Shape from Specular Flow

Yair Adato, *Student Member*, *IEEE*, Yuriy Vasilyev, *Student Member*, *IEEE*, Todd Zickler, *Member*, *IEEE*, and Ohad Ben-Shahar, *Member*, *IEEE*

Abstract—An image of a specular (mirror-like) object is nothing but a distorted reflection of its environment. When the environment is unknown, reconstructing shape from such an image can be very difficult. This reconstruction task can be made tractable when, instead of a single image, one observes relative motion between the specular object and its environment, and therefore, a motion field—or *specular flow*—in the image plane. In this paper, we study the shape from specular flow problem and show that observable specular flow is directly related to surface shape through a nonlinear partial differential equation. This equation has the key property of depending only on the relative motion of the environment while being independent of its content. We take first steps toward understanding and exploiting this PDE, and we examine its qualitative properties in relation to shape geometry. We analyze several cases in which the surface shape can be recovered in closed form, and we show that, under certain conditions, specular shape can be reconstructed when both the relative motion and the content of the environment are unknown. We discuss numerical issues related to the proposed reconstruction algorithms, and we validate our findings using both real and synthetic data.

Index Terms—Specular objects, specular flow, shape reconstruction, environment motion field, Gaussian curvature, parabolic points, specular curvature.

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 if and how 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. This is leveraged, for example, in some forms of anamorphic art (Fig. 1).

Despite the intrinsic difficulties, the human visual system seems quite adept at inferring specular shape in unknown environments, even when few other shape cues are available (Fig. 1). Computational solutions have proven elusive, however, and existing methods are limited to recovering only sparse or qualitative shape information, considering a limited class of surfaces, or requiring calibrated conditions in which the environment structure is known.

In contrast to the existing methods, this paper presents a theory for dense and quantitative specular surface reconstruction that specifically considers 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. This model has two key features as follows:

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 simplifies the reconstruction problem.
2. The camera is set fixed relative to the object and it observes a relative motion between the specular surface and the environment. This induces a *specular flow* [1], [2] on the image plane, which 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 and assuming no interreflections, 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 in closed form. Here, we explore such cases in detail.

We begin our analysis in a two-dimensional world (Section 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 by solving a separable nonlinear ordinary differential equation (ODE) with initial conditions provided by an occlusion boundary.

We then consider the three-dimensional world (Section 4), where the specular object is a surface, the environment is a spherical map, and the environment motion is described by a spherical vector field. Fundamental to our analysis is the derivation of a pair of coupled nonlinear partial differential equations (PDEs) that relate specular flow to the surface shape. While these equations may be difficult to solve in general, we show that, for a special class of environment rotations, they can be reduced to a system of decoupled linear equations which can be solved to recover the shape in

- *Y. Adato and O. Ben-Shahar are with the Department of Computer Science, Ben-Gurion University of the Negev, PO Box 653, Beer-Sheva 84105, Israel. E-mail: {adato, ben-shahar}@cs.bgu.ac.il.*
- *Y. Vasilyev and T. Zickler are with the School of Engineering and Applied Science, Harvard University, 33 Oxford Street, Cambridge, MA 02138. E-mail: vasilyev@fas.harvard.edu, zickler@eecs.harvard.edu.*

Manuscript received 24 Sept. 2008; revised 2 Apr. 2009; accepted 20 June 2009; published online 28 June 2010.
Recommended for acceptance by K. Kutulakos.
For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2008-09-0641.
Digital Object Identifier no. 10.1109/TPAMI.2010.126.

0162-8828/10/$26.00 © 2010 IEEE Published by the IEEE Computer Society

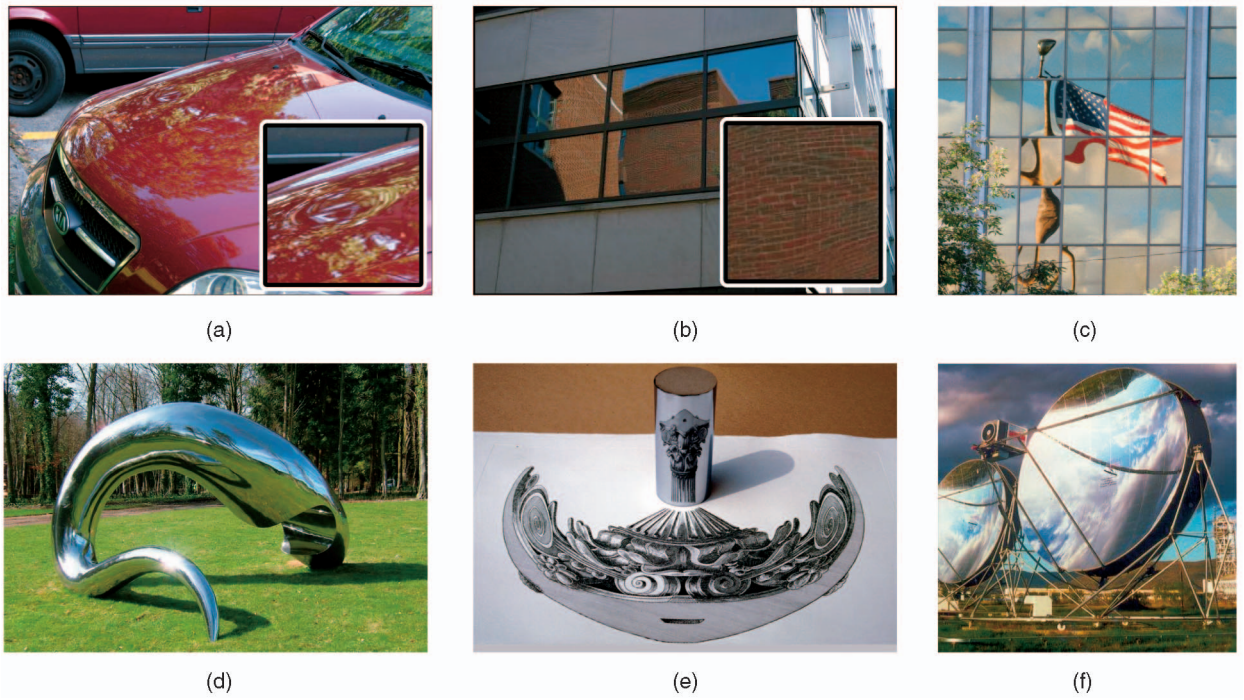


Fig. 1. Specular surfaces under dense, (approximately) far-field illumination are frequently encountered in everyday life (a)-(c), in art (d), (e), and in technological applications (f). These images often convey useful shape information, such as the dent in the car (a), or the imperfections in a building's window (b), (c). Can this information be extracted computationally when the environment is unknown? The fact that one can manipulate the environment to create almost any image from any specular object (e) casts doubt on the ability to do so.

closed form when sufficient initial conditions are available (Section 5).

Armed with a solution to this special case, we turn to the general case of *arbitrary* environment rotations in three dimensions (Section 6) and show that closed-form solutions can be obtained in this case as well if one combines observations from multiple, distinct environment rotations. Somewhat surprisingly, we find that this is true even when the environment rotations are unknown. This yields an "autocalibrating" reconstruction procedure that requires virtually no knowledge about the environment content or its axes of rotation.

It is important to note that the focus of this paper is the recovery of shape once specular flow is known, and for this reason, we do not directly address the problem of estimating a specular flow field from image data. Although many existing optical flow techniques exhibit excellent performance in classical cases (see [3] for an extensive comparison), the unique characteristics and singularities of specular flows suggest that their estimation from image data requires a significantly different approach. Whether or not this task can be achieved reliably and what novel techniques would facilitate this computation are important directions for future work (see Section 7) and we expect that the analysis presented in this paper will be useful for this problem as well.

2 Related Work

Most studies of the relationship between specular reflections and surface shape avoid the complications of natural lighting by considering environments that contain a single point light source. In these cases, one observes a small number of isolated "specularities" and, provided they can be accurately localized, each of these induces a constraint between a surface point, its normal, and its local view and illumination directions. When the camera and source positions are known, for example, the surface normal and the surface depth at a specular point are both determined up to a one-parameter family. Constraints of this type have been used computationally for different tasks, including recognition (e.g., [4]) and surface reconstruction (e.g., [5], [6], [7]). In a similar vein, complex illumination environments that are known and controlled have been used to obtain higher order surface information (e.g., curvature) [8], [9], [10], [11], [12], [13], [14], [15] and to extract shape from multiple specular bounces [16].

When two (stereo) views of a specular highlight are available and illumination consists of a known point source, the apparent shift of the specular highlight relative to the surface constrains the surface curvature [17], [18]. While less information is available when the position of the source is unknown, two views of a specular reflection are still sufficient to determine whether a surface is locally concave or convex [19], [20].

Not unrelated to these "specular stereo" approaches are methods that seek to exploit the *motion* of sparse specularities. Qualitatively, it is known that as the observer moves, specularities are created and annihilated in pairs at (or, in the near-field case [21], close to) parabolic surface points [22], [23]. More quantitatively, theory suggests that one can recover a complete surface curve by observing the specular motion induced by continuous camera motion [19], although the existing techniques for doing so are restricted to convex (or concave) surfaces without parabolic points [24]. Dense 3D shape for more general surfaces can be obtained by observing rotating objects under properly configured extended light sources, and by integrating the profiles from tracked specularities [25], [26].

All of the methods mentioned above rely on either point source environments or complex environments that are known and calibrated. The inference of specular shape in natural, uncontrolled illumination environments has received significantly less attention. Such "real-world" environments induce *dense* specular reflections (Fig. 1) 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 [27]) infer accurate shape, even when the illumination environment and bounding contour of the surface are unknown [28]. 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 [28], an observation that has been exploited in the extraction of local shape features and the nonphotorealistic rendering of specular objects [29] and the detection of specular surfaces in natural images [30].

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 the object, the viewer, and/or the environment. In a qualitative analysis, Waldon and Dyer [1] 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. (This is directly related to the pairwise creation and destruction of sparse specularities described above.) More quantitatively, Roth and Black [2] present an optical flow algorithm that estimates a specular flow field from an image sequence and simultaneously identifies a surface. Unlike our work, their system considers flow estimation as well as reconstruction, and it can handle surfaces with mixtures of specular and diffuse materials. The main limitation of their approach is the set of surfaces it considers, because the reconstruction amounts to choosing a member from a parameterized family of convex implicit surfaces (e.g., spheres of varying radii).

In contrast to the previous work, we seek quantitative shape recovery for general free-form surfaces that are not constrained to be convex/concave or of any particular parametric form. We consider unknown, natural illumination environments, and surfaces that are free of diffuse texture that could otherwise assist in the reconstruction process (e.g., [2]). The main contribution of our work is to show that recovering specular shape in these conditions is tractable.

As a final note, it is worth recognizing the connection between our work and the problem of nonparametric camera calibration [31], [32], which also benefits from observations of small motions. The difference of our work is that we explicitly interpret the nonparametric image plane "warp" as being caused by reflection from a single continuous specular surface.

3 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. This also allows us to introduce our theoretical framework in simplified terms. 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).

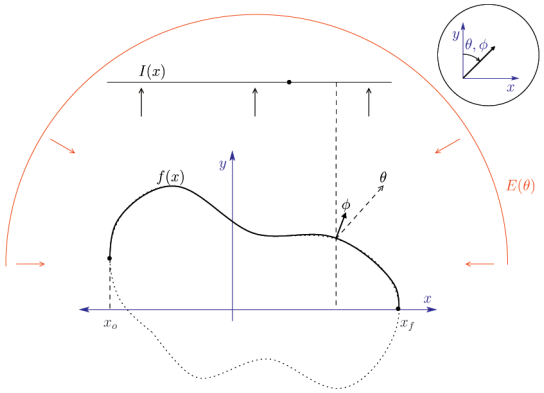


Fig. 2. The specular shape reconstruction problem in two dimensions. A surface profile $f(x)$ (a plane curve) is illuminated by a far-field illumination environment $E(\theta)$, and is viewed orthographically to produce a 1D image $I(x)$. The sign convention for the angular dimension is shown in the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

Here, we show that, under the conditions of our model (i.e., far-field illumination and observer), one can analytically recover an arbitrary continuous 2D "surface" from the observed specular flow.

As 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 $\theta = 0$, 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 $u = dx/dt$ on the image plane. From the sign convention in Fig. 2, it follows that

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

Taking temporal derivatives of this last 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 of the curve 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 exemplified in Fig. 3, projections of these points behave as either "sources" or "sinks" of specular flow, in accordance with the pairwise specular "birth" and "death" that is expected at parabolic lines in three dimensions [22], [23]. In this example, where the environment rotates in a counterclockwise manner, the flow is divergent (expands "outward" 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

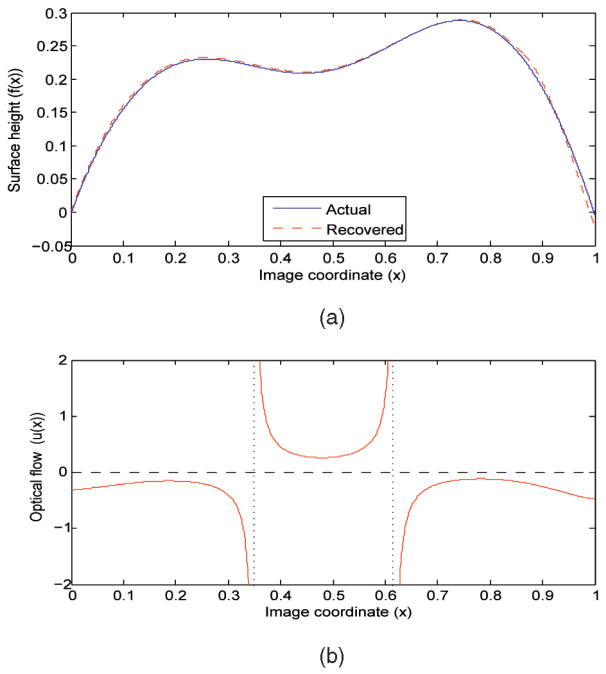


Fig. 3. Recovering a surface profile from specular flow. A surface profile $f(x)$ (blue solid curve in (a)) is viewed under a rotating environment as depicted in Fig. 2. This induces a specular flow (b) that is singular at inflection points. Using this flow, we recover the surface by solving (4) using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

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 (1) to obtain a Riccati equation:

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

This second order nonlinear 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 integration 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.

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 counterclockwise direction. While the content of the environment is irrelevant, for this experiment we extracted it from a great circle of the captured "St. Peter's" environment map [33]. 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 (2) using the leftmost point of the curve as an initial condition.

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 parametrization 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 reparameterize the surface derivative f_x using a stereographic projection [34], according to which we define the following:

$$\begin{aligned} q &= \frac{2f_x}{1+\sqrt{f_x^2+1}}, \\ f_x &= \frac{4q}{4-q^2}. \end{aligned} \tag{3}$$

Substituting in (2) yields an equation of the same form:

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

whose solution is

$$q(x) = 2\tan\left(-\frac{\omega}{4}\int_{x_o}^{x}\frac{d\lambda}{u(\lambda)} + C\right). \tag{5}$$

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

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

To recover the surface $f(x)$, the solution $q(x)$ is transformed via (3) and then integrated as before. This entire solution approach is demonstrated in Fig. 4, where, unlike in Fig. 3, the initial conditions were extracted automatically from the observable occluding points of the unknown curve.

3.2 Observations

Since it enables the recovery of surface shape, we refer to (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 ("texture"), 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. However, surface reconstruction requires the integration of the *inverse flow*, which is well defined everywhere. We note that this ability to handle general free-form surfaces differentiates our approach from that of Oren and Nayar [24], which represents 2D curves

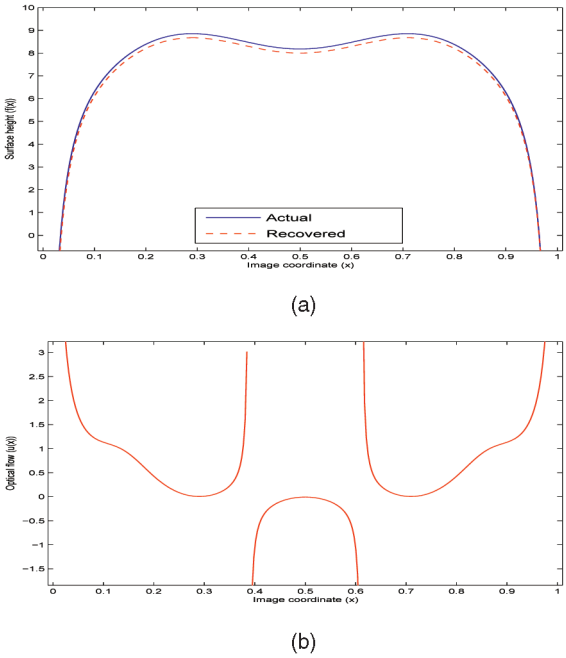


Fig. 4. 2D shape from specular flow with (5) and initial condition extracted automatically from an observable occluding boundary point. All graphs are color-coded as in Fig. 3.

using a support function, and therefore, requires them to be convex.

The 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 Section 5, many of these desirable properties carry over to the three-dimensional case as well.

4 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 (bivariate) function. As before, the surface is viewed orthographically from above and illuminated by a far-field environment (see Fig. 5).

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 samples 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, with the zenith measured from the viewing direction). In particular, we represent reflection directions as (α, β) and normal directions as (θ, ϕ), both under the sign convention shown in Fig. 5. 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:

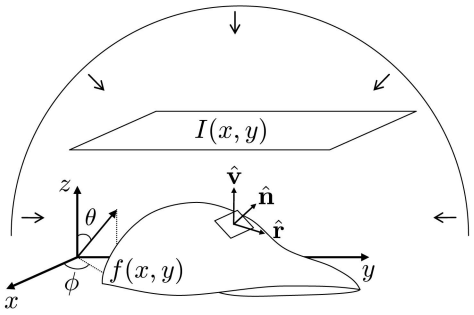


Fig. 5. The specular shape reconstruction problem in three dimensions. A surface $f(x,y)$ is illuminated by a far-field illumination environment and is viewed orthographically to produce a 2D image $I(x,y)$. The illumination sphere is parameterized using spherical coordinates (θ, ϕ).

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

From this relationship, we can deduce that

$$\begin{aligned}\tan\alpha &= \frac{2\|\nabla f\|}{1 - \|\nabla f\|^2},\\ \tan\beta &= \frac{f_y}{f_x}.\end{aligned} \tag{6}$$

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* (EMF)[1]:

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

This is a general representation that can describe both "rigid" motion (i.e., an environment that rotates around some fixed axis relative to the object) and an arbitrary "nonrigid" motion. In the rigid case, which is the focus of our analysis, any small motion corresponds to the environment rotating about an axis, say,

$$\hat{\mathbf{a}} = (\sin\alpha_\circ\cos\beta_\circ, \sin\alpha_\circ\sin\beta_\circ, \cos\alpha_\circ)$$

with angular velocity ω. In this case, the EMF can be shown to take the following analytic form:

$$\begin{aligned}\omega_\alpha(\alpha,\beta) &= \omega\sin\alpha_\circ\sin(\beta_\circ - \beta),\\ \omega_\beta(\alpha,\beta) &= \omega(\cos\alpha_\circ - \sin\alpha_\circ\cos(\beta - \beta_\circ)\cot\alpha).\end{aligned} \tag{7}$$

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

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

1. In the absence of interreflections, the law of specular reflection guarantees that the coordinates (α, β) of a surface reflection vector will be equal to those of the corresponding point in the environment sphere. Thus, we use the same notation for both.

is related to the environment motion through the Jacobian:

$$\boldsymbol{\omega} = \frac{d(\alpha,\beta)}{dt} = \frac{\partial(\alpha,\beta)}{\partial(x,y)}\,\frac{d(x,y)}{dt} = \mathbf{J}\mathbf{u}. \tag{8}$$

The Jacobian $\mathbf{J}$ can be expressed in terms of surface shape by taking temporal derivatives of (6), which yields

$$\mathbf{J} \triangleq \frac{\partial(\alpha,\beta)}{\partial(x,y)} = \begin{pmatrix} \dfrac{f_x f_{xx} + f_y f_{xy}}{\|\nabla f\| \cdot (1+\|\nabla f\|^2)} & \dfrac{f_x f_{xy} + f_y f_{yy}}{\|\nabla f\| \cdot (1+\|\nabla f\|^2)} \\ \dfrac{f_x f_{xy} - f_y f_{xx}}{2\|\nabla f\|^2} & \dfrac{f_x f_{yy} - f_y f_{xy}}{2\|\nabla f\|^2} \end{pmatrix}. \tag{9}$$

Implicit in (8) is the fact that the environment direction reflected by each point is determined by the surface geometry, i.e.,

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

Hence, when the environment motion field $\boldsymbol{\omega}$ is known and the induced specular flow $\mathbf{u}$ is observed (and measured) on the image plane, (8) becomes a system of nonlinear PDEs in the unknown shape $f(x,y)$. We refer to this equation as the SFSF equation in three dimensions. Solving this equation is equivalent to reconstructing the unknown specular shape.

4.1 Behavior at Parabolic Points

Equation (8) may be used to solve for an unknown shape $f(x,y)$ from a known specular flow (u,v), but we can also invert 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 $\boldsymbol{\omega}$:

$$\mathbf{u} = \mathbf{J}^{-1}\boldsymbol{\omega}. \tag{10}$$

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

$$Det(\mathbf{J}) = \frac{2K(1+\|\nabla f\|^2)}{\|\nabla f\|}, \tag{11}$$

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

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

Equation (11) 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.[2] This is 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 [1], [10], where the flow singularities can drift away from parabolic points.) A synthetic example of this phenomenon, with the specular flow computed using (10), is demonstrated in Fig. 6.

The unique behavior of specular flow near parabolic points must be considered when one seeks to estimate specular flow from image data. Like standard optical flow, the aperture problem means that purely local estimation is underconstrained and that regularization is required to accomplish the task. Most existing regularization schemes for optical flow are unsuitable, however, because they are designed to handle piecewise smooth flows with finite magnitude. While not considered in this paper, it is possible that the differential geometric analysis provided above will prove useful in designing more suitable regularization schemes.

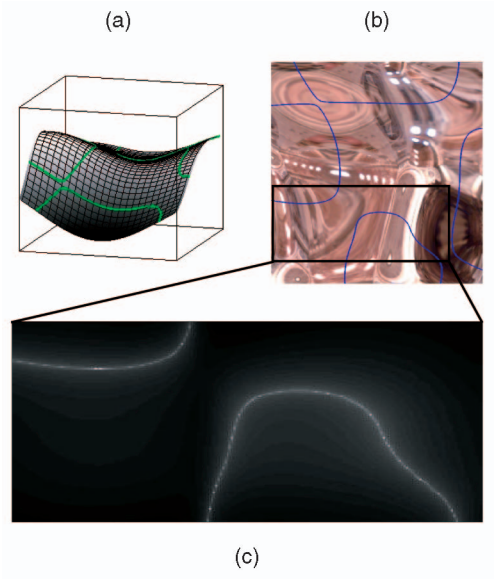


Fig. 6. The magnitude of the specular flow induced by a specular surface under environment motion. (a) The shape of a synthetic specular surface with parabolic lines superimposed. (b) One frame of a rendered image sequence under a rotating environment. Superimposed are the projection of the parabolic curves. (c) A zoomed image of the log-magnitude of the induced specular flow in the marked region of interest. Note how the flow magnitude is incomparably larger (even in the logarithmic scale) near parabolic lines than in any other surface point.

5 Environment Motion about the View Direction

Solving the 3D SFSF (8) in general may be a very difficult task as is usually the case with nonlinear PDEs. One special case in which this equation 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 $\hat{\mathbf{v}}$ (i.e., $\hat{\mathbf{a}} = (0,0,1)$). In our spherical coordinate system, environment rotation about the view direction induces the motion field

$$\begin{aligned} \omega_\alpha(\alpha,\beta) &= 0, \\ \omega_\beta(\alpha,\beta) &= \omega, \end{aligned} \tag{12}$$

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:

2. Singularities also occur at points where $\|\nabla f\| = 0$, which is caused by our choice of spherical coordinates. In practice, these points are likely to 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.

$$\begin{aligned} h(x,y) &\triangleq f_x^2 + f_y^2, \\ k(x,y) &\triangleq \tan^{-1}(f_y/f_x). \end{aligned} \tag{13}$$

The SFSF equation then reduces to a particularly simple form:

$$\begin{pmatrix} h_x & h_y \\ k_x & k_y \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} 0 \\ 2\omega \end{pmatrix}, \tag{14}$$

with the derivatives of h and k derived directly from (13):

$$\begin{aligned} h_x &= 2(f_x f_{xx} + f_y f_{xy}), \quad & h_y &= 2(f_x f_{xy} + f_y f_{yy}), \\ k_x &= \frac{f_x f_{xy} - f_y f_{xx}}{\|\nabla f\|^2}, & k_y &= \frac{f_x f_{yy} - f_y f_{xy}}{\|\nabla f\|^2}. \end{aligned} \tag{15}$$

For convenience, we define

$$\tilde{\mathbf{J}} \triangleq \begin{pmatrix} h_x & h_y \\ k_x & k_y \end{pmatrix}, \tag{16}$$

in which case (14) can be written as

$$\tilde{\mathbf{J}}\mathbf{u} = (0, 2\omega).$$

Equation (14) represents a pair of *decoupled*, *linear* PDEs in h and k. Hence, by observing the specular flow $\mathbf{u} = (u, v)$, one can solve each of these equations in closed form using the method of characteristics. The characteristic curves for both equations are the same, and they 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., ∇f is known at one or more points along each integral curve of the specular flow. Once $h(x,y)$ and $k(x,y)$ are known, the surface gradient can be recovered by inverting the expressions in (13):

$$\begin{aligned} f_x &= \sqrt{h}\cos k, \\ f_y &= \sqrt{h}\sin k, \end{aligned} \tag{17}$$

and derivative functions can be integrated to recover the surface $f(x,y)$. 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.

5.1 Shape Recovery under *Unknown* Angular Velocity

In addition to facilitating a shape recovery procedure, (14) has a straightforward and useful geometric interpretation: Its two linear equations constrain the directional derivatives of the magnitude ($h(x,y)$) and orientation ($k(x,y)$) of the surface gradient along integral curves of specular flow.

In particular, consider the second equation that constrains the behavior of $k(x,y)$. That equation can be rewritten in terms of the *unit* specular flow $\hat{\mathbf{u}}$:

$$\hat{\mathbf{u}} \cdot \nabla k = \frac{\mathbf{u}}{\|\mathbf{u}\|} \cdot \nabla k = \frac{2\omega}{\|\mathbf{u}\|}, \tag{18}$$

which implies that the rate of change in the surface gradient orientation along each arc-length-parameterized characteristic must be proportional to the angular velocity ω and inversely proportional to the flow magnitude $\|\mathbf{u}\|$. Since k is an angle, 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., [35]). Given the specular flow $\mathbf{u}$ induced by environment rotation about $\hat{\mathbf{v}}$, we can therefore define its *specular curvature* as the covariant derivative $\nabla_{\hat{\mathbf{u}}} k$, i.e.,

$$\kappa_s \triangleq \nabla_{\hat{\mathbf{u}}} k = \nabla k \cdot \hat{\mathbf{u}} = \frac{2\omega}{\|\mathbf{u}\|}. \tag{19}$$

Moreover, since the characteristics (the integral curves of the specular flow) are generically closed curves, integrating κ_s along each such curve must yield $2N\pi$, where N is an integer, in order to satisfy surface integrability. While there are cases where N takes arbitrary integral value, generically it is either 0 or 1, depending on the local shape of the surface patch and whether or not the characteristic intersects parabolic curves on the surface (see Fig. 7). While a complete account of this issue is outside the scope of this paper, we prove the following weaker proposition, which is sufficient for our purposes.

Proposition 1. *Let p be an elliptic extremum point of a bounded Monge patch $(x, y, f(x,y))$. Let $\mathbf{u}$ be the observed specular flow under a view-axis EMF. Then, there exists a closed characteristic curve $C(s), s \in [0, l]$ of the SFSF (14) that goes around p and satisfies*

$$\oint_C \kappa_s ds = \oint_C \frac{2\omega}{\|\mathbf{u}\|} ds = 2\pi.$$

That is, $C(s)$ is a characteristic for which $k(x,y)$ completes exactly one revolution ($N = 1$).

Proof. Assume, without loss of generality, that $f(x,y)$ is given such that p is the origin and the frame axes coincide with the principal directions of the surface at p (or, if p is umbilical, assume arbitrary frame). A second order Taylor expansion of $f(x,y)$ in this coordinate frame is

$$f(x,y) = \frac{1}{2}\left(\kappa_1 x^2 + \kappa_2 y^2\right) + R$$

where κ_1 and κ_2 are the principal curvatures at p and R includes higher order terms. The surface gradient and the function $h(x,y)$ (13) are therefore approximated by

$$\begin{aligned} \nabla f &= (\kappa_1 x, \kappa_2 y), \\ h(x,y) &= \kappa_1^2 x^2 + \kappa_2^2 y^2, \end{aligned}$$

while the gradient of h (15) is approximated by

$$\nabla h(x,y) = \left(2\kappa_1^2 x, 2\kappa_2^2 y\right).$$

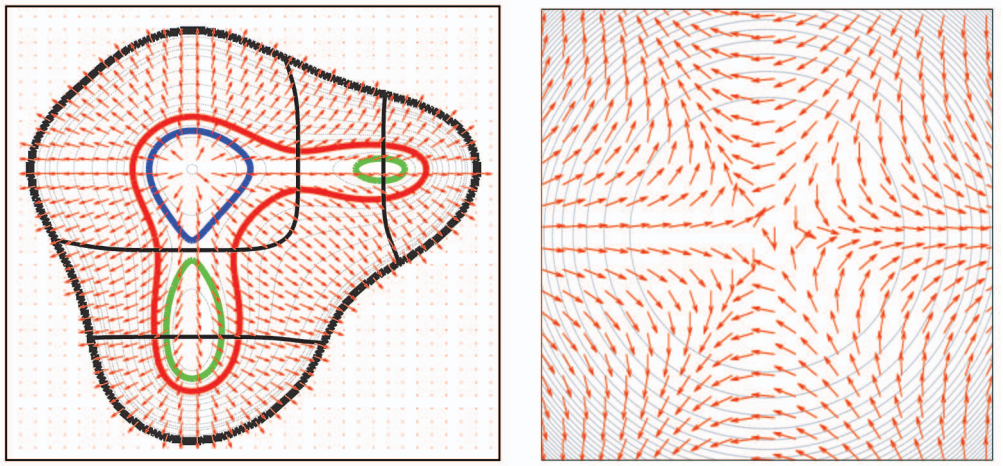
Since $\nabla h \cdot \mathbf{u} = 0$ (see 14), the gradient of h is always perpendicular to the specular flow, i.e., ∇h is everywhere normal to $C(s)$. Also note that $k(x,y)$ can be expressed in unit vector form $\hat{k}(x,y)$ as

$$\hat{k} = \frac{\nabla f}{\|\nabla f\|}.$$

Given all of the above, consider the angle (expressed as a dot product) between $\hat{k}$ and ∇h along a characteristic $C(s)$ that is sufficiently close to p:

$$\hat{k} \cdot \frac{\nabla h}{\|\nabla h\|} = \frac{2}{\|\nabla f\| \cdot \|\nabla h\|}\left(\kappa_1^3 x^2 + \kappa_3^2 y^2\right).$$

Since p is elliptic, κ_1 and κ_2 have the same sign and are both nonzero, which implies that the dot product never vanishes. Thus, $\hat{k}$ is nowhere perpendicular to ∇h and is



(a) (b)

Fig. 7. The specular curvature κ_s of the auxiliary function $k(x,y)$ (the surface gradient orientation field) integrates to $2N\pi$ along any closed characteristics, where generically, but not always, $N \in \{0,1\}$. (a) Several characteristics (in gray), the k function (depicted as a unit length vector field in red), and the parabolic lines (in black) of a particular shape (Fig. 12a). Several characteristics are highlighted by different colors. Note that k completes one full revolution ($N=1$) along the blue characteristic and along the red characteristic, but zero revolutions ($N=0$) along the green characteristics. (b) Degenerate surfaces can produce specular flows for which the specular curvature along closed characteristics of the SFSF equation integrates to $2N\pi$ with $N>1$. Shown is the flow field induced by the monkey saddle surface $f(x,y)=x^3-3xy^2$ for which the function k completes two full revolutions (i.e., κ_s integrates to 4π) along any characteristic curve.

nowhere tangent to the specular flow $\mathbf{u}$. Since this is equivalent to being nowhere tangent to the integral curve $C(s)$, $\hat{k}$ points either inside or outside $C(s)$ along $C(s)$'s entire length. This guarantees that $\hat{k}$ completes the same number of revolutions as the tangent of $C(s)$, and, following the theorem of turning tangents [36, p. 37], we conclude that

$$\oint_C \kappa_s ds = \oint_C \kappa_c ds = 2\pi,$$

where κ_c is the curvature of the curve $C(s)$. □

Given Proposition 1, suppose that a bounded specular surface is viewed under a view-axis EMF, but unlike the previous section, suppose the angular velocity ω is unknown. We claim that the specular shape can still be reconstructed since ω can be recovered automatically from the specular flow field by exploiting this proposition. Indeed, such bounded surfaces generically have at least one elliptic extremum point (the one closer to the viewer) having a normal that coincides with the viewing direction. Such points produce zero specular flow and can be identified in the flow field (see also footnote 2). A characteristic $C(s)$ around such a point that does not intersect a parabolic curve (which can also be identified in the flow field) satisfies Proposition 1, and by computing

$$\omega = \frac{\pi}{\oint_C \|\mathbf{u}\|^{-1} ds}, \tag{20}$$

the angular velocity ω can be recovered.

In summary, the following procedure provides dense surface shape from specular flow induced by a view-axis rotation of the environment with arbitrary and unknown angular velocity:

Reconstruction Algorithm I: View-axis Rotation
1) Recover ω using (20).
2) Solve (14) for $h(x,y)$ and $k(x,y)$.
3) Recover f_x and f_y using (17).
4) Integrate f_x and f_y to obtain $f(x,y)$.

5.2 Numerical Considerations

Carrying out the reconstruction procedure numerically requires special care. While integration can be accomplished using standard numerical tools, such as Runge-Kutta methods (e.g., Matlab's `ode45` PDE solver), doing so naively fails at the singularities—both in magnitude and orientation—that occur along parabolic curves. To insure numerical stability, we must preprocess the observed flows and rewrite the SFSF equation.

Throughout the paper, we have assumed that a specular flow field is given as input. Of course, any practical system that estimates such a flow field from image data will necessarily identify parabolic curves (as flow singularities) during the estimation process. Thus, if we assume that the flow fields are given as input, it is natural to assume that the locations of the parabolic curves are given as well. Since parabolic curves determine the Gaussian curvature function up to a binary choice, this means that we have access to a function g up to sign, where g is defined by

$$g(x,y) \triangleq \begin{cases} 1, & K(x,y) \geq 0, \\ -1, & K(x,y) < 0, \end{cases}$$

and $K(x,y)$ is the Gaussian curvature function.

The function $\pm g$ can be used to multiply both sides of (14) to effectively "flip" the flow direction in regions of positive (or negative) Gaussian curvature and remove orientation singularities. Coupled with this, we can scale the equations by the inverse flow magnitude to remove instabilities caused by large flow magnitudes. This yields the equations

$$\begin{aligned} \frac{u \cdot g}{\sqrt{u^2+v^2}} h_x + \frac{v \cdot g}{\sqrt{u^2+v^2}} h_y &= 0, \\ \frac{u \cdot g}{\sqrt{u^2+v^2}} k_x + \frac{v \cdot g}{\sqrt{u^2+v^2}} k_y &= \frac{2\omega \cdot g}{\sqrt{u^2+v^2}}, \end{aligned} \tag{21}$$

which have the same characteristic curves as those of (14) but can be solved numerically using standard tools. Since this manipulation flips the flow direction in some regions on the image plane, we refer to the function g as the "flipping function" in the sequel.

5.3 Experimental Results for View-Axis EMF

We have tested the theory and reconstruction procedures discussed above on both synthetic and real data. Synthetic data were generated by placing virtual specular surfaces in real far-field illumination environments [33] and rendering them under orthographic projection. This type of evaluation provides access to ground truth shape for comparison, and it also enables us to compute specular flow fields (using (10)) instead of measuring them from image data. One such example is shown in Fig. 8 and it includes both the measured and flipped specular flow, the flipping function, the computed characteristic curves versus the function h level sets (which represent the ground truth computation), and the reconstructed surface. To demonstrate the essential role of the flipping procedure, we also show the result of integrating selected characteristics *without* it (Fig. 8f). To compare the recovered surface against the ground truth, we report both the absolute reconstruction error (Fig. 8k) and the relative reconstruction error in percent of surface height at each point (Fig. 8l). Initial conditions for this example were supplied along the two main axes.

In addition to synthetic results, we also applied our reconstruction procedure to real image data acquired using the system depicted in Fig. 9a. A camera (Canon 10D, EF 75-300 mm lens, EF 25 II extension tube) was placed 1 m from a 5 cm-diameter chrome sphere. Both the camera and the object were placed in an unknown, far-field illumination environment (a large lecture hall), and rotated as a fixed pair about an axis parallel and close to the camera's optical axis. One frame of the captured image sequence is shown in Figs. 9b and 9c.

Specular flow was recovered from this image sequence using the Horn and Schunck algorithm [37]. Based on this flow, the surface was recovered as described in the previous section, with the important exception that the angular velocity was given as input ($\omega = 0.5^\circ$/frame) as opposed to being estimated from data. (This experimental setup precludes the existence of the closed integral curves required for estimating ω.) Initial conditions were provided manually by specifying the surface gradient along the red curve shown in the figure. The RMS error in the reconstruction as a fraction of the sphere radius was approximately 1 percent.

While a specular sphere is a relatively simple object, it provides the benefit of readily available ground truth shape for the purpose of quantitative performance evaluation. Unfortunately, no such data were available to us for more complex specular objects (and devising ways to construct such objects is part of our future work). Still, to obtain additional confirmation of the reconstruction abilities of our method for more general shapes, we have applied it to nontrivial, real specular objects with no ground truth shape and compared the results visually. One such example is shown in Fig. 10. As with the sphere, the corresponding specular flow was measured using the Horn and Schunck algorithm [37] and initial data were measured manually (and therefore, inevitably, inaccurately).

We reiterate that measuring specular flows reliably and correctly from images is a difficult problem in itself, which is not addressed in this paper. Standard approaches succeed in Fig. 9 due to the absence of parabolics, but to our knowledge, no existing optical flow technique can handle the unique singularities that occur when surfaces are nonconvex. Developing suitable new optical flow techniques is an important direction of future work (see Section 7).

In addition to these results on synthetic and captured data, we present a *theoretical* example that emphasizes the ability to obtain analytical solutions whenever the measured specular flow can be expressed analytically (either explicitly or via interpolation of discrete measurements). As before, consider the specular flow observed for a sphere under view-axis rotation. One can show that the resulting specular flow field is

$$u(x,y) = -\omega y,$$
$$v(x,y) = \omega x.$$

This expression for the specular flow can be used in (14) to obtain the Cauchy problem

$$-\omega y h_x + \omega x h_y = 0,$$
$$-\omega y k_x + \omega x k_y = \omega,$$

where the initial conditions are

$$\Gamma_h = \left\{ x = 0,\ y = \tau,\ h = \frac{\tau^2}{r - \tau^2} \right\},$$
$$\Gamma_k = \left\{ x = 0,\ y = \tau,\ k = \frac{\pi}{2} \right\}.$$

The first step toward solving this problem is to find its characteristic curves. Given their similar structure, both equations have the same characteristic curves

$$\frac{dx}{-y} = \frac{dy}{x} \quad \Rightarrow \quad -x dx = y dy \quad \Rightarrow \quad x^2 + y^2 = c.$$

Using the characteristic curves and Γ_h, we solve for h in a straightforward way

$$c = \tau^2 \quad \Rightarrow \quad x^2 + y^2 = \tau^2 \quad \Rightarrow \quad h = \frac{x^2 + y^2}{r - x^2 - y^2}.$$

Since the second equation is inhomogeneous, we require an additional step. We first note that

$$\frac{dy}{x} = dk \quad \Rightarrow \quad \frac{dy}{\sqrt{\tau^2 - y^2}} = dk \quad \Rightarrow \quad \arcsin\frac{y}{\tau} = k + c,$$

and, using the initial condition Γ_k, we conclude that $c = 0$ and

$$k = \arcsin\frac{y}{\sqrt{x^2 - y^2}}.$$

Finally, the surface derivatives are

$$f_x = \sqrt{h}\cos k = \frac{\sqrt{x^2 + y^2}}{\sqrt{r - x^2 - y^2}} \frac{x}{\sqrt{x^2 - y^2}} = \frac{x}{\sqrt{r - x^2 - y^2}},$$
$$f_y = \sqrt{h}\sin k = \frac{\sqrt{x^2 + y^2}}{\sqrt{r - x^2 - y^2}} \frac{y}{\sqrt{x^2 - y^2}} = \frac{y}{\sqrt{r - x^2 - y^2}},$$

and by integrating these derivative functions, the sphere is recovered as a height function

$$f = \sqrt{r - x^2 - y^2}.$$

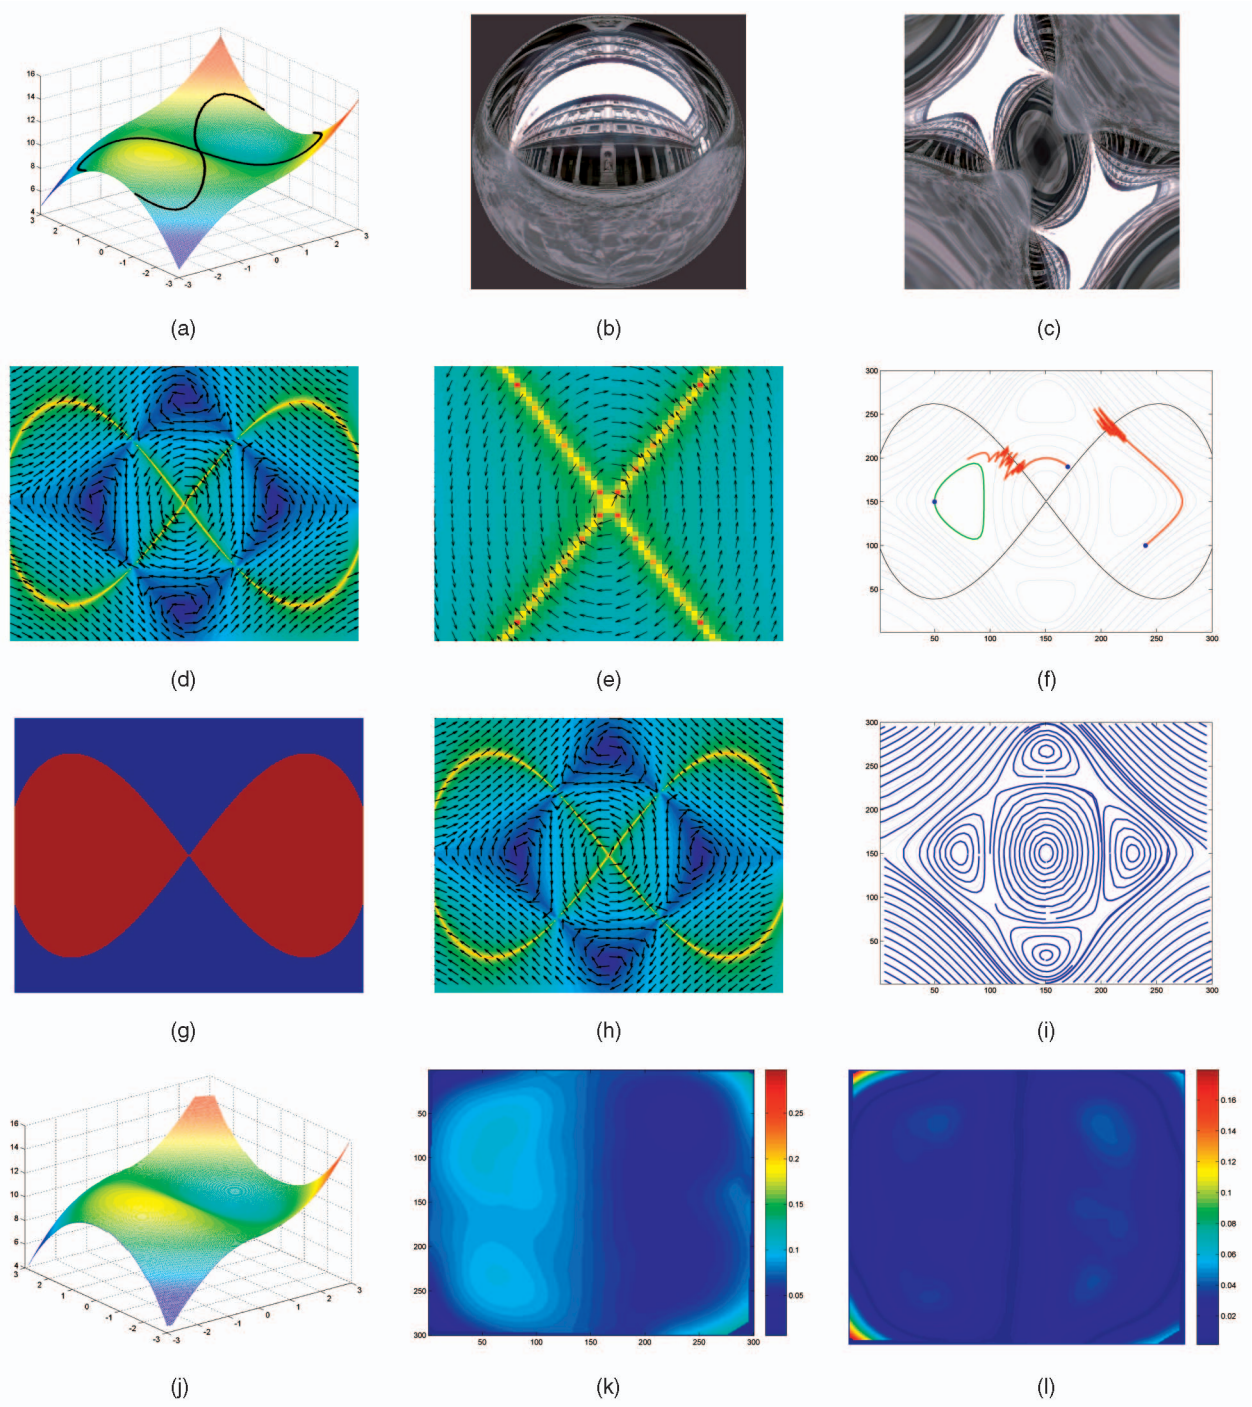


Fig. 8. Example of specular shape reconstruction of a general surface observed under environment rotation about the viewing direction. (a) The original surface $f(x, y) = 0.2ab^2 - 1.1\sin(a)$. This surface includes multiple extrema and several elliptic and hyperbolic regions. The parabolic lines that separate these regions are marked here in black. (b) The illumination environment (mapped on a disk of directions) in which the specular surface was observed [33]. (c) One image from the image sequence obtained under view-axis environment motion. (d) The observed specular flow field under this EMF. Flow magnitude is color-coded and orientation is shown by the superimposed unit-length vector field. (e) Detail of singularities in the flow field. Note the singularities in both magnitude and orientation. (f) Solving the SFSF equation in the presence of specular flow singularities is prone to fail. Note how the red integrations that begin from the blue points fail to follow the true level sets of the h function (show in light gray) as soon as they need to cross a singularity. Compare to the green curve that keeps off these singularities. (g) The flipping mask. Compare this mask to the specular flow magnitude in panel (c). (h) The flipped specular flow. Note how the orientation singularities were eliminated. (i) Solving the normalized SFSF equation provides robust solutions. Shown here are characteristics of the function h based on (21). (j) The reconstructed surface. (k) Absolute reconstruction error, expressed as percentage of surface relief. (l) Relative reconstruction error, expressed as percentage of surface height at each point.

6 General Environment Motion

The special case described above—the case of rotation about the view direction—is important because it enables closed-form solutions. However, it is quite "accidental." In this section, we relax this restriction on â while maintaining the ability to find closed-form solutions. We show that

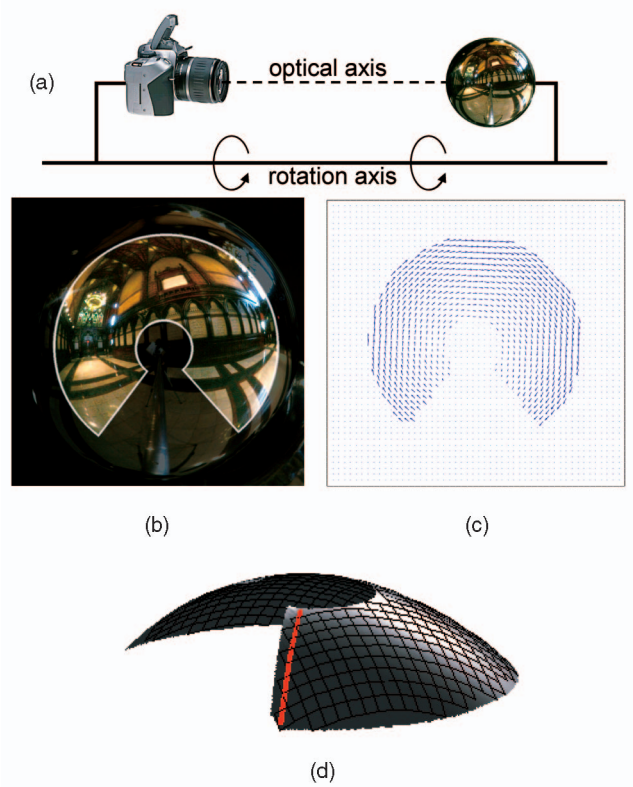


Fig. 9. Recovering general shape from specular flow in three dimensions. (a) Under far-field illumination, a camera and object rotate as a fixed pair around a line parallel to the optical axis. (b) One image from a captured sequence. (c) Estimated specular flow. (d) Shape recovered by solving the 3D shape from specular flow equation as described in Section 5. The surface gradient is assumed known along the red curve, which provides the necessary initial conditions.

this can be done for *arbitrary* rotation axes $\hat{\mathbf{a}}$ in the case when the object is observed under two or more known environment motions, each inducing a distinct specular flow. Furthermore, we show that with multiple specular flows, one can recover the shape even without knowing the environment motions under which they were observed. This leads to an "autocalibrating" reconstruction procedure that requires no knowledge about the environment content or its motion.

We begin by expressing the general EMF in (7) explicitly in terms of the surface and its derivatives. For this, we first note that (see also (6))

$$\begin{aligned}\cot\alpha &= \frac{1-\|\nabla f\|^2}{2\cdot\|\nabla f\|},\\ \sin\beta &= \frac{-f_y}{\|\nabla f\|},\\ \cos\beta &= \frac{-f_x}{\|\nabla f\|},\end{aligned} \tag{22}$$

and, substituting these expressions into (7), we obtain

$$\begin{aligned}\omega_\alpha(\alpha,\beta) &= \frac{\omega\sin\alpha_\circ}{\|\nabla f\|}(f_y\cos\beta_\circ - f_x\sin\beta_\circ),\\ \omega_\beta(\alpha,\beta) &= \omega\cos\alpha_\circ\\ &\quad + \omega\sin\alpha_\circ\frac{1-\|\nabla f\|^2}{2\cdot\|\nabla f\|^2}(f_x\cos\beta_\circ + f_y\sin\beta_\circ).\end{aligned} \tag{23}$$

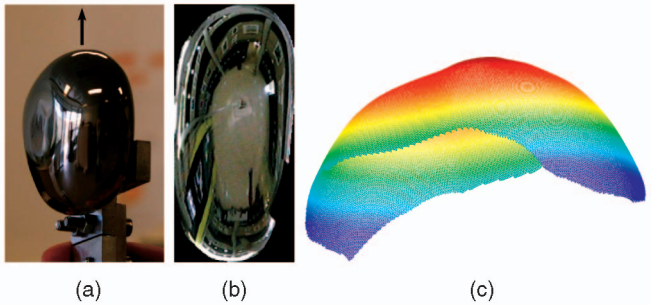


Fig. 10. Reconstruction of a nontrivial real specular object. (a) A blob-like specular object (mounted on a stand). The black arrow shows the viewing direction. (b) A single frame from the specular flow sequence captured by the camera. (c) The reconstructed shape. Compare (qualitatively, since no ground truth shape data were available) to the perceived shape in panel a.

Using this in conjunction with the auxiliary functions from (13), we can therefore rewrite the SFSF as follows:

$$\tilde{\mathbf{J}}\mathbf{u} = \omega\sin\alpha_\circ\begin{pmatrix}2(1+h)(f_y\cos\beta_\circ - f_x\sin\beta_\circ)\\ 2\cot\alpha_0 + \dfrac{1-h}{h}(f_x\cos\beta_\circ + f_y\sin\beta_\circ)\end{pmatrix}. \tag{24}$$

At first sight, this expression (24) seems to offer little advantage over the original SFSF equation ((8) and (9)). Indeed, while it has been reduced from second order to first order, the equations are still coupled and nonlinear. Even worse, it now involves three unknown and coupled functions (h, k, and f). As we will see, however, this form is the key to the multiple flow approach discussed next.

6.1 Two Equal-Azimuth Motions

Consider a specular object viewed under two different EMFs defined by the two different rotation axes,

$$\begin{aligned}\hat{\mathbf{a}}_1 &= (\alpha_1,\beta_0),\\ \hat{\mathbf{a}}_2 &= (\alpha_2,\beta_0),\end{aligned} \tag{25}$$

and angular velocities $\omega_1 \neq 0$ and $\omega_2 \neq 0$, respectively. Assume that $\hat{\mathbf{a}}_1$, $\hat{\mathbf{a}}_2$, and the viewing direction are all in the same plane, i.e., the two rotation axes have identical azimuth angle, as depicted graphically in Fig. 11. Note that if either $\alpha_1 = 0$ or $\alpha_2 = 0$, the corresponding EMF reduces to the special solvable case of rotation around the viewing direction. We therefore assume that neither zenith angle vanishes.

Let $\mathbf{u}_1 = (u_1, v_1)$ and $\mathbf{u}_2 = (u_2, v_2)$ be the specular flows obtained due to the first and second EMFs, respectively, and consider the following "normalized" specular flows $\tilde{\mathbf{u}}_i$ for $i \in \{1,2\}$:

$$\tilde{\mathbf{u}}_i = \begin{pmatrix}\tilde{u}_i\\ \tilde{v}_i\end{pmatrix} \triangleq \frac{\mathbf{u}_i}{\omega_i\sin\alpha_i} = \left(\frac{u_i}{\omega_i\sin\alpha_i}, \frac{v_i}{\omega_i\sin\alpha_i}\right). \tag{26}$$

Since $\omega_i \neq 0$ and $\alpha_i \neq 0$, both normalized flows are well defined. By rewriting (24) in terms of $\tilde{\mathbf{u}}_i$, we now obtain the following two SFSF equations for $i \in \{1,2\}$:

$$\tilde{\mathbf{J}}\tilde{\mathbf{u}}_i = \begin{pmatrix}2(1+h)(f_y\cos\beta_\circ - f_x\sin\beta_\circ)\\ 2\cot\alpha_i + \dfrac{1-h}{h}(f_x\cos\beta_\circ + f_y\sin\beta_\circ)\end{pmatrix}.$$

This equation is significant because it is linear in the arbitrary azimuthal angles α_1 and α_2. Hence, by subtracting the equations due to the first and second EMFs, we can obtain

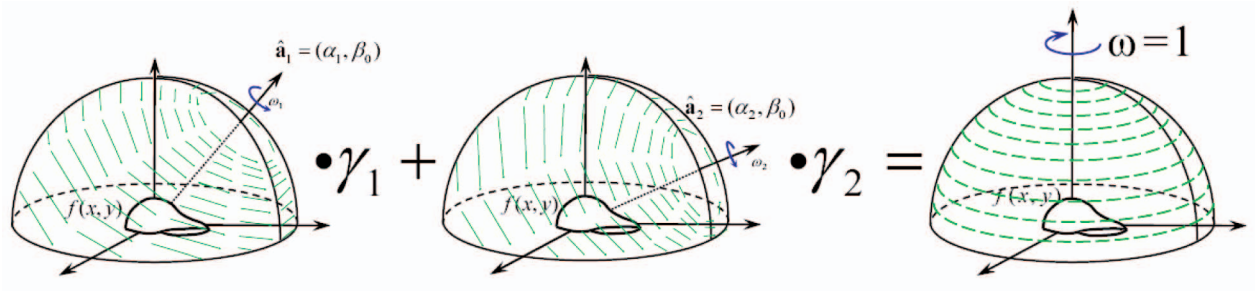


Fig. 11. Two EMFs with rotation axes of equal azimuthal angle. A linear combination of these two vector fields is equal to an EMF around the viewing direction. Since the SFSF is solvable in the latter case, the specular shape can be recovered after being observed under two such EMFs.

$$\tilde{\mathbf{J}}(\tilde{\mathbf{u}}_1 - \tilde{\mathbf{u}}_2) = \begin{pmatrix} 0 \\ 2\cot\alpha_1 - 2\cot\alpha_2 \end{pmatrix}. \tag{27}$$

Since $\tilde{\mathbf{u}}_1 - \tilde{\mathbf{u}}_2$ is easily computed from the measured pair of specular flows, (27) is a pair of decoupled linear PDEs in the unknown h and k and the known EMFs. This equation is in exactly the same form as that of a single, view-axis EMF, and it is therefore solvable using the technique discussed in Section 5.

6.2 Combining EMFs Linearly

Equation (27) implies that a linear combination of two measured specular flows (under equal-azimuth motions) gives an SFSF equation of the type obtained for a single EMF around the viewing direction. This should not come as a surprise. In fact, (7) already suggests that if $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$ are two equal-azimuth EMFs (cf., 25), then there exist two constants γ_1 and γ_2 such that their linear combination

$$\gamma_1\boldsymbol{\omega}_1 + \gamma_2\boldsymbol{\omega}_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \tag{28}$$

corresponds to a unit-speed, view-axis EMF (Fig. 11). Applying this requirement first to $\omega_\alpha(\alpha, \beta)$ from (7) implies that γ_1 and γ_2 must satisfy

$$\gamma_1\omega_1 \sin\alpha_1 + \gamma_2\omega_2 \sin\alpha_2 = 0,$$

a constraint which can be trivially satisfied by the normalization factor used in (26), i.e.,

$$\gamma_1 = \frac{1}{\omega_1 \sin\alpha_1}, \ \gamma_2 = -\frac{1}{\omega_2 \sin\alpha_2}.$$

Next, by dividing both γ_1 and γ_2 by the common factor $2\cot\alpha_1 - 2\cot\alpha_2$ that appears in (27), we also satisfy the requirement on $\omega_\beta(\alpha, \beta)$ and find the final γ_1 and γ_2.

Exploiting the observation above, we can now rewrite the SFSF for the special case of rotation around the viewing direction (and with angular velocity $\omega = 1$) as a linear combination of the general SFSF equations due to two equal-azimuth EMFs

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \sum_{i=1}^{2} \gamma_i\boldsymbol{\omega}_i = \mathbf{J}\sum_{i=1}^{2} \gamma_i\mathbf{u}_i = \tilde{\mathbf{J}}\sum_{i=1}^{2} \frac{\gamma_i\mathbf{u}_i}{2}, \tag{29}$$

which clearly reduce to (27) developed above.

Generalizing the insights discussed above, we note that, given $n \geq 3$ EMFs $\boldsymbol{\omega}_i$, at least three of which have distinct rotation axes $\hat{\mathbf{a}}_i$, one can find a set of constants γ_i such that the linear combination

$$\sum_{i=1}^{n} \gamma_i\boldsymbol{\omega}_i = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \tag{30}$$

is a view-axis EMF with unit angular speed. If the rotation axes (α_i, β_i) and angular velocities ω_i are known, one can use (7) to solve for γ_i and express them in terms of these parameters. Then, as in the two-EMF case, the view-axis SFSF can be rewritten as a linear combination of general SFSF equations due to the general EMFs,

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \sum_{i=1}^{n} \gamma_i\boldsymbol{\omega}_i = \mathbf{J}\sum_{i=1}^{n} \gamma_i\mathbf{u}_i = \tilde{\mathbf{J}}\sum_{i=1}^{n} \frac{\gamma_i\mathbf{u}_i}{2}, \tag{31}$$

which can be solved using the technique discussed in Section 5.

6.3 Three Arbitrary Motions

Section 6.1 showed that the shape from specular flow problem can be solved for two *general* EMFs as long as they obey one restriction—their rotation axes must be coplanar with the view axis. Equation (31) suggests, however, that this restriction can be eliminated once we use three (or more) EMFs. Indeed, consider a specular object viewed under three different EMFs defined by the three different and arbitrary rotation axes

$$\hat{\mathbf{a}}_i = (\alpha_i, \beta_i) \quad i \in \{1, 2, 3\}$$

and angular velocities $\omega_i \neq 0$, $i \in \{1, 2, 3\}$, respectively. Let $\mathbf{u}_i = (u_i, v_i), i \in \{1, 2, 3\}$, be the corresponding specular flows observed in the image plane for each EMF and consider the following "normalized" specular flows $\tilde{\mathbf{u}}_i$ for $i \in \{1, 2, 3\}$:

$$\tilde{\mathbf{u}}_i \triangleq \frac{\mathbf{u}_i}{\omega_i \sin\alpha_i \cos\beta_i} = \left(\frac{u_i}{\omega_i \sin\alpha_i \cos\beta_i}, \frac{v_i}{\omega_i \sin\alpha_i \cos\beta_i}\right).$$

By rewriting (24) in terms of these definitions, we now get the following three equations:

$$\tilde{\mathbf{J}}\tilde{\mathbf{u}}_i = \begin{pmatrix} 2(1+h)(f_y - f_x\tan\beta_i) \\ 2\dfrac{\cot\alpha_i}{\cos\beta_i} + \dfrac{1-h}{h}(f_x + f_y\tan\beta_i) \end{pmatrix} \quad i \in \{1, 2, 3\}.$$

Then, by subtracting proper weighted combinations of pairs of equations, we obtain

$$\tilde{\mathbf{J}}\begin{pmatrix} \dfrac{\tilde{u}_1 - \tilde{u}_2}{\tan\beta_2 - \tan\beta_1} - \dfrac{\tilde{u}_1 - \tilde{u}_3}{\tan\beta_3 - \tan\beta_1} \\ \dfrac{\tilde{v}_1 - \tilde{v}_2}{\tan\beta_2 - \tan\beta_1} - \dfrac{\tilde{v}_1 - \tilde{v}_3}{\tan\beta_3 - \tan\beta_1} \end{pmatrix} = \begin{pmatrix} 0 \\ \lambda \end{pmatrix}, \tag{32}$$

where λ is the following scalar defined from the three rotation axes:

$$\lambda = \frac{\cot\alpha_1\cos\beta_2 - \cot\alpha_2\cos\beta_1}{\cos\beta_1\sin\beta_2 - \cos\beta_2\sin\beta_1} - \frac{\cot\alpha_1\cos\beta_3 - \cot\alpha_3\cos\beta_1}{\cos\beta_1\sin\beta_3 - \cos\beta_3\sin\beta_1}.$$

Notably, dividing (32) by λ reduces it to (31), as expected.

Since the weighted combination of specular flows is easily computed from the measured specular flows and the known EMFs, (32) is a set of two decoupled linear PDEs in the unknown h and k, and once again, we arrived at an equation that resembles in structure the case of a single view-axis EMF. Hence, as in Section 6.1, (32) is solvable using the same technique discussed in Section 5.

6.4 Shape Recovery under *Unknown* Arbitrary EMFs

The previous section demonstrates that the shape of a specular surface can be recovered in closed form given its specular flows under two or more known environment motion fields. A natural question is whether or not shape can still be recovered in the "uncalibrated case," i.e., when neither the angular velocities nor the rotation axes of the input EMFs are known. In this section, we answer this question in the affirmative. Somewhat surprisingly, this requires no additional information or assumptions beyond those used so far.

In order to address surface reconstruction under multiple *unknown* EMFs, recall first that we need to linearly combine several EMFs to yield a solvable view-axis SFSF equation (cf., (31) or (32)). When the EMF parameters (i.e., their rotation axes $\hat{\mathbf{a}}_i = (\alpha_i, \beta_i)$ and angular velocities ω_i) are known, the coefficients γ_i of this linear combination can be computed analytically. However, as we argue next, these coefficients can be recovered even when the EMF parameters are unknown.

Consider again (14), which represents the SFSF problem for EMFs about the viewing direction. Section 5.1 exploited a geometrical interpretation of this equation with regard to the behavior of the function $k(x, y)$ in order to automatically recover the angular velocity of the EMF. In the same spirit, consider now the equation in $h(x, y)$ from (14), i.e., the equation

$$\mathbf{u}\cdot\nabla h = 0.$$

This equation dictates that, for a view-axis rotation, the specular flow must be perpendicular to the gradient of $h(x, y)$ and therefore parallel to its isocontours (see also Proposition 1). Since $h(x, y)$ is defined as the square magnitude of the surface gradient (13), the integral curves of specular flow must therefore be curves of constant gradient magnitude. One useful curve for which this holds is the shape's occluding contour. While the flow is zero at the contour itself, there are always curves of isogradient magnitude arbitrarily close to it in the object domain, and when the environment rotates about the view axis, the flow direction along such a curve must be tangent to the occluding contour. Assuming that the occluding contour can be detected in an image, this tangency condition can be used to determine the coefficients γ_i.

More formally, if $\alpha(s)$ is an arc-length-parameterized curve representing the object's occluding contour, then the specular flow $\mathbf{u}$ due to an EMF about the view axis must satisfy

$$\mathbf{u}(\alpha(s)) = v(s)\alpha'(s), \tag{33}$$

where $v(s)$ is some scalar function of the arc length. Given multiple observed specular flows $\mathbf{u}_i$ due to EMFs $\boldsymbol{\omega}_i$, the coefficients γ_i that bring the linear combination $\sum_i \gamma_i\boldsymbol{\omega}_i$ to a view-axis rotation are those for which a combined flow $\bar{\mathbf{u}} = \sum_i \gamma_i\mathbf{u}_i$ satisfies the tangency condition in (33). Computationally, these coefficients are given by

$$\underset{\gamma_i}{\operatorname{argmin}} \int_\alpha \mathbf{E}[\theta(\alpha'(s)), \theta(\bar{\mathbf{u}}(\alpha(s)))]ds, \tag{34}$$

where $\theta(\cdot)$ is a function returning the orientation of its vector argument and $\mathbf{E}$ is any suitable penalty metric between two orientations. Although in our experiments we used least squares, other penalty functions can be used as well.

Finally, we note that the solution to (34) may not be unique if the number of linearly independent input flows exceeds 3. Note also that the tangency constraint only determines the coefficients up to scale, but since this scale factor is simply the angular velocity, it can be recovered using the procedure described in Section 5.1. In summary, the following procedure provides dense surface shape from multiple specular flows induced by arbitrary and unknown rotations of the environment:

> Reconstruction Algorithm II: Multiple unknown arbitrary EMFs
>
> 1) Use (34) to find γ_i such that $\bar{\mathbf{u}} = \sum_i \gamma_i\mathbf{u}_i$ satisfies the tangency condition along the occluding boundary.
> 2) Apply Reconstruction Algorithm I (page 8) on $\bar{\mathbf{u}}$ to reconstruct $f(x, y)$.

6.5 Experimental Results

The procedure for specular shape reconstruction from multiple specular flows induced by arbitrary and unknown EMFs was evaluated using data similar to that of Section 5.3. Synthetic surfaces were placed in captured illumination environments and used to generate multiple distinct specular flow fields corresponding to different rotation axes. Given multiple specular flows for different (two or three) EMFs, the "normalized" linear combination that corresponds to a view-axis EMF was computed by minimizing the least square error of the tangency condition. Given the recovered coefficients of the combination, the corresponding view-axis SFSF equation (31) was derived to yield the PDEs for $h(x, y)$ and $k(x, y)$ as described in Section 6. The resulting final equation was then solved for $h(x, y)$ and $k(x, y)$ as discussed in Section 5, using manually provided initial conditions. Two examples of this entire process are shown in Figs. 12 and 13. Note that the numerical considerations described in Section 5.2 apply here as well.

7 Summary and 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

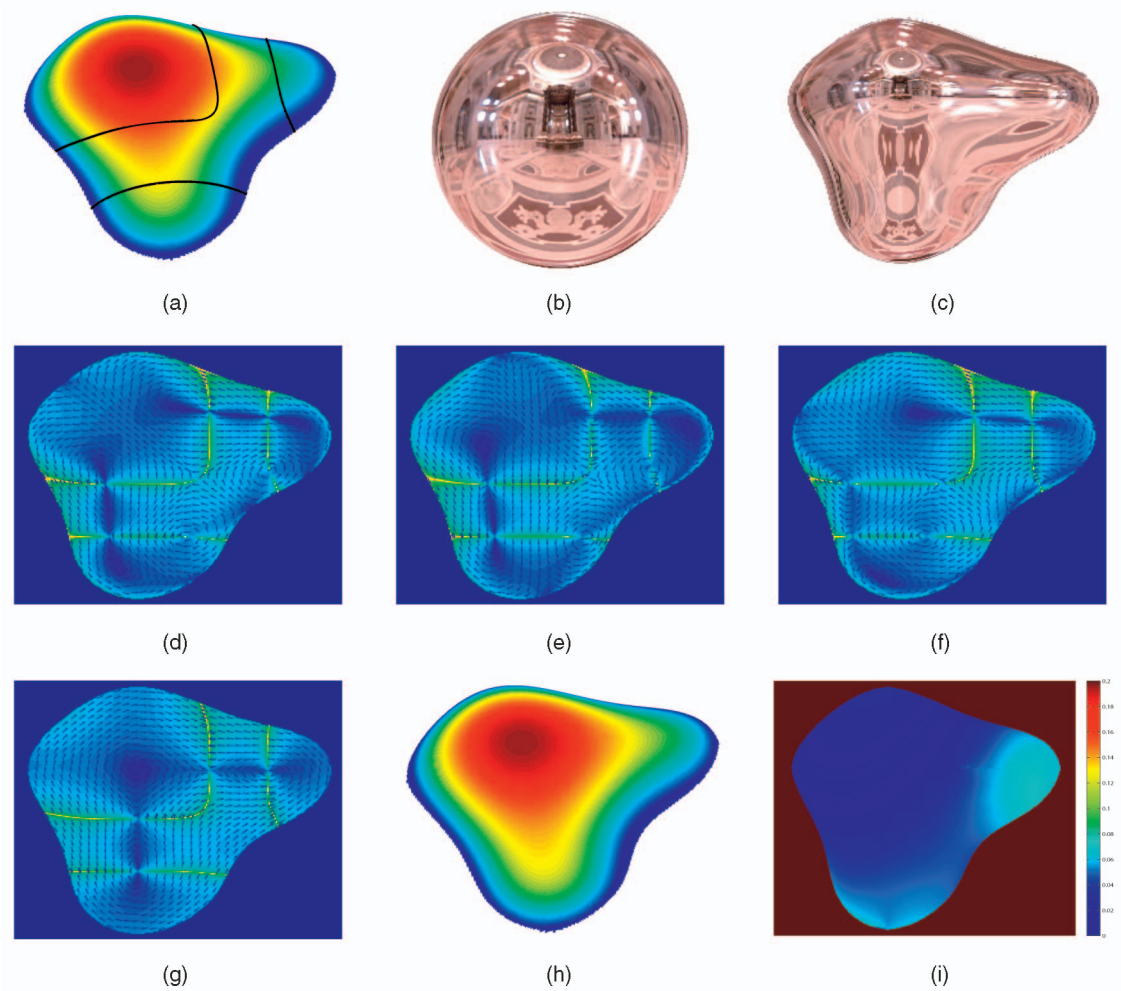


Fig. 12. An example of specular shape reconstruction of a general surface observed under three arbitrary and unknown EMFs. (a) The original surface $f(x,y) = \sqrt{4 - x^2 - y^2 - \cos(2x - 2) - \sin 2y}$. As in the previous examples, this surface includes multiple maxima and several elliptic and hyperbolic regions separated by parabolic lines, shown here in black. (c) One snapshot of the specular surface obtained under the first EMF. (d)-(f) Specular flow obtained from the three different EMFs. ($\hat{\mathbf{a}}_1 = (180°, -45°)$, $\hat{\mathbf{a}}_2 = (-30°, -18°)$, $\hat{\mathbf{a}}_3 = (45°, -60°)$). Although the EMFs were known, this information was *not* used in the reconstruction. (g) The weighted specular flow obtained by a linear combination of the three different EMFs. The coefficients of the combinations were recovered based on (34) and least squares as the optimality criterion. (h) The reconstructed surface. (i) Absolute reconstruction error, reported as percent of maximum input surface height.

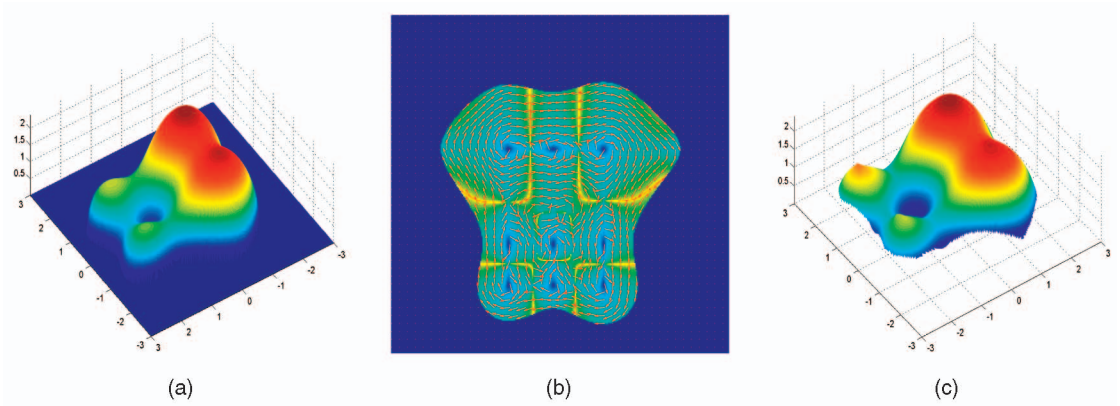


Fig. 13. A second example using the surface $f(x,y) = \sqrt{4 - x^2 - y^2 - \cos(3x - 6) - 2\sin 2y}$. Note that this surface is significantly more complex than the previous one in the sense of having more parabolic lines as well as multiple maxima, minima, and saddle points. (a) The original surface. (b) The "corrected" weighted specular flow used in the integration. (c) The recovered surface.

relative object/environment motion, we analyze the relationship between observed specular flow and specular shape (assuming no interreflections), and 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.

Armed with a computational procedure for solving the SFSF problem for environment rotations about the view direction, we then show that specular shape can be reconstructed from specular flow under arbitrary rotations once the same object is observed under multiple distinct motions. Furthermore, we show that with multiple specular flows, one can recover the shape without knowing the environment rotations under which they were observed.

Several extensions and research directions arise naturally from our results. In particular, the problem of measuring the specular flow itself was not discussed in this paper. This problem shares many features with optical flow estimation, but it is also fundamentally different due to the unique structures and singularities that specular flows exhibit, both in terms of orientation and their magnitude, especially along parabolic lines. While many optical flow techniques do exist, it is likely that reliable specular flow estimation, if possible, would require novel schemes and regularization approaches. As mentioned before, it is possible that the differential geometric analysis we provide in this paper may be constructive for this aspect of the problem as well (perhaps using the framework described in [38]). 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 and interreflections). 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 the required initial conditions for the solution of the SFSF equation and replace the present need to provide these data by other means (for example, by spatial sampling or using calibrated illumination sources). It can also assist in determining the occluding boundaries of specular objects (e.g., for the application of the tangency condition from Section 6.4), a task that could become nontrivial, for example, in complex specular of hybrid scenes with multiple objects that occlude each other.

Another issue that deserves attention regarding the practical reconstruction of shape from specular flow is the numerical sensitivity of the proposed methods. For example, even if specular flows could be measured from image data, one should expect imperfect measurements and flow fields corrupted with noise. It is therefore natural to ask whether the proposed shape reconstruction methods exhibit graceful degradation in the presence of noise, or if they collapse completely. Preliminary analysis can be conducted by adding noise to synthetic specular flows (Table 1 and Fig. 14), and such an analysis suggests that the proposed methods may indeed be robust. But an in-depth investigation of numerical sensitivity as well as experiments on real data is an important future research direction. As part of this research thread, it would be beneficial to design and build an acquisition system that generalizes the system in Fig. 9 by providing the ability to induce arbitrary rotation axes. Until these issues are studied

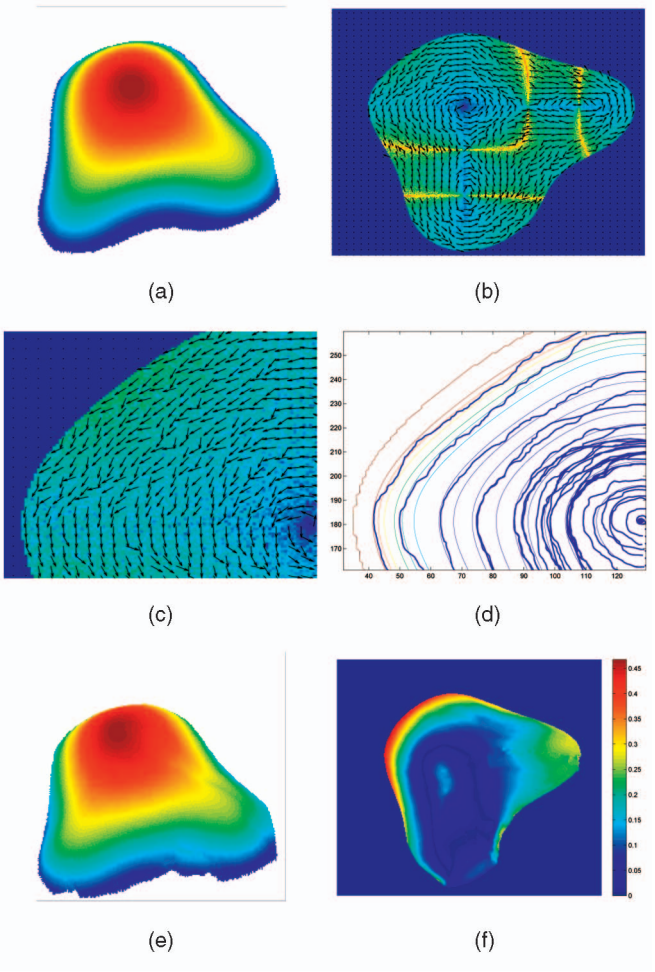


Fig. 14. Qualitative sensitivity to specular flow noise. (a) Another reconstruction of the same object from Fig. 12 from a single view-axis specular flow $\mathbf{u} = (u, v)$. (b) A noisy version $\hat{\mathbf{u}}$ of $\mathbf{u}$, with i.i.d. Gaussian noise corrupting both the u and v coordinate functions. (c) A close-up of one section of $\hat{\mathbf{u}}$. (d) Selected characteristic curves recovered from $\hat{\mathbf{u}}$ (shown in thick blue). Note that while the noise in the specular flow translates to inaccurate integral curves, the latter maintains the qualitative structure of the true curves (shown in thin). (e) The full reconstructed surface from $\hat{\mathbf{u}}$. Compare to panel a. (f) The reconstruction error.

TABLE 1
Quantitative Sensitivity to Specular Flow Noise: Reconstruction Errors for Multiple Runs of the Noise Experiment in Fig. 14

Noise level, σ		0	$\frac{\|\mathbf{u}(x,y)\|}{1000}$	$\frac{\|\mathbf{u}(x,y)\|}{100}$	$\frac{\|\mathbf{u}(x,y)\|}{10}$
ω given	mean AAE (degrees)	0.45	0.51	0.54	1.52
	rms height error (%)	4.20	4.20	4.10	3.69
ω estimated	mean AAE (degrees)	0.61	0.71	2.08	10.73
	rms height error (%)	4.34	4.14	5.02	11.74
mean absolute difference from true speed $\omega = 1^\circ/s$		0.0017	0.0018	0.0192	0.1526

Algorithm 5.1 was applied to flows with added spatially varying, zero-mean Gaussian noise having standard deviation σ proportional to local flow magnitude. Reported are average errors over five runs with angular speed ω both given and estimated from the data. Rows 2 and 4: Mean absolute angular error in the recovered normal field. Rows 3 and 5: RMS surface height error as a percentage of maximum surface height. Row 6: Mean absolute difference between estimated angular speed and true value, $\omega = 1^\circ/s$.

closely, even the appropriate noise model that applies to specular flows remains an open question.

While our analysis shows that shape can be recovered from a single specular flow observed under view-axis rotation and from multiple specular flows observed under arbitrary rotations, it is still an open question whether the problem can be solved from a single specular flow under an arbitrary motion. Similarly, reframing the SFSF problem for perspective projection is considered a key issue in our forthcoming research.

In addition to addressing these computational questions, our analysis may aid the 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 [39], whether flow is used to recover shape is less certain [40]. 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 does utilize specular flow. In particular, the distinguished behavior at parabolic points would seem to play a role in this process.

Finally, this paper considers the case in which the camera and object move as a fixed pair relative to a distant environment. Another important direction to pursue in the future is to consider the more natural case in which the object (resp. viewer) moves relative to a fixed camera (resp. object) and environment. This is effectively what a human does when holding a specular surface in their hand and investigating it under small rotations. The analysis presented in this paper is likely relevant to this case as well, and it suggests that when shape is the desired output, observations of an object under multiple distinct rotations may be beneficial.

Acknowledgments

This research is supported by the Israel Science Foundation under grant no. 1245/08 and the US National Science Foundation (NSF) under grant no. IIS-0712956. Ohad Ben-Shahar and Yair Adato also thank the generous support of the Frankel Fund and the Paul Ivanier Robotics Center at Ben-Gurion University. Additional funding for Todd Zickler and Yriy Vasilyev was provided by the NSF under CAREER Award IIS-0546408.

References

[1] S. Waldon and C. Dyer, "Dynamic Shading, Motion Parallax and Qualitative Shape," *Proc. IEEE Workshop Qualitative Vision,* pp. 61-70, 1993.

[2] S. Roth and M. Black, "Specular Flow and the Recovery of Surface Structure," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* pp. 1869-1876, 2006.

[3] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. Black, and R. Szeliski, "A Database and Evaluation Methodology for Optical Flow," *Proc. IEEE Int'l Conf. Computer Vision,* 2007.

[4] M. Osadchy, D. Jacobs, and R. Ramamoorthi, "Using Specularities for Recognition," *Proc. IEEE Int'l Conf. Computer Vision,* pp. 1512-1519, 2003.

[5] O. Drbohlav and R. Sára, "Specularities Reduce Ambiguity of Uncalibrated Photometric Stereo," *Proc. European Conf. Computer Vision,* pp. 46-60, 2000.

[6] O. Drbohlav and M. Chantler, "Can Two Specular Pixels Calibrate Photometric Stereo?" *Proc. IEEE Int'l Conf. Computer Vision,* pp. 1850-1857, 2005.

[7] J. Solem, H. Aanaes, and A. Heyden, "A Variational Analysis of Shape from Specularities Using Sparse Data," *Proc. Conf. 3D Data Processing, Visualization, and Transmission,* pp. 26-33, 2004.

[8] K. Ikeuchi, "Determining Surface Orientations of Specular Surfaces by Using the Photometric Stereo Method," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 3, no. 6, pp. 661-669, Nov. 1981.

[9] S. Nayar, K. Ikeuchi, and T. Kanade, "Determining Shape and Reflectance of Hybrid Surfaces by Photometric Sampling," *IEEE Trans. Robotics and Automation,* vol. 6, no. 4, pp. 418-431, Aug. 1990.

[10] S. Savarese, M. Chen, and P. Perona, "Local Shape from Mirror Reflections," *Int'l J. Computer Vision,* vol. 64, no. 1, pp. 31-67, 2005.

[11] T. Bonfort, P. Sturm, and P. Gargallo, "General Specular Surface Triangulation," *Proc. Asian Conf. Computer Vision,* pp. 872-881, 2006.

[12] M. Knauer, J. Kaminski, and G. Häusler, "Phase Measuring Deflectometry: A New Approach to Measure Specular Free-Form Surfaces," *Proc. Conf. SPIE,* pp. 366-376, 2004.

[13] M. Tarini, H. Lensch, M. Goesele, and H. Seidel, "3D Acquisition of Mirroring Objects Using Striped Patterns," *Graphical Models,* vol. 67, pp. 233-259, 2005.

[14] S. Rozenfeld, I. Shimshoni, and M. Lindenbaum, "Dense Mirroring Surface Recovery from 1D Homographies and Sparse Correspondence," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2007.

[15] J. Lellmann, J. Blazer, A. Reider, and J. Beyerer, "Shape from Specular Reflection and Optical Flow," *Int'l J. Computer Vision,* vol. 80, no. 2, pp. 226-241, Nov. 2008.

[16] K. Kutulakos and E. Steger, "A Theory of Refractive and Specular 3D Shape by Light-Path Triangulation," *Proc. IEEE Int'l Conf. Computer Vision,* pp. 1448-1455, 2005.

[17] A. Blake, "Specular Stereo," *Proc. Int'l Joint Conf. Artificial Intelligence,* pp. 973-976, 1985.

[18] A. Blake and G. Brelstaff, "Geometry from Specularities," *Proc. IEEE Int'l Conf. Computer Vision,* 1988.

[19] A. Zisserman, P. Giblin, and A. Blake, "The Information Available to a Moving Observer from Specularities," *Image and Vision Computing,* vol. 7, pp. 38-42, 1989.

[20] A. Blake and H. Bülthoff, "Does the Brain Know the Physics of Specular Reflection?" *Nature,* vol. 343, pp. 165-168, 1990.

[21] A. Blake and H. Bülthoff, "Shape from Specularities: Computation and Psychophysics," *Philosophical Trans. Biological Sciences,* vol. 331, no. 1260, pp. 237-252, 1991.

[22] J. Koenderink and A. van Doorn, "Photometric Invariants Related to Solid Shape," *Optical Acta,* vol. 27, no. 7, pp. 981-996, 1980.

[23] M. Longuet-Higgins, "Reflection and Refraction at a Random Moving Surface. I. Pattern and Paths of Specular Points," *J. Optical Soc. Am.,* vol. 50, no. 9, pp. 838-844, 1960.

[24] M. Oren and S. Nayar, "A Theory of Specular Surface Geometry," *Int'l J. Computer Vision,* vol. 24, no. 2, pp. 105-124, 1997.

[25] J. Zheng, Y. Fukagawa, and N. Abe, "Shape and Model from Specular Motion," *Proc. IEEE Int'l Conf. Computer Vision,* pp. 72-79, 1995.

[26] J. Zheng and A. Murata, "Acquiring a Complete 3D Model from Specular Motion under the Illumination of Circular-Shaped Light Sources," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 22, no. 8, pp. 913-920, Aug. 2000.

[27] S. Savarese, F. Li, and P. Perona, "What Do Reflections Tell Us about the Shape of a Mirror?" *Proc. Symp. Applied Perception in Graphics and Visualization,* 2004.

[28] R. Fleming, A. Torralba, and E. Adelson, "Specular Reflections and the Perception of Shape," *J. Vision,* vol. 4, pp. 798-820, 2004.

[29] U. Weidenbacher, P. Bayerl, H. Neumann, and R. Fleming, "Sketching Shiny Surfaces: 3D Shape Extraction and Depiction of Specular Surfaces," *ACM Trans. Applied Perception,* vol. 3, no. 3, pp. 262-285, 2006.

[30] A. DelPPozo and S. Savarese, "Detecting Specular Surfaces on Natural Images," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2007.

[31] D. Nister, H. Stewenius, and E. Grossmann, "Non-Parametric Self-Calibration," *Proc. IEEE Int'l Conf. Computer Vision,* vol. 1, 2005.

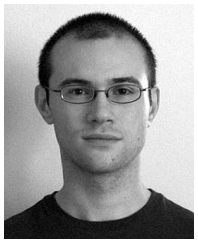
[32] S. Ramalingam, P. Sturm, and S. Lodha, "Towards Generic Self-Calibration of Central Cameras," *Proc. Workshop Omnidirectional Vision,* pp. 20-27, 2005.

[33] "Light Probe Gallery," http://www.debevec.org/Probes/, 2008.

[34] K. Ikeuchi and B. Horn, "Numerical Shape from Shading and Occluding Boundaries," *Artificial Intelligence,* vol. 17, pp. 141-184, 1981.
[35] O. Ben-Shahar and S. Zucker, "The Perceptual Organization of Texture Flows: A Contextual Inference Approach," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 25, no. 4, pp. 401-417, Apr. 2003.
[36] M. do Carmo, *Differential Geometry of Curves and Surfaces.* Prentice-Hall, Inc., 1976.
[37] B. Horn and B. Schunck, "Determining Optical Flow," *Artificial Intelligence,* vol. 17, pp. 185-203, 1981.
[38] S. Agarwal, S. Mallick, D. Kriegman, and S. Belongie, "On Refractive Optical Flow," *Proc. European Conf. Computer Vision,* pp. 483-494, 2004.
[39] B. Hartung and D. Kersten, "Distinguishing Shiny from Matte," *J. Vision,* vol. 2, no. 7, p. 551, 2002.
[40] S. Roth, F. Domini, and M. Black, "Specular Flow and the Perception of Surface Reflectance," *J. Vision,* vol. 3, no. 9 p. 413, 2003.



Yair Adato received the BSc degree in mathematics and computer science and the MSc degree in computer science from Ben Gurion University of the Negev, Israel, in 2005 and 2007, respectively. He is currently working toward the PhD degree in the Computer Science Department at the same university, under the supervision of Dr. Ohad Ben-Shahar. He is the winner of the 2008 Friedman Award for excellence in research. His research interests include computational and biological vision. He is a student member of the IEEE and the IEEE Computer Society.



Yuriy Vasilyev received the BSE degree in computer science from the University of Pennsylvania in 2005, where he was also the winner of the 2005 A. Atwater Kent Prize in electrical engineering. He received the MSc degree in computer science from Harvard University in 2007 and is currently working toward the PhD degree in the School of Engineering and Applied Science at Harvard University, where professor Todd Zickler is his advisor. His research interests include computer vision and computer graphics. He is a student member of the IEEE and the IEEE Computer Society.



Todd Zickler received the BEng degree (with honors) in electrical engineering from McGill University in 1996 and the PhD degree in electrical engineering from Yale University in 2004, under the direction of Peter Belhumeur. He joined the School of Engineering and Applied Sciences at Harvard University as an assistant professor in 2004 and was appointed a John L. Loeb associate professor of the natural sciences in 2008. He is the director of the Harvard Computer Vision Laboratory. His research is focused on modeling the interaction between light and materials and developing algorithms to extract scene information from visual data. His work is motivated by applications in face, object, and scene recognition; image-based rendering; content-based image retrieval; image and video compression; robotics; and human-computer interfaces. He is a recipient of the US National Science Foundation (NSF) Career Award and a Research Fellowship from the Alfred P. Sloan Foundation. His research is funded by the NSF, the US Army Research Office, the US Office of Naval Research, and the Sloan Foundation. He is a member of the IEEE and the IEEE Computer Society.



Ohad Ben-Shahar received the BSc and MSc degrees in computer science from the Technion, Israel Institute of Technology, and the PhD degree in computer science from Yale University, Connecticut, in 2003. He is a senior lecturer in the Department of Computer Science at Ben Gurion University (BGU) of the Negev and the director of the BGU Computational Vision Lab. Before pursuing the PhD degree, he was employed for several years as an R&D scientist in the Israeli high-tech industry. Before joining BGU in 2004, he was a postdoctoral researcher at Yale University. His main area of research is in computational vision and image analysis, where he is focusing primarily on issues related to the differential geometrical foundations of perceptual organization and early vision. His work is explicitly multidisciplinary and his computational research is often endowed by investigations into human perception, visual psychophysics, and computational neuroscience of biological vision. He is the recipient of the 2007 Psychobiology Young Investigator Award and his research is funded by the Israel Science Foundation (ISF), the US Air Force Office of Scientific Research (AFOSR), the Deutsche Forschungsgemeinschaft (DFG) in Germany, the US National Science Foundation (NSF), and the US National Institute for Psychobiology. He is a member of the IEEE and the IEEE Computer Society.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.