Toward a Theory of Shape from Specular Flow

Yair Adato[1]   Yuriy Vasilyev[2]   Ohad Ben-Shahar[1]   Todd Zickler[2]

adato@cs.bgu.ac.il   vasilyev@fas.harvard.edu   ben-shahar@cs.bgu.ac.il   zickler@seas.harvard.edu

[1]Computer Science Department, Ben-Gurion University, Beer Sheva, Israel
[2]Harvard School of Engineering and Applied Sciences, Cambridge, USA

Abstract

The image of a curved, specular (mirror-like) surface is a distorted reflection of the environment. The goal of our work is to develop a framework for recovering general shape from such distortions when the environment is neither calibrated nor known. To achieve this goal we consider far-field illumination, where the object-environment distance is relatively large, and we examine the dense specular flow that is induced on the image plane through relative object-environment motion. We show that under these very practical conditions the observed specular flow can be related to surface shape through a pair of coupled nonlinear partial differential equations. Importantly, this relationship depends only on the environment's relative motion and not its content. We examine the qualitative properties of these equations, present analytic methods for recovery of the shape in several special cases, and empirically validate our results using captured data. We also discuss the relevance to both computer vision and human perception.

1. Introduction

An image of a specular (mirror-like) surface is a distortion of the surrounding environment. Since this distortion depends on surface shape, it is natural to ask how and if surface structure can be recovered from such an image. Like most vision problems this one is ill-posed; without knowledge of the environment, veridical shape information is theoretically inaccessible. Indeed, as has often been observed, it is possible to create any given image from any given specular surface by suitably manipulating the environment.

In spite of this difficulty, the human visual system is quite adept at inferring specular shape in unknown environments, even when no other shape cues are available (Fig. 1). Computationally, however, the recovery of specular shape in such general conditions has proven illusive, and existing methods have been limited to recovering only sparse or qualitative shape information, considering limited class of surfaces, or requiring calibrated conditions where the environment structure is known.

Figure 1. Images of specular surfaces under dense, (approximately) far-field illumination. As shown here, specular reflections can convey useful shape information—in this case revealing a dent in the car, the imperfections in the building's window, and the word that is pressed into a specular sheet. Can this information be extracted computationally when the environment is unknown?

ronment structure is known.

In contrast to previous work, this paper presents an approach to specular surface reconstruction that specifically targets *general* surfaces in *unknown* real-world environments. Our approach is built on an image formation model that is complex enough to be practical but simple enough for tractable analysis. The model has two essential features:

1. The environment and observer are far from the specular surface relative to the surface relief. This implies a parallel-projection camera and a reduced, two-dimensional plenoptic function (i.e., an environment map), which simplify the reconstruction problem.

2. The camera observes relative motion between the object and the environment which induces a *specular flow* [21] on the image plane. As we show, this flow provides direct access to surface shape since it depends only on the motion of the environment, not its content.

We show that based on this model one can derive differential equations relating observed specular flow to the environment motion and surface shape. In some cases these

equations can be solved (analytically) to yield dense surface shape, and here we explore two such cases in detail. We begin our exploration in a two-dimensional world (Sec. 3) in which the specular object is a plane curve, the image plane is a line, and the surrounding environment is a function defined on the unit circle. In this case, one can uniquely recover the surface (convex or not) by solving a separable non-linear ODE with initial conditions provided by an occlusion boundary. We then consider a three-dimensional world (Sec. 4) where the specular object is a surface, and we derive a coupled pair of non-linear PDEs that relate specular flow to surface shape. We show how singularities of these PDEs relate directly to the parabolic lines of the shape (where the specular flow generically grows unbounded) and how analytic reconstruction is feasible under a specific class of environment motions. Based on the analytic approach we demonstrate numerical shape recovery using both 2D and 3D experimental data.

2. Related work

Most studies of the relationship between specular reflections and surface shape consider environments that contain a single point light source. In these cases, one observes a small number of 'specularities', each of which induces a constraint between a surface point, its normal, and its local view and illumination directions. In addition, small changes in viewpoint induce specular motion, and by observing this specular motion relative to the motion of fixed surface texture, one can make local inferences about the sign of the Gaussian curvature [4, 5, 28, 6]. In order to obtain more quantitative surface information from sparse specular observations, however, one must employ significant regularization [25].

More information regarding surface shape can be obtained by observing the motion of sparse specularities over extended motion sequences. Qualitatively, it is known that as the observer moves, specularities are created and annihilated in pairs at (or in the near-field case, close to [7]) parabolic surface points [17, 19]. More quantitatively, theory suggests that one can recover a complete surface profile (i.e., a curve) by observing the specular motion induced by continuous camera motion [28]. Practical methods for doing so, however, have been developed for convex (or concave) surfaces and do not allow parabolic points [20].

Specular shape inference in natural, uncontrolled environments has received significantly less attention. Since curved specular surfaces reflect illumination from all directions, real-world environments induce *dense* specular reflections that are qualitatively very different from the sparse specularities described above. For still images of this type, it has been observed that humans often (but not always [24]) infer accurate shape, even when the illumination environment and bounding contour of the surface are un-

known [11]. While the exact mechanisms underlying these results are not yet known, it has been suggested that humans exploit the fact that image gradient directions are often correlated with second derivatives of the surface [11].

Computationally, the inference of shape in such general conditions is severely ill-posed. One can obtain additional constraints, however, through observations of dense specular flow induced by relative motion of an object, viewer, and/or environment. In a qualitative analysis, Walden and Dyer [26], show that specular flow is singular along parabolic curves when either the environment or viewer is far from the surface, and that singularities can drift from parabolic curves when both are nearby. More quantitatively, Roth and Black [21] present an optical flow algorithm that estimates a specular flow field and simultaneously identifies a surface from a parametric family of implicit functions (e.g., spheres of varying radii).

It should be mentioned that previous work on the recovery of specular shape also include '3D scanning' systems that use *calibrated* environments to obtain shape information. Example configurations include extended light sources with object or source motion [27, 15], and one or more views of a fixed object under one or more 'grid-like' environments [23, 12, 8, 9].

In contrast to previous work, we seek quantitative shape recovery for general surfaces that are not constrained to be convex or of a particular parametric form. We consider completely unknown, dense illumination environments and surfaces that are absent of diffuse texture that could otherwise assist in the reconstruction process. The main contribution of our work is to show that shape can be recovered under these conditions.

3. Specular shape from specular flow in 2D

Before addressing the general three-dimensional problem, important insights can be gained from analyzing the inference of specular shape in two dimensions (i.e., surface profiles). In this case, surfaces are reduced to plane curves, images and specular flow fields are one-dimensional, and the space of illumination directions is a circle (Fig. 2). Here we show that under the conditions of our model (i.e., far-field illumination and observer), one can *analytically* recover an *arbitrary* continuous surface profile from the observed specular flow.

As is shown in Fig. 2, the visible part of a smooth surface profile is assumed to be the graph of a function $f(x)$, and the far-field illumination environment $E(\theta)$ describes the incident radiance, which is independent of x. At a point x on the image plane, we observe the radiance reflected from a point on the surface having normal orientation $\phi(x)$, and the radiance measured at $I(x)$ is simply the value of the illumination environment $E(\theta(x))$ in the mirror-reflected direction $\theta(x)$. Since the viewing direction is aligned with

Figure 2. The specular shape reconstruction problem in two dimensions. A surface profile $f(x)$ (a plane curve) is illuminated by a far-field illumination environment $E(\theta)$, and is viewed orthographically to produce a 1D image $I(x)$. The sign convention for the angular dimension is shown in the inset.
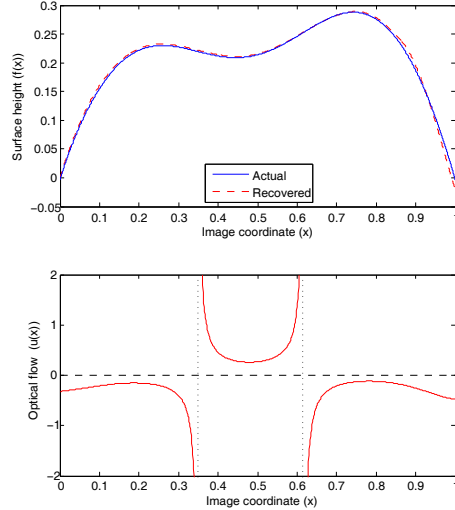

Figure 3. Recovering a surface profile from specular flow. A surface profile $f(x)$ (top, blue solid curve) is viewed under a rotating environment as depicted in Fig. 2. This induces a specular flow (bottom) that is singular at inflection points. Using this flow, we recover the surface by solving Eq. 4 using the left-hand surface boundary as an initial condition. The surface is recovered (top, red dashed curve) despite the singularities in specular flow because reconstruction relies on the integration of *inverse* flow, which is well-defined everywhere.

$\theta = 0$, it follows that $\theta(x) = 2\phi(x)$.

To recover shape from specular flow, we seek a relationship between motion of the environment $\omega = d\theta/dt$ and the induced motion field—or specular flow—on the image plane $u = dx/dt$. From the sign convention in Fig. 2 it follows that

$$\tan(\theta(x)/2) = -f_x(x).$$

Taking temporal derivatives of this expression and using the fact that $\sec^2(\theta/2) = 1 + f_x^2$ we obtain the desired relationship:

$$u(x) = \frac{-\omega}{2\kappa(x)\sqrt{1 + f_x(x)^2}}, \tag{1}$$

where $\kappa(x) = f_{xx}/(1 + f_x^2)^{3/2}$ is the curvature at point x.

Equation 1 is a generative equation for specular flow, and it shows that specular flow is well defined everywhere except at the projections of surface points having zero curvature. Furthermore, as is exemplified in Fig. 3, projections of these points behave as either 'sources' or 'sinks' of specular flow, in accordance with the pair-wise specular 'birth' and 'death' that is expected at parabolic lines in three dimensions [17, 19]. In this example, where the environment rotates in a counter-clockwise manner, the flow is divergent (expands 'outwards' in both directions) at the left inflection point. This point behaves as a flow source—a point where new regions of the environment come into view on the image plane. By the same reasoning, the right inflection point is a sink because the flow is convergent there. Their roles would change if one were to reverse the direction of environment rotation.

In order to recover the surface from the observed specular flow, we rearrange Eq. 1 to obtain a Riccati equation,

$$2u(x)f_{xx} + \omega f_x^2 + \omega = 0. \tag{2}$$

This second order non-linear ODE can be reduced to a separable first-order equation by making the substitution

$v = f_x$, and it has a relatively simple analytic solution. We first obtain the first derivative of the surface using

$$f_x(x) = \tan\left(-\frac{\omega}{2}\int_{x_i}^{x}\frac{d\lambda}{u(\lambda)} + C\right),$$

where λ is a dummy variable, x_i is some initial point on the image plane, and C is an arbitrary constant that can be determined using an initial condition, $C = \tan^{-1}(f_x(x_i))$. Once the first derivative is known, the surface f can be recovered through integration. This introduces another arbitrary constant which determines the absolute depth of the surface and can be set to zero.

3.1. Object boundaries as initial conditions

When the surface profile is a smooth closed curve, object boundaries occur where the surface normal is orthogonal to the viewing direction, and the derivative is therefore known at these points (see x_0 and x_f in Fig. 2). Thus, object boundaries provide a convenient source for initial conditions. Using the surface parameterization proposed above, however, the initial conditions at the left and right object boundaries are $f_x(x_o) \to \infty$ and $f_x(x_f) \to -\infty$, which are inconvenient for numerical purposes.

To get around this, we can re-parameterize the surface derivative f_x using a stereographic projection [16], accord-

ing to which we define:

$$q = \frac{2f_x}{1 + \sqrt{f_x^2 + 1}}$$

$$f_x = \frac{4q}{4 - q^2}. \qquad (3)$$

Substituting in Eq. 2 yields an equation of the same form:

$$u(x)q_x + \frac{\omega}{8}q^2 + \frac{\omega}{2} = 0, \qquad (4)$$

whose solution is

$$q(x) = 2\tan\left(-\frac{\omega}{4}\int_{x_o}^{x}\frac{d\lambda}{u(\lambda)} + C\right).$$

Here, C is an arbitrary constant that can now be determined using the initial condition provided by one of the object boundaries

$$C = \tan^{-1}\left(q(x_o)\right) = \tan^{-1}\left(-2\right).$$

To recover the surface $f(x)$, the solution $q(x)$ is transformed via Eq. 3 and then integrated as before.

A demonstration of this procedure is shown in Fig. 3. Here, a sequence of 1D images is rendered under an environment that rotates in a counter-clockwise direction. The environment is extracted from a great circle of the captured "St. Peter's" environment map [18]. Flow is estimated numerically and independently at each pixel using the optical flow equation $I_x u + I_t$=0, and the surface is recovered by solving Eq. 4 using the left-most point as an initial condition.

3.2. Observations

Since it enables the recovery of surface shape, we refer to Eq. 4 as the *shape-from-specular-flow* (SFSF) equation in two dimensions. It has a number of notable properties. The ODE can be solved analytically given an analytic expression for the specular flow, and a unique solution can be readily obtained using an 'occluding contour' (or any other point at which the first derivative is known) as a boundary condition. Since there is no aperture problem in two dimensions, specular flow can be estimated independently at every image point from as few as two images. Thus, provided that the illumination environment exhibits sufficient angular radiance variation, we are able to completely recover a two-dimensional surface profile from as few as two frames.

Another important property is that the 2D SFSF equation enables the recovery of arbitrary smooth surfaces, including those with points of zero curvature. As noted above, the specular flow approaches $\pm\infty$ at the projection of an inflection point. Surface reconstruction requires the integration of the *inverse flow*, however, which is well defined everywhere.


Figure 4. The specular shape reconstruction problem in three dimensions. A surface $f(x, y)$ is illuminated by a far-field illumination environment and is viewed orthographically to produce a 2D image $I(x, y)$. The illumination sphere is parameterized using spherical coordinates (θ, ϕ).

These nice properties of the 2D SFSF equation follow directly from an image formation model that includes a far-field viewer and environment and relative object/environment motion. As we show in Sect. 4.2, many of these desirable properties carry over to the three-dimensional case as well.

4. Specular shape from specular flow in 3D

Much like the two-dimensional case described in the previous section, we begin the three-dimensional analysis by considering a surface $\mathbf{S}(x, y) = (x, y, f(x, y))$ that is the graph of a (bi-variate) function. As before the surface is viewed orthographically from above and illuminated by a far-field environment (see Fig. 4).

Let $\hat{\mathbf{v}} = (0, 0, 1)$ be the viewing direction, $\hat{\mathbf{n}}(x, y)$ the surface normal at surface point $(x, y, f(x, y))$, and $\hat{\mathbf{r}}(x, y)$ the mirror-reflection direction at the same point. An image of $\mathbf{S}(x, y)$ on the orthographic image plane constitutes radiance values of the distant illumination environment. In the 3D case this environment constitutes a *sphere* of directions, which we parameterize with two spherical angles (zenith and azimuth). In particular, we represent reflection directions as (α, β) and normal directions as (θ, ϕ), both under the usual sign convention shown in Fig. 4. As in the previous section, these directions are related by

$$\begin{aligned} \alpha(x, y) &= 2\theta(x, y) \\ \beta(x, y) &= \phi(x, y) \end{aligned}.$$

In order to relate displacements on the image plane to those on the illumination sphere, we note that the reflection vector at each point can be expressed both in terms of surface derivatives and spherical coordinates;

$$\begin{aligned} \hat{\mathbf{r}} &= (\sin\alpha\cos\beta, \sin\alpha\sin\beta, \cos\alpha) \\ &= \frac{(-2f_x, -2f_y, 1 - f_x^2 - f_y^2)}{f_x^2 + f_y^2 + 1}. \end{aligned}$$

4

From this relationship we can deduce that

$$
\begin{aligned}
\tan \alpha &= \frac{2\|\nabla f\|}{1-\|\nabla f\|^2} \\
\tan \beta &= \frac{f_y}{f_x} \quad .
\end{aligned}
\tag{5}
$$

As in the two-dimensional case, we are interested in the effects of environment motion. In three dimensions, the angular motion of a far-field environment can be represented as a vector field on the unit sphere. We use the following notation to describe this *environment motion field*:

$$
\boldsymbol{\omega}(\alpha, \beta) = (\omega_\alpha(\alpha, \beta), \omega_\beta(\alpha, \beta))) = \left(\frac{d\alpha}{dt}, \frac{d\beta}{dt} \right).
$$

This is a general representation that can describe both 'rigid' motion (i.e., an environment that rotates around some fixed axis) and an arbitrary 'non-rigid' motion. In the rigid case, when the environment rotates about axis $\hat{\mathbf{a}} = (\sin \alpha_\circ \cos \beta_\circ, \sin \alpha_\circ \sin \beta_\circ, \cos \alpha_\circ)$ with angular velocity ω, the environment motion field is [1]

$$
\begin{aligned}
\omega_\alpha(\alpha, \beta) &= \omega \sin \alpha_\circ \sin(\beta_\circ - \beta) \\
\omega_\beta(\alpha, \beta) &= \omega \left(\cos \alpha_\circ - \sin \alpha_\circ \cos(\beta - \beta_\circ) \cot \alpha \right).
\end{aligned}
\tag{6}
$$

Environment motion induces a motion field, or a specular flow, on the image plane. This flow, represented as

$$
\mathbf{u} = (u(x, y), v(x, y)) = \left(\frac{dx}{dt}, \frac{dy}{dt} \right),
$$

is related to the environment motion through the Jacobian:

$$
\boldsymbol{\omega} = \frac{d(\alpha, \beta)}{dt} = \frac{\partial(\alpha, \beta)}{\partial(x, y)} \frac{d(x, y)}{dt} = \mathbf{J}\mathbf{u}.
\tag{7}
$$

The Jacobian $\mathbf{J}$ can be expressed in terms of surface shape by taking temporal derivatives of Eq. 5, which yields:

$$
\mathbf{J} \triangleq \frac{\partial(\alpha, \beta)}{\partial(x, y)} = \begin{pmatrix} \frac{f_x f_{xx} + f_y f_{xy}}{\|\nabla f\| \cdot (1 + \|\nabla f\|^2)} & \frac{f_x f_{xy} + f_y f_{yy}}{\|\nabla f\| \cdot (1 + \|\nabla f\|^2)} \\ \frac{f_x f_{xy} - f_y f_{xx}}{2\|\nabla f\|^2} & \frac{f_x f_{yy} - f_y f_{xy}}{2\|\nabla f\|^2} \end{pmatrix}.
\tag{8}
$$

Implicit in Eq. 7 is the fact that the environment direction reflected by each point is determined by the surface geometry. That is,

$$
\boldsymbol{\omega} = \boldsymbol{\omega}(\alpha(f_x, f_y), \beta(f_x, f_y)) = \tilde{\boldsymbol{\omega}}(f_x, f_y) .
$$

When the environment motion field is known and we observe the induced specular flow $\mathbf{u}$ on the image plane, one hopes to recover the shape by solving Eq. 7, which represents a system of non-linear PDEs in $f(x, y)$. Thus, we refer to this equation as the *shape from specular flow* (SFSF) equation in three dimensions.

4.1. Behavior at parabolic points

While Eq. 7 may be used to solve for an unknown shape $f(x, y)$ from a known specular flow (u, v), it can be rearranged, through inversion of the Jacobian $\mathbf{J}$, to derive a generative equation for an unknown specular flow $\mathbf{u}$ induced by a known surface $f(x, y)$ under a given environment motion $\boldsymbol{\omega}$

$$
\mathbf{u} = \mathbf{J}^{-1} \boldsymbol{\omega} .
\tag{9}
$$

Important insight into this relationship is revealed from the determinant of $\mathbf{J}$, which can be written as

$$
Det(\mathbf{J}) = \frac{2K(1 + \|\nabla f\|^2)}{\|\nabla f\|},
\tag{10}
$$

where K is the Gaussian curvature of the surface, i.e.,

$$
K = (f_{xx} f_{yy} - f_{xy}^2)/(1 + \|\nabla f\|^2)^2.
$$

Eq. 10 tells us that the environment motion field and specular flow are related by an isomorphism at all surface points except parabolic points, where the Gaussian curvature vanishes[1]. This is directly analogous to the two-dimensional case in which the specular flow is infinite at inflection points. In the three-dimensional case, the magnitude of the specular flow generically grows unbounded at the image projection of a parabolic point on the surface. (Note that this is different from the near-field case [26], where the flow singularities can drift away from parabolic points.) One synthetic example of this phenomenon, with the specular flow computed using Eq, 9, is demonstrated in Fig. 5

4.2. Environment motion around the view direction

One special case in which Eq. 7 assumes a simple form that can be solved analytically occurs when the axis of environment rotation $\hat{\mathbf{a}}$ is aligned with the view direction $\hat{\mathbf{v}}$. In our spherical coordinate system, environment rotation about the view direction induces the motion field

$$
\begin{aligned}
\omega_\alpha(\alpha, \beta) &= 0 \\
\omega_\beta(\alpha, \beta) &= \omega
\end{aligned}
\tag{11}
$$

with ω being the scalar angular velocity.

To exploit the reduced complexity, we define two auxiliary functions corresponding to the surface gradient magnitude and orientation:

$$
\begin{aligned}
h(x, y) &\triangleq f_x^2 + f_y^2 \\
k(x, y) &\triangleq \tan^{-1}(f_y/f_x),
\end{aligned}
$$

$$
\tag{12}
$$
$$
\tag{13}
$$

---

[1]Another singular case in this context are points with $\|\nabla f\| = 0$, i.e., fronto-parallel surface points. Note that these points reflect the observer who, unlike the rest of the environment, does not move relative to the surface and hence provides little in the way of shape information.

Figure 5. The magnitude of the specular flow induced by a specular surface under environment motion. **Top left:** The shape of a synthetic specular surface with parabolic lines superimposed. **Top right:** One frame of a rendered image sequence under a rotating environment. Superimposed are the projection of the parabolic curves. **Bottom:** A zoomed image of the log-magnitude of the induced specular flow in the marked region of interest. Note how the flow magnitude is generally very large near parabolic lines.

where $\tan^{-1}$ is the four quadrant arctangent of its argument. Given these definitions and the flow from Eq. 11, the two coupled non-linear equations in Eq. 7 reduce to two *linear* PDEs in h and k:

$$u(x,y)h_x(x,y) + v(x,y)h_y(x,y) = 0 \qquad (14)$$
$$u(x,y)k_x(x,y) + v(x,y)k_y(x,y) = 2\omega . \quad (15)$$

These equations immediately suggest the following reconstruction procedure:

Reconstruction algorithm:

1. Use observed specular flow (u, v) to solve Eq. 14 for $h(x,y)$ and Eq. 15 for $k(x,y)$.

2. Recover f_x and f_y from $h(x,y)$ and $k(x,y)$ using definitions in Eqs. 12 and 13.

3. Integrate f_x and f_y to obtain $f(x,y)$.

Both Eqs. 14 and 15 can be solved using the method of characteristics, and in both cases the characteristics correspond to the integral curves of the specular flow. Of course, the surface will be recovered through this procedure provided that requisite initial conditions are available, i.e., that ∇f is known at one or more points along each integral curve of the specular flow. Clearly, the absolute depth of the surface cannot be recovered unless the absolute depth of one or more surface points is given by some other means.

4.3. Observations

In addition to facilitating a shape recovery procedure, Equations 14 and 15 have straight-forward and useful geometric interpretations: they constrain the directional derivatives of h and k in the direction of the specular flow.

Eq. 14 dictates that h be constant along each characteristic. Thus, iso-contours of h necessarily correspond to integral curves of the specular flow. Similarly, since Eq. 15 can be written in terms of the *unit* specular flow $\hat{\mathbf{u}}$

$$\hat{\mathbf{u}} \cdot \nabla k = \frac{\mathbf{u}}{\|\mathbf{u}\|} \cdot \nabla k = \frac{2\omega}{\|\mathbf{u}\|} , \qquad (16)$$

it implies that the rate of change in k along each arc-length parameterized characteristic must be proportional to ω and inversely proportional to the flow magnitude $\|\mathbf{u}\|$. This observation becomes particularly interesting once we realize that k is an angle (i.e., an orientation) and its rate of change can be interpreted as a *curvature* measure (in the spirit proposed in studies of oriented patterns and visual flows, e.g. [3]). Given the specular flow $\mathbf{u}$ of a specular object under environment rotation around $\hat{\mathbf{v}}$, we can therefore define its *specular curvature* based on Eq. 16 as

$$\kappa_s \triangleq \frac{2\omega}{\|\mathbf{u}\|} . \qquad (17)$$

Moreover, since the characteristics (or, equivalently, the integral curves of the specular flow) are generically closed curves, integrating κ_s along each such curve must yield a multiple of 2π (or else the system violates integrability). Although not exploited in our paper, this observation could be used to recover the angular velocity ω and hence to facilitate specular shape recovery even when the environment angular velocity is unknown.

4.4. Experimental results

Although the contribution of this paper is primarily theoretical, as a proof of concept, the approach developed above was applied to image data acquired using the system depicted in the top of Fig. 6. A camera (Canon 10D, EF 75-300mm lens, EF 25 II extension tube) was placed 1m from a 2" diameter chrome sphere. In an unknown, far-field illumination environment, both the camera and sphere rotated as a fixed pair about an axis parallel to the camera's optical axis. One frame of the captured image sequence is shown in the middle of Fig. 6.

Given an image sequence captured at a known angular velocity of $0.5°$/frame, specular flow was recovered using the Horn and Schunck algorithm [14]. Based on this flow, the surface was recovered as described in the previous section. Initial conditions were provided manually by specifying the surface gradient along the red curve shown in the figure. The integral curves of the specular flow were determined by numerically integrating the flow field, and these curves served as characteristics for the numerical integration of Eqs. 14 and 15 to recover h and k. With h and k known at each point, the surface derivatives f_x and f_y are

Figure 6. Recovering general shape from specular flow in three dimensions. **Top:** Under far-field illumination, a camera and object rotate as a fixed pair around a line parallel to the optical axis. **Middle-left:** One image from a captured sequence. **Middle-right:** estimated specular flow. **Bottom:** Shape recovered by solving the 3D shape-from-specular-flow equation as described in Sec. 4.2. The surface gradient is assumed known along the red curve, which provides the necessary initial conditions.

easily computed with

$$f_x = \sqrt{h} \cos k$$
$$f_y = \sqrt{h} \sin k$$

and then integrated to yield the surface shown in the bottom of Fig. 6. The RMS error in the reconstruction as a fraction of the sphere radius was found to be 1.2%.

5. Discussion

This paper introduces a novel theoretical framework for the reconstruction of smooth specular shapes from observed motion in natural, unknown, and uncontrolled environments. Using far-field view and illumination conditions and relative object/environment motion, we analyze the relationship between observed specular flow and specular shape, and we derive a system of coupled nonlinear PDEs that can be solved for reconstruction. We show that in the particular case of environment rotation around the viewing direction, this system can be reduced to a pair of linear PDEs and solved either analytically or numerically.

Several extensions and research directions arise naturally from our results. First, we show that the shape from specular flow (SFSF) equation can be solved in three dimensions when the environment rotates about the view direction, but

is likely that one can extend this analysis to include more general rotation directions, either through a change of angular coordinates applied to the Z-axis equations (Eqs. 14 and 15), or by attempting to solve the SFSF equation (Eq. 7) with ω defined via Eq. 6 rather than Eq. 11. Second, although we restricted our attention to the case of environment rotation in this paper, it is likely that small relative object/observer motion can be described in a similar manner. Third, while we focus on a distant observer, it may be possible to extend our analysis to a near-field viewer by replacing the constant $\hat{\mathbf{v}}$ vector with a vector field corresponding to perspective projection.

The problem of recovering the specular flow itself was not discussed in this paper, but it is possible that the differential geometric analysis we provide may be useful for this aspect of the problem as well (perhaps using the framework described in [2]). In this context, the unique behavior of specular flow around parabolic curves may be useful for identifying them and distinguishing them from other types of flow singularities(e.g., due to surface discontinuities). This, in turn, could facilitate the reconstruction of *piecewise* smooth specular shape as well. In any case, a study of the specular flow itself should also assist in recovering or estimating initial conditions for the solution of the SFSF equation, either from the regular parts of the flow, or from its singularities either around parabolic lines, or in the vicinity of the object's occluding contour.

In addition to these computational issues, this analysis may also aid our understanding of human perception of specular shape. While the human visual system seems to exploit specular flow in distinguishing between surfaces that are specular and diffuse [13], whether flow is used to recover shape is less certain [22]. But since we often move our heads while inspecting specular objects (e.g., when locating a dent in a car), it seems plausible that shape recovery might utilize specular flow. In particular, the distinguished behavior at parabolic points would seem to play a role in this process.

Finally, similar analysis to that presented here might be used to recover shape from *still* images under natural lighting as well. The relationship between the image gradient ∇I in a still image and the angular derivatives ∇E at the corresponding point of an illumination sphere is $\nabla I = \mathbf{J}^\top \nabla E$, with $\mathbf{J}$ being the same as in the SFSF equation (Eq. 7). Analysis of this type could lead to statistical methods for shape recovery that exploit the structure of natural lighting [10], and it could help to elucidate the mechanisms underlying the human ability to recover specular shape from such images [11].

Acknowledgments

References

[1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Shape recovery from specular flow. BGU Technical Report 07-10, Ben-Gurion University, Computer Science Department, March 2007.

[2] S. Agarwal, S. Mallick, D. Kriegman, and S. Belongie. On refractive optical flow. pages 483–494, 2004.

[3] O. Ben-Shahar and S. Zucker. The perceptual organization of texture flows: A contextual inference approach. 25(4):401–417, 2003.

[4] A. Blake. Specular stereo. pages 973–976, 1985.

[5] A. Blake and G. Brelstaff. Geometry from specularities. 1988.

[6] A. Blake and H. Bülthoff. Does the brain know the physics of specular reflection? *Nature*, 343:165–168, 1990.

[7] A. Blake and H. Bülthoff. Shape from specularities: Computation and psychophysics. *Philosophical Transactions: Biological Sciences*, 331(1260):237–252, 1991.

[8] T. Bonfort and P. Sturm. Voxel carving for specular surfaces. pages 591–596, 2003.

[9] T. Bonfort, P. Sturm, and P. Gargallo. General specular surface triangulation. pages 872–881, 2006.

[10] R. Dror, T. Leung, E. Adelson, and A. Willsky. Statistics of real-world illumination. 2001.

[11] R. Fleming, A. Torralba, and E. Adelson. Specular reflections and the perception of shape. 4:798–820, 2004.

[12] M. Halstead, B. Barsky, S. Klein, and R. Mandell. Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. volume 1, pages 335–342, 1996.

[13] B. Hartung and D. Kersten. Distinguishing shiny from matte. 2(7):551–551, 2002.

[14] B. Horn and B. Schunck. Determining optical flow. 17:185–203, 1981.

[15] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. 3(6):661–669, 1981.

[16] K. Ikeuchi and B. Horn. Numerical shape from shading and occluding boundaries. 17:141–184, 1981.

[17] J. Koenderink and A. van Doorn. Photometric invariants related to solid shape. *Optical Acta*, 27(7):981–996, 1980.

[18] Light probe gallery. http://www.debevec.org/Probes/.

[19] M. Longuet-Higgins. Reflection and refraction at a random moving surface. I. pattern and paths of specular points. 50(9):838–844, 1960.

[20] M. Oren and S. Nayar. A theory of specular surface geometry. 24(2):105–124, 1997.

[21] S. Roth and M. Black. Specular flow and the recovery of surface structure. pages 1869–1876, 2006.

[22] S. Roth, F. Domini, and M. Black. Specular flow and the perception of surface reflectance. 3(9):413–413, 2003.

[23] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. 64(1):31–67, 2005.

[24] S. Savarese, F. Li, and P. Perona. What do reflections tell us about the shape of a mirror? 2004.

[25] J. Solem, H. Aanaes, and A. Heyden. A variational analysis of shape from specularities using sparse data. In *Proc. 3D Data Processing, Visualization, and Transmission*, pages 26–33, 2004.

[26] S. Waldon and C. Dyer. Dynamic shading, motion parallax and qualitative shape. In *Proc. IEEE Workshop on Qualitative Vision*, pages 61–70, New York City, NY, USA, 1993.

[27] J. Zheng and A. Murata. Acquiring a complete 3D model from specular motion under the illumination of circular-shaped light sources. 22(8):913–920, 2000.

[28] A. Zisserman, P. Giblin, and A. Blake. The information available to a moving observer from specularities. 7:38–42, 1989.