Shape from Specular Flow: Is One Flow Enough?

Yuriy Vasilyev¹ Todd Zickler¹ ¹ Harvard School of Engineering and Applied Sciences Cambridge, USA

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

Abstract

Specular flow is the motion field induced on the image plane by the movement of points reflected by a curved, mirror-like surface. This flow provides information about surface shape, and when the camera and surface move as a fixed pair, shape can be recovered by solving linear differential equations along integral curves of flow. Previous analysis has shown that two distinct motions (i.e., two flow fields) are generally sufficient to guarantee a unique solution without externally-provided initial conditions. In this work, we show that we can often succeed with only one flow. The key idea is to exploit the fact that smooth surfaces induce integrability constraints on the surface normal field. We show that this induces a new differential equation that facilitates the propagation of shape information between integral curves of flow, and that combining this equation with known methods often permits the recovery of unique shape from a single specular flow given only a single seed point.

1. Introduction

A curved, mirror-like surface presents the observer with a distortion of the surrounding environment. Such distortions provide information about the shape of the mirror. Reconstruction from this cue is difficult, however, because the problem is ill-posed when the environment is unknown. One way to simplify the reconstruction task is to exploit motion, and there is a significant body of work devoted to analyzing the motion of specular reflections and their use in surface reconstruction. Much of this work focuses on the motion of isolated specular highlights, such as those induced by compact light sources, either differentially in two nearby images [4, 20] or as tracked over an extended image sequence [10, 12, 19]. In these cases, one can recover surface information at a point, or as demonstrated by Oren and Nayar [12], along a surface curve.

More related to our work are analyses of the dense image motion induced by specular reflections in complex lighting environments. In these cases, relative motion between the observer, mirror object, and/or environment induces a 2D Steven Gortler¹ Ohad Ben-Shahar²

² Computer Science Department, Ben-Gurion University Beer Sheva, Israel ben-shahar@cs.bgu.ac.il

motion field on the image plane. This motion field—termed specular flow [14]—provides direct information about the mirror surface as a whole [1, 2, 6, 14, 15, 16, 18].

Of particular interest is the analysis that expresses the relationship between specular flow and surface shape as a partial differential equation that can be solved for the surface shape [1, 2, 6, 16]. Such an equation has been formulated in the case that the mirror surface is observed orthographically, the object and camera move as a fixed pair under a distant lighting environment (Fig. 1, left), and occlusions and interreflections are absent. In this case, since the camera's view direction $\hat{\mathbf{v}}$ is fixed relative to the object, one can represent its visible shape by its field of *reflection vectors* $\hat{\mathbf{r}}(\mathbf{x}) \in \mathbb{S}^2$, which is related to the surface normal field, $\hat{\mathbf{n}}(\mathbf{x})$, by the law of specular reflection: $\hat{\mathbf{n}}(\mathbf{x}) \propto \hat{\mathbf{v}} + \hat{\mathbf{r}}(\mathbf{x})$. The surface shape (represented by $\hat{\mathbf{r}}(\mathbf{x})$) must then satisfy [6]

$$\hat{\mathbf{r}}_{\mathbf{u}} = \frac{1}{\|\mathbf{u}\|} \left[\boldsymbol{\omega}\right]_{\times} \hat{\mathbf{r}},\tag{1}$$

where $\boldsymbol{\omega} \in \mathbb{R}^3$ is the angular velocity vector of the environment motion, $[\boldsymbol{\omega}]_{\times}$ its skew matrix, and $\mathbf{u}(\mathbf{x}) \in \mathbb{R}^2$ is the observed specular flow field. In previous work this equation has been called the *shape from specular flow equation* (SFSF). For consistency we use the same terminology here.

This PDE formulation allows characterizing the conditions for unique and complete reconstruction. For example, given knowledge of the surface normal at a single point (represented by its reflection vector), one can integrate the SFSF equation along an integral curve of the observed specular flow (henceforth termed a *flow curve*) to obtain the surface normal at every point along that curve. The entire surface can be recovered, then, as long as one has knowledge of the surface normal at one or more points along each flow curve (see blue arrows, right of Fig. 1) [1, 2, 6]. Alternatively, it has been shown that a second specular flow field can be used to propagate the normal field between the flow curves of the first flow. Solving the SFSF equation along the flow curves of the two flows together, then, uniquely determines the surface [6].

The existing analysis is incomplete, however, because it



Figure 1. Left: A mirror-like object is observed by an orthographic camera as they move as a fixed pair under a distant illumination environment. This induces a motion field on the image plane (specular flow; middle) that encodes information about the shape. Previous analysis describes how to integrate shape information *along* integral curves of the flow field (blue arrow; right). In this paper, we derive an equation that allows integration *across* the curves as well. This allows recovering an entire surface from a single seed point (red arrow; right). This approach exploits the curvature of the flow field, which is depicted in the top-right as a moving frame.

effectively treats each motion field as a collection of independent 1D flow curves-not too different from the 1D trajectories of Oren and Nayar [12]-and ignores higher-order properties, such as curvature, of the observed motion field. One might guess that this higher order structure embeds additional shape information, and in this paper we show this intuition to be correct. Beginning with the condition that the underlying normal field corresponds to a continuous surface and must therefore be integrable [9], we derive a second PDE that complements the SFSF equation. This new equation naturally incorporates the measured curvature of the specular flow field, and allows integrating surface shape information orthogonal to the flow curves. By combining it with the SFSF equation, we can recover the entire surface from a single initial point and a single observed specular flow, as depicted by the red integration path in Fig. 1.

After deriving this new equation in section 3, in section 4 we describe a surface property, which we call reconstructibility, and show that the entire reflection field, and therefore the shape, of surfaces that have this property can be recovered from the knowledge of the reflection vector at a single seed point. Then, in section 4.3, we outline a proof, based on the notion of transversality, of the fact that reconstructibility is a generic property of surfaces. We note that presently our proof cannot be easily completed, as it requires finding the dimension of a variety that is described by polynomials of which one, by our estimate, has well over a trillion terms. This makes finding the dimension of that variety analytically prohibitive. Instead we have performed numerical tests that are encouraging in that they indicate that the proof may be successfully completed. We conclude by confirming the correctness of our mathematical analysis with proof-of-concept results on synthetic data.

2. Construction

Following Canas et al. [6], we assume the input to be a specular flow over a closed connected region U of the image plane. We assume that U is the orthographic projection of a

closed smooth manifold surface M in the view direction $\hat{\mathbf{v}}$ as shown in Fig. 1, and use \dot{U} and ∂U to represent the interior and boundary of U, respectively. We represent M as a height field $f(\mathbf{x})$ over \dot{U} and further assume the following:

Assumption 1.

- a) M does not self-reflect
- b) M contains neither developable regions nor flat points
- c) The third derivatives of the height field are continuous

Note that assumption c) is a necessary and sufficient condition for the equality of the mixed partials of the reflection field ($\hat{\mathbf{r}}_{xy} = \hat{\mathbf{r}}_{yx}$) and is exactly the same as the assumption made in [6] to obtain a reconstruction in the case of unknown environment rotations.

The *flow curves* were defined in [6] as level sets $\{\mathbf{x} : \boldsymbol{\omega}^{\mathsf{T}} \hat{\mathbf{r}}(\mathbf{x}) = C\}, C \in [-\|\boldsymbol{\omega}\|, \|\boldsymbol{\omega}\|]$. Furthermore, the authors showed that under the above assumptions, almost all flow curves are smooth nonintersecting curves. The only exceptions are those flow curves that intersect the boundary ∂U at isolated points. Additionally they showed that the flow curves cover all of \dot{U} except for a finite number of points. Such *exceptional points* are points \mathbf{x} of U where the Jacobian $\frac{\partial \hat{\mathbf{r}}(\mathbf{x})}{\partial \mathbf{x}}$ satisfies $\boldsymbol{\omega}^{\mathsf{T}} \frac{\partial \hat{\mathbf{r}}(\mathbf{x})}{\partial \mathbf{x}} = 0$; that is, points where the one dimensional column span of $\frac{\partial \hat{\mathbf{r}}(\mathbf{x})}{\partial \mathbf{x}}$ happens to be orthogonal to $\boldsymbol{\omega}$. Since these exceptional points are isolated, we ignore them in the rest of this paper.

There is another set of curves related to the flow curves the integral curves of the gradient field of the scalar function $\omega^{\mathsf{T}} \hat{\mathbf{r}}(\mathbf{x})$ —that also cover \dot{U} . This gradient field is everywhere orthogonal to the specular flow, and is not well defined only at the same exceptional points as described above for the flow curves. We call the integral curves of this gradient field *ortho curves* (for orthogonal curves).

We have already seen that the SFSF equation (Eq. 1) specifies the change in the reflection field along flow curves, $\hat{\mathbf{r}}_{\mathbf{u}}$. We would like to derive a complementary differential equation that specifies the change in the reflection field

along the ortho curves. Together, the two equations would specify the entire reflection field Jacobian. The knowledge of the full Jacobian would allow us to propagate the reflection field information in any direction, not solely along flow curves.

The derivation of the equation that specifies the change in the reflection field along ortho curves, which follows in the next section, is simpler to present and easier to interpret if we perform it at a specific point \mathbf{x}_0 and in the coordinate system that is aligned with the specular flow at \mathbf{x}_0 . Specifically, we take \mathbf{x}_0 to be the origin, $\mathbf{u}(\mathbf{x}_0)$ to be in the direction of the positive x-axis, $\hat{\mathbf{v}}$ to be in the direction of the positive z-axis, and the positive y-axis to be that direction that makes this a right handed coordinate system. Note that in this coordinate system the x and y axes are tangents to the specular flow and ortho curves at \mathbf{x}_0 , respectively.

The reflection field Jacobian can be represented as a 3×2 matrix $\mathbf{J} = \begin{bmatrix} \hat{\mathbf{r}}_x & \hat{\mathbf{r}}_y \end{bmatrix}$. Because at the point \mathbf{x}_0 the x and y directions are aligned with the flow and ortho curves, $\hat{\mathbf{r}}_x(\mathbf{x}_0)$ and $\hat{\mathbf{r}}_{u}(\mathbf{x}_{0})$ are the changes of the unit length reflection field along flow and ortho curves respectively. For this reason we call $\hat{\mathbf{r}}_x(\mathbf{x}_0)$ the *tangent derivative* and $\hat{\mathbf{r}}_y(\mathbf{x}_0)$ the *ortho derivative*. In this coordinate system $\hat{\mathbf{r}}_x(\mathbf{x}_0) = \hat{\mathbf{r}}_u(\mathbf{x}_0)$ and is specified by the SFSF equation (Eq. 1) in terms of the rotation axis, the reflection field, and the flow at x_0 . In the next section we show that by imposing surface integrability we can also obtain an equation for the ortho derivative in terms of known and measured quantities. Note that, to avoid clutter, we will omit x_0 from the expressions in our derivation. It is important to keep in mind, however, that the derivation is performed at \mathbf{x}_0 , and, since the x and y directions are aligned with the flow and ortho directions only at \mathbf{x}_0 , is only valid there.

3. Integrability and Reflection Field Jacobian

One constraint on the ortho derivative $\hat{\mathbf{r}}_y$ can be obtained from the fact that the reflection field is unit length, and therefore its derivative in the direction of the ortho curve (and in fact any direction) must satisfy:

$$\hat{\mathbf{r}}^{\mathsf{T}}\hat{\mathbf{r}}_{y} = 0. \tag{2}$$

When the reflection vector $\hat{\mathbf{r}}$ is known, this is a single linear constraint on the three components of the ortho derivative.

Another constraint comes from our assumption of surface integrability. The integrability constraint is a constraint on the surface height field, $f(\mathbf{x})$. On an integrable surface, the total change in height along any closed path is zero, which is equivalent to requiring that at every point: $f_{xy} = f_{yx}$. This last condition holds so long as the height field has continuous second derivatives. In [6] it was shown that this constraint can be expressed at every image point in terms of the components of the surface reflection field $\hat{\mathbf{r}} = (r_1, r_2, r_3)$:

$$\left(\frac{r_1}{r_3+1}\right)_y = \left(\frac{r_2}{r_3+1}\right)_x.$$
(3)

If the reflection vector $\hat{\mathbf{r}}$ is known, this expression reduces to a linear constraint on the directional derivatives $\hat{\mathbf{r}}_x = (r_{1x}, r_{2x}, r_{3x})$ and $\hat{\mathbf{r}}_y = (r_{1y}, r_{2y}, r_{3y})$. Namely,

$$-(r_3+1)r_{1y}+r_1r_{3y}=r_2r_{3x}-(r_3+1)r_{2x}.$$
 (4)

In that case we can also use the SFSF equation (Eq. 1) to determine the tangent derivative $\hat{\mathbf{r}}_x$ at \mathbf{x}_0 , and Eq. 4 is a linear constraint on the ortho derivative $\hat{\mathbf{r}}_y$ at \mathbf{x}_0 .

To summarize, if the scene rotation, the flow, and the reflection vector are known, then Eq. 1 specifies the tangent derivative $\hat{\mathbf{r}}_x$ and the unit length and integrability constraints in Eqs. 2 and 4 are two linear constraints on the three components of the ortho derivative $\hat{\mathbf{r}}_y$. To obtain the full Jacobian we need another linear constraint on the ortho derivative. In the next section we show how such a constraint can be obtained from a linear combination of the derivatives of the specular flow and integrability constraints.

3.1. Linear system for the ortho derivative

To further constrain the ortho derivative $\hat{\mathbf{r}}_y$, we observe that both the SFSF equation (Eq. 1) and the integrability constraint must hold at every point of \dot{U} . Therefore, the derivatives of both sides of each of these constraints must be equal. These derivatives specify how the constraints change between points, and depend on both the curvature κ of the flow curve and the curvature κ_{\perp} of the ortho curve at \mathbf{x}_0 . These curvatures are new measurements that we are taking advantage of to gain new information about the surface. They are exactly the "tangential curvature" and "normal curvature" that have been investigated in the context of texture flows in [3]. However, these higher order properties of the specular flow field have not been utilized by previous analyses of shape from specular flow.

By differentiating the SFSF equation (Eq. 1) in the y direction, the direction perpendicular to the flow, we get a sense of how this constraint changes as we move away from the flow curve through x_0 . This derivative must satisfy:

$$\nabla_{y}\left(\hat{\mathbf{r}}_{\mathbf{u}}\right) = \nabla_{y}\left(\mathbf{J}\mathbf{u}\right) = \nabla_{y}\left(\frac{1}{\|\mathbf{u}\|}[\boldsymbol{\omega}]\times\hat{\mathbf{r}}\right).$$
(5)

It can be shown that by expanding the derivative with product and chain rules and by using Assumption 1 to equate the mixed partial derivatives of the reflection vector field, we obtain a differential equation which specifies the change in $\hat{\mathbf{r}}_y$ along the *x* direction in terms of $\hat{\mathbf{r}}, \hat{\mathbf{r}}_x$, the flow, \mathbf{u} , and the rotation $\boldsymbol{\omega}$ (see [17]). That is, we obtain a differential equation that defines the change in $\hat{\mathbf{r}}_y$ in the flow direction at \mathbf{x}_0 in terms of the known vector $\mathbf{g} \triangleq \left(\frac{1}{\|\mathbf{u}\|}\right)_u [\boldsymbol{\omega}]_{\times} \hat{\mathbf{r}}$:

$$\hat{\mathbf{r}}_{xy} = \frac{1}{\|\mathbf{u}\|} [\boldsymbol{\omega}]_{\times} \hat{\mathbf{r}}_y - \kappa_{\perp} \hat{\mathbf{r}}_y + \mathbf{g}.$$
 (6)

The integrability constraint (Eq. 4) also holds at every point. Its derivative in the x direction can be written as (see [17]):

$$\begin{pmatrix} -(r_3+1) \\ 0 \\ r_1 \end{pmatrix}^{\mathsf{T}} \hat{\mathbf{r}}_{xy} + \begin{pmatrix} -r_{3x} \\ -(r_3+1)\kappa \\ r_{1x}+r_2\kappa \end{pmatrix}^{\mathsf{T}} \hat{\mathbf{r}}_y = b(\hat{\mathbf{r}}, \hat{\mathbf{r}}_x),$$
(7)

where the scalar function b can be expressed at x_0 as:

$$b \triangleq r_2 \left(\left(\frac{1}{\|\mathbf{u}\|} \right)_x (\omega_1 r_2 - \omega_2 r_1) - \frac{1}{\|\mathbf{u}\|} (\omega_1 r_{2x} - \omega_2 r_{1x}) \right) + \\ - (r_3 + 1) \left(\left(\frac{1}{\|\mathbf{u}\|} \right)_x (\omega_3 r_1 - \omega_1 r_3) + \frac{1}{\|\mathbf{u}\|} (\omega_3 r_{1x} - \omega_1 r_{3x}) \right)$$

We can now obtain a third linear constraint on the ortho derivative $\hat{\mathbf{r}}_y$ by eliminating the $\hat{\mathbf{r}}_{xy}$ terms in Eq. 7 using Eq. 6. Together with the unit length and integrability constraints in Eqs. 2 and 4, this gives a linear system for the ortho derivative $\hat{\mathbf{r}}_y$ at the point \mathbf{x}_0 :

$$\begin{bmatrix} r_1 & r_2 & r_3 \\ -(r_3+1) & 0 & r_1 \\ b_1 & b_2 & b_3 \end{bmatrix} \hat{\mathbf{r}}_y = \begin{pmatrix} 0 \\ c \\ b+(r_3+1)g_1 - r_1g_3 \end{pmatrix},$$
(8)

where the third row is the constraint we just derived:

$$b_1 \triangleq -r_{3x} + (r_3 + 1)\kappa_{\perp} - r_1 \frac{\omega_2}{\|\mathbf{u}\|}$$

$$b_2 \triangleq (r_3 + 1) \left(\frac{\omega_3}{\|\mathbf{u}\|} - \kappa\right) + r_1 \frac{\omega_1}{\|\mathbf{u}\|}$$

$$b_3 \triangleq r_{1x} + \kappa r_2 - (r_3 + 1) \frac{\omega_2}{\|\mathbf{u}\|} - r_1 \kappa_{\perp}$$

and g_1, g_3 are the components of the vector g.

In this section we have shown that the tangent and ortho derivatives of the reflection field can be obtained in terms of the scene rotation, the reflection field, and the flow at the point \mathbf{x}_0 , but, of course, the equations can be derived at any point \mathbf{x} , and the full Jacobian can be obtained whenever the linear system for the ortho derivative can be inverted.

3.2. Degenerate Points

It can be shown (see [17]) that the determinant of the linear system in Eq. 8 is zero at those image points x at which

$$\kappa \|\mathbf{u}\| = \frac{\boldsymbol{\omega}^{\mathsf{T}} \hat{\mathbf{r}} + \omega_3}{(r_3 + 1)}.$$
(9)

Note that each quantity in this equation is independent of the coordinates chosen for the image plane. This fact makes this equation particularly useful for identifying and reasoning about the points at which the linear system cannot be inverted. Without considering the exceptional points of section 2, we will refer to the points in U at which the linear system in Eq. 8 cannot be formed or is singular as *degenerate points*. Note that this includes the occluding boundary points and those points in \dot{U} where Eq. 9 holds.

In the next section we describe a surface reconstruction procedure that involves solving the linear system in Eq. 8 to

reconstruct the reflection vector field along directions perpendicular to the flow curves. Since this system cannot be solved at degenerate points, the set of degenerate points is of great interest to us. Experimentally, we will show that for almost all rotations and surfaces degenerate points lie on curves in the image plane that are distinct from flow curves, and that in this case the entire surface can be reconstructed from a single seed point. This observation is also confirmed mathematically for the specific case of the sphere (see figure 2 for examples and [17] for the proof).

4. Surface Reconstruction

In the previous section we showed that it is possible to obtain the full reflection field Jacobian at every nondegenerate point. Therefore, if we let $\alpha(t)$ be any connected, smooth segment of an ortho curve that has no degenerate points and is arc-length parameterized by t, then using Eq. 8 we can write a continuous ODE in t:

$$K(\hat{\mathbf{r}}(t),\boldsymbol{\omega},\mathbf{u}(t),\kappa(t),\kappa_{\perp}(t))\,\hat{\mathbf{r}}'(t) = h\left(\hat{\mathbf{r}}(t),\boldsymbol{\omega},\mathbf{u}(t)\right).$$
 (10)

Thus, a single initial point on such a curve segment constrains the reflection field everywhere on this segment. From [6] we know that the same is true for flow curves. Therefore, if we consider any open subset W of U in which there are no degenerate points, then starting from any $\mathbf{x} \in W$ at which the reflection vector is known, we can, by solving Eq. 1 along specular flow curves and Eq. 10 along ortho curves, obtain the reflection vector field at all points of W. This is true because both sets of curves cover all of W, and thus there exists a path from x to every other point of W that is composed of segments of flow and ortho curves along which the reflection field can be propagated from x. In fact, even if there are isolated degenerate points in W we can still propagate the reflection field information to every point from x, since there is always some path from x to each other point of W that passes through the degenerate points along flow curves and not along ortho curves.

In general, however, we find that degenerate points are not isolated but lie on curves, which we call *degenerate curves*, and for certain surfaces and rotation axes may even cover all of U. In the latter case the differential equation in Eq. 10 is not defined anywhere, and therefore the reflection vector cannot be propagated orthogonally to any flow curve.

4.1. Reconstruction and Degenerate Curves

We now look at the most common scenario: the case in which degenerate points lie on curves, which we call *degenerate curves*. To simplify this discussion we define the following property of surfaces:

Definition 1. For a given scene rotation ω , a surface M is said to be *reconstructible* or to possess the *reconstructibility* property if none of the flow curves is a subset of the degenerate point set.

Note that this is indeed a property of the surface for a given rotation, since both the flow and the set of degenerate points depend only on the height field and the scene rotation.

We claim that if the surface is reconstructible, then its reflection field (and thus surface shape) can be reconstructed everywhere from a single seed point.

Lemma 1. If the surface M satisfies Assumption 1 and is reconstructible, then its reflection field can be reconstructed over \dot{U} from a single point where it is known.

Proof. Let \mathbf{x}_1 , \mathbf{x}_2 be two points in \dot{U} . If M satisfies Assumption 1, then the reflection vector field can be propagated along all segments of flow curves using Eq. 1 and along those segments of ortho curves that do not contain a degenerate point using Eq. 10. Thus, if there is a path between \mathbf{x}_1 and \mathbf{x}_2 composed of flow and ortho curve segments such that degenerate points are contained on flow curve segments only, then the reflection field can be propagated from \mathbf{x}_1 to \mathbf{x}_2 (see for example the black curve segment in Fig. 3(a)).

Suppose that every possible path in U from x_1 to x_2 along ortho and flow curves passes through a degenerate point. Thus, there exists a connected open subset W that contains x_1 , does not contain any degenerate points, and whose boundary ∂W is composed entirely of degenerate points and is a set of closed curves in U. For simplicity, assume that there are paths from x_1 to x_2 that cross only a single closed degenerate curve which is part of ∂W . We can repeat the argument multiple times for pairs of points that are separated by multiple degenerate curves.

Now the only way that there is no path between x_1 and x_2 along flow and ortho curves such that all degenerate points lie on segments of flow curves only is if it is impossible to cross ∂W along a flow curve. However, because ∂W is a set of closed curves and the flow curves are closed curves that do not intersect in \dot{U} and cover all of \dot{U} (see [6]), this is possible only if some flow curve is fully contained in ∂W . We call such a flow curve, one on which every point is a degenerate point, a *degenerate flow curve*. Note that if x_1 is itself degenerate, then we can replace it with any non-degenerate point on the flow curve through x_1 . Such a point must exist since M has the reconstructibility property.

Since \mathbf{x}_1 and \mathbf{x}_2 were arbitrary, this implies that if M satisfies Assumption 1 and has the reconstructibility property, then there is a path between any two points of \dot{U} such that all degenerate points lie on segments of flow curves only. This means that the reflection field of M can be propagated to all points of \dot{U} from a single point where it is known. \Box

In the next few sections we provide evidence that for most scene rotation and surface pairs, degenerate points lie on curves in the image plane and that flow curves are generally not contained in the set of degenerate points.

4.2. Degenerate Curve Examples

In the previous section we saw that reconstruction is possible in the case that no flow curve is composed entirely of degenerate points. In this section we summarize our experiments indicating that for most combinations of rotations and surfaces such degenerate flow curves do not exist.

In Figs. 2 and 3 we show the specular flow curves (in blue) and the degenerate curves (in red) for two surfaces, the sphere and a more complex shape, observed under a variety of scene rotations. In (a) the scene rotation axis is parallel to the view direction. In (f) it is about a horizontal axis, and in (b)-(e) it is at some position in between.

For the sphere, in the case of the view axis rotation, Fig. 2 (a), every point is degenerate and no reconstruction is possible. (In fact one can prove that in the case of view axis rotation this is true of any surface of revolution about the view axis.) A degenerate flow curve also exists when the scene rotation is about an axis on the equator, (f). In this case the specular flow curve through the center of the sphere is degenerate. Since it divides the sphere in half, a full reconstruction can only obtained if at least one reflection vector is known in each half. For no other scene rotations, examples shown in (b)-(e), is there a flow curve composed entirely of degenerate points. We show in [17] that the sphere is indeed reconstructible for all rotations except the ones whose axis lies on the equator or is parallel to the view direction.

Figure 3 (a)-(f) shows that for most rotations the situation is similar for a more complex surface-one with elliptic, parabolic, and hyperbolic points. In the figure the parabolic points are shown in green. At these points the reflection field Jacobian cannot be computed, so they are degenerate. Note, however, that the flow is generally transverse to the parabolic curves. It was shown in [6] that parabolic points at which the flow is parallel to the parabolic curve are isolated. Thus, using Eq. 1 one can always propagate the reflection field across a parabolic curve. Note that unlike the case of the sphere, the surface in Fig. 3 (a) has no degenerate flow curves in the case of view axis rotation. Empirically, we find this to be true for surfaces that are not surfaces of revolution about the view axis. For other rotation axes, the results are the same as for the sphere-there is a degenerate flow curve when the scene rotation axis is on the equator, but not in any other cases we have seen.

4.3. Reconstruction in the General Case

Our experiments indicate that the cases shown in Figs. 2 (a), (f) and 3 (f) are uncommon. For most surfaces and rotation axes, degenerate points lie on curves and there are no degenerate flow curves. In this section we attempt to provide some mathematical evidence for these observations.

Rather than work directly with the reconstructibility property we defined earlier, we find it easier to investigate a "stronger" property of surfaces.



Figure 2. (a)-(f) Flow curves (in blue) and degenerate curves (in red) for the unit sphere for scene rotations ranging from parallel to view axis in (a) to perpendicular to the view axis in (f). The rotation axes are in the xz plane inclined at angles [0 30 50 70 85 90] degrees measured from the view axis. (a) In the case of view axis rotation every point on the sphere is degenerate. (f) If the rotation is perpendicular to the view axis, the flow curve through the center of the sphere is degenerate. (b)-(e) In all other cases there are flow curves that cross the degenerate curves.

Definition 2. For a given scene rotation ω , a surface M is said to be *strongly reconstructible* or to possess the *strong reconstructibility property* if none of the flow curves are simultaneously tangent to and have the same curvature as a degenerate curve.

Note that this definition implies that if degenerate points form a region of the image plane rather than a curve, then the strong reconstructibility property does not hold.

The strong reconstructibility property is stronger in the sense that if it holds for a given surface, it implies that no flow curve can be a subset of the degenerate point set, and hence that reconstructibility holds as well. We would now like to provide evidence that for a given scene rotation strong reconstructibility is a generic property of surfaces. Because we know by Lemma 1 that reconstructibility implies that the entire surface reflection field can be reconstructed from a single point, this will also give evidence for the fact that a single specular flow is sufficient for reconstruction in the general case. Due to space constraints we limit our discussion, which we base on the analysis of generic properties of curves in [5], to a summary of the key concepts involved in showing that strong reconstructibility is a generic property. See [17] for more details.



Figure 3. (a)-(f) Flow curves (in blue), parabolic curves (in green), and degenerate curves (in red) for a more complex surface. The rotation axes are along the horizon, inclined at angles [0 30 50 70 85 90] degrees measured from the view axis. (a) Even in the case of view axis rotation degenerate points lie on curves. A path connecting two points along flow and ortho curves that passes through degenerate points along flow curves only is shown in black. (f) If the rotation is perpendicular to the view axis, the flow curve through the center is degenerate. (b)-(e) In all other cases there are flow curves that cross the degenerate curves.

Following [5], we say a property is *generic* if it holds for an open and dense set of surfaces. Intuitively, denseness means that for any given surface, M, and rotation, ω , there is an arbitrarily close surface which is strongly reconstructible, and openness means that strong reconstructibility holds in an open neighborhood of surfaces about a strongly reconstructible surface.

A useful tool for working with arbitrary surfaces is the Monge-Taylor map (see for example [5]). This map associates to each domain point \mathbf{x} the coefficients of the k-order Taylor approximation of the height field. Thus an arbitrary smooth surface is represented by $\{(a_1(\mathbf{x}), ..., a_n(\mathbf{x})) : \mathbf{x} \in U\}$, an immersion of the two dimensional surface into \mathbb{R}^n (here the $a_i(\mathbf{x})$ are the Taylor coefficients). Our goal is to show that the Monge-Taylor representation of an open and dense set of smooth surfaces does not intersect the subset of \mathbb{R}^n in which strong reconstructibility does not hold.

To begin we express in terms of the a_i the conditions which must be satisfied in order for strong reconstructibility not to hold. Specifically we have three *bad conditions* that a point $(a_1, ..., a_n)$ must satisfy. These conditions specify that: 1) the point lies on a degenerate region or curve; 2) the tangent space to this region or curve contains the tangent space to the flow curve through this point; and 3) the curvatures of the two curves are equal (here we must ensure that the expression we use is satisfied in the case of a degenerate region). These conditions can be written in a way that defines a variety in \mathbb{R}^n , and if the dimension of this variety is less than n - 2, then transversality [5] guarantees that this variety does not intersect an open and dense set of smooth surfaces, and therefore strong reconstructibility is a generic property (see [17] for more details).

Since it is possible to express the bad conditions as polynomials, one may be able to analytically find the dimension of this variety by computing its Gröbner basis. Unfortunately, we have found that one of these polynomials (the one corresponding to the equality of curvatures condition) has well over a trillion terms, which makes this computation intractable with current means.

Another way to investigate the dimension of this variety is to determine the rank of the Jacobian of the functions that express the bad conditions at test points on each irreducible subvariety of this variety. The dimension of each such subvariety is less than or equal to n minus the greatest rank of the Jacobian at the test points on that subvariety. The greatest dimension of these irreducible subvarieties is then an upper bound on the dimension of the variety.

To perform this investigation, we express Eq. 9 in terms of the height field and its derivatives and specify the variety of bad conditions as the intersection in \mathbb{R}^n of zero level sets of three functions - one for each condition. Note that because Eq. 9 can be expressed in terms of the height field and its derivatives to third order and we only need to represent the tangent and curvature of this condition, it suffices to choose k = 5, and hence n = 20. For a given rotation ω we find a Monge patch whose center point satisfies each of the three conditions and compute the rank of the 3×20 Jacobian matrix at that point. Fig. 4 shows sample fifth order Monge patches with the image centers satisfying the bad conditions. The flow curves are in blue, the flow curve through the center in yellow, and the degenerate points in red. As was the case in all our experiments, the rank of the Jacobian at each test point is three.

Now because the variety of the bad conditions may be reducible, we cannot say for certain that its dimension is less than or equal to seventeen. Since we have not proved this variety to be irreducible and clearly cannot run an exhaustive search, we might not have tested any points on a subvariety of dimension greater than seventeen. However,



Figure 4. Example monge patches whose centers satisfy the bad conditions (image center). The flow curve through the image center (in yellow, others shown in blue) and degenerate curve (in red) meet with parallel tangents and the same curvature. While for any given ω (here $\omega = \frac{(.8,0,1)}{\|(.8,0,1)\|}$) surfaces with such points seem to exist, we believe that for general surfaces this is not the case.

we find it encouraging that a random search in \mathbb{R}^n never yielded a point on such a higher dimensional subvariety.

5. Synthetic Reconstruction

To verify the correctness of our derivation in Section 3, we implemented a simple version of a propagative reconstruction described in 4. Given the rotation ω , the reflection vector at a single point, and the flow over a discretized image plane, we propagate the reflection field information from the seed point in perpendicular directions until the reflection field is obtained everywhere. A sample reconstruction of an spheroid with $\omega = \frac{(.2,0,1)}{\|(.2,0,1)\|}$ is shown in Fig. 5. Note that a reconstruction of this surface of revolution about the view axis is possible for this scene rotation. In (a) and (b) we show the true and reconstructed surface reflection fields, respectively, with the reflection vector components coded as RGB. Fig. 5(c) shows the absolute angular error of our reconstruction in degrees. In (d) we show a wire mesh of the ground truth shape (in blue) overlaid with the result of integrating the surface derivatives that were obtained from the reconstructed reflection field in (b) (shown in red). The true and the recovered surfaces are almost identical.

6. Conclusion

The analysis in this paper suggests that an observation of a single specular flow field, along with knowledge of the motion axis and the surface normal at a single interior point, is sufficient to recover a complete surface in the general case. The argument proceeds by expressing the integrability condition on the underlying normals as a partial differential equation, combining it with the known flow derivative equation, and analyzing conditions for a unique solution.

Our analysis suggests a few theoretical directions worth considering. First, at the core of our argument is the exploitation of surface smoothness through surface integrability. Since integrability is a very general notion of smoothness, we are providing a conservative description of what is achievable. For example, if one were to use a spline-based surface representation, one would expect stronger reconstruction guarantees. A second consideration is the poten-



Figure 5. (a): the "ground truth" surface reflection field, with xyz-components color coded as RGB. (b): the estimated surface reflection field. (c): absolute angular surface reflection field error in degrees. (d): wire meshes of ground truth (in blue) and surface reconstructed from the estimated reflection field (in red).

tial utility of the occluding contour, along which the surface normals are known. While almost all flow curves do not intersect the contour [6], it is possible that appropriate use of surface smoothness will allow one to propagate surface constraints from the contour, and allow unique reconstruction *without* the knowledge of an interior surface normal or the motion axis. Additionally, because of the close similarity between the specular shape from flow we have considered here and the problem of flow-based calibration of generic central projection cameras, see for example [7, 11, 13], the theory we have presented in this paper can likely be applied to show that in many cases a single flow is sufficient for calibration of such cameras.

The theory presented in this paper also has practical implications. It has recently been shown that a spline-based surface representation allows one to recover surface shape from a single pair of images in the special case of a viewaxis rotation [15]. Our analysis suggests that the same should be true for *any* rotation, and that practical systems for producing specular surfaces directly from image sequences are within reach. Moreover, because of the close connection between the shape from specular flow equation and the relation between image and environment intensity gradients [1], our analysis provides hope for systems that seek to match the human ability to infer specular surface shape from a single still image [8].

Acknowledgements

This project is funded by the US National Science Foundation under grant IIS-0712956. Additional support for O.B.S was provided by the Israel Science Foundation under grant No. 1245/08, the Frankel Fund, and the Paul Ivanier Robotics Center at Ben-Gurion University.

References

- Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Toward a theory of shape from specular flow. *ICCV*, 2007.
- [2] Y. Adato, Y. Vasilyev, T. Zickler, and O. Ben-Shahar. Shape from specular flow. *PAMI*, 32(11), 2010.
- [3] O. Ben-Shahar and S. Zucker. The perceptual organization of pexture flow: a contextual inference approach. *PAMI*, 25(4), 2003.
- [4] A. Blake and H. Bülthoff. Shape from specularities: Computation and psychophysics. *Phil. Trans. R. Soc. B*, 331, 1991.
- [5] J. W. Bruce and P. J. Giblin. *Curves and Singularities. Second Edition.* Cambridge University Press, 1992.
- [6] G. Canas, Y. Vasilyev, Y. Adato, T. Zickler, S. Gortler, and O. Ben-Shahar. A linear formulation of shape from specular flow. *ICCV*, 2009.
- [7] F. Espuny. A closed-form solution for the generic selfcalibration of central cameras from two rotational flows. *Proc. VISAPP*, 1:26–31, 2007.
- [8] R. Fleming, A. Torralba, and E. Adelson. Specular reflections and the perception of shape. J Vis, 4(9), 2004.
- [9] B. Horn and M. Brooks. The variational approach to shape from shading. *CVGIP*, 33(2), 1986.
- [10] 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.
- [11] D. Nistér, H. Stewénius, and E. Grossmann. Non-parametric self-calibration. *ICCV*, 2005.
- [12] M. Oren and S. Nayar. A theory of specular surface geometry. *IJCV*, 24(2):105–124, 1997.
- [13] S. Ramalingam, P. Sturm, and S. Lodha. Towards generic self-calibration of central cameras. In *OMNIVIS*, 2005.
- [14] S. Roth and M. Black. Specular flow and the recovery of surface structure. In *Proc. CVPR*, 2006.
- [15] A. Sankarnarayanan, A. Veeraraghavan, O. Tuzel, and A. Agrawal. Specular surface reconstruction from sparse reflection correspondences. In *Proc. CVPR*, 2010.
- [16] Y. Vasilyev, Y. Adato, T. Zicker, and O. Ben-Shahar. Dense specular shape from multiple specular flows. *CVPR*, 2008.
- [17] Y. Vasilyev, T. Zickler, S. Gortler, and O. Ben-Shahar. Toward shape from a single specular flow. *Tech. Rep. TR-04-11*, *Harvard School of Engineering and Applied Sciences*, 2011. ftp://ftp.deas.harvard.edu/techreports/tr-04-11.pdf.
- [18] S. Waldon and C. Dyer. Dynamic shading, motion parallax and qualitative shape. In *Proc. IEEE Workshop on Qualitative Vision.*, pages 61–70, 1993.
- [19] J. Zheng and A. Murata. Acquiring a complete 3d model from specular motion under the illumination of circularshaped light sources. *PAMI*, 22(8), 2000.
- [20] A. Zisserman, P. Giblin, and A. Blake. The information available to a moving observer from specularities. *Image* and Vision Computing, 7(1):38–42, 1989.