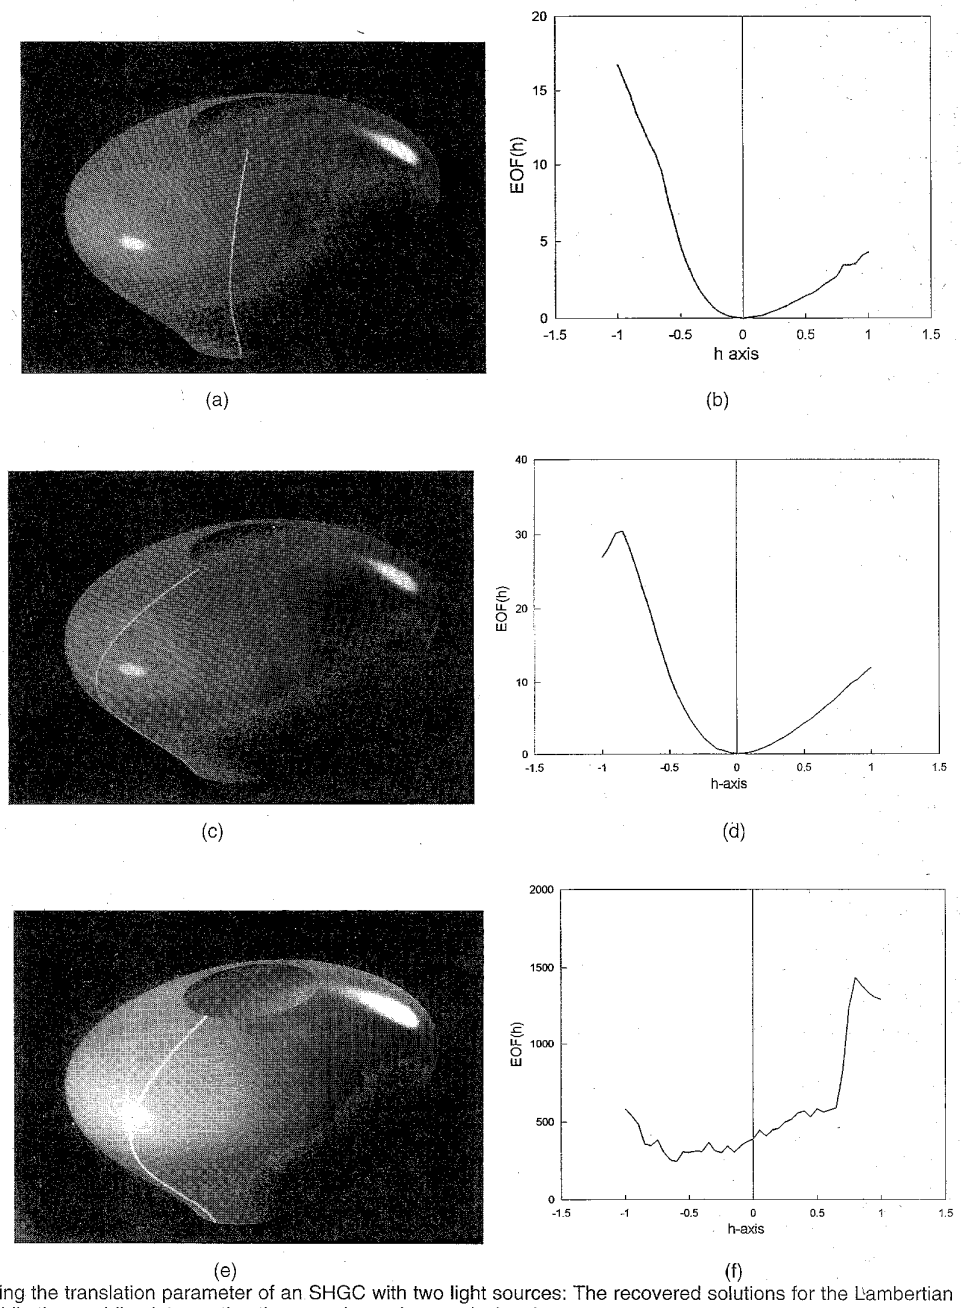

Fig. 11. Recovering the translation parameter of an SHGC with two light sources: The recovered solutions for the Lambertian meridians are quite accurate (a-d), while the meridian intersecting the specular region results in a less accurate recovery.

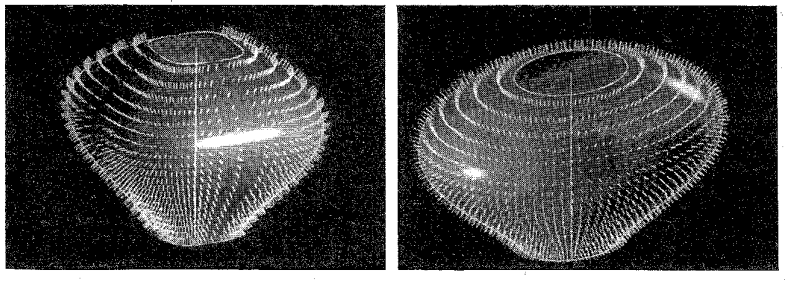

Fig. 12. The two recovered SHGCs are shown.

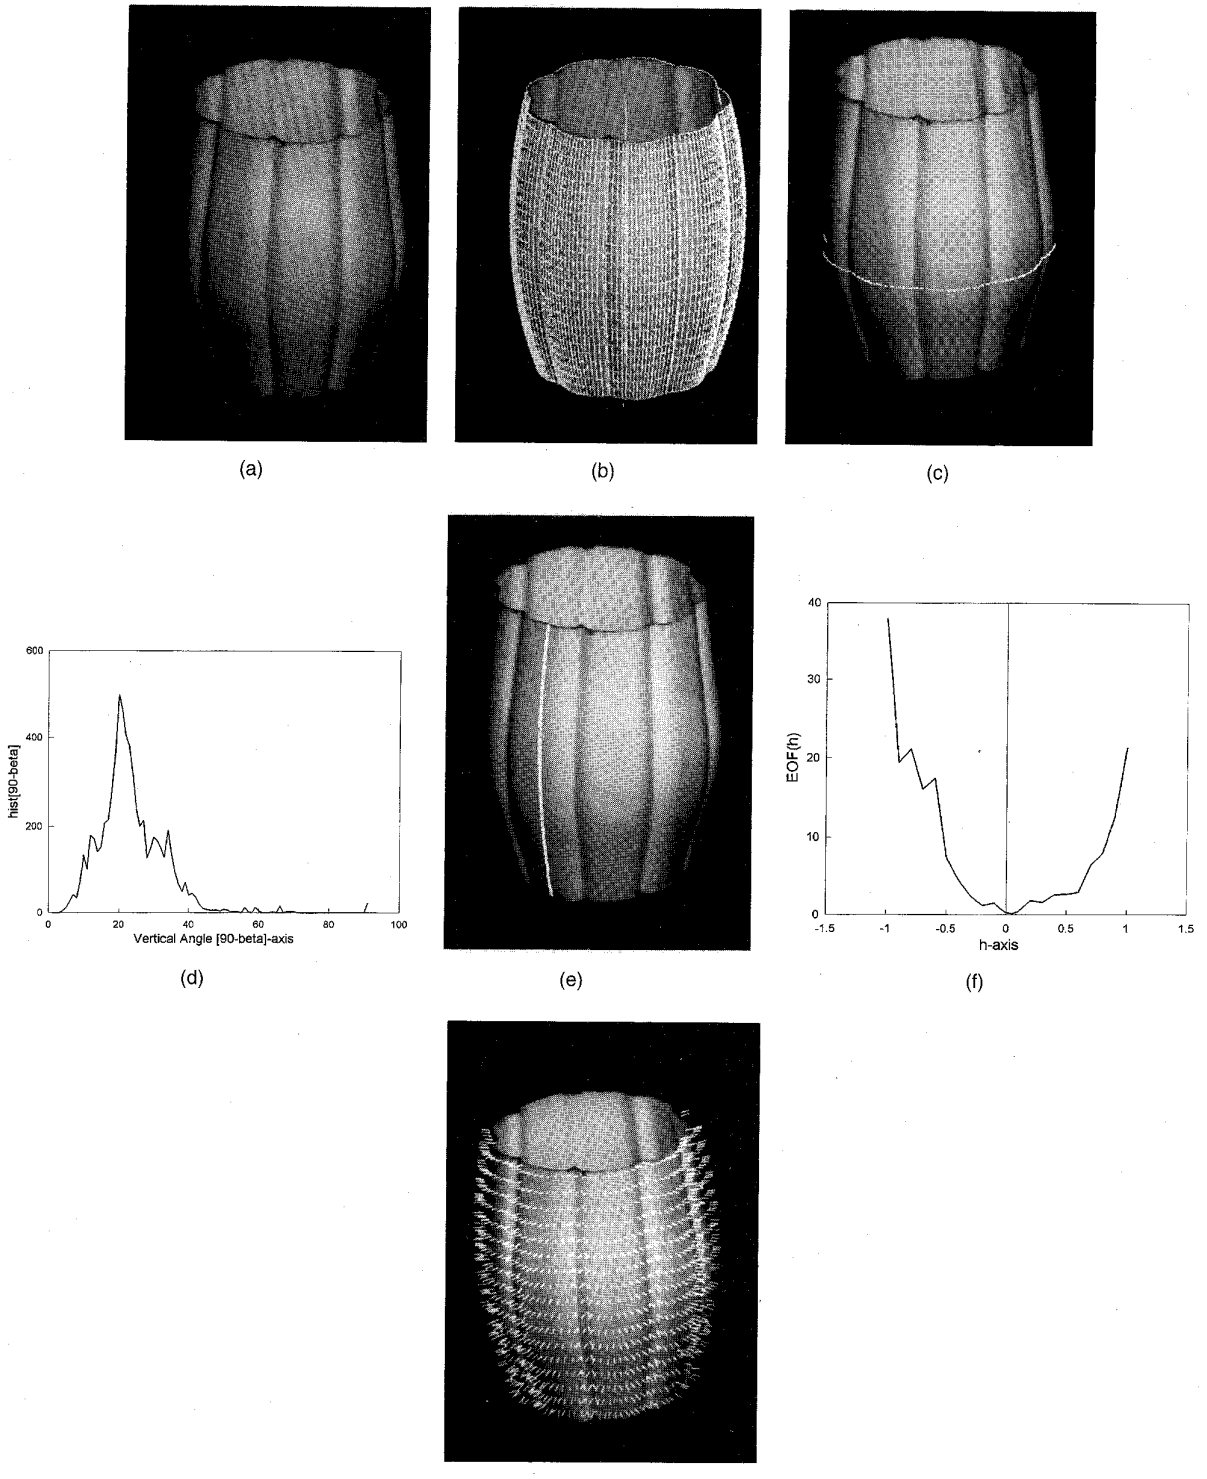

Fig. 13. Recovering a real SHGC: (a) the SHGC intensity image, (b) the recovered rulings, (c) and image parallel extrema, (d) the computed histogram in β, (e) an image meridian and (f) its recovered *EOF* graph, and (g) the recovered SHGC surface.

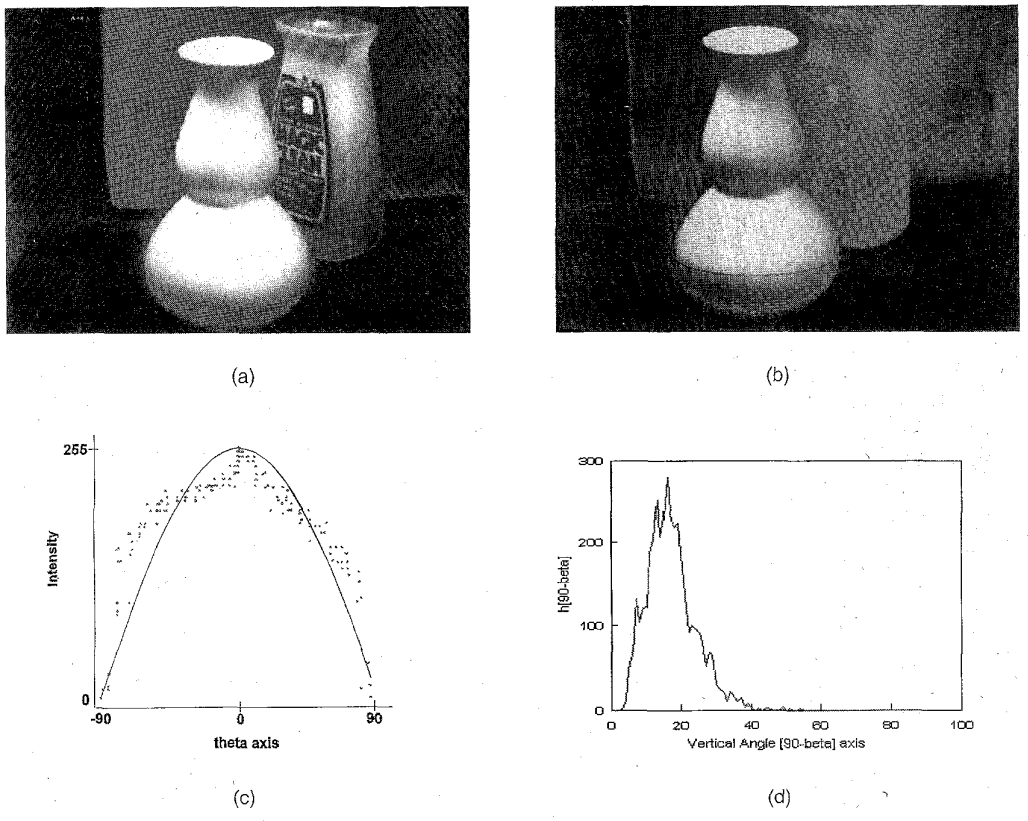

Fig. 14. Recovering a real "burnt" SHGC (not strictly Lambertian): (a) SHGC intensity image, (b) a parallel extema, (c) the intensity plot demonstrates that the curve is not strictly Lambertian, and (d) the recovered histogram for the rotation parameter, accurate to within 1° of the correct solution.

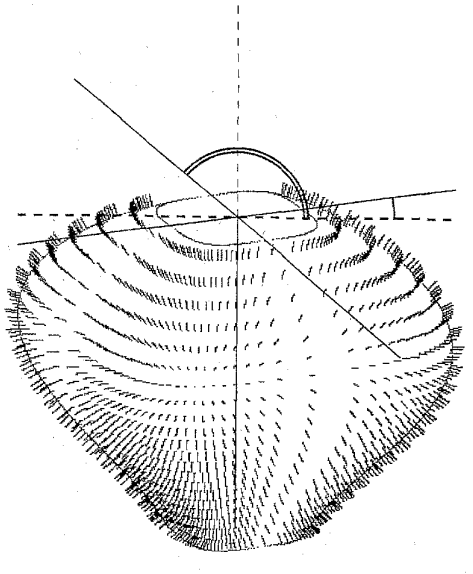

Fig. 15. Recovering the parameters of an SHGC using contour alone: the skew symmetry of the projected cross-section curve can be used to recover rotation angle β while the nonaccidentalness constraint can be used to recover translation parameter h.

We now try out the algorithm on a real SHGC image, as shown in Fig. 13a. First, the image contour is computed and the image limbs are separated from the image edges. Next, the image axis is recovered. Taking the top edge to be the top cross-section curve, the SHGC ruling is then computed as shown in Fig. 13b. From the SHGC ruling, a parallel extrema is found, see Fig. 13c. Using the intensity values along the extrema, a histogram is computed for the vertical rotation angle, as shown in Fig. 13d. The histogram attains a maximum at 19° from vertical alignment, while the actual value (which can be determined from the rotational symmetry) is approximately 20°. Next, an image meridian is selected, as shown in Fig. 13e. The intensity-based translation recovery algorithm presented in Section 5.3 is used on this meridian and the graph of $EOF(h)$ is shown in Fig. 13f. The minima for the axis translation EOF is obtained at $h = 0.05$, i.e., the axis of the SHGC roughly intersects the centroid of the cross-section curve, which is consistent with the image axis intersecting the projected center of the cross-section curve. The recovered SHGC is shown in Fig. 13g.

Next, we are interested in testing the robustness of the algorithm to instances where the reflectance function is not purely Lambertian, but more closely resembles the "weak Lambertian function" given in [3]. In this example, the object of interest shown in Fig. 14a is known to be a surface of revolution (aka SOR) by virtue of its bilateral symmetry. In addition, the object seems to be somewhat "burnt" so that we do not expect it to follow a strict cosine rule. The projected parallel extrema is shown in Fig. 14b and corresponds to a vertical rotation of approximately 17°. Taking the brightest point on the parallel extrema as the direction of the light source, we graph the intensity along the parallel extrema as a function of the cosine angle between the light source and the surface normal. In the graph, shown in Fig. 14c, it appears that the observed image intensity is approximately a monotonically de-

creasing function of the cosine angle between the surface normal and the light source. The histogram of the recovered solution for the vertical rotation angle is shown in Fig. 14d, where the peak of the histogram is at 16°, within a degree of the correct solution.

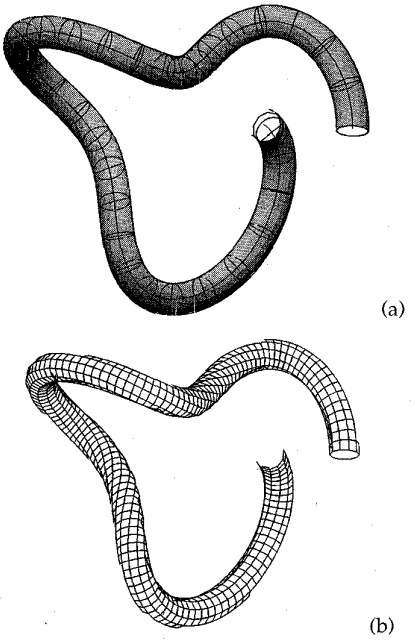

Fig. 16. The intensity-based recovery algorithm applied to a tube surface, where heuristic methods cannot readily be applied: (a) Discs are recovered along the tube axis using the algorithm given for recovering SHGC rotation, and (b) the recovered discs are combined to form a tube mesh.

Finally, we are interested in demonstrating the applicability of the intensity-based methods presented in this paper to object classes other than SHGCs. For most SHGCs, the contour is often sufficient to recover the 3D shape (modulo scale) as demonstrated in [14]. For example, the contour image of the SHGC in Fig. 7 is shown in Fig. 15. Since the projected cross-section curve is skew-symmetric, it seems reasonable to assume that the 3D cross-section curve is symmetric. Moreover, since the projected cross-sectional axes of symmetry and the SHGC image axis intersect at a point, we can safely assume (using the non-accidentalness criterion described in [14]) that these axes intersect in 3D and form an orthogonal coordinate system. But since the projection of three orthogonal axes in 3D has a unique solution (modulo a mirror reversal), both contour-unconstrained SHGC parameters can be determined. Thus, if the SHGC is assumed to be both rotationally and bilaterally symmetric, its 3D shape can often be recovered without using the intensity image. The tools developed in Section 5, however, can also be used to recover classes of surfaces that are *provably unrecoverable from contour information alone* (see [15]). For example, the 3D shape of the tube shown in Fig. 1 cannot be recovered from contour alone. As shown in [15], the intensity image of the tube in Fig. 1a has a unique solution, while the contour image in Fig. 1b has one degree of freedom corresponding to the slant of the disc at each point on the tube axis. Using the intensity-based slant recovery algorithm of Section 5.2, the orientation of discs along the tube can be recovered, as shown in Fig. 16a. This can be then be used to generate the image mesh (Fig. 16b) which uniquely determines (modulo scale) the 3D shape of the underlying tube surface.