

A Linear Formulation of Shape from Specular Flow

Guillermo D. Canas¹, Yuriy Vasilyev¹, Yair Adato², Todd Zickler¹, Steven Gortler¹, Ohad Ben-Shahar²

¹ Harvard School of Engineering and Applied Sciences
Cambridge, USA

{gdiez, yuriy, zickler, sjg}@seas.harvard.edu

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

{adato, ben-shahar}@cs.bgu.ac.il

Abstract

When a curved mirror-like surface moves relative to its environment, it induces a motion field—or specular flow—on the image plane that observes it. This specular flow is related to the mirror’s shape through a non-linear partial differential equation, and there is interest in understanding when and how this equation can be solved for surface shape. Existing analyses of this ‘shape from specular flow equation’ have focused on closed-form solutions, and while they have yielded insight, their critical reliance on externally-provided initial conditions and/or specific motions makes them difficult to apply in practice. This paper resolves these issues. We show that a suitable reparameterization leads to a linear formulation of the shape from specular flow equation. This formulation radically simplifies the reconstruction process and allows, for example, both motion and shape to be recovered from as few as two specular flows even when no externally-provided initial conditions are available. Our analysis moves us closer to a practical method for recovering shape from specular flow that operates under arbitrary, unknown motions in unknown illumination environments and does not require additional shape information from other sources.

1. Introduction

An image of a curved, mirror-like surface presents an observer with a distortion of its surrounding environment, and there is interest in understanding when and how a mirror’s shape can be recovered from these distortions. Of particular interest are cases like those in Fig. 1, where the surface is viewed in a natural lighting environment, and little or no information about the environment is available from other image cues. This reconstruction problem is difficult for two reasons: 1) it is ill-posed, and 2) the relationship between a mirror’s shape and its image is quite complex.

One promising approach for coping with these difficulties is to exploit motion [16, 19, 1, 18]. When a specular surface moves relative to its environment, it induces a motion field on the image plane, and this motion field—



Figure 1. Specular surfaces under natural lighting present the observer with a distortion of the surrounding lighting environment. We explore when and how surface shape can be recovered from these observed distortions. (Image credits: Flickr users *spud* and @DS.)

termed *specular flow* [16]—provides valuable information about surface shape. In particular, it allows one to avoid reasoning about the content of the unknown environment by reasoning about the (simpler) relative motion instead. Such reasoning is facilitated by the basic relationship between a moving specular surface and the specular flow it induces, as has been recently derived by Adato *et al.* [1]. According to their formulation, when the surface and the observer move as a fixed pair relative to a distant environment, the observed specular flow field and the object motion yield a non-linear partial differential equation (PDE) in terms of the surface shape. This equation has been referred to as the *shape from specular flow* (SFSF) equation, and previous analyses have explored closed-form solutions for shape when one or more observed flows are available [1, 18].

Despite the insight it affords, reconstruction according to the existing formulation of the SFSF equation has two severe limitations. First, due to the complexity of the non-linear PDE, solutions thus far have only been determined for a very specific class of motions [1]. Second, even in these restricted cases, the solution cannot be obtained unless significant initial conditions (*e.g.*, multiple known surface curves) are provided by an external source [1, 18].

This paper resolves both of these limitations. We develop an alternative formulation of the shape from specular flow problem that produces a *linear* differential equation instead of a non-linear one. This linear formulation is intuitive and provides additional insights regarding the qualitative struc-

ture of specular flow fields. It also radically simplifies the reconstruction problem, providing a foundation for deeper studies as well as practical slutions for shape from specular flow.

Using this new formulation we prove two main results:

1. Specular shape can be reconstructed from a single flow under *general* known motion with adequate externally-provided initial conditions; and
2. Specular shape can be reconstructed from two or more specular flows under general *unknown* motions with *no* externally-provided initial conditions.

The latter result is particularly powerful because it means that reconstruction can succeed even when no additional shape information is available from other image cues.

2. Background and related work

While there have been numerous studies of specular surfaces under point-source lighting (e.g., [3, 14, 20]) and calibrated and/or active lighting environments (e.g., [4, 10, 12, 17]), there have been comparatively fewer studies involving complex natural lighting. Most notable is the work of Waldon and Dyer [19], who analyze the qualitative structure of specular flow fields; Roth and Black [16], who study parametric surfaces; and the aforementioned [1, 18] which consider general surfaces.

Following Adato *et al.* [1], and as shown in Fig. 2, the problem of shape from specular flow is formulated as follows. A specular surface is viewed from direction \hat{v} by an orthographic observer, and this surface is illuminated by a far-field illumination environment (an environment map). The surface is assumed to be the graph of a twice-differentiable function $f(\mathbf{x})$ defined on a region U of the image plane, and we define the *reflection vector* $\hat{r}(\mathbf{x})$ at each image point \mathbf{x} to be the mirror-reflection of the view direction \hat{v} about the surface normal $\hat{n}(\mathbf{x})$ at the point's back-projection. Image formation can be imagined as shooting a ray in direction \hat{v} through each image point \mathbf{x} , computing the corresponding reflection vector $\hat{r}(\mathbf{x})$, and then sampling the illumination environment in that direction. Note that self-reflections of the surface are not allowed.

When the surface and observer move as a fixed pair relative to the environment, the motion induces a motion field $\mathbf{u}(\mathbf{x})$ on the image plane. This observed specular flow provides access to surface shape information through the *shape from specular flow* equation [1]:

$$\mathbf{w}(f_x, f_y) = \mathbf{J}(f_x, f_y, f_{xx}, f_{xy}, f_{yy})\mathbf{u}(\mathbf{x}), \quad (1)$$

where \mathbf{w} is a spherical vector field on the environment sphere that describes the relative motion, and \mathbf{J} is a 2×2 matrix with non-linear entries in the first and second derivatives of the unknown height field f . This equation represents a coupled pair of non-linear second-order PDEs in




Figure 2. A mirror-like surface is illuminated by a far-field environment and viewed orthographically to produce an image.

terms of the unknown surface; and given an observed specular flow \mathbf{u} and known motion field \mathbf{w} , one can theoretically reconstruct the surface by solving these equations.

As stated earlier, what makes this approach attractive is that it provides shape information in terms of the environment *motion* instead of environment *content*, hence facilitating reconstruction in uncontrolled settings. Due to the complexity of the non-linear PDE, however, solutions have only been determined for the very specific class of motions in which the environment rotates about the view direction [1]. (In this case, Eq. 1 reduces to a pair of linear first-order PDEs.) Additionally, a solution cannot be obtained unless significant initial conditions are provided as input. This usually includes knowledge of the surface gradient (f_x, f_y) along one or more surface curves [1, 18].

3. Linear shape from specular flow

To begin, we depart from the original formulation [1] and derive the shape from specular flow equation in terms of the reflection vector field defined on the surface. Similar to the previous section, the surface we seek is assumed to be a smooth closed manifold surface M , and we assume that we observe a closed connected region $U \subset \mathbb{R}^2$ that is the orthogonal projection of M in view direction \hat{v} and onto the image plane. (See Fig. 2.) We denote by \dot{U} and ∂U the interior and (occluding) boundary of U , respectively. We further assume the following:

Assumption 1.

- a) M does not self-reflect
- b) M contains no developable regions
- c) M contains no flat points

Note that (b) and (c) are true for generic smooth surfaces.

Consider a lighting direction $\hat{c}(0) \in \mathbb{S}^2$ as a point on the environment sphere. Environment motion can be considered a time-varying rotation of this sphere, and such a rotation induces an angular trajectory $\hat{c}(t)$ of the lighting

direction, written

$$\hat{\mathbf{c}}(t) = \mathbf{R}(t)\hat{\mathbf{c}}(0).$$

As this environment point moves, it is reflected from the surface onto the image plane, and assuming the Gaussian curvature of the surface does not vanish (to be relaxed momentarily), this reflection induces a smooth trajectory $\mathbf{x}(t)$ on this plane. Corresponding to the trajectory on the image plane is a surface curve $\mathbf{p}(t)$, and if we denote the surface normal and reflection vectors along this curve by $\hat{\mathbf{n}}(t)$ and $\hat{\mathbf{r}}(t)$, respectively, then the law of specular reflection guarantees that $\hat{\mathbf{r}}(t) = \hat{\mathbf{c}}(t)$ for all t .

Having the environment trajectory $\hat{\mathbf{c}}(t)$, its derivative at time $t = 0$ is

$$\hat{\mathbf{c}}'(0) = \boldsymbol{\omega} \times \hat{\mathbf{c}}(0), \quad (2)$$

where $\boldsymbol{\omega}$ is the angular velocity vector at that initial time ($\hat{\boldsymbol{\omega}} \triangleq \boldsymbol{\omega}/\|\boldsymbol{\omega}\|$ is the rotation axis and $\|\boldsymbol{\omega}\|$ is the angular speed). At the same time, the chain rule tells us

$$\hat{\mathbf{r}}'(0) = (D\hat{\mathbf{r}}(0)) \mathbf{u}(\mathbf{x}(0)), \quad (3)$$

where $\mathbf{u}(\mathbf{x}) \triangleq d\mathbf{x}/dt$ denotes the specular flow field, and $D\hat{\mathbf{r}} \triangleq \partial\hat{\mathbf{r}}/\partial\mathbf{x}$ is the 2×3 Jacobian of the reflection map $\hat{\mathbf{r}}: U \rightarrow \mathbb{S}^2$ that assigns a reflection vector to each image point. Equations 2 and 3 are equal by the law of specular reflection, and clearly this is true for all initial environment points $\hat{\mathbf{c}}(0)$ and thus all $\mathbf{x} \in \dot{U}$. This leads us our desired relation,

$$(D\hat{\mathbf{r}}(\mathbf{x}))\mathbf{u}(\mathbf{x}) = \boldsymbol{\omega} \times \hat{\mathbf{r}}(\mathbf{x}), \quad (4)$$

which is valid for all non-parabolic $\mathbf{x} \in \dot{U}$.

Equation 4 is an alternative expression for the shape from specular flow equation derived previously [1], and it expresses the relation in terms of the surface’s reflection vector field $\hat{\mathbf{r}}(\mathbf{x})$ instead of the surface itself. This has the important advantage of producing a linear PDE instead of a non-linear one, and since the reflection field uniquely determines the surface when the view vector is known (*i.e.*, $\hat{\mathbf{n}} \propto \hat{\mathbf{v}} + \hat{\mathbf{r}}$), it provides equivalent shape information.

The remainder of this paper analyzes the linear SFSF PDE (Eq. 4) to understand when and how it can be solved for $\hat{\mathbf{r}}(\mathbf{x})$ and surface shape. We begin with the case in which we observe a single flow \mathbf{u} under a known environment motion $\boldsymbol{\omega}$, and we study the uniqueness of solutions by generalizing the equation to handle parabolic surface points in a natural way. We then move on to consider two flows under motions that may or may not be known. For brevity, results are stated without proof in this document. The proofs can be found in an associated technical report [6].

4. Shape from one specular flow: uniqueness

Suppose we are given as input a specular flow field $\mathbf{u}(\mathbf{x})$ captured under known environment rotation $\boldsymbol{\omega}$ and from




Figure 3. ℓ -curves (right) induced by a smooth closed specular surface observed from direction $\hat{\mathbf{v}}$ under a far-field illumination environment (left) rotating about axis $\boldsymbol{\omega}$. (This “blob” surface (from [18]) is shown with gray iso-height contours and black parabolic curves.) ℓ -curves on the image plane (in blue) generically form smooth closed curves, corresponding to the “reflection images” of concentric circles on the environment sphere. An exception is the reflection image of the concentric circle that passes through the anti-view $-\hat{\mathbf{v}}$; this curve intersects the boundary (at red points). Zeros of flow (in magenta) occur at surface points where the surface reflection vector is aligned with $\pm\boldsymbol{\omega}$.

known viewpoint $\hat{\mathbf{v}}$. Throughout this document, we will assume that this flow field has been estimated by an “ideal” optical flow algorithm and is identical to the true two-dimensional motion field. As has been noted previously [1, 19], specular flow fields exhibit unique structure at parabolic surface points. In particular, as one approaches the parabolic line, the flow magnitude grows unbounded, the orientation is aligned with the flat direction on the surface, and as one crosses the parabolic line the orientation undergoes a 180° jump (see, for example, Fig. 5c). It is thus safe to assume that if the flow \mathbf{u} is known, the parabolic points are known, and the flat direction at each parabolic point is also known.

Our immediate goal is to understand uniqueness of solutions for shape, particularly in the presence of parabolic surface points. Doing so requires generalizing Eq. 4, and to this end, we define the ℓ -curves: $\{\mathbf{x} \mid \hat{\boldsymbol{\omega}}^\top \hat{\mathbf{r}}(\mathbf{x}) = C\}$ for each constant $C \in [-1, 1]$. As depicted in Fig. 3, these are curves on the image plane where we observe, in reflection, an “ $\boldsymbol{\omega}$ -centered latitude” circle of the environment.

For generic smooth surfaces, almost all of such ℓ -curves are smooth closed curves (blue in Fig. 3). The Inverse Function Theorem tells us that a point \mathbf{x} is a smooth point of an ℓ -curve if and only if the derivative $D(\hat{\boldsymbol{\omega}}^\top \hat{\mathbf{r}}(\mathbf{x})) = \hat{\boldsymbol{\omega}}^\top (D\hat{\mathbf{r}}(\mathbf{x}))$ exists and is non-zero. “Exceptional points” appear where this derivative vanishes. This happens at isolated points where we see the reflection of the rotation axis: $\hat{\boldsymbol{\omega}}^\top \hat{\mathbf{r}}(\mathbf{x}) = \pm 1$ (magenta points in Fig. 3). It also happens

at isolated parabolic points where the one-dimensional column span of $D(\hat{\mathbf{r}}(\mathbf{x}))$ happens to be orthogonal to ω (green in Fig. 3). (While space prohibits a detailed discussion, one can show that the structure of the ℓ -curves near these points is similar to the structure of isophotes near illuminance critical points of diffusely shaded surfaces [11].) Additional isolated non-smooth points occur at the boundary. Since $\hat{\mathbf{r}}(\mathbf{x}) = -\hat{\mathbf{v}}, \forall \mathbf{x} \in \partial U$, the boundary is the pre-image of a single point on the environment sphere, and ℓ -curves with $C = -\hat{\omega}^\top \hat{\mathbf{v}}$ can intersect this boundary (red points in Fig. 3). Since all of these exceptional points are isolated points, we ignore them for the remainder of this section.

The following lemma relates the flow field \mathbf{u} to the ℓ -curves, and shows that a globally smooth orientation can be assigned to this data.

Lemma 1. *Given a flow field $\mathbf{u}(\mathbf{x})$ corresponding to rotation ω , there exists a unit-length vector field $\hat{\mathbf{d}}(\mathbf{x})$ that has the following four properties.*

1. $\hat{\mathbf{d}}(\mathbf{x})$ is smooth away from the isolated exceptional points.
2. The field lines of $\hat{\mathbf{d}}(\mathbf{x})$ are the ℓ -curves.
3. For all non-parabolic \mathbf{x} , $\hat{\mathbf{d}}(\mathbf{x}) = \text{sgn}(\mathbf{x})\mathbf{u}(\mathbf{x})/\|\mathbf{u}(\mathbf{x})\|$ where $\text{sgn}(\mathbf{x})$ is the sign of Gaussian curvature function.
4. For all parabolic \mathbf{x} , $\hat{\mathbf{d}}(\mathbf{x})$ is aligned with the flat direction at \mathbf{x} .

With this Lemma in hand, one can show that the reflection field $\hat{\mathbf{r}}(\mathbf{x})$ is a solution to the PDE

$$(D\hat{\mathbf{r}}(\mathbf{x}))\hat{\mathbf{d}}(\mathbf{x}) = (\omega \times \hat{\mathbf{r}}(\mathbf{x})) \text{sgn}(\mathbf{x})\|\mathbf{u}(\mathbf{x})\|^{-1}, \quad (5)$$

which naturally extends Eq. 4 to handle parabolic points. In particular, at non-parabolic points the two equations are equivalent by Property 3 in Lemma 1. At parabolic points, $\|\mathbf{u}\|^{-1} = 0$ and Eq. 5 becomes $(D\hat{\mathbf{r}}(\mathbf{x}))\hat{\mathbf{d}}(\mathbf{x}) = 0$, which is true by Property 4. Moreover, the right hand side of Eq. 5 is continuous because the discontinuities in $\text{sgn}(\mathbf{x})$ only occur at parabolic points, where $\|\mathbf{u}\|^{-1} = 0$.

Let us now select a connected smooth segment of an ℓ -curve, say $\gamma(s)$, that is arc-length parameterized by s . We can plug this curve into equation 5 to obtain:

$$(D\hat{\mathbf{r}}(\gamma(s)))\gamma'(s) = \omega \times \hat{\mathbf{r}}(\gamma(s)) \text{sgn}(s)\|\mathbf{u}(\gamma(s))\|^{-1}.$$

This is a continuous ordinary differential equation in s , and we can thus conclude that knowledge of the reflection vector $\hat{\mathbf{r}}$ at any single point on an ℓ -curve constrains $\hat{\mathbf{r}}$ along the entire connected component of that ℓ -curve in any solution to Eq. 5. We summarize this with the following corollary, which simply follows from the fact that two ℓ -curves $L(C), L(C')$ are disjoint for $C \neq C'$:

Corollary 1. *Given a closed, connected subset $U \subset \mathbb{R}^2$ of the image plane that is the orthogonal projection in direction $\hat{\mathbf{v}}$ of a manifold surface M (satisfying Assumptions 1a–c), away from exceptional points, the solution $\hat{\mathbf{r}}(\mathbf{x})$ to Eq. 5 is unique if and only if we are given at least one constraint point $\hat{\mathbf{r}}(\mathbf{x}_i)$ on each connected smooth component of each ℓ -curve.*

This result, which shows that the surface can be recovered under arbitrary rotations ω , represents a significant advance over the existing single-flow approach [1], which is restricted to the special case in which $\omega \propto \hat{\mathbf{v}}$. Our analysis does not, however, alleviate the need for initial conditions. On the contrary, it shows that with a single flow, the need for externally-provided initial conditions is unavoidable.

5. Shape from two specular flows

We have shown that a complete surface cannot be reconstructed from a single flow without additional input. That said, there are isolated exceptional points at which the shape is known. These are the points where one observes the rotation axis in reflection (magenta points in Fig. 3), and at these points the reflection vector—and therefore the surface normal—is known up to sign: $\hat{\mathbf{r}} = \pm\hat{\omega}$. These points are detectable because they are points of zero flow, but since there are no continuous ℓ -curves that pass through them, we cannot propagate their shape information to reconstruct larger regions.

The key idea exploited in the remainder of this paper is that shape information at these “zero-flow points” can be propagated when two or more distinct flows are observed. In fact, in Sect. 6 we will show that the surface can be recovered even when the environment motions (the ω_i ’s) are not known *a priori*.

A depiction of two distinct specular flows is shown in Fig. 4. Here, the red and blue curves represent ℓ -curves of two linearly independent flow fields induced by the rotations ω_1 and ω_2 under view $\hat{\mathbf{v}}$. As in the single-flow case of the previous section, the ℓ -curves are reflection-images of “latitude” circles on the environment sphere. Another significant landmark in this figure is the green circle on the environment sphere, which represents the great circle spanned by ω_1 and ω_2 . This circle has the special property that at every point on its reflection image (green curves in U), the red and blue flow directions are co-linear. These image points are rare, however, and the two flow directions are distinct over most of the surface.

To begin our analysis, we first show that one will always observe within \hat{U} at least one zero-flow point where the reflection vector is known and collinear with the rotation axis ω . Doing so requires additional assumptions about surface M , so we supplement Assumptions 1a–c, by further assuming that the surface does not self-occlude, that points of zero




Figure 4. l -curves induced by distinct rotations ω_1, ω_2 (red and blue). As in the one-flow case, l -curves are reflection-images of concentric circles on the environment sphere. The reflection vector $\hat{\mathbf{r}}$ is known at a point of zero flow (blue point), and as depicted by the black curve, this can be used as a seed to recover the entire surface by integrating along segments of l -curves.

flow are not parabolic points, and that there are no closed cusp curves (i.e., parabolic curves where the flat direction is everywhere tangent to the parabolic curve).

Assumption 2.

- a) $\forall \mathbf{x} \in \dot{U}, \hat{\mathbf{n}}(\mathbf{x})^\top \hat{\mathbf{v}} > 0$
- b) $\hat{\mathbf{r}}(\mathbf{x}) = \pm \omega_i / \|\omega_i\| \Rightarrow |D\hat{\mathbf{r}}(\mathbf{x})| \neq 0$
- c) There are no closed cusp curves

Note that (c) is true for generic smooth surfaces. With these assumptions in hand, one can indeed show the following.

Lemma 2. *Given a closed, connected subset $U \subset \mathbb{R}^2$ of the image plane that is the orthogonal projection of a manifold surface M (smooth, closed, with no flat points), and a rotation ω with associated specular flow field $\mathbf{u}(\mathbf{x})$ defined over U , we can compute a single constrained point $\hat{\mathbf{r}}(\mathbf{x}) = \pm \omega, \mathbf{x} \in \dot{U}$.*

Our goal now is to use a second flow field to propagate this shape information and reconstruct the entire surface.

Definition 1. We say that two points $\mathbf{x}_1, \mathbf{x}_2 \in \dot{U}$ are l -connected under flows $\{\mathbf{u}_i\}$ if there is a piecewise-smooth path between them in \dot{U} composed of a sequence of pieces of l -curves of the SFSF equations (Eq. 5) from any of the input flows $\{\mathbf{u}_i\}$ and their corresponding $\{\omega_i\}$. We call one such path an l -path of the two points under $\{\mathbf{u}_i\}$. We say that an open set \dot{U} is l -connected under $\{\mathbf{u}_i\}$ if $\forall \mathbf{x}_1, \mathbf{x}_2 \in \dot{U}, \mathbf{x}_1, \mathbf{x}_2$ are l -connected under $\{\mathbf{u}_i\}$.

In Fig. 4, one l -path is shown in black, and it connects a point at which the reflection vector is known (a blue point of zero flow) to another point on the surface. As one might imagine, this connection can be exploited to recover the reflection vector at the second point:

Lemma 3. *Given two l -connected points $\mathbf{x}_1, \mathbf{x}_2 \in \dot{U}$, if $\hat{\mathbf{r}}(\mathbf{x}_1)$ is known, then $\hat{\mathbf{r}}(\mathbf{x}_2)$ is uniquely determined in any solution to a pair of SFSF equations Eq. 5 using the two flows.*

Finally, we can prove that the set \dot{U} is l -connected, and this yields the following.

Theorem 1. *Given a closed, connected subset $U \subset \mathbb{R}^2$ of the image plane that is the orthogonal projection in direction $\hat{\mathbf{v}}$ of a manifold surface M (satisfying Assumptions 1a-c and 2a-c), along with one constrained point $\hat{\mathbf{r}}(\mathbf{x}), \mathbf{x} \in \dot{U}$, and two linearly independent rotations ω_1, ω_2 with associated specular flow fields $\mathbf{u}_1(\mathbf{x}), \mathbf{u}_2(\mathbf{x})$ defined over U , there is a unique solution to the associated pairs of SFSF equations Eq. 5.*

Taken together, Lemma 2 and Theorem 1 tell us that we can reconstruct the reflection field $\hat{\mathbf{r}}$ (and hence the surface) over U given two rotations ω_1, ω_2 and flows $\mathbf{u}_1, \mathbf{u}_2$. Importantly, externally-provided initial conditions are not required for this reconstruction, because there are sufficient constraints that are “intrinsic” to the observed flows.

5.1. A reconstruction algorithm

Theorem 1 suggests that when two linearly independent flows are available, we can use a relatively simple reconstruction algorithm. Given two flows $\mathbf{u}_1(\mathbf{x})$ and $\mathbf{u}_2(\mathbf{x})$ captured under known rotations ω_1 and ω_2 , we build an overconstrained linear system

$$\mathbf{A}\mathbf{r} = \mathbf{b}, \tag{6}$$

where \mathbf{r} is a vector that contains the three components of all of the reflection vectors in the domain of interest. The rows of \mathbf{A} result from discretizing the linear SFSF PDE using a first-order finite differences scheme for both $\{u_1, \omega_1\}$ and $\{u_2, \omega_2\}$, and the corresponding entries in \mathbf{b} are zero. Since a discrete image does not generically include exact samples of parabolic surface points, it is sufficient to discretize

$$(D\hat{\mathbf{r}}(\mathbf{x}))\hat{\mathbf{u}}(\mathbf{x}) = (\omega \times \hat{\mathbf{r}}(\mathbf{x})) \|\mathbf{u}(\mathbf{x})\|^{-1}, \tag{7}$$

which is numerically stable and equivalent to Eq. 5 at non-parabolic points. In addition, we augment these linear constraints on \mathbf{r} by including constraints from one or more points at which the reflection vectors are known. As shown in Lemma 2, under the given conditions, it is always possible to find a constraint using the zeros of flow. At the points where the flow \mathbf{u}_i vanishes, we know that $\hat{\mathbf{r}} = \pm \omega_i / \|\omega_i\|$. By Assumption 2b, we know that these points are isolated and, in practice, we find them as the local minima of the flow magnitude. Once we choose a sign for \mathbf{r} at the constrained point, the resulting linear system is overconstrained and can be solved in the least squares sense to obtain a reflection vector field $\hat{\mathbf{r}}(\mathbf{x})$ (and a surface). Note that this reconstruction is up to a binary choice.

6. Motion and shape from two specular flows

The previous section showed that two distinct specular flows carry enough information to recover the surface (up to a binary choice) when the environment motions ω_1, ω_2 are known. Here we show that these motions can also be recovered from the specular flow fields. The result is an ‘auto-calibrating’ reconstruction procedure that simultaneously recovers shape and motion from two (or more) flows.

We show that shape and motion can be recovered from two specular flows in three stages. First, borrowing results from a related problem, we show that the unknown environment motions ω_1 and ω_2 can be recovered up to an orthogonal transform. Second, using the ‘calibrated’ reconstruction results from the previous section, we show that these motion estimates determine the reflection field $\hat{\mathbf{r}}$ up to this same orthogonal transform. Finally, we show that enforcing integrability of the specular surface resolves the orthogonal ambiguity, and thus that shape can be uniquely recovered.

6.1. Recovering environment motions

When the SFSF equation is expressed in terms of the reflection field, the resulting equation (Eq. 4) is the same as that arising in flow-based calibration of generic central-projection cameras [13, 15, 7]. In that problem, one seeks to recover a map $g: U \in \mathbb{R}^2 \rightarrow \mathbb{S}^2$ that assigns a camera ray to each image point. The only difference in the present case is that the rays being assigned to each image point are interpreted as reflection vectors on a surface (*i.e.*, $g(\mathbf{x}) = \hat{\mathbf{r}}(\mathbf{x})$).

The similarity between the two problems allows us to borrow the following result related to recovering the unknown motions ω_1 and ω_2 .

Result 1. [Espuny] *Given two linearly independent flows $\mathbf{u}_1, \mathbf{u}_2$ induced by unknown but linearly independent angular velocities ω_1, ω_2 , the angular velocities can be determined up to an orthogonal transformation of \mathbb{R}^3 .*

Espuny’s result [7] comes in the form of a clever linear solution for the unknown motions. The algorithm translates to our problem as follows. Up to a per-point sign choice, the reflection field $\hat{\mathbf{r}}$ can be conveniently represented in a coordinate system formed by (the unknown) ω_1, ω_2 :

$$\hat{\mathbf{r}} \propto \bar{\mathbf{r}} = \alpha_1 \omega_1 + \alpha_2 \omega_2 + (\omega_1 \times \omega_2),$$

where the coefficients α_i depend only on the components of the two flows and their first derivatives. Specifically, denoting the components of each flow as $\mathbf{u}_i = (u_i, v_i)$; letting $\mathbf{V}(\mathbf{x}) = [\mathbf{u}_1(\mathbf{x}) \mid \mathbf{u}_2(\mathbf{x})]$ be the 2×2 matrix formed by stacking two flow vectors column-wise; and letting $\mathbf{V}_x, \mathbf{V}_y$ be derivatives of \mathbf{V} with respect to the the image coordinates, these coefficients can be written [7]:

$$\begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} = \begin{pmatrix} \frac{\partial u_2}{\partial x} + \frac{\partial v_2}{\partial y} - \frac{\det \mathbf{V}_x}{\det \mathbf{V}} u_2 - \frac{\det \mathbf{V}_y}{\det \mathbf{V}} v_2 \\ -\frac{\partial u_1}{\partial x} - \frac{\partial v_1}{\partial y} + \frac{\det \mathbf{V}_x}{\det \mathbf{V}} u_1 + \frac{\det \mathbf{V}_y}{\det \mathbf{V}} v_1 \end{pmatrix}.$$

Then, provided that the flows \mathbf{u}_1 and \mathbf{u}_2 are linearly independent, the Gram matrix of the two rotation vectors can be determined using

$$\begin{pmatrix} \omega_1^\top \omega_1 & \omega_1^\top \omega_2 \\ \omega_2^\top \omega_1 & \omega_2^\top \omega_2 \end{pmatrix} = \begin{pmatrix} \alpha_2 \\ -\alpha_1 \end{pmatrix} (-\alpha_2, \alpha_1) - \begin{pmatrix} (D\alpha_2)^\top \\ -(D\alpha_1)^\top \end{pmatrix} V, \quad (8)$$

where $D\alpha_i$ is the spatial gradient of the coefficient function.

Given two flow fields \mathbf{u}_1 and \mathbf{u}_2 , we can write a linear constraint on the Gram matrix (Eq. 8) for each image point \mathbf{x} at which $\mathbf{u}_1(\mathbf{x})$ and $\mathbf{u}_2(\mathbf{x})$ are not collinear. As described in Sect. 5, while points with collinear flows always exist, the flow directions at most points are distinct. Thus, from a typical pair of flows, we can reliably estimate the Gram matrix by combining a large number of independent estimates (see discussion in [8]). This is important because individual instances of Eq. 8 requires computing second derivatives of the flow field and are likely to be noisy.

Once the Gram matrix is determined, the rotation vectors ω_1 and ω_2 are determined up to an orthogonal transform.

6.2. Recovering the reflection field

The results of Sect. 5 tell us that surface shape can be recovered from as few as two known rotations ω_1 and ω_2 . Since $D(\mathbf{T}\hat{\mathbf{r}}) = \mathbf{T}\omega \times \mathbf{T}\hat{\mathbf{r}}$ for any orthogonal matrix \mathbf{T} , similar results will hold when the rotations are modified by an orthogonal transform. The only difference is that the surface will be determined up to the same orthogonal transform, with an additional binary choice.

So, given two flow fields $\mathbf{u}_1(\mathbf{x})$ and $\mathbf{u}_2(\mathbf{x})$, we can use the results above to recover initial estimates of the rotation vectors ω_1 and ω_2 , and the algorithm described in Sect. 5.1 to recover initial estimates of the reflection field $\hat{\mathbf{r}}(\mathbf{x})$. This initial estimate, $\hat{\mathbf{r}}^*(\mathbf{x})$ say, will differ from the true reflection field by an orthogonal transform.

6.3. Recovering the surface

Since the surface we seek is a smooth manifold, its normal field $\hat{\mathbf{n}}(\mathbf{x})$ must be integrable [9], and this places constraints on the reflection field $\hat{\mathbf{r}}(\mathbf{x})$. Here we show that these constraints are sufficient to resolve the orthogonal ambiguity and uniquely determine the surface.

Let the components of a surface normal vector be written $\hat{\mathbf{n}} = (n_1, n_2, n_3)$. The normal field corresponding to an integrable surface must satisfy [2]

$$\begin{pmatrix} n_1 \\ n_3 \end{pmatrix}_y = \begin{pmatrix} n_2 \\ n_3 \end{pmatrix}_x.$$

The law of specular reflection implies $\hat{\mathbf{n}} \propto \hat{\mathbf{r}} + \hat{\mathbf{v}}$, and by assuming that the view direction is $\hat{\mathbf{v}} = (0, 0, 1)^\top$, we can write this constraint in terms of the components of the unit




Figure 5. Results with synthetic data. (a): Image of a specular surface rendered with the St. Peter’s environment map. (b): Specular flow induced by an environment rotation. Flow magnitude is color coded while orientation is shown by the superimposed unit-length vector field. The green circle indicates a constraint point (a vanishing point of flow). (c): Ground-truth surface normal field, with xyz -components color in RGB. (d): Surface normal field estimated from two distinct specular flows. (e): Absolute angular error in the surface normal estimates, measured in degrees.

reflection field, $\hat{\mathbf{r}} = (r_1, r_2, r_3)$:

$$\begin{aligned} \left(\frac{r_1}{r_3 + 1} \right)_y &= \left(\frac{r_2}{r_3 + 1} \right)_x \\ \Leftrightarrow (r_3 + 1)(r_{2x} - r_{1y}) + r_1 r_{3y} - r_2 r_{3x} &= 0. \end{aligned} \quad (9)$$

Theorem 2. *If $\hat{\mathbf{r}}(\mathbf{x})$ is the unit-length field of reflection vectors from a smooth (integrable) surface under view $\hat{\mathbf{v}} = (0, 0, 1)^\top$, the set of orthogonal transformations $\hat{\mathbf{r}}^*(\mathbf{x}) = \mathbf{T}\hat{\mathbf{r}}(\mathbf{x})$ that satisfy the integrability constraint of Eq. 9 are the identity matrix and the diagonal matrix with diagonal entries $(-1, -1, 1)$.*

In Sects. 6.1 and 6.2 we showed how to estimate the reflection vector field up to an orthogonal transform. Theorem 2 implies that there are only two unit-length reflection fields that correspond to integrable surfaces, and it is unlikely that the estimated reflection field $\hat{\mathbf{r}}^*(\mathbf{x})$ is one of them. Furthermore, it implies that we can find the two “correct” surfaces by searching for an orthogonal transform that maps $\hat{\mathbf{r}}^*(\mathbf{x})$ to a reflection field that satisfies the integrability constraint of Eq. 9.

7. Experimental Results

As proof of concept, we first test our approach on synthetic data obtained by placing a virtual specular surface in a captured far-field illumination environment and rendering an image for a far-field observer (Fig. 5a). Two specular flows, $\mathbf{u}_1, \mathbf{u}_2$, corresponding to scene rotations about two independent axes, ω_1, ω_2 are computed by a generative equation for specular flow [1] (one of these flows is shown in Fig. 5c). The angular velocity vectors are first estimated up to an orthogonal transformation by computing the Gram matrix as an average of the values computed from Eq. 8 at every pixel. At this point, we are free to choose any rotation vectors ω_1^* and ω_2^* that are consistent with the Gram matrix, so we select them to be along the x -axis, and in the xy -plane, respectively, in the image coordinate system.

The given flows $\mathbf{u}_1, \mathbf{u}_2$ and the initial angular velocity vectors ω_1^*, ω_2^* are then used as input to the algorithm of

Sect. 5.1, whose output is a reflection vector field $\hat{\mathbf{r}}^*(\mathbf{x})$. From Sect. 6.2, we know that this reflection vector field is correct up to an unknown orthogonal transformation. We find this transformation, \mathbf{T} , by finding a (local) minima of a constrained non-linear system of equation based on Eq. 9. The constraints guarantee the solution \mathbf{T} to be orthogonal, and in all cases we have considered, we have succeeded by using the 3×3 identity matrix as initialization.

Applying transformation \mathbf{T} to the estimated reflection field resolves the orthogonal ambiguity and allows us to compute the normal vector field. Figures 5c and 5d show the x, y, z components of the actual and estimated normal fields, respectively, represented as RGB colors. Except near the boundary, where the flow becomes degenerate, the angular error in our estimate is smaller than 0.1° (Fig. 5e). Even near the boundary the error is less than one degree.

In addition to these synthetic results, we have evaluated our approach using data acquired with a custom device that replicates the imaging process of Fig. 2. This device, shown in the left of Fig. 6, provides an object-camera distance of approximately 1m and allows them to be rotated as a fixed pair about an arbitrary axis. We acquired two image sequences for the object shown in the middle of Fig. 6, with one rotation axis approximately aligned with the view direction and the other approximately orthogonal to it. Specular flow fields were estimated from the image data using [5], and these flows were used to recover the motions and the normal field as described above. The right of Fig. 6 shows the surface that was recovered by integrating this normal field as well as a comparable view of the actual object.

8. Conclusion

This paper shows that the shape from specular flow equation, which is non-linear in the unknown surface, can be re-formulated as a linear PDE. The enabling idea is to represent the surface by its field of reflection vectors. The resulting linear PDE provides insight into the qualitative structure of specular flow fields and produces some powerful reconstruction algorithms.



Figure 6. Results with captured data. **Left:** Acquisition system that enables camera/object rotation about arbitrary axes. **Middle:** Object under study. **Right:** Shape recovered from two specular flows, shown with a comparable view of the object.

As in previous work, this paper focusses on the task of recovering shape when specular flow is known, and it does not address the problem of estimating specular flow from image data. Accurately recovering specular flow in the presence of singularities induced by parabolic surface points is an interesting open problem, and it is one that is likely to benefit from the linear formulation presented here.

In fact, our linear formulation of shape from specular flow should enable deeper studies of the problem in general. Directions to consider include an expanded use of the integrability constraint, which yields significant power and may even be sufficient to allow single-flow reconstructions. Another is the case in which a specular object moves relative to the observer, which is more natural for a human or robotic observer.

Acknowledgements

The authors thank Peter Sturm for suggesting the link between shape from specular flow and generic camera calibration. This project is funded by the US National Science Foundation under grant IIS-0712956 and the Israel Science Foundation under grant No. 1245/08. Additional funding for T.Z. and Y.V. was provided by a Sloan Research Fellowship and NSF Career Award IIS-0546408. 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.

References

- [1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Toward a theory of shape from specular flow. *ICCV*, 2007.
- [2] P. Belhumeur, D. Kriegman, and A. Yuille. The bas-relief ambiguity. *IJCV*, 35(1):33–44, 1999.
- [3] A. Blake. Specular stereo. In *Proc. of the Int. Joint Conf. on Artif. Intell.*, pages 973–976, 1985.
- [4] T. Bonfort, P. Sturm, and P. Gargallo. General specular surface triangulation. In *Proc. ACCV*, pages 872–881, 2006.
- [5] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *ECCV*, 2004.
- [6] G. D. Canas, Y. Vasilyev, Y. Adato, T. Zickler, S. Gortler, and O. Ben-Shahar. Unique specular shape from two specular flows. Technical Report XXXX, Harvard University, 2009.
- [7] F. Espuny. A closed-form solution for the generic self-calibration of central cameras from two rotational flows. *Proc. VISAPP*, 1:26–31, 2007.
- [8] F. Espuny. Generic self-calibration of central cameras from two “real” rotational flows. In *Proc. OMNIVIS*, 2008.
- [9] B. Horn and M. Brooks. The variational approach to shape from shading. *CVGIP*, 33:174–208, 1986.
- [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. Illuminance critical points on generic smooth surfaces. *JOSA A*, 10(5):844–854, 1993.
- [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] D. Nistér, H. Stewénus, and E. Grossmann. Non-parametric self-calibration. *ICCV*, 2005.
- [14] M. Oren and S. Nayar. A theory of specular surface geometry. *Int. J. Comput. Vision*, 24(2):105–124, 1997.
- [15] S. Ramalingam, P. Sturm, and S. Lodha. Towards generic self-calibration of central cameras. In *Proc. OMNIVIS*, 2005.
- [16] S. Roth and M. Black. Specular flow and the recovery of surface structure. In *CVPR*, 2006.
- [17] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. *Int. J. Comput. Vision*, 64(1):31–67, 2005.
- [18] Y. Vasilyev, Y. Adato, T. Zickler, and O. Ben-Shahar. Dense specular shape from multiple specular flows. *CVPR*, 2008.
- [19] S. Waldon and C. Dyer. Dynamic shading, motion parallax and qualitative shape. In *Proc. IEEE Workshop on Qualitative Vision.*, pages 61–70, 1993.
- [20] A. Zisserman, P. Giblin, and A. Blake. The information available to a moving observer from specularities. *Image and Vision Computing*, 7:38–42, 1989.