

Dense Specular Shape from Multiple Specular Flows

Yuriy Vasilyev¹

Yair Adato²

Todd Zickler¹

Ohad Ben-Shahar²

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

¹Harvard School of Engineering and Applied Sciences, Cambridge, USA

²Computer Science Department, Ben-Gurion University, Beer Sheva, Israel

Abstract

The inference of specular (mirror-like) shape is a particularly difficult problem because an image of a specular object is nothing but a distortion of the surrounding environment. Consequently, when the environment is unknown, such an image would seem to convey little information about the shape itself. It has recently been suggested (Adato et al., ICCV 2007) that observations of relative motion between a specular object and its environment can dramatically simplify the inference problem and allow one to recover shape without explicit knowledge of the environment content. However, this approach requires solving a non-linear PDE (the ‘shape from specular flow equation’) and analytic solutions are only known to exist for very constrained motions.

In this paper, we consider the recovery of shape from specular flow under general motions. We show that while the ‘shape from specular flow’ PDE for a single motion is non-linear, we can combine observations of multiple specular flows from distinct relative motions to yield a linear set of equations. We derive necessary conditions for this procedure, discuss several numerical issues with their solution, and validate our results quantitatively using image data.

1. Introduction

Shape inference is one of the hallmarks of computer vision, and the automatic inference of shape has been a key research problem in image understanding since the inception of the field. One important visual cue that is exploited by humans is *specular reflection*. Specular (mirror-like) surfaces are abundant in both natural and man-made environments, and specular reflections often provide powerful information about surface shape (Fig. 1).

Despite the apparent utility of specular reflections, computational analysis of specular objects has received limited attention. This is likely due to the seemingly arbitrary relationship between an object’s shape and the image it induces.

Indeed, since the induced image is a distorted reflection of the environment, it is possible to create essentially *any* image from a given specular surface by placing it in a suitably manipulated setting. This is leveraged, for example, in what is known as *anamorphic art*—visual art that must be observed through special mirrors to make sense (Fig. 1d.)

Due to the difficulty of recovering specular shape in unknown environments, most existing work has: 1) focussed on recovering only sparse or qualitative shape; 2) considered limited classes of surfaces; or 3) required calibrated conditions in which environment structure is known. In contrast, this paper presents an approach for quantitative reconstruction that specifically targets general surfaces in unknown real-world environments. It is motivated by recent work [1, 17] that seeks to recover 3D shape from *specular flow*—the optical flow that is induced by relative motion between a specular object, an observer and their environment.

As shown by Adato et al. [1], observing specular flow can simplify the reconstruction problem because it decouples the shape and relative motion of the object from the content of its environment. In fact, one can derive a non-linear PDE—termed the *shape from specular flow (SFSF) equation*—that explicitly relates an observed specular flow field to the object shape and motion. Then, inferring 3D shape amounts to solving this equation [1].

The SFSF equation is highly nonlinear, and analytic solutions are known to exist only when the motion is known and only if it follows a very restricted form [1]. In this paper, we extend this analysis by deriving analytic solutions for motions that are *general* and *unknown*. We show that the non-linear SFSF equation can be reduced to a set of linear equations by combining the specular flow fields induced by several distinct motions. Conveniently, this can be achieved without having knowledge of the motions *a priori*. The result of our work is therefore an ‘auto-calibrating’ approach for recovering dense specular shape from general, unknown motions in natural, unknown environments.




Figure 1. Specular surfaces under dense, (approximately) far-field illumination are frequently encountered in everyday life, in technical applications, and in the art. The human visual system is often able to infer specular shape under these general conditions, even when the illumination environment is unknown. Can this task be achieved computationally also?

2. Related work

Inferring specular shape is often simplified by assuming simple illumination environments, such as a single point light source. In these conditions, each isolated specular-ity provides strong constraints on the camera position, the source position, and the surface normal at the observed point. When the camera and source positions are known, for example, the surface normal and the surface depth at the 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., [16]) and surface reconstruction (e.g. [8, 7, 20]). 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 [2, 3]. 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 [24, 4].

Not unrelated to “specular stereo” are methods that seek additional surface information from motion under simple lighting. Relative motion between a specular surface, the viewer, and the environment induces motion of the observed specularities, and these specularities are known to be created and annihilated in pairs at (or, in the near-field case [5], close to) parabolic surface points [11, 13]. In theory, this type of motion can be used facilitate the recovery of a surface profile [24], but existing methods are limited to convex (or concave) surfaces that are void of parabolics [15].

In addition to these methods that exploit simple lighting, a separate class of techniques relies on complex lighting that is known and controlled. Using this approach, one can often obtain higher order surface information such as curvature [10, 14, 18, 12, 6]. Controlled environments can also be combined with object motion to yield complete 3D reconstructions [22, 23].

Real-world illumination environments that are complex and unknown (Figs. 1a,c) pose additional challenges. Un-

der these conditions, observed specular reflections are dense and qualitatively different from the sparse or controlled cases discussed above. While the capability of the human visual system to handle such complex lighting is still controversial (c.f., see [9] vs. [19]), the computational problem is clearly ill-posed.

A small number of computational studies have shown that motion can be exploited to simplify the inference of specular shapes in unknown complex lighting. Under these conditions, motion induces a *specular flow* on the image plane. Qualitatively, specular flow is known to exhibit singularities along (or in the near-field case, close to) parabolic curves [21], and specular flow can also be used to identify a unique convex surface from a small parametric family [17]. Quantitatively, Adato et al. [1] show that general surfaces can be recovered from specular flow, but their analytic solution is only applicable when the motion is known and of a very particular form. In this paper we build on this by considering closed-form solutions when the motion is both unknown and relatively unconstrained.

3. Background: Shape from specular flow

This section introduces notation and summarizes the relevant findings of Adato et al. [1], who analyze the relationship between shape, relative object/environment motion, and the specular flow they induce.

As shown in Fig. 2, it is assumed that both the viewer and the illumination environment are far from the observed surface relative to the surface relief. Hence, the camera is approximately orthographic and illumination depends only on angular direction (i.e., the plenoptic function is bivariate at any instant of time). It is also assumed that the viewer is fixed relative to the object and that a specular flow is induced by a motion of the environment. The visible region of the surface is represented by the graph of a function $S(x, y) = (x, y, f(x, y))$ defined over the image plane XY , and the viewing direction is $\hat{\mathbf{v}} = (0, 0, 1)$. Let $\hat{\mathbf{n}}(x, y)$ be the surface normal at surface point $(x, y, f(x, y))$, and




Figure 2. The SFSF problem involves a surface $f(x, y)$ illuminated by a far-field illumination environment and viewed orthographically to produce an image $I(x, y)$. The illumination sphere, parameterized by spherical coordinates (θ, ϕ) , moves relative to the camera and object, and this induces a specular flow on the image plane. (Figure reproduced from [1].)

$\hat{\mathbf{r}}(x, y)$ be the mirror-reflection direction at the same point (see Fig. 2). Also, the far-field illumination environment is represented by an *unknown* non-negative radiance function defined on the sphere of directions, which is parameterized by spherical angles θ (elevation) and ϕ (azimuth).

If (θ, ϕ) are the spherical coordinates of the normal vector at a point on the surface, then the spherical coordinates of the reflection vector $\hat{\mathbf{r}}$ at that point are simply $(\alpha, \beta) = (2\theta, \phi)$. Since $\hat{\mathbf{r}}$ can also be expressed in terms of surface properties, it is possible to derive explicit relationships between the reflection direction and the first-order derivatives of the surface:

$$\tan \alpha = \frac{2\|\nabla f\|}{1 - \|\nabla f\|^2}, \quad \tan \beta = \frac{f_y}{f_x}. \quad (1)$$

As discussed above, the surface is observed under motion of the environment, and this *environment motion field* (EMF) can be represented as a vector field on the unit sphere:

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

While an EMF could be arbitrary, a reasonable class of motions to consider is the space of “rigid” rotations, where the entire environment rotates at an angular velocity $\omega \neq 0$ around an axis $\hat{\mathbf{a}}$ defined by spherical angles (α_o, β_o) , i.e.,

$$\hat{\mathbf{a}} = (\sin \alpha_o \cos \beta_o, \sin \alpha_o \sin \beta_o, \cos \alpha_o).$$

Following [1], such an EMF can be written as follows:

$$\begin{aligned} \omega_\alpha(\alpha, \beta) &= \omega \sin \alpha_o \sin(\beta_o - \beta) \\ \omega_\beta(\alpha, \beta) &= \omega \cos \alpha_o - \omega \sin \alpha_o \cos(\beta - \beta_o) \cot \alpha. \end{aligned} \quad (2)$$

Observing the specular object under an EMF induces a motion field – termed the specular flow – in the image plane. When this specular flow is represented as a vector field

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

it can be related to the EMF through the chain rule

$$\omega = \frac{d(\alpha, \beta)}{dt} = \frac{\partial(\alpha, \beta)}{\partial(x, y)} \frac{d(x, y)}{dt} = \mathbf{J} \mathbf{u}, \quad (3)$$

where the entries of the Jacobian \mathbf{J} can be expressed in terms of the surface by differentiating Eq. 1. This yields:

$$\mathbf{J} \triangleq \frac{\partial(\alpha, \beta)}{\partial(x, y)} = \begin{pmatrix} \frac{f_x f_{xx} + f_y f_{xy}}{\|\nabla f\| \cdot (1 + \|\nabla f\|^2)} & \frac{f_x f_{xy} + f_y f_{yy}}{\|\nabla f\| \cdot (1 + \|\nabla f\|^2)} \\ \frac{f_x f_{xy} - f_y f_{xx}}{2\|\nabla f\|^2} & \frac{f_x f_{yy} - f_y f_{xy}}{2\|\nabla f\|^2} \end{pmatrix}. \quad (4)$$

As discussed by Adato et al. [1], Eq. 3 provides a coupled pair of non-linear, second-order PDEs in terms of the unknown surface. Thus, given an observed specular flow \mathbf{u} and the known parameters of an EMF ($\hat{\mathbf{a}}$ and ω) one can theoretically solve this equation to recover the surface. For this reason, Eq. 3 has been termed the *shape from specular flow* (SFSF) equation [1]. Importantly, since the SFSF equation does not depend on the content of the environment, it provides the ability to recover (dense) surface shape without any knowledge of the scene illumination.

A limitation of the SFSF approach from [1], however, is that it requires the solution of Eq. 3, which is highly non-trivial. For this reason, the authors only consider a very special class of environment rotations: those in which the rotation axis $\hat{\mathbf{a}}$ is aligned with the viewing direction. In this case, the EMF rotation axis becomes $\hat{\mathbf{a}} = (\alpha_o, \beta_o) = (0, 0)$ and the general EMF (Eq. 2) reduces to a constant spherical vector field $\omega = (0, \omega)$. If one now defines two auxiliary functions that represent the (unknown) surface gradient magnitude and orientation,

$$\begin{aligned} h(x, y) &\triangleq f_x^2 + f_y^2 \\ k(x, y) &\triangleq \tan^{-1}(f_y/f_x), \end{aligned} \quad (5)$$

the SFSF reduces to a particularly simple form:

$$\tilde{\mathbf{J}} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} 0 \\ 2\omega \end{pmatrix}. \quad (6)$$

where, for notational convenience, we have defined

$$\tilde{\mathbf{J}} \triangleq \begin{pmatrix} h_x & h_y \\ k_x & k_y \end{pmatrix} \quad (7)$$

and the derivatives of h and k are derived directly from Eq. 5

$$\begin{aligned} h_x &= 2(f_x f_{xx} + f_y f_{xy}) & h_y &= 2(f_x f_{xy} + f_y f_{yy}) \\ k_x &= \frac{f_x f_{xy} - f_y f_{xx}}{\|\nabla f\|^2} & k_y &= \frac{f_x f_{yy} - f_y f_{xy}}{\|\nabla f\|^2} \end{aligned} \quad (8)$$

Eq. 6 is a pair of *decoupled, linear* PDEs in h and k . Hence, given the specular flow $\mathbf{u} = (u, v)$ one can solve each of these equations in closed form using the method of characteristics. The characteristics for both equations are the same, and they correspond to the integral curves of the specular flow. Once $h(x, y)$ and $k(x, y)$ are known, the surface gradient can be recovered by inverting Eq. 5, and this can be integrated to recover the surface $f(x, y)$.

4. General Environment Motion

The special case described above – the case of rotation about the view direction – is important because it enables closed form solutions. That said, it is extremely ‘accidental’. One of the main contributions of our paper is to relax this restriction on $\hat{\mathbf{a}}$ while maintaining the ability to find closed-form solutions. Our key result is that this can be done for *arbitrary* $\hat{\mathbf{a}}$ in the case when the object is observed under two or more distinct environment motions (EMFs).

We begin by expressing the general EMF in Eq. 2 explicitly in terms of the surface and its derivatives. For this we first note that

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

and substituting these expressions into Eq. 2 we obtain

$$\begin{aligned}\omega_\alpha(\alpha, \beta) &= \frac{\omega \sin \alpha_o}{\|\nabla f\|} (f_y \cos \beta_o - f_x \sin \beta_o) \\ \omega_\beta(\alpha, \beta) &= \omega \cos \alpha_o \\ &+ \omega \sin \alpha_o \frac{1 - \|\nabla f\|^2}{2 \cdot \|\nabla f\|^2} (f_x \cos \beta_o + f_y \sin \beta_o).\end{aligned}\quad (9)$$

Using this in conjunction with the auxiliary functions from Eqs. 5 and 8, we can therefore re-write the SFSF equation as follows

$$\tilde{\mathbf{J}}\mathbf{u} = \omega \sin \alpha_o \left(\frac{2(1+h)(f_y \cos \beta_o - f_x \sin \beta_o)}{2 \cot \alpha_o + \frac{1-h}{h}(f_x \cos \beta_o + f_y \sin \beta_o)} \right). \quad (10)$$

At first sight, this expression (Eq. 10) seem to offer little advantage over the original SFSF equation (Eq. 3). Indeed, while it has been reduced from second order to first order, the equations are still coupled and non-linear. 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.

4.1. Two equal-azimuth motions

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

$$\begin{aligned}\hat{\mathbf{a}}_1 &= (\alpha_1, \beta_0) \\ \hat{\mathbf{a}}_2 &= (\alpha_2, \beta_0)\end{aligned}\quad (11)$$

and angular velocities $\omega_1 \neq 0$ and $\omega_2 \neq 0$, respectively. Note that $\hat{\mathbf{a}}_1, \hat{\mathbf{a}}_2$ and the viewing direction are all in the same plane, i.e., that the two rotation axes have identical azimuth angle β_0 , as is also depicted graphically in Fig. 3. Note that if either $\alpha_1 = 0$ or $\alpha_2 = 0$ then 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 $i \in \{1, 2\}$

$$\tilde{\mathbf{u}}_i = \begin{pmatrix} \tilde{u}_i \\ \tilde{v}_i \end{pmatrix} \triangleq \frac{\mathbf{u}_i}{\omega_i \sin \alpha_i} = \left(\frac{u_i}{\omega_i \sin \alpha_i}, \frac{v_i}{\omega_i \sin \alpha_i} \right) \quad (12)$$

Since $\omega_i \neq 0$ and $\alpha_i \neq 0$, both normalized flows are well defined. Rewriting Eq. 10 in terms of $\tilde{\mathbf{u}}_i$ we now obtain the following two SFSF equations for $i \in \{1, 2\}$

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

This equation is significant because most of it does not depend on the arbitrary azimuthal angles α_1 and α_2 . Hence, by subtracting the equations due to the first and second EMFs we get

$$\tilde{\mathbf{J}}(\tilde{\mathbf{u}}_1 - \tilde{\mathbf{u}}_2) = \begin{pmatrix} 0 \\ 2 \cot \alpha_1 - 2 \cot \alpha_2 \end{pmatrix}. \quad (13)$$

Since $\tilde{\mathbf{u}}_1 - \tilde{\mathbf{u}}_2$ is easily computed from the measured pair of specular flows, Eq. 13 is a set of two decoupled linear PDEs in the unknown h and k and the known EMFs. Since this equation resembles the case of a single EMF around the viewing direction, it is solvable using the same technique discussed in Sec. 3.

Eq. 13 implies that a linear combination of the two measured specular flow gives a SFSF equation of the type obtained for a single EMF around the viewing direction. This result should not come as a surprise. In fact, Eq. 2 already suggests that if ω_1 and ω_2 are two equal-azimuth EMFs (cf. Eq. 11) then there exist two constants γ_1 and γ_2 such that

$$\gamma_1 \omega_1 + \gamma_2 \omega_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad (14)$$

i.e., that this linear combination is an EMF around the viewing direction (Fig. 3). Indeed, applying this requirement first to $\omega_\alpha(\alpha, \beta)$ from Eq. 2 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 Eq. 12, i.e.,

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

Satisfying also the requirement on $\omega_\beta(\alpha, \beta)$ is then achieved by further dividing both γ_1 and γ_2 by the common factor $(2 \cot \alpha_1 - 2 \cot \alpha_2)$ which appears in Eq. 13.

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

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \sum_{i=1}^2 \gamma_i \omega_i = \mathbf{J} \sum_{i=1}^2 \gamma_i \mathbf{u}_i = \tilde{\mathbf{J}} \sum_{i=1}^2 \frac{\gamma_i \mathbf{u}_i}{2}, \quad (15)$$

which clearly reduce to Eq. 13 which was developed above.




Figure 3. 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 observed under two such EMFs.

4.2. Three arbitrary motions

In the last section we showed that the shape from specular flow problem can be solved for two general EMFs that obey one restriction—their azimuthal angle must be the same. Here we show that this restriction can be eliminated by using three EMFs.

Consider a specular surface viewed under three different EMFs defined by the three different 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 $i \in \{1, 2, 3\}$.

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

Rewriting Eq. 10 in terms of these definitions we now get the following three equations

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

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

$$\tilde{\mathbf{J}} \begin{pmatrix} \frac{\tilde{u}_1 - \tilde{u}_2}{\tan \beta_2 - \tan \beta_1} - \frac{\tilde{u}_1 - \tilde{u}_3}{\tan \beta_3 - \tan \beta_1} \\ \frac{\tilde{v}_1 - \tilde{v}_2}{\tan \beta_2 - \tan \beta_1} - \frac{\tilde{v}_1 - \tilde{v}_3}{\tan \beta_3 - \tan \beta_1} \end{pmatrix} = \begin{pmatrix} 0 \\ \lambda \end{pmatrix} \quad (16)$$

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

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

Since the weighted combinations of specular flows is easily computed from measured data and known EMFs, Eq. 16 is a set of decoupled linear PDEs in h and k . Once

again, we arrive at an equation that resembles in structure the case of a single EMF around the viewing direction. Hence, as in the previous case, this equation can be solved using the same technique discussed in Sec. 3.

Not unlike the discussion in the case of two equal-azimuth EMFs, here too the result should not come as a surprise. Generalizing the discussion from Sec. 4.1, we note that given any number of EMFs ω_i , at least three of which have mutually distinct rotation axes $\hat{\mathbf{a}}_i$ that are unaligned with the view direction, one can find a set of constants γ_i such that the weighted sum

$$\sum_i \gamma_i \omega_i = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (17)$$

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 Eq. 2 to solve for γ_i . 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

$$\begin{pmatrix} 0 \\ 1 \end{pmatrix} = \sum_i \gamma_i \omega_i = \mathbf{J} \sum_i \gamma_i \mathbf{u}_i = \tilde{\mathbf{J}} \sum_i \frac{\gamma_i \mathbf{u}_i}{2}. \quad (18)$$

This equation then trivially reduces to Eq. 16 and can be solved using the technique discussed in Sec. 3.

5. Shape recovery under *unknown* EMFs

The previous section demonstrates that shape 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 rotations axes of the input EMFs are known. In this section we answer this question in the affirmative. Somewhat surprisingly, we next show that this requires no additional assumptions beyond those made so far.

Suppose first that the specular object is observed under an EMF around the viewing direction, but unlike Sec. 3, assume that the angular velocity ω is unknown. While they do not make use of it, Adato et al. [1] point out that for




Figure 4. The auxiliary function $k(x, y)$, i.e., the surface gradient orientation field, integrates to $2N\pi$ along any closed characteristic, where $N \in \{0, 1\}$ is determined by the number of times the characteristic intersects parabolic curves on the surface. Shown are several characteristics (in gray), the k function (red unit-length vector field), and the parabolic curves (in black) of a particular shape (Fig. 5a). Several characteristics are highlighted in different colors. k completes one full revolution along the blue and red characteristics (zero and eight intersections with parabolic lines, respectively), but zero revolutions along the green one (two intersections).

integrable surfaces, the SFSF equation provides a constraint on ω . They note that the linear equation in k shown in Eq. 6 can be written in terms of the *unit* specular flow field as

$$\hat{\mathbf{u}} \cdot \nabla k = \frac{2\omega}{\|\mathbf{u}\|}, \quad (19)$$

and since k is an orientation, the left hand side of this expression can be interpreted as a curvature measure, or what they termed the *specular curvature*. Integrating this curvature along a (closed) characteristic of the function must yield $2N\pi$ for some integer N in order to satisfy surface integrability [1]. Here we further observe that generically, N would be either 0 or 1, which depends on the number of times the characteristic intersects parabolic lines. While a proof is omitted for space considerations, Fig. 4 exemplifies that if the characteristic crosses (transversely) parabolic lines $2M$ times¹ then $N = 1 - (M \bmod 2)$.

The above observation translates to an algorithm for recovering the surface under a view-axis EMF with *unknown* velocity since the latter can be computed by

$$\omega = \frac{\pi}{\oint_C \|\mathbf{u}\|^{-1} ds} \quad (20)$$

along a closed characteristic C that crosses a parabolic line $2M$ times, with M odd.

Armed with the ability to recover the angular velocity of a view-axis EMF from specular flow, we proceed to address surface reconstruction under multiple unknown EMFs. Toward this goal, recall that several of EMFs can be linearly combined to yield a solvable SFSF of the sort derived in Eqs. 15 and 18. The key observation is that we can recover

¹The number of transversal intersections must be even since for a smooth surface both curves are closed.

the coefficients γ_i from the observed specular flows themselves, without knowledge of the EMFs. To do so we exploit the behavior of specular flow fields near the occluding contour of a smooth surface. Eq. 6 (in fact, its top row) suggests that when the specular flow is induced by a rotation about the view direction, its integral curves correspond to iso-contours of gradient magnitude (i.e., $h(x, y)$ from Eq. 5). It follows, then, that under such an EMF, the occluding contour of the surface is an integral curve of flow, or equivalently, that the specular flow field is everywhere tangent to the occluding contour. Thus, given a set of two or more observed specular flows \mathbf{u}_i , the desired coefficients γ_i are those that produce a combined flow $\mathbf{u} = \sum_i \gamma_i \mathbf{u}_i$ which satisfies this tangency condition. This 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 earlier in this section.

6. Experimental results

As proof of concept, we applied our theoretical results to data obtained synthetically by placing specular surfaces (e.g., Fig. 5a) in far-field illumination environments (Fig. 5b) and rendering them for a far-field viewer (Fig. 5c). This provides access to ground truth shape for comparison, and it also enables us to compute specular flow fields (using the generative equation of [1]) instead of measuring them from image data. An example is shown in Fig. 5d.

Given multiple specular flows from (two or three) distinct EMFs, the “normalized” linear combination was computed (Fig. 5e) by: 1) minimizing the least square error relative to the tangency condition; and 2) determining the view-axis angular velocity as described in Sec. 5. The recovered angular velocity ω and the coefficients γ_i completely determine a view-axis specular flow field. Using this flow field as input, the view-axis SFSF equation (Eq. 18) was solved for $h(x, y)$ and $k(x, y)$ (and thus the surface) using manually-provided initial conditions.

Carrying out this reconstruction procedure numerically requires care. While integration can be accomplished using standard numerical tools such as Runge-Kutta methods (e.g., Matlab’s `ode45` PDE solver), some preprocessing is necessary to deal with the singularities in specular flow that are exhibited along parabolic curves. At these points, specular flow exhibits singularities in *both* magnitude and orientation. Singularities in flow magnitude do not cause any difficulty when solving Eq. 18. (The equation in h is homogeneous and is therefore independent of flow magnitude; the equation in k depends only on inverse magnitude, which is well-defined everywhere.) To handle singularities in orientation, we note that parabolic curves can be localized from the observed flow field because of the unique behavior there; and once these curves are detected, one can remove the orientation singularities by ap-

appropriately ‘flipping’ the flow direction in regions between pairs of parabolic curves. As example of this procedure is shown in Figs. 5(e–h), where we have incorporated a computational step that creates a ‘flipping mask’ (Fig. 5g) that reflects the flow orientation in appropriate regions (Fig. 5h). By “correcting” the specular flow field in this way, one can then solve for $h(x, y)$ and $k(x, y)$ using standard numerical tools. A surface recovered using this technique is shown in Fig. 5i, and to compare this to ground truth, we report the absolute relative error as a percent of maximum ground truth surface height. A second example is presented in Fig. 6.

7. Summary and discussion

This paper addresses the reconstruction of specular surfaces in unknown, complex environments based on observations of relative object/environment motion. Environment motion induces a specular flow on the image plane, and this is related to shape through a non-linear PDE (the shape from specular flow equation). Whereas previous work has considered a limited class of environment motions (rotations about the view direction), we relax these restrictions and demonstrate that closed-form solutions can be found for *arbitrary* rotations, provided that one observes multiple specular flows corresponding to distinct environment motions. We also present a reconstruction procedure that is ‘auto-calibrating’ in the sense that no *a priori* information is required about the environment motions. The combined theory allows the recovery of general, non-convex surfaces in natural environments with content and relative motion that are both completely unknown.

As in the work of Adato et al. [1], this paper assumes that specular flow is given as input, and it does not address the problem of estimating specular flow from image data. This is an important direction of future work, and it is likely that some of the numerical tools developed here may prove useful for this task as well.

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 applicable 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

Funds for this project are being provided by the US National Science Foundation under grant IIS-0712956. O.B.S

and Y.A. also thank the generous support of the Frankel Fund and the Paul Ivanier Robotics Center at Ben-Gurion University. Additional funding for T.Z. and Y.V. was provided by the US National Science Foundation under CAREER award IIS-0546408.

References

- [1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Toward a theory of shape from specular flow. In *ICCV*, 2007.
- [2] A. Blake. Specular stereo. In *Proc. of the Int. Joint Conf. on Artif. Intell.*, pages 973–976, 1985.
- [3] A. Blake and G. Brelstaff. Geometry from specularities. In *ICCV*, 1988.
- [4] A. Blake and H. Bülthoff. Does the brain know the physics of specular reflection? *Nature*, 343:165–168, 1990.
- [5] A. Blake and H. Bülthoff. Shape from specularities: Computation and psychophysics. *Philosophical Transactions: Biological Sciences*, 331(1260):237–252, 1991.
- [6] T. Bonfort, P. Sturm, and P. Gargallo. General specular surface triangulation. In *Proc. ACCV*, pages 872–881, 2006.
- [7] O. Drbohlav and M. Chantler. Can two specular pixels calibrate photometric stereo? In *ICCV*, pages 1850–1857, 2005.
- [8] O. Drbohlav and R. Šára. Specularities reduce ambiguity of uncalibrated photometric stereo. In *ECCV*, pages 46–60, 2000.
- [9] R. Fleming, A. Torralba, and E. Adelson. Specular reflections and the perception of shape. *Journal of Vision*, 4:798–820, 2004.
- [10] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Trans. Pattern Anal. Mach. Intell.*, 3(6):661–669, 1981.
- [11] J. Koenderink and A. van Doorn. Photometric invariants related to solid shape. *Optical Acta*, 27(7):981–996, 1980.
- [12] K. Kutulakos and E. Steger. A theory of refractive and specular 3D shape by light-path triangulation. In *ICCV*, pages 1448–1455, 2005.
- [13] M. Longuet-Higgins. Reflection and refraction at a random moving surface. I. pattern and paths of specular points. *J. Opt. Soc. Am.*, 50(9):838–844, 1960.
- [14] S. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Trans. Robotics and Automation*, 6(4):418–431, 1990.
- [15] M. Oren and S. Nayar. A theory of specular surface geometry. *Int. J. Comput. Vision*, 24(2):105–124, 1997.
- [16] M. Osadchy, D. Jacobs, and R. Ramamoorthi. Using specularities for recognition. In *ICCV*, pages 1512–1519, 2003.
- [17] S. Roth and M. Black. Specular flow and the recovery of surface structure. In *Proc. CVPR*, pages 1869–1876, 2006.
- [18] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. *Int. J. Comput. Vision*, 64(1):31–67, 2005.
- [19] S. Savarese, F. Li, and P. Perona. What do reflections tell us about the shape of a mirror? In *Symp. Applied Percept. Graph. Visualiz.*, 2004.
- [20] 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.
- [21] 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.
- [22] J. Zheng, Y. Fukagawa, and N. Abe. Shape and model from specular motion. In *ICCV*, pages 72–79, 1995.
- [23] J. Zheng and A. Murata. Acquiring a complete 3D model from specular motion under the illumination of circular-shaped light sources. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):913–920, 2000.
- [24] A. Zisserman, P. Giblin, and A. Blake. The information available to a moving observer from specularities. *Image and Vision Computing*, 7:38–42, 1989.

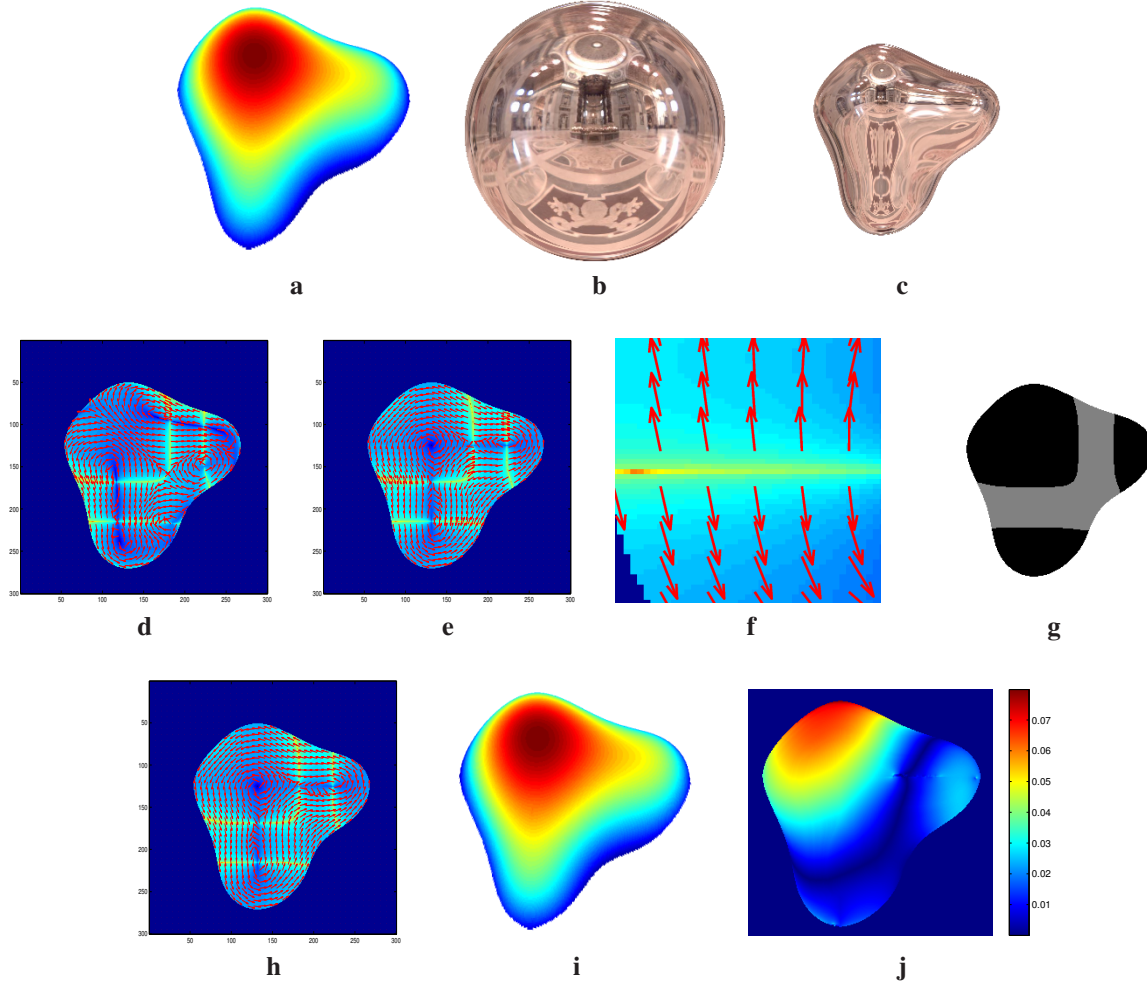


Figure 5. Example of specular shape reconstruction of a general specular surface observed under three arbitrary EMFs defined by the rotation axes $\hat{\mathbf{a}}_1 = (120^\circ, -66^\circ)$, $\hat{\mathbf{a}}_2 = (22.5^\circ, 16^\circ)$, $\hat{\mathbf{a}}_3 = (30^\circ, 36^\circ)$. (a): The original surface $f(x, y) = \sqrt{4 - x^2 - y^2} - \cos(2x - 2) - \sin 2y$. This surface includes multiple maxima and several elliptic and hyperbolic regions separated by parabolic curves. (b): The illumination environment (mapped to a disk of directions) in which the specular surface was observed. (c): One snapshot of the specular surface obtained under the first EMF. (d): A typical specular flow obtained from one of the EMFs. Flow magnitude is color coded while direction is shown by the superimposed unit-length vector field. (e): The weighted specular flow obtained from three different EMFs as specified in Eq. 16. (f): Detail of the weighted specular flow near a parabolic curve. Note the singularities in both magnitude and direction. (g): Mask used to eliminate orientation singularities in the weighted specular flow. Flow in the gray region was direction-reversed. (h): ‘Corrected’ weighted specular flow after direction flipping. (i): Reconstructed surface. (j): Absolute relative reconstruction error. Reported values are percent of maximum input surface height.




Figure 6. 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 from 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.