

Specular Surface Reconstruction from Sparse Reflection Correspondences

Aswin Sankaranarayanan, Ashok Veeraraghavan, Oncel Tuzel, Amit Agrawal

TR2010-044 July 2010

Abstract

We present a practical approach for surface reconstruction of smooth mirror-like objects using sparse reflection correspondences (RCs). Assuming finite object motion with a fixed camera and un-calibrated environment, we derive the relationship between RC and the surface shape. We show that by locally modeling the surface as a quadric, the relationship between the RCs and unknown surface parameters becomes linear. We develop a simple surface reconstruction algorithm that amounts to solving either an eigen-value problem or a second order cone program (SOCP). Ours is the first method that allows for reconstruction of mirror surfaces from sparse RCs, obtained from standard algorithms such as SIFT. Our approach overcomes the practical issues in shape from specular flow (SFSF) such as the requirement of dense optical flow and un-defined/infinite flow at parabolic points. We also show how to incorporate auxiliary information such as sparse surface normals into our framework. Experiments, both real and synthetic are shown that validate the theory presented.

CVPR 2010

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Specular Surface Reconstruction from Sparse Reflection Correspondences

Aswin C. Sankaranarayanan[†], Ashok Veeraraghavan[‡], Oncel Tuzel[‡] and Amit Agrawal[‡]
[†]Rice University, Houston, TX
[‡]Mitsubishi Electric Research Labs, Cambridge, MA

Abstract

We present a practical approach for surface reconstruction of smooth mirror-like objects using sparse reflection correspondences (RCs). Assuming finite object motion with a fixed camera and un-calibrated environment, we derive the relationship between RC and the surface shape. We show that by locally modeling the surface as a quadric, the relationship between the RCs and unknown surface parameters becomes linear. We develop a simple surface reconstruction algorithm that amounts to solving either an eigenvalue problem or a second order cone program (SOCP). Ours is the first method that allows for reconstruction of mirror surfaces from sparse RCs, obtained from standard algorithms such as SIFT. Our approach overcomes the practical issues in shape from specular flow (SFSF) such as the requirement of dense optical flow and undefined/infinite flow at parabolic points. We also show how to incorporate auxiliary information such as sparse surface normals into our framework. Experiments, both real and synthetic are shown that validate the theory presented.

1. Introduction

Objects that exhibit mirror reflectance have no appearance of their own, but rather distort the surrounding environment. Traditional shape recovery methods designed for Lambertian surfaces such as structure from motion (SfM), stereo or multi-view stereo can not be directly used for such objects. Shape recovery of highly specular and mirror-like objects was first studied as an extension of the photometric stereo by Ikeuchi [10]. Since then methods have been developed for a wide range of imaging conditions, including known motion or scene patterns and active illumination.

Recently, there [1, 13] has been significant progress in shape from specular flow (SFSF). SFSF explores surface estimation by measuring dense optical flow of environment features as observed on the mirror under a known motion of environment/mirror/camera. Much of the prior work in SFSF assumes infinitesimal rotation of the environment, where in the forward flow equations linking the motion field, surface parameters and specular flow (SF) can be expressed as a partial differential equation. The elegance of




Figure 1. Effect of finite motion: For infinitesimal environmental motion, points P_1 and P_2 having the same normal and curvature would exhibit identical specular flow. But, for finite motion, the observed flow is different since it also depends on the normal and curvature of the neighborhood.




Figure 2. Two images of the same mirror rotated with a few reflection correspondences highlighted.

existing approaches do not extend to finite or large motion since large displacements can not be incorporated into the pde framework. Further, SFSF requires dense optical flow which is difficult to obtain for specular objects since SF exhibits certain undesirable properties such as undefined/infinite flow and one-to-many mappings.

Specular Flow vs Reflection Correspondences: Figure 1 shows points (P_1 and P_2) on two different surfaces, where the local normal and curvature are identical. Since the specular flow forward equations depend only on local normal and curvature, for the same infinitesimal environment motion, the flow will be identical for these two points. However, in case of a finite/large environment motion as shown, the observed flow depends not only on the local normal and curvature, but also on properties of the neighborhood. Notice that the actual flow observed at these two surface points is radically different because of the difference in their neighborhood. This dependence is not easily incorporated into the specular flow framework.

In this paper, we develop a theory of specular surface reconstruction under finite motion using *reflection correspondences* (RCs). Any two points in one or more images of a specular surface which observe the same environmental feature are denoted as RCs (refer Figure 2). We derive the relationship between RCs and the shape of the mirror for the case of finite motion and use a locally quadric surface parameterization to develop an efficient algorithm for surface reconstruction. By using correspondences, we avoid the undesirable properties of specular flow and design a practical solution. To our knowledge, this is the first method that allows for reconstruction of mirror surfaces from sparse RCs as would be obtained when using a feature matching algorithm such as SIFT.

1.1. Contributions

The specific technical contributions of this paper are

- We formulate the problem of recovering surface shape from RCs for the case of finite motion under uncalibrated environment.
- We model the surface as a locally quadric leading to a linear formulation solvable using efficient algorithms.
- We avoid practical issues in SFSF requiring only sparse RCs, leading to a practical workable solution.
- We show how to incorporate auxiliary information such as sparse surface normals to improve the reconstruction.

1.2. Prior Work

Qualitative Properties: Zisserman et al. [19] show that local surface properties such as concave/convexity can be determined under motion of the observer without knowledge of the lighting. Blake [3] analyzes stereoscopic images of specular highlights and shows that the disparity of highlights is related to its convexity/concavity. Fleming et al. [8] discuss human perception of shape from images of specular objects even when the environment is unknown.

Active Illumination: Ikeuchi [10] present the idea of estimating the structure of specular objects via photometric stereo. Oren and Nayar [12] use the notion of caustics to determine if an image feature is real or a reflection. Surface recovery is done by tracking an unknown scene point. Chen et al. [7] use this property to estimate surface mesostructures at high resolution. Hertzmann and Seitz [9] suggest using probes of known surface and reflectance in order to determine the properties of an unknown surface.

Calibrated Environment: Surface reconstruction under known environment has been studied in great detail. Savarse et al. [15] show that local properties of a smooth surface can be determined upto third order by the observed projection of two calibrated lines. Bonford and Sturm [4] use multiple images of a calibrated pattern to reconstruct the shape of a mirror. Nayar et al. [11] use structured highlighting to obtain shape of specular objects. An extension for a stereo based imaging system is proposed in [14].

Uncalibrated Environments: Walden and Dyer [18] show that *specular flow* can be significantly different from motion field especially near points of parabolic curvature. Chen and Arvo [6] developed specular path perturbation theory for fast rendering of multiple inter-reflections. This theory is applied to shape from specular flow by Roth and Black [13]. A variational approach for interpolating sparse set of normals is proposed in Solen et al. [16].

Recently, significant headway into the SFSF problem has been made in a series of papers [1, 17, 5]. Adato et al. [1] solve for the structure of the object from an infinitesimal rotation of the environment around the camera’s optical axis. Vaillyev et al. [17] show that the specular flow for rotation about the optical axis can be computed for two known or three unknown rotations about arbitrary axes. Both of these approaches require a significant number of initial conditions (in terms of normals) to solve for the shape of the mirror. Canas et al. [5] show that under a different parametrization of the surface, the estimation problem becomes linear. This is used to estimate the shape of the mirror from multiple specular flows *without* any additional information. But, the SFSF formulation is inherently differential and does not extend to large motion. Here, we develop the problem of surface reconstruction from RCs that are obtained from images in which the environment/object motion is large.

2. Reflection Correspondences

Consider a camera looking at an object with mirror reflectance. For the sake of simplicity, we first describe the method as applicable to an orthographic camera. We assume that the environment is at infinity, which ensures that the environment feature imaged at a pixel is dependent only on the normal of the surface element observed at the pixel and not on its location. Finally, we use the terms *object* and *mirror* inter-changeably in rest of the paper to denote the object/mirror whose surface we seek to estimate.

Let $I(\mathbf{x})$ be the intensity observed at pixel \mathbf{x} , and $E(\Theta)$ be the environment intensity map parameterized by spherical coordinates $\Theta = (\theta, \phi)$. We model the mirror surface in its Monge form $z = f(\mathbf{x}) = f(x, y)$. Under the imaging setup assumed, the forward imaging equation can be written as,

$$I(\mathbf{x}) = E(\Theta(\mathbf{x})) = E(\theta(\mathbf{x}), \phi(\mathbf{x}))$$

$$\tan \theta(\mathbf{x}) = \frac{2\|\nabla f\|}{1-\|\nabla f\|^2}, \tan \phi(\mathbf{x}) = \frac{f_y(\mathbf{x})}{f_x(\mathbf{x})} \quad (1)$$

and $\nabla f(\mathbf{x}) = (f_x(\mathbf{x}), f_y(\mathbf{x}))^T$ are the surface gradients at \mathbf{x} . For rest of the paper, we model the surface in terms of its gradient field $\nabla f(\mathbf{x})$. Our goal is to estimate the shape $f(\mathbf{x})$ from multiple images of the object under known motion of the camera/object/environment.

2.1. Influence of Motion on Observed Features

For an orthographic camera, translation of the mirror does not change the surface normals in a camera coordinate system, and produces just a translation of the image. Thus,

it does not provide any new information. Similarly, translation of the object gives no additional information since the environment is assumed to be at infinity. Rotation of the camera about the optical axis (in-plane rotation) is a remapping of pixels and produces no new information. For the same reason, rotation of the camera about its optical axis is just a remapping of the pixels and produces no new information. A rotation of environment is equivalent to placing the object on the camera’s optical axis and rotating the camera-object pair in tandem. Thus, ignoring pixel remapping, environment rotation is equivalent to object rotation about the optical axis of the camera. For the rest of the paper, we only consider object rotation since this changes the normal field observed by the camera, thereby producing images not seen earlier.

Due to assumption of environment at infinity, observed environment features depend only on the surface normal of the object. Under object rotation, its image changes in a structured way. The environment feature observed at a pixel could reappear at locations where the surface normal reappears. It is also possible that the environment feature does not reappear after object rotation, or reappear at multiple locations.

In this paper, we focus *only* on rotation of the object around the camera’s optical axis. This particular choice leads to a linear formulation in the surface gradients given the RCs, and as a consequence, enables the use of well-known optimization tools. Extensions to more general motion form the focus of future research.

2.2. Infinitesimal Rotation

For completeness, we first show that RCs are equivalent to specular flow under infinitesimal object rotation (c.f. [1]). Let the mirror surface be $z = f(\mathbf{x})$ in the camera coordinate system at $t = 0$. After rotation about the optical axis of the camera, the mirror surface at time t is given by given as $z(t, \mathbf{x}) = f(t, \mathbf{x}) = f(R(t)^T \mathbf{x})$, where $R(t) = \exp(t\omega)$ is the 2×2 rotation matrix defined by the skew-symmetric matrix ω . The surface gradients are given as,

$$\nabla f(t, \mathbf{x}) = R(t) \nabla f(R(t)^T \mathbf{x}). \quad (2)$$

Consider a environment feature that moves along $\mathbf{x}(t) = \mathbf{x}_0 + \boldsymbol{\nu}(t)$ as the object rotates. Under the modeling assumptions, $\nabla f(t, \mathbf{x}_0 + \boldsymbol{\nu}(t)) = \nabla f(0, \mathbf{x}_0) = \nabla f(\mathbf{x}_0)$. This implies that,

$$\frac{d\nabla f(t, \mathbf{x}(t))}{dt} = 0 \quad (3)$$

Evaluating this expression at $t = 0$, with $R(0) = \mathbb{I}$, we get

$$\omega \nabla f(\mathbf{x}_0) - \left[\frac{\partial^2 f}{\partial \mathbf{x}^2} \right] (\omega \mathbf{x}_0 - \dot{\boldsymbol{\nu}}(\mathbf{x}_0)) = 0 \quad (4)$$

This gives the forward flow equations relating the infinitesimal specular flow $\dot{\boldsymbol{\nu}}$ to the local surface parameters and the

known rotation encoded in ω . Given specular flow, the surface gradient field can be recovered by solving the partial differential equation defined in (4)[1]. The second order properties of the surface are encoded in the Hessian matrix $[\partial^2 f / \partial \mathbf{x}^2]$ play an important role in determining the specular flow. When this matrix is singular specular flow goes to infinity. Parabolic curvature points have singular Hessian matrix and exhibit infinite flow.

Infinitesimal to Finite Rotation: While infinitesimal motion leads to an elegant theoretical formulation, in practice it is approximated by a small motion. For Lambertian objects, small motion is often a good approximation to infinitesimal motion as the flow values are well-defined. However, the presence of infinite flow at parabolic curvature points indicates that no matter how small the motion is, the specular correspondences will be undefined in the neighborhood of parabolic curvature points. The size of this neighborhood depends on the magnitude of the rotation. Hence, it becomes hard to estimate specular flow field. This observation motivates modeling matches between images in terms of *sparse* RCs as opposed to a *dense* flow field. Sparse correspondence allows us to incorporate regions/pixels that have no flow associated with them as well as those that exhibit one-to-many mappings across images. More importantly, one to many mappings, or disappearance of correspondences (similar to occlusion/disocclusion for stereo) can be handled as well.

2.3. Finite Motion

Now we consider finite rotation R of the object about the camera’s optical axis. Let $\{\mathbf{x}_i^A \leftrightarrow \mathbf{x}_i^B; i = 1, \dots, N\}$ be a set of correspondences on a image pair such that environment feature observed at \mathbf{x}_i^A on image A matches with that in \mathbf{x}_i^B on image B. This induces the following relationships on the surface gradients of the object,

$$\nabla f_A(\mathbf{x}_i^A) = \nabla f_B(\mathbf{x}_i^B), i = 1, \dots, N \quad (5)$$

Further, $f_B(\mathbf{x}) = f_A(R^T \mathbf{x})$, giving us a constraint purely in terms of the object at location A.

$$\nabla f_A(\mathbf{x}_i^A) = R \nabla f_A(R^T \mathbf{x}_i^B) \quad (6)$$

Given multiple images under varying (known) rotation, we can use correspondences obtained between *any* image pair by suitably rewriting the relationships in terms of ∇f_A . Thus, the key idea is that, given object rotation, RCs provide constraints on the surface gradients. To recover a dense surface, we need additional assumptions on the surface. Towards this end, we propose the use of locally quadric model for the surface.

3. Shape from Sparse Correspondences

In this section we describe modeling the surface as a locally quadric and algorithms for surface reconstruction.

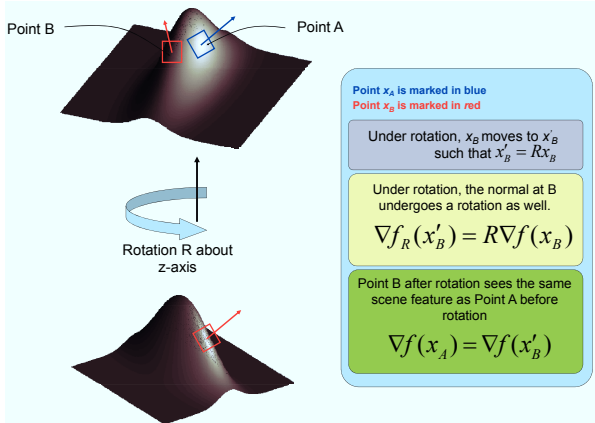


Figure 3. Rotation of a mirror causes its appearance to change. Under rotation of the mirror, its corresponding normal field changes. Environment features reappear wherever its corresponding surface normal reappears. Hence, a RC implies that the normal at the corresponding locations are the same.

3.1. Surface Approximation using Quadrics

We approximate a general smooth surface using local quadric models over non-overlapping neighborhoods. A surface $z = f(\mathbf{x}) = f(x, y)$ is approximated as,

$$z = f(x, y) \approx \frac{1}{2} \mathbf{x}^T H_i \mathbf{x} + J_i \mathbf{x} + \alpha_i, \mathbf{x} \in \mathcal{N}_i \quad (7)$$

H_i, J_i and α_i are the quadric parameters defining the neighborhood \mathcal{N}_i . The surface gradients are defined as,

$$\nabla f(\mathbf{x}) = J_i + H_i \mathbf{x}, \mathbf{x} \in \mathcal{N}_i \quad (8)$$

Note that H_i is a symmetric matrix and α_i is unobservable in an orthographic camera setup.

The choice of quadric model for local surface parametrization is motivated primarily from the fact the associated surface gradient field is *affine* in \mathbf{x} . This, as we show later, leads to linear relationships in the unknown surface parameters (J_i and H_i) given RCs. It is entirely possible to use alternate surface parameterizations at the cost of solving a harder inference problem.

Note that in our imaging setup (environment at infinity with an orthographic camera), lower order models such as piecewise planar produce far inferior approximations as the entire neighborhood images the same environment feature leading to highly coarse imagery (equivalently, implying that the assumptions are valid for extremely small neighborhoods). The quadric model serves as a good approximation for a wide range of smooth surfaces.

3.2. Quadric Reflectance Correspondences

We can now re-interpret the surface gradient relationship of (6) under the local quadric surface model. Note that (6) relates the surface gradients at points \mathbf{x}^A and $R^T(\mathbf{x}^B)$. Equation (6) can be rewritten in terms of the surface parameters of the neighborhoods to which the two point belong.

Let $\mathbf{x} \in \mathcal{N}_k$ and $R^T(\mathbf{x}^B) \in \mathcal{N}_m$. Then, substituting the expressions for the surface gradients (8) in (6), we get:

$$R(J_m + H_m R^T(\mathbf{x}^B)) = J_k + H_k \mathbf{x}^A \quad (9)$$

When the rotation R is known, (9) is linear in the unknown surface parameters. Equation (9) also hints at a scale ambiguity in the solution space. A solution for the quadric parameters can be scaled to obtain another solution which satisfies the equations as well. This implies that any solution to the quadric parameters at best suffers from a global scale ambiguity. However, the problem can be more severe when cliques of neighborhood are formed, with correspondences only between points in the same clique. This could potentially lead to each clique having its own scale ambiguity thereby leading to a far more severe problem globally. To resolve this, we enforce smoothness of the surface gradients at the boundary between two adjacent neighborhoods.

3.3. Smoothness Constraints

As mentioned earlier, enforcing the RCs in (9) alone can potentially lead to multiple scale ambiguities that can be very difficult to resolve. Further, the local quadric model does not necessarily lead to smooth surfaces, in that, there could be discontinuities at the boundary of the neighborhoods. Both of these problems can be resolved if we further enforce the gradients at the boundary of two patches to be the same on either side of the patch. Let \mathbf{x} be a pixel at the boundary of two neighborhoods \mathcal{N}_m and \mathcal{N}_n . We can ensure that the surface gradient at \mathbf{x} as obtained by the quadric model of \mathcal{N}_m and \mathcal{N}_n are the same.

$$J_m + H_m \mathbf{x} = J_n + H_n \mathbf{x}, \mathbf{x} \in \mathcal{B}(\mathcal{N}_n) \cap \mathcal{B}(\mathcal{N}_m) \quad (10)$$

As with the relationships in (9), the boundary smoothness constraint is linear in the surface parameters and exhibits a scale ambiguity. This further reinforces the point that the global scale ambiguity cannot be resolved without an absolute information of a few surface normals or gradients. Such information can potentially be obtained from occluding contours or knowledge of normals at few distinct points. We later show how such constraints can be incorporated into our framework.

3.4. Solving for the Surface

In order to solve for the surface, the neighborhood size of the local quadric must be defined. A smaller neighborhood allows for a better approximation of the surface, but leads to more unknowns in the surface estimation problem. Without a priori knowledge of the surface, it is impossible to choose neighborhoods that well approximates the surface and also reduces the number of parameters. We chose non-overlapping square neighborhoods over the image plane, whose size depends on the size of the problem we can solve efficiently.

Given a set of RCs $\{\mathbf{x}_i^A \leftrightarrow \mathbf{x}_i^B\}$ and the rotation undergone by the surface R , we can now solve for the quadric surface parameters. For each available correspondence, we formulate the relationships as given in (9). Next, between boundary of two patches, we enforce the boundary smoothness constraints of (10). Stacking these together, we get a system of linear system of equations $A\Theta = 0$, where Θ is a vector consisting of all the surface parameters $\{J_i, H_i, i = 1, \dots, N\}$.

Solution can be obtained by formulating the solution as the smallest eigenvalue of the matrix $A^T A$.

$$\hat{\Theta} = \arg \min_{\|\theta\|=1} \theta^T A^T A \theta \quad (11)$$

Alternatively, we can formulate it as a second order cone program (SOCP):

$$\hat{\Theta} = \arg \min_{\Theta} \|A\Theta\|_2 \text{ s.t } \mathbf{1}^T \Theta \geq 1 \quad (12)$$

Solving as an eigenvalue problem is often more accurate in terms of numerical precision. However, for very large systems it is not possible to solve the eigenvalue problem due to resource constraints. In such cases, it might still be possible to solve it as a convex program. The convex formulation also allows to incorporate additional constraints that might be available. Using the surface parameters, the gradient field is recovered. The surface is then estimated by integrating the gradient field by solving a Poisson equation [2]. The various steps involved in our surface reconstruction algorithm are shown graphically in Figure 4.




Figure 4. Block Diagram of the various steps involved in surface reconstruction from sparse RCs.

3.5. Additional Normal Constraints

In addition to RCs, surface normal information might be available at a sparse set of points on the surface. This information might come either from observing a known scene point, or from observing the reflection of the camera, or from the occluding contour. In any of these cases, we can reformulate the problem such that this additional information is seamlessly integrated into the reconstruction algorithm. The convex optimization framework for estimating the shape of the mirror is more flexible in its ability to incorporate additional constraints. One such constraint is that of the knowledge of surface normals or gradients at a few locations on the mirror. Given a set of such surface gradients, $\{\mathbf{x}_l^n, \nabla f^n(\mathbf{x}_l), l = 1, \dots, Q_n\}$, we can enforce these by adding additional constraints to the optimization problem. Each of the known gradients induce a linear constraint on the quadric parameters.

$$J_{n_l} + H_{n_l} \mathbf{x}_l^n = \nabla f^n(\mathbf{x}_l), \mathbf{x}_l^n \in \mathcal{N}_{n_l}, l = 1, \dots, Q_n. \quad (13)$$

We can now solve an optimization problem defined as follows,

$$\begin{aligned} & \min_{J_k, H_k} \|A\Theta\|_2 \\ & \text{subject to } \mathbf{1}^T \Theta \geq 1 \\ & J_{n_l} + H_{n_l} \mathbf{x}_l^n = \nabla f^n(\mathbf{x}_l), \mathbf{x}_l^n \in \mathcal{N}_{n_l}, l = 1, \dots, Q_n. \end{aligned} \quad (14)$$

The basic nature of the optimization problem does not change, in that, it remains a SOCP.

In practice, the ability to incorporate such additional information to estimate the surface of the mirror is extremely useful especially when there are not sufficient RCs to reconstruct the surface reliably. Having a few known surface gradients/normals goes a long way in regularizing the surface estimate.

4. Experiments

We perform a thorough evaluation of our algorithm on several challenging surfaces: a) Quadric b) Gaussian c) Mixture of Gaussians d) Test Surface 1 e) Test Surface 2 f) Gravy Ladle and g) Ice cream scoop. The analytical equations for Test surface 1 and 2 are $z = f(x, y) = \sqrt{4 - x^2 - y^2 - \cos(2x - 2) - \sin(2y)}$ and $z = f(x, y) = \sqrt{4 - x^2 - y^2 - \cos(3x - 6) - 2 \sin(2y)}$ respectively. Some of the surfaces on which we have tested the surface reconstruction algorithm are shown in Figure 5.



Figure 5. Our approach can handle mirror surfaces with different shapes and sizes as shown. The leftmost object was machined using a CNC milling machine. All other objects were bought from local departmental store.

4.1. Simulations:

Known RCs: In order to test the reconstruction algorithm, we first computed true RCs by matching normals before and after rotation. Using these ground truth RCs, we can reconstruct the surfaces. Figure 6 shows reconstructions for surfaces of varying complexity. This highlights the connection between the complexity of the surface and the number of correspondences required to reconstruct it reliably.

Reflection Correspondences from Images: Figure 7 shows reconstructions from rendered images. We used POV-Ray to render mirrors of various shapes, obtaining

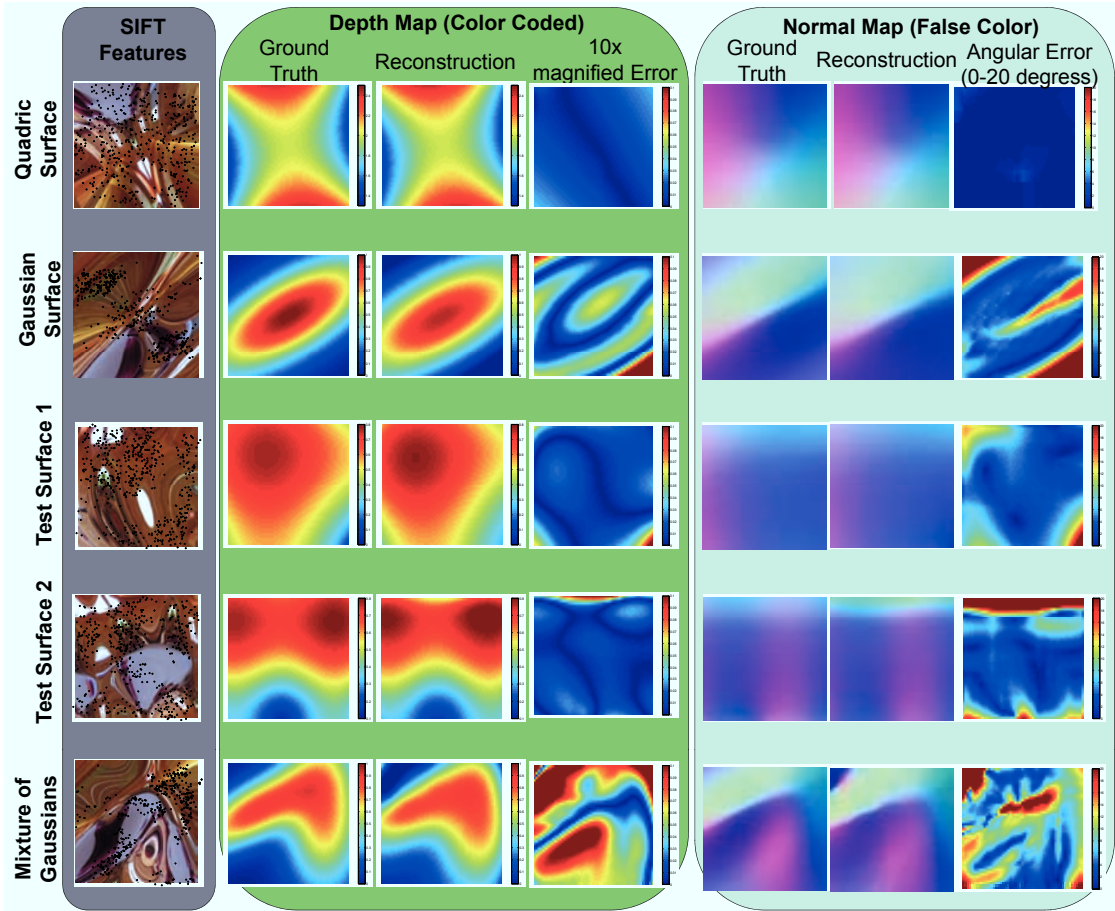


Figure 7. Reconstruction of various surfaces from SIFT matches obtained from rendered images. The reconstructions were obtained from images directly without any additional information. For each surface, multiple images of the mirror were obtained by rotating the environment in PovRay. Reflection correspondences were obtained using SIFT, and subsequently processed by the proposed surface reconstruction algorithm. Shown are reconstructed normal fields and normal mismatch error in degrees.




Figure 6. Three surfaces and the minimum number of RCs required to reliably reconstruct them. Reflection correspondences are assumed to have an uniform spatial spread. The more complex the surface, the more the number of RCs required to reconstruct it.

multiple images by rotating the mirror, and imaging using an ideal orthographic camera. Figure 7 shows reconstruction results obtained from 5 images of each mirror. RCs were obtained using SIFT. All surfaces were reconstructed by partitioning the region of interest into 10×10 non-overlapping neighborhoods and modeling the surface

locally as a quadric in each. Finally, outliers in the SIFT matches were manually removed before running the surface reconstruction algorithm. As expected the quality of reconstruction depends on the inherent complexity of the surface.

When very few images of a mirror are available, invariably SIFT matches do not produce a dense sampling on correspondences on the image plane. This results in features in the mirror being missed out. In such cases, having additional information in terms of actual normal values on the surface of the mirror helps in regularizing the solution. Figure 8 shows reconstruction results when surface normals at 16 locations were provided to the algorithm.

4.2. Experiments on Real Data

To test the algorithm with real data, we built an experimental rig as shown in Figure 9. We used a Point Gray Dragonfly camera with a 35mm lens. The mirror was mounted on a turn-table/screw about 85 cm from the camera. Motion of the mirror was estimated using markers stuck on the mirror and on the turn-table. It is noteworthy that modeling assumptions such as orthographic camera and

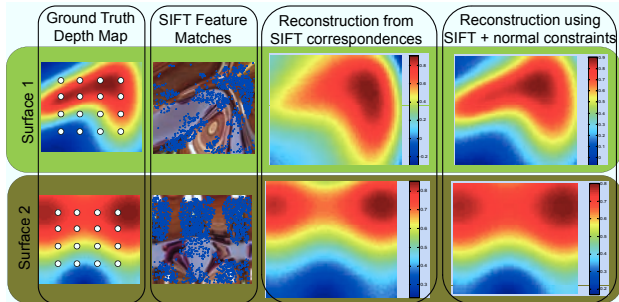


Figure 8. Additional information on the surface in terms of normals known at few location can be easily incorporated in the estimation process. This is especially useful when SIFT matching does not give a spatially dense covering. In this case, having just two images of the mirror do not produce enough SIFT matches to reconstruct the mirror faithfully. Additional knowledge of 16 normals at marked location improves the fidelity of the reconstruction.

scene at infinity are violated in this imaging setup.




Figure 9. Our experimental setup consists of mirror object on a turn table with the camera looking down on the mirror. The camera is calibrated and positioned such that its optical axis is parallel to the rotation axis of the turn-table/mirror. Markers on the mirror and turn-table allow for the recovery of object motion.

To analyze the accuracy of our approach, we used a quadric surface machined using a CNC milling machine, giving us precise ground truth information. Our dataset for this surface comprised of 38 images obtained by rotating the object. Markers placed on the mirror were used to estimate the object motion. Figure 10 shows a surface reconstruction result on this surface. Outliers were removed using RANSAC to enforce local affinity in the RCs. Here, a local affine model on the RCs was used motivated by the affine relationship between the correspondence pair x^A and x^B in (9). To test the robustness of our estimation, we reconstructed the mirror surface by choosing random pairs of images in which the mirror was less than 30 degrees apart. In all, our average reconstruction error for the depth map was 1.2 mm while the depth variation in the mirror was 27 mm.

Comparison with 3D scanner: We also show reconstruction results on a gravy-ladle in Figure 11. The gravy-ladle object was painted with a diffuse white paint and scanned using a commercial \$3000 NextEngine desktop 3D

scanner. Figure 11 compares the scanned point cloud with our reconstruction. We use the standard ICP algorithm to register the two surfaces. The difference between our reconstruction and scanned point cloud is less than 0.5 mm for 70% point matches. Real dimensions of the object was 67mm x 44mm x 25mm.




Figure 10. Surface estimation results from real data for the quadric surface. Shown are reconstruction results from using 2 images from the dataset. Average absolute error of the reconstruction is 1.2 mm over a depth variation of 27 mm.

Database: In the course of experiments, we collected a database consisting of two parts. For a total of six synthetic surfaces, we provide the depth map, several POV-Ray rendered images of the surface with known motion between the images. The database also consists of four real mirror surfaces, for which we provide a depth map (using a commercial laser based structured light 3D sensor), several real images of the surface with known rotation between these images. The dataset and our code for surface reconstruction from RCs are available through the authors' websites.

5. Discussions and Conclusions

It is worthwhile discussing the relative merits and demerits of our proposition. The theory described in [5] establish both uniqueness and existence of solution from multiple *dense* specular flows. In contrast, we show that it is possible to reconstruct mirror surfaces from a *sparse* set of RCs corresponding to a *single* rotation axis provided we enforce additional constraints on the surface smoothness. However, our method does rely on rotation about the optical axis of the camera, and hence, fails for objects exhibit rotational symmetry about this axis. Finally, Roth and Black [13] use a global surface parametrization which is more restrictive than the local patch-wise model assumed in this paper.




Figure 11. Comparison of reconstructed surface with 3D laser scan of object obtained by coating it with matte paint. Registration error was around 0.5 mm with about 70% match.

Sparsity of Correspondences: All of the information of the surface shape in our estimation comes from RCs. It is noteworthy that, in the absence of RCs, the matrix A in (11) has a 5-dimensional null space that spans the space of quadrics. The constraints generated by the RC shift the solution from a global quadric to the surface we are interested in. In this sense, the density of RCs determine the granularity with which features on the surface are resolved. In particular, capturing fine nuances in the surface require RCs that capture the subtle variations in the shape. In addition to this, to capture these variations, it is also important to make the neighborhoods of quadric approximation adaptive to the local density of RCs.

Complexity of Mirror Surface: The complexity of the mirror shape affects the reconstruction process in two ways. As the surface complexity increases, the size of the neighborhoods over which the quadric assumptions are made needs to be decreased. Complex surfaces affect the algorithm in an indirect way by introducing points/curves of parabolic curvature. As mentioned earlier in the paper, under finite motion, features appear/disappear in regions close to parabolic curvature. As a consequence, there is a possibility of obtaining no SIFT features in these regions (see Figure 7). However, parabolic curves do provide strong surface priors: local cues such as points and curves of zero Gaussian curvature, as well as global cues by dividing the surface of the mirror into regions of elliptic and hyperbolic curvature. Detecting parabolic curves from one or more images of the mirror and using such information appropriately will, potentially, enable reconstructions of arbitrarily complex surfaces.

Outliers: Robustness to outliers remain one of the main challenges to the surface estimation procedure presented in this paper. For the experiments in this paper, the RCs were pruned using RANSAC to enforce local affinity of features. However, a rigorous outlier handling procedure remains an important direction for future research.

Conclusions: In this paper, we proposed using sparse RCs for specular surface reconstruction. We showed that using RC avoids practical difficulties with previous SFSF methods. We show that modeling the mirror surface as a locally quadric makes the relationship between RC and surface linear. This allows us to reconstruct the surface using an efficient algorithm. We highlight the ability of the algorithm by working with real images of mirror and use off-the-shelf feature matching algorithms such as SIFT for estimating the shape of the mirror. In this regard, we believe that this paper makes significant progress towards making surface recovery of mirrors practical.

Acknowledgement

The authors thank John Barnell for help with the hardware, and Prof. Rama Chellappa for his encouragement. Thanks to Ming-Yu Liu, Prof. Todd Zickler and Dr. Jay Thornton for discussions on the manuscript. A.S. thanks Prof. Richard Baraniuk for his

support and encouragement.

References

- [1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Toward a Theory of Shape from Specular Flow. In *ICCV*, pages 1–8, 2007. 1, 2, 3
- [2] A. Agrawal, R. Raskar, and R. Chellappa. Surface reconstruction from gradient fields via gradient transformations. *IJCV*, 2009. 5
- [3] A. Blake. Specular stereo. In *Proc. Int'l Conf. Artificial Intell.*, pages 973–976, 1985. 2
- [4] T. Bonfort and P. Sturm. Voxel carving for specular surfaces. In *ICCV*, pages 591–596, 2003. 2
- [5] G. D. Canas, Y. Vasilyev, Y. Adato, T. Zickler, S. Gortler, and O. Ben-Shahar. A Linear Formulation of Shape from Specular Flow. In *ICCV*, pages 1–8, 2009. 2, 7
- [6] M. Chen and J. Arvo. Theory and application of specular path perturbation. *ACM Trans. Graph.*, 19(1):246–278, 2000. 2
- [7] T. Chen, M. Goesele, and H. Seidel. Mesostructure from specularity. In *CVPR*, volume 2, 2006. 2
- [8] R. Fleming, A. Torralba, and E. Adelson. Specular reflections and the perception of shape. *J. Vision*, 4(9):798–820, 2004. 2
- [9] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *PAMI*, 27(8):1254–1264, 2005. 2
- [10] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE PAMI*, 3(6):661–669, 1981. 1, 2
- [11] S. K. Nayar, A. C. Sanderson, L. E. Weiss, and D. A. Simon. Specular surface inspection using structured highlight and Gaussian images. *IEEE Trans. Robotics and Automation*, 6(2):208–218, 1990. 2
- [12] M. Oren and S. Nayar. A theory of specular surface geometry. *IJCV*, 24(2):105–124, 1997. 2
- [13] S. Roth and M. Black. Specular flow and the recovery of surface structure. In *CVPR*, volume 2, 2006. 1, 2, 7
- [14] A. Sanderson, L. Weiss, and S. Nayar. Structured highlight inspection of specular surfaces. *IEEE PAMI*, 10(1):44–55, 1988. 2
- [15] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. *IJCV*, 64(1):31–67, 2005. 2
- [16] J. Solem, H. Aanaes, and A. Heyden. A variational analysis of shape from specularities using sparse data. In *3DPVT*, pages 26–33, 2004. 2
- [17] Y. Vasilyev, Y. Adato, T. Zickler, and O. Ben-Shahar. Dense specular shape from multiple specular flows. In *CVPR*, pages 1–8, 2008. 2
- [18] S. Waldon and C. Dyer. Dynamic shading, motion parallax and qualitative shape. In *Proc. IEEE Workshop on Qualitative Vision*, pages 61–70, 1993. 2
- [19] 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. 2