

Mirror Surface Reconstruction from a Single Image

Miaomiao Liu, Richard Hartley, Mathieu Salzmann
NICTA* & ANU

{miaomiao.liu, mathieu.salzmann}@nicta.com.au, richard.hartley@anu.edu.au

Abstract

This paper tackles the problem of reconstructing the shape of a smooth mirror surface from a single image. In particular, we consider the case where the camera is observing the reflection of a static reference target in the unknown mirror. We first study the reconstruction problem given dense correspondences between 3D points on the reference target and image locations. In such conditions, our differential geometry analysis provides a theoretical proof that the shape of the mirror surface can be uniquely recovered if the pose of the reference target is known. We then relax our assumptions by considering the case where only sparse correspondences are available. In this scenario, we formulate reconstruction as an optimization problem, which can be solved using a nonlinear least-squares method. We demonstrate the effectiveness of our method on both synthetic and real images.

1. Introduction

In this paper, we tackle the problem of mirror surface reconstruction from a single image. Traditional 3D reconstruction methods typically perform poorly on mirror objects, since the information captured by the camera comes from the objects' surroundings rather than from the objects themselves. Methods specifically designed to handle mirror surfaces have been introduced, but usually exploit motion and thus do not apply to the single image scenario [16, 8, 13, 20]. Solutions to the single image case have nonetheless been proposed [9, 17, 15, 18]. However, existing approaches do not offer theoretical guarantees of the uniqueness of the reconstructed surface.

Here, we introduce an approach to reconstructing a mirror surface from a single image with a provably unique solution. To this end, we consider the scenario where the camera observes the reflection of a static reference plane¹

*NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

¹Although we consider a plane, any object of known shape can be used.




Figure 1. Real reconstruction setup and sample input images.

with known pose in the unknown mirror. Furthermore, we assume that *reflection correspondences* between 3D points on the reference plane and 2D image locations are given. With dense correspondences, a differential geometry analysis reveals that, for a *smooth* mirror surface without inter-reflections, reconstruction reduces to solving an initial value problem (IVP) with two partial differential equations (PDEs). We derive a theoretical proof of uniqueness of the solution to this IVP. Furthermore, studying the order of integration of the two PDEs yields a generally unique solution for the starting point of the IVP. This therefore implies uniqueness of the mirror surface reconstruction.

To address the more realistic scenario where only sparse reflection correspondences are available, we parametrize the depths of points on the mirror surface as a uniform cubic B-spline. We then formulate reconstruction as an optimization problem that minimizes the 3D error between the points on the reference plane and the image correspondences back-projected to the reference plane via the mirror.

In summary, the key contributions of this paper are

- A solution to the problem of reconstructing a smooth mirror surface from a single image given dense correspondences between the image and a reference plane with known pose.
- A theoretical proof of uniqueness of this solution.
- A practical optimization formulation of the reconstruction problem given sparse reflection correspondences.

We demonstrate the effectiveness of our reconstruction method from both dense and sparse correspondences on synthetic and real images, such as those depicted in Fig. 1.

2. Related Work

Most existing methods that tackle mirror surface reconstruction exploit temporal information, such as the motion of the camera [13, 20], or that of the environment [8, 16]. Within this class of methods, shape from specular flow has become a popular approach [14, 1, 5, 19]. As an alternative, shape recovery can be performed by exploiting multiple reference planes with known pose relative to the camera. This can be achieved either by utilizing multiple views of the object with a reference plane fixed relative to the camera [3, 12, 2], or with a static camera, but a moving reference plane [4, 11]. In [10], it was shown that the surface shape can be recovered from a single viewpoint when two 3D reference points on the light path are known, which is similar to using a moving reference plane. In this paper, we consider the problem of reconstructing a smooth mirror surface from a single image, and therefore cannot exploit motion. As such, our approach is most related to [9, 17, 15, 18].

In [9], a method to recover the shape of the human cornea from a fixed camera and a static reference plane with known pose was introduced. The shape was modeled as a uniform biquintic B-spline, thus bearing similarities with our cubic B-spline formulation. However, the approach in [9] requires the 3D location of one point on the surface to be known. Furthermore, it comes with no theoretical guarantees of convergence, or uniqueness. Here, we provide a theoretical proof of uniqueness of the mirror shape that does not require knowing the position of any surface point.

A differential geometry analysis of surface patches was proposed in [17]. While the reconstruction in [17] was limited to local patches, it was extended to modeling a global surface shape in [15]. To this end, local patches were linked by a Delaunay triangulation, and this rough shape estimate was refined by constrained interpolation. While we also base our analysis on differential geometry, our formulation naturally extends to the entire mirror surface without requiring stitching patches together.

In [18], the shape of the surface was recovered by iteratively estimating the normal at a point from its depth, and the depth of a neighboring point from this normal. The initial depth was obtained by searching for a value that gave coherent depth and normal fields. Our method can be thought of as an analytical formulation of this approach. Our analytical formulation allows us to establish a theoretical proof of uniqueness of the solution of both the IVP that corresponds to the iterative procedure described above and the starting depth required to solve this IVP.

3. Notation and Setup

Objects and scalars are denoted by italic letters (e.g., λ , P). Points and vectors are denoted by bold letters (e.g., $\mathbf{v} = (x, y, 1)^\top$) with the exception of the rotation matrix and translation vector, which we denote by R and T , respectively.




Figure 2. **Mirror surface reconstruction setup.** A pinhole camera centred at O is observing a mirror surface P that reflects a reference plane Q in the image I . A point \mathbf{m} on Q is reflected to the image point \mathbf{v} on I via the 3D mirror point \mathbf{p} on P . We refer to \mathbf{m} and \mathbf{v} as reflection correspondences. The reflected ray l is determined by \mathbf{m} and \mathbf{p} . We denote \mathbf{i} as the incident ray for image point \mathbf{v} and \mathbf{n} as the normal to P at \mathbf{p} . R and T denote the pose of the reference plane w.r.t. the camera.

Elements of a vector are referenced with subscripts that indicate their position in the vector (e.g., $\mathbf{n} = (n_x, n_y, n_z)^\top$). Vector-valued functions are denoted by bold letters (e.g., $\mathbf{m}(x, y)$). The L_2 norm of a vector is denoted $\|\cdot\|$, and $\langle \cdot, \cdot \rangle$ denotes the inner product of two vectors.

Mirror surface representation. Fig. 2 depicts our setup for mirror surface reconstruction. Without loss of generality, we suppose that the camera is centred at the origin of the coordinate system and oriented so that its principal axis is aligned with the positive z -axis. Furthermore, we assume that the focal length is 1. A point (x, y) in the image therefore corresponds to a point $\mathbf{v}(x, y) = (x, y, 1)^\top$ in an image plane at distance 1 in front of the camera. We focus on a rectangular region of interest (ROI) $I_x \times I_y$ ² in which the mirror P is visible, non-tangentially.

Thus, for points $(x, y) \in I_x \times I_y$, the ray through point $\mathbf{v}(x, y) = (x, y, 1)^\top$ meets the mirror simply (not tangentially) at a point $\mathbf{p}(x, y) = s(x, y)\mathbf{v}$, where $s(x, y)$ is referred to as the depth of the mirror at this point. The function $s(x, y)$ therefore determines the shape of the mirror, and finding this function is equivalent to finding the shape of the part of the mirror that lies within the ROI. We assume that the function $s(x, y)$ is at least twice continuously differentiable.

4. Differential Geometry Analysis

In this section, we present our approach to mirror surface reconstruction given dense reflection correspondences, as well as our proof of uniqueness of a solution. Our analysis of the reconstruction problem relies on the normal \mathbf{n} to the mirror at \mathbf{p} . Therefore, we start by deriving an analytical expression for \mathbf{n} . Let \mathbf{m} be a point on the reference plane

² I_x and I_y are closed sets.

Q that reflects to \mathbf{v} on the image plane I via \mathbf{p} . Based on the geometry of reflection, \mathbf{n} bisects the angle between the incident ray $\mathbf{i} = \mathbf{v}/\|\mathbf{v}\|$ and the reflected ray $\mathbf{l} = (\mathbf{m} - s\mathbf{v})/\|\mathbf{m} - s\mathbf{v}\|$. This lets us write the unnormalized normal to the mirror surface P at \mathbf{p} as

$$\mathbf{n} = \|\mathbf{v}\|(\mathbf{m} - s\mathbf{v}) - \|\mathbf{m} - s\mathbf{v}\|\mathbf{v}. \quad (1)$$

Note: It is important to remember that all the quantities \mathbf{n} , \mathbf{v} , \mathbf{m} , and s are functions of (x, y) . To simplify notation, the explicit dependence on (x, y) will henceforth usually be omitted.

Our formulation relies on differential geometry. To this end, we assume that s , \mathbf{m} , and therefore \mathbf{p} and \mathbf{n} , are sufficiently *smooth* functions of the image location (x, y) . This will certainly be true if the mirror is smooth, and with no occluding contours. Under this assumption, the normal to the surface can also be expressed as $\mathbf{n} = \partial\mathbf{p}/\partial x \times \partial\mathbf{p}/\partial y$. This implies that the surface normal must be orthogonal to the partial derivatives $\partial\mathbf{p}/\partial x$ and $\partial\mathbf{p}/\partial y$. Since $\mathbf{p} = s\mathbf{v}$, this can be expressed, for x , as

$$\begin{aligned} 0 &= \langle \mathbf{n}, \partial\mathbf{p}/\partial x \rangle \\ &= \langle \mathbf{n}, s \partial\mathbf{v}/\partial x + \partial s/\partial x \mathbf{v} \rangle \\ &= s \langle \mathbf{n}, \partial\mathbf{v}/\partial x \rangle + \partial s/\partial x \langle \mathbf{n}, \mathbf{v} \rangle. \end{aligned} \quad (2)$$

Since $\mathbf{v} = (x, y, 1)^\top$, its x -derivative is $(1, 0, 0)$, and Eq. (2) can be re-written as

$$\frac{\partial s}{\partial x} = \frac{-n_x s}{\langle \mathbf{n}, \mathbf{v} \rangle}, \quad (3)$$

where n_x denotes the x coordinate of \mathbf{n} . Similarly, for y , we can derive

$$\frac{\partial s}{\partial y} = \frac{-n_y s}{\langle \mathbf{n}, \mathbf{v} \rangle}. \quad (4)$$

Suppose that the reflection correspondence \mathbf{m} is known for each point in the ROI (either by dense matching, or modeling) and can therefore be written as a function $\mathbf{m}(x, y)$. Eq. (1) expresses \mathbf{n} as a function of x , y , s and $\mathbf{m}(x, y)$. Substituting this expression into Eqs. (3) and (4) yields the Initial Value Problem (IVP)

$$\begin{aligned} \frac{\partial s}{\partial x} &= f_x(s, x, y, \mathbf{m}(x, y)) \\ \frac{\partial s}{\partial y} &= f_y(s, x, y, \mathbf{m}(x, y)) \\ s(x_0, y_0) &= s_0 \end{aligned} \quad (5)$$

where $(x_0, y_0) \in I_x \times I_y$, s_0 is an initial value for the function s at this point, and the particular form of f_x and f_y is given by Eqs. (3), (4) and (1). Assuming that $\mathbf{m}(x, y)$ is known for points in $I_x \times I_y$, both f_x and f_y are ultimately functions of x , y and s only.

The important point here is that given the function $\mathbf{m}(x, y)$, the shape of the mirror, determined by $s(x, y)$, must be a solution to this IVP.

4.1. Uniqueness Results

The goal of this section is to state certain uniqueness results for the mirror shape, based on the formulation of the problem as the IVP (5). Proofs will follow in later sections.

Uniqueness of the solution to the IVP. Assuming that the function $\mathbf{m}(x, y)$ is derived from an actual instantiation of the model, without noise, then there is at least one solution to the IVP (5) for a particular value of s_0 , namely the true solution. However, if $\mathbf{m}(x, y)$ is a given arbitrary function, even smooth, and s_0 is some given value, then it cannot be expected that the IVP has a solution. We therefore concentrate on the uniqueness of the solution, which would imply that the true shape of the mirror is the only solution to the IVP.

Theorem 4.1. *Suppose that for each point (x, y) in a region of interest $I_x \times I_y$ in an image, the corresponding ray meets a mirror non-tangentially. Given the data $\mathbf{m}(x, y)$ and an initial value s_0 , the IVP (5) has (at most) a unique solution, defined on the whole of $I_x \times I_y$.*

In Section 4.2, this result will be seen to follow from a standard uniqueness result in Ordinary Differential Equations (ODEs), namely the Picard Lindelöf Theorem [6].

Uniqueness of the starting point. From Theorem 4.1 it may appear that there is a family of solutions to the IVP (5), one for each value of s_0 . However, studying the order of integration of the PDEs in the IVP reveals that, for generic mirror shapes, there is a single valid s_0 . This will be shown in Section 4.3.

4.2. Solution to the IVP

Rather than considering the more difficult case of PDEs, we start by fixing $y = y_0$, thus reducing the PDEs in (5) to a single ordinary differential equation (ODE) of the form

$$f(x, s) = \frac{ds}{dx} = \frac{-s n_x(x, s)}{\langle \mathbf{n}(x, s), \mathbf{v}(x) \rangle}. \quad (6)$$

Any solution to the PDEs restricted to the scan line $y = y_0$ is a solution to the ODE (6).

Existence and uniqueness of solutions to ODEs are much simpler problems than for PDEs. In particular, the Picard Lindelöf Theorem gives the required existence and uniqueness conditions. In our context, it can be stated as follows:

Theorem 4.2 Picard Lindelöf Theorem. *Let I_x and $I_s = [s_0 - \beta, s_0 + \beta]$ be closed intervals of the real line, and let $f : I_x \times I_s \rightarrow \mathbb{R}$ be a continuous function bounded by M . Suppose that $f(x, s)$ satisfies a Lipschitz condition in s , namely*

$$|f(x, s_1) - f(x, s_2)| \leq L|s_1 - s_2|$$

for all $(x, s_1), (x, s_2) \in I_x \times I_s$ and for some constant L (not depending on x). Then the initial value problem

$$\frac{ds}{dx} = f(x, s), \quad s(x_0) = s_0$$

for $x_0 \in I_x$ has a unique solution in $I_x \cap [x_0 - b, x_0 + b]$ where $b = \beta/M$.

The general statement, as well as the proof of Theorem 4.2 are given in [6]. To show that the solution to our problem is unique, we now show that the conditions of this theorem are satisfied in our context.

To apply Theorem 4.2, let $s_0 > 0$ and $I_s = [0, 2s_0]$ (i.e., $\beta = s_0$), and observe that $f(x, s)$ in Eq. (6) is defined for $x \in I_x$ and $s \in I_s$. Observe in particular that when $s = 0$ and \mathbf{n} is given by Eq. (1), $\langle \mathbf{n}, \mathbf{v} \rangle$ is nonzero (assuming that \mathbf{m} is not at the camera center). We now need to verify that $f(x, s)$ is bounded and satisfies a Lipschitz condition in s .

Bound for $f(x, s)$. Since f is continuous on a closed bounded set, it is bounded. However, we need to find the bound for $f(x, s)$. We first observe that $f(x, s)$ is unchanged if \mathbf{n} is multiplied by a constant factor, since it appears both in the numerator and denominator of Eq. (6). Therefore, we may assume that \mathbf{n} has unit norm. In this case, we have $n_x(x, s) \leq 1$ and thus $|f(x, s)| \leq s / \langle \mathbf{n}(x, s), \mathbf{v}(x) \rangle$. Now, since both $\mathbf{n}(x, s)$ and $\mathbf{v}(x)$ are continuous on $I_s \times I_x$, so is $|\langle \mathbf{n}(x, s), \mathbf{v}(x) \rangle|$; hence it is bounded below, and since it is nonzero, by the assumption that rays meet the mirror non-tangentially, $|\langle \mathbf{n}(x, s), \mathbf{v}(x) \rangle| \geq \delta$ for some value δ . For $s \in [0, 2s_0]$ and $\beta = s_0$, it then follows that

$$|f(x, s)| \leq M = 2s_0/\delta. \quad (7)$$

Lipschitz condition in s . By a standard result, if a function has bounded derivative, $|f'(s)| < L$, then f is Lipschitz with constant L . Here, $f(s)$ has the form $p(s)/q(s) = s / \langle \mathbf{n}(x, s), \mathbf{v}(x) \rangle$. Its derivative is continuous whenever $q(s)$ is non-zero. Being defined on a closed bounded domain, it is hence bounded. This applies to the present function for which the derivative is bounded by a constant independent of x , since $|\langle \mathbf{n}(x, s), \mathbf{v}(x) \rangle| \geq \delta$. Hence f satisfies a Lipschitz condition in s , as required.

Extension in one dimension. Having shown that $f(x, s)$ is bounded and satisfies a Lipschitz condition in s , Theorem 4.2 implies that there is a unique solution to the ODE on the interval $I_x \cap [x_0 - b, x_0 + b]$, where $b = \beta/M = s_0/M = \delta/2$. Hence there exists a unique solution $s(x)$ on the interval $[x_0 - \delta/2, x_0 + \delta/2]$ to the extent that this interval lies within the domain I_x of x . Since δ is a fixed value, independent of s_0 , this solution can be extended further by a distance of $\delta/2$ in either direction, and hence ultimately

to the whole interval I_x , by iteratively extending the previous solution. In particular, let $x_1 = x_0 + \delta/4$ be the new starting point with depth $s_1 = s(x_0 + \delta/4)$. By the previous argument, there exists a unique solution defined on the interval $[x_1 - \delta/2, x_1 + \delta/2]$, which must be an extension of the first solution, since they coincide at x_1 . To reach this conclusion, it is necessary only to show that $s_1 > 0$. Assuming that $s_1 \leq 0$, by the intermediate value theorem there exists a point x in the interval $[x_0, x_1]$ such that $|ds/dx| = |f(x, s(x))| = |s_1 - s_0|/|x_1 - x_0| \geq 4s_0/\delta = 2M$. This is contrary to hypothesis that $|f(x, s)| \leq M$, where $x \in [x_0 - \delta/2, x_0 + \delta/2]$, which thus shows that $s_1 > 0$. This implies that there is a unique solution to the IVP on the interval $[x_1 - \delta/2, x_1 + \delta/2]$. This extends the original solution up to $x_1 + \delta/2 = x_0 + 3\delta/4$. By iterating over this process, the solution can therefore be extended to the complete interval I_x .

Extending in both dimensions. The previous discussion shows that, by holding y_0 fixed, one can propagate in the x direction, to find a unique solution $s(x, y_0)$ that satisfies Eq. (3) on $I_x \times \{y_0\}$. Subsequently, starting at each point $s(x, y_0)$, for all $x \in I_x$, a unique solution to Eq. (4) can be obtained. This yields $s(x, y)$ for all $(x, y) \in I_x \times I_y$, which is uniquely determined by the initial condition $s(x_0, y_0) = s_0$. Therefore, the entire visible mirror surface can be reconstructed uniquely given the depth of one starting point.

Note that the solution obtained by this method is not guaranteed to satisfy both Eqs. (3) and (4) everywhere. It satisfies Eq. (4) everywhere, but Eq. (3) only along the line $y = y_0$. The essential point, however, is that if a full solution to the IVP (5) exists, then the described method must find it and it must therefore be unique. Alternatively, it is possible to propagate a solution in the y direction first, followed by x . If this gives a different result than propagating in x first, then no solution can exist for the full IVP (5). On the other hand, if both methods give the same solution, $s(x, y)$, then this solution satisfies both Eqs. (3) and (4) everywhere and hence is a solution to the IVP.

Numerically, this method can be implemented using a method for solving ODEs, such as a Runge-Kutta method [7], resulting in a solution if one exists.

4.3. Computing a Starting Depth

Although guaranteed to be unique (if it exists), the solution derived in the previous section relies on knowing the depth of one point. Hence, there is at most a one-parameter family of solutions to this problem, given known positions of the points $\mathbf{m}(x, y)$. In this section, we show that not all depths s_0 give a valid solution, and more specifically, that the valid depth of a starting point can be obtained uniquely.

Potentially, propagating a solution first in the x direction then y may give different results from propagating first in the y direction then x . In this case no solution exists to

satisfy the PDEs everywhere. One may speculate that this observation may rule out some values of the starting depth s_0 . This observation suggests an additional constraint to find the depth of the starting point: We want these two orders of integration (i.e., first x then y , or first y then x) to yield the same result. This motivates a consideration of the constraint

$$\frac{\partial^2 s}{\partial x \partial y} = \frac{\partial^2 s}{\partial y \partial x}. \quad (8)$$

Swapping the order of partial derivation is valid for any C^2 (twice continuously differentiable) function, which s is, by hypothesis. Given Eqs. (3) and (4), Eq. (8) can be further written as

$$\frac{\partial}{\partial x} f_y(x, y, s(x, y)) = \frac{\partial}{\partial y} f_x(x, y, s(x, y)). \quad (9)$$

In general, for arbitrary functions f_x and f_y , this equality will not hold. However, it does hold if f_x and f_y are C^1 continuous functions (i.e., the mirror is a C^2 continuous surface), and if $s(x, y)$ satisfies the PDEs in Eqs. (3) and (4). Note that if $s(x, y)$ satisfies Eq. 9, then it does **not** follow that s satisfies the PDEs; it is only a necessary condition.

Eq. (9) allows us to derive an equation involving only x , y , s , and \mathbf{m} . More specifically, given $f_x = \frac{-n_x s}{\langle \mathbf{v}, \mathbf{n} \rangle}$, we can write

$$\begin{aligned} \frac{\partial f_x}{\partial y} &= \frac{-\langle \mathbf{n}, \mathbf{v} \rangle (s \partial n_x / \partial y + n_x \partial s / \partial y)}{\langle \mathbf{n}, \mathbf{v} \rangle^2} \\ &+ \frac{s n_x (\langle \partial \mathbf{v} / \partial y, \mathbf{n} \rangle + \langle \mathbf{v}, \partial \mathbf{n} / \partial y \rangle)}{\langle \mathbf{n}, \mathbf{v} \rangle^2}. \end{aligned} \quad (10)$$

Continuing with only the numerator yields

$$-\langle \mathbf{n}, \mathbf{v} \rangle (s \partial n_x / \partial y + n_x \partial s / \partial y) + s n_x (n_y + \langle \mathbf{v}, \partial \mathbf{n} / \partial y \rangle).$$

Substituting $\partial s / \partial y$ according to Eq. (4) gives

$$-\langle \mathbf{n}, \mathbf{v} \rangle (s \partial n_x / \partial y - s n_x n_y / \langle \mathbf{n}, \mathbf{v} \rangle) + s n_x (n_y + \langle \mathbf{v}, \partial \mathbf{n} / \partial y \rangle).$$

Equating with the numerator of $\frac{\partial f_y}{\partial x}$ (because the denominators are the same) and cancelling s which is a common factor yields

$$\left\langle -\mathbf{n} \frac{\partial n_x}{\partial y} + n_x \frac{\partial \mathbf{n}}{\partial y} + \mathbf{n} \frac{\partial n_y}{\partial x} - n_y \frac{\partial \mathbf{n}}{\partial x}, \mathbf{v} \right\rangle = 0. \quad (11)$$

We substitute \mathbf{n} and its partial derivatives in Eq. (11) according to its definition in Eq. (1), which gives a polynomial of the form

$$(As^2 + Bs + C) + (as + b)\sqrt{ps^2 + qs + t} = 0, \quad (12)$$

where the square root is introduced by the term $\|\mathbf{m} - s\mathbf{v}\|$ and all the coefficients rely on the values (x, y) and $\mathbf{m}(x, y)$ at a given point and do not involve the unknown parameter

s . To obtain a solution for s , we multiply Eq. (12) with $(As^2 + Bs + C) - (as + b)\sqrt{ps^2 + qs + t}$ and simplify the resulting equation symbolically with the *Matlab Symbolic Toolbox*. To our surprise, the coefficients of the degree 4 and 3 terms disappear, which thus leaves us with a degree 2 polynomial equation, whose coefficients depend on x , y , \mathbf{m} and its partial derivatives w.r.t. x and y . Although this degree 2 polynomial has two solutions, for generic surfaces, only one of them satisfies Eq. (12). This is due to the fact that our strategy to remove the square root in (12) introduces an additional, invalid solution. Therefore, for generic surfaces, there exists a single valid s_0 , which, combined with the proof of uniqueness of the solution to the IVP, implies that the mirror surface can be reconstructed uniquely.

To conclude, this gives us two ways of computing the shape of the mirror: We can solve the polynomial equation at one image point, and then solve the two PDEs of the IVP (5) sequentially, or we can solve the polynomial equation at each image point. Note that both methods require dense reflection correspondences, both for integration purposes and to compute accurate partial derivatives of the reflection correspondences.

5. Shape Recovery from Sparse Measurements

While the previous section describes two possible ways of reconstructing the mirror surface, the solution to the polynomial equation, as well as the integration of the PDEs strongly rely on dense and noise-free correspondences. In a more realistic scenario, reflection correspondences will be sparse and noisy. In this section, we introduce an approach to performing reconstruction in these challenging conditions. We first present our surface parametrization, and then describe our reconstruction framework.

5.1. Surface Representation

In the formulation of Section 4, we directly modeled the surface in terms of the depth of image points. Since our proof of uniqueness relies on the surface being C^2 continuous, we can make use of other parametrizations that encode such a smoothness. In particular, here, we parametrize the depth s using a uniform cubic B-spline (UCBS). This implicitly satisfies the geometric constraint that surface points lie on their respective visual rays.

More specifically, let $s : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a function mapping the image point at $z = 1$ to a scalar value. For each point (x, y) on the image plane, s is defined as

$$s = \sum_{j=1}^{u_x} \sum_{k=1}^{u_y} C_{jk} N_j(x) N_k(y), \quad (13)$$

where C_{jk} are the unknown depth values of the *control points*, which are organized in a $u_x \times u_y$ grid over the image, and $N_j(x)$, $N_k(y)$ are the UCBS basis functions. For a

point (x, y) on the image plane, $N_j(x)$ and $N_k(y)$ are fully determined. We can re-write Eq. (13) as

$$s = \mathbf{w}\mathbf{c}, \quad (14)$$

where \mathbf{w} is the $1 \times u_x u_y$ vector of known basis functions, and \mathbf{c} is the $u_x u_y \times 1$ vector of unknowns. Therefore, a point on the surface $\mathbf{p} = \mathbf{w}\mathbf{c}(x, y, 1)^\top$. With this parametrization, we can derive the normal at \mathbf{p} from differential geometry as $\mathbf{n} = \frac{\partial \mathbf{p}}{\partial x} \times \frac{\partial \mathbf{p}}{\partial y}$, where

$$\begin{cases} \frac{\partial \mathbf{p}}{\partial x} = \frac{\partial s}{\partial x}(x, y, 1)^\top + s(1, 0, 0)^\top, \\ \frac{\partial \mathbf{p}}{\partial y} = \frac{\partial s}{\partial y}(x, y, 1)^\top + s(0, 1, 0)^\top. \end{cases} \quad (15)$$

This finally yields

$$\mathbf{n} = \left(\frac{\partial \mathbf{w}}{\partial x} \mathbf{c}, \frac{\partial \mathbf{w}}{\partial y} \mathbf{c}, -x \frac{\partial \mathbf{w}}{\partial x} \mathbf{c} - y \frac{\partial \mathbf{w}}{\partial y} \mathbf{c} - s \right)^\top, \quad (16)$$

where $\frac{\partial \mathbf{w}}{\partial x}, \frac{\partial \mathbf{w}}{\partial y}$ are derived from the UCBS basis functions.

5.2. Shape Recovery as an Optimization Problem

Given our parametric representation of the mirror surface, shape recovery reduces to estimating the depth of the control points \mathbf{c} . Given a set of sparse reflection correspondences between image points $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m\}$ and points on the reference plane $\{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_m\}$, we formulate reconstruction as the problem of finding the \mathbf{c} that minimizes the distance between \mathbf{m}_i and the backprojection of \mathbf{v}_i to the reference plane. In the remainder of this section, we first explain the backprojection process and then derive the resulting optimization problem.

As mentioned earlier, a 3D point on the mirror corresponding to image point (x_i, y_i) can be expressed as $\mathbf{p}_i = \mathbf{w}_i \mathbf{c}(x_i, y_i, 1)^\top$, with its normal given by Eq. (16). The corresponding reflected ray \mathbf{l}_i can be computed as $\mathbf{l}_i = \mathbf{i}_i - 2 \langle \tilde{\mathbf{n}}_i, \mathbf{i}_i \rangle \tilde{\mathbf{n}}_i$, where $\mathbf{i}_i = (x_i, y_i, 1)^\top / \|(x_i, y_i, 1)\|$ and $\tilde{\mathbf{n}}_i = \mathbf{n}_i / \|\mathbf{n}_i\|$. Points along this reflected ray can be represented as $\mathbf{a}_i = \mathbf{p}_i + \lambda_i \mathbf{l}_i$. The pose of the reference plane relative to the camera is determined by the rotation matrix \mathbf{R} and the translation vector \mathbf{T} . Let $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3)$, where $\mathbf{r}_j, j \in \{1, 2, 3\}$, denotes the j -th column of \mathbf{R} . The reference plane can be represented by the vector $\mathbf{q} = (\mathbf{r}_3^\top, -\mathbf{r}_3^\top \mathbf{T})^\top$, such that $\langle \mathbf{q}, (\mathbf{a}^\top, 1) \rangle = 0$ for any point \mathbf{a} on the reference plane. Backprojection is achieved by computing the intersection of the reflected ray with the reference plane. This intersection can be computed as $\hat{\mathbf{m}}_i = \mathbf{p}_i - \frac{\langle \mathbf{r}_3, \mathbf{p}_i \rangle - \mathbf{r}_3^\top \mathbf{T}}{\langle \mathbf{r}_3, \mathbf{l}_i \rangle} \mathbf{l}_i$.

We estimate the parameter \mathbf{c} of the surface by solving a non-linear least-squares problem that minimizes the error between our backprojections to the plane and the real points on the plane. This can be written as

$$\min_{\mathbf{c}} \sum_{i=1}^m \|\hat{\mathbf{m}}_i(\mathbf{c}) - \mathbf{m}_i\|. \quad (17)$$




Figure 3. Initialization of the unknown mirror.

Note that Section 4 shows that only the correct mirror surface corresponds to reflected rays that intersect the reference plane at the observed points. The problem in Eq. (17) can be solved using an iterative scheme such as the Levenberg-Marquardt method.

Initialization. The optimization problem in Eq. (17) is non-convex and therefore requires initialization. In our experiments, we use the following strategy, depicted by Fig. 3: We initialize the unknown mirror P as a plane, and seek for its pose such that the camera can best see the reflection of the reference plane Q . With an actual planar mirror, P is the bisector of the angle between Q and its virtual image Q' , which is related to Q by a rigid transformation. To achieve full visibility of Q , P should be such that Q' is parallel to the image plane. We therefore search for a translation T_v of Q' that minimizes the squared distance between the correspondence points \mathbf{m}'_i on Q' and the intersection of the visual rays with Q' , which can be expressed as the least-squares solution to a linear system. Given this translation, we compute the pose of P by bisecting the angle between Q and Q' . In practice, the mirror is not planar. Since the divergence, resp. convergence, of the reflection rays with a convex, resp. concave, mirror, T_v defines an upper bound for a convex mirror and a lower bound for a concave one. Therefore, we search for a translation of Q' within the bounds $[0 T_v]$, or $[T_v -3T_v]$, and take the mirror shape P that gives the smallest energy value in Eq. (17) after optimization.

6. Experiments

We now demonstrate the effectiveness of our approach on synthetic and real surfaces.

6.1. Synthetic Data

For our synthetic experiments, we used a UCBS, an ellipsoid and a sphere as mirror surfaces. The reflection correspondences were obtained by backprojecting all image pixels to the reference plane. More than $4M$ reflection correspondences were used in our synthetic experiments. Recall that we have 3 possible ways of reconstructing the surface: With dense correspondences, we can either solve a degree 2 polynomial equation for each pixel, or solve this equation for a single pixel and solve the IVP (5). With sparse corre-




Figure 4. **Shapes obtained by solving polynomials and by integration.** Red dots denote the surface reconstructed by solving polynomial equations, Cyan dots the surface obtained by solving the PDEs in Order A, and Magenta dots in Order B. Blue dots denote ground truth.




Figure 5. **Reconstruction and self-consistency errors as a function of the density of noisy data:** We show the average errors, as well as standard deviations, over 50 runs with different random noise.

spondences, we solve the optimization problem in Eq. (17). We now present results for these 3 methods. Due to space limitation, results for the sphere, as well as additional results for the UCBS and ellipsoid are provided in supplementary material.

Solving Polynomials and Integrating PDEs. Fig. 4 compares with ground truth the shapes obtained by solving either a polynomial equation at each pixel independently, or the IVP. As mentioned in Section 4.2, the PDEs can be solved in two different orders. *Order A*: Solve for x with $y = y_0$, then solve for y for fixed values of x , and *Order B*: Solve in the opposite order. Therefore, in Fig. 4, we show the results obtained with these 2 orders. Note that all the shapes in Fig. 4 are essentially identical. This was to be expected, since with dense, noise-free correspondences, the shape is unique. We then added zero-mean Gaussian noise with one pixel standard deviation. Fig. 5 depicts the reconstruction and self-consistency errors as a function of the percentage of correspondences used for reconstruction for the approach by solving IVP. The reconstruction error is computed as the mean 3D point-to-point distance between the reconstructed shape and ground-truth. The self-consistency error is defined as the mean 3D point-to-point distance between the shapes obtained with order A and order B. These errors were averaged over 50 runs with different random noise. We observed from experiments that the approach by solving polynomial equations is more sensitive to noise and data density than that by solving the IVP.

Solving an Optimization Problem. Fig. 6 depicts our results obtained by solving the optimization problem of Eq. (17) with sparse, noisy data. As before, we used one pixel for the noise standard deviation, and only kept 10% of the correspondences. Note that our reconstructions are accurate. In the case of the ellipsoid, representing the surface as a UCBS introduces approximation errors, depending on the number of control points used. Fig. 7 shows the influence of the number of control points on the objective




Figure 6. **Reconstructions by optimization with noisy data.** Red dots show our reconstructions, and Blue dots ground truth.

function value, the reconstruction error and the runtime. In general, the approximation error decreases as the number of control points increases, but the computational complexity increases accordingly. Note, however, that there is a range of values for which the approximation error is small and the runtime reasonable.

6.2. Experiments on Real Data

To evaluate our approach on real surfaces, we used the stainless steel spoon and gravy boat depicted in Fig. 1. We acquired images with a Sony 1920×1080 HDR-XR200 digital camera and used an LCD monitor as reference plane, on which we displayed patterns that let us extract correspondences following the strategy of [11]. This yielded 1919 and 2029 correspondences for the spoon and gravy boat, respectively. The pose of the reference plane relative to the camera was calibrated with the Matlab Calibration Toolbox.

Since we only had sparse correspondences, we relied on our optimization framework to perform reconstruction. The spoon and gravy boat were approximated by a UCBS with 20×20 and 10×10 control points, respectively. In both cases, the mirror was initialized using the same strategy as for synthetic experiments. In Fig. 8, we show our reconstructions, as well as histograms of objective function values. Note that the reconstructed shapes correspond to those of the original objects shown in Fig. 1, and that most of the objective function values are small, suggesting accurate reconstruction.

7. Conclusion

In this paper, we have addressed the problem of mirror surface reconstruction from a single image, and provided a




Figure 7. **Influence of control points for ellipsoid reconstruction.** From left to right, mean objective function value over the pixels, reconstruction error and runtime as a function of the number of control points used to approximate the ellipsoid. To really evaluate the approximation error, these results were obtained with noise-free data. Standard deviations were computed over the image pixels.

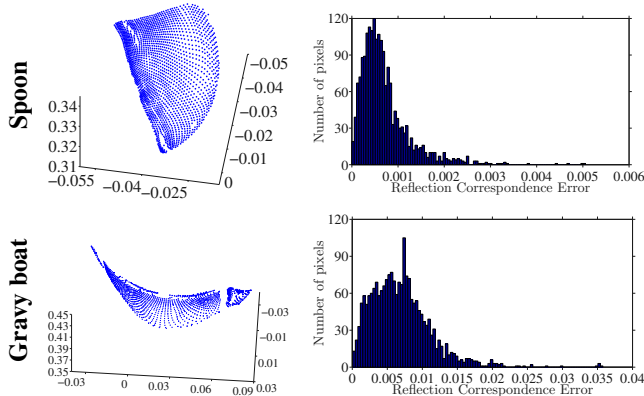


Figure 8. **Reconstructions from real images.** Reconstruction of the spoon and gravy boat depicted in Fig. 1. The estimated surfaces visually match the shape of the objects. The majority of low objective function values in the histograms suggests accurate reconstruction.

theoretical proof of uniqueness of the solution in the presence of dense reflection correspondences. Furthermore, we have introduced an optimization framework to reconstruct the mirror surface when only sparse correspondences are available. Currently, our approach requires the pose of the reference plane to be known. In the future, we will study the case where this pose is unknown, and optimized in the problem of Eq. (17). Our initial experiments in this direction have shown that, for generic surfaces, the Jacobian matrix of the energy around a local minimum has full rank. This indicates that only a finite number of discrete solutions exist, as opposed to a continuous family of solutions. Finding the global minimum among these discrete solutions, as well as studying the impact of our shape representation on the energy will be the focus of our future research.

References

- [1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Toward a theory of shape from specular flow. In *ICCV*, 2007.
- [2] J. Balzer, S. Höfer, and J. Beyerer. Multiview specular stereo reconstruction of large mirror surfaces. In *CVPR*, 2011.
- [3] T. Bonfort and P. Sturm. Voxel carving for specular surfaces. In *ICCV*, 2003.
- [4] T. Bonfort, P. Sturm, and P. Gargallo. General Specular Surface Triangulation. In *ACCV*, 2006.
- [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 *Proc. ICCV*, 2009.
- [6] N. L. Earl A. Coddington. *Theory of ordinary differential equations*. McGraw-Hill, 1955.
- [7] D. F. M. Endre Sli. *An Introduction to Numerical Analysis*. Cambridge University Press, 2003.
- [8] T. A. Fleming, R. W. and E. H. Adelson. Specular reflections and the perception of shape. *Journal of Vision*, 2004.
- [9] M. A. Halstead, B. A. Barsky, S. A. Klein, and R. B. Mandell. Reconstructing curved surfaces from specular reflection patterns using spline surface fitting of normals. In *SIGGRAPH*, 1996.
- [10] K. N. Kutulakos and E. Steger. A theory of refractive and specular 3d shape by light-path triangulation. *IJCV*, 2008.
- [11] M. Liu, K.-Y. K. Wong, Z. Dai, and Z. Chen. Pose estimation from reflections for specular surface recovery. In *ICCV*, 2011.
- [12] D. Nehab, T. Weyrich, and S. Rusinkiewicz. Dense 3d reconstruction from specular consistency. In *CVPR*, 2008.
- [13] M. Oren and S. K. Nayar. A theory of specular surface geometry. *IJCV*, 1997.
- [14] S. Roth and M. J. Black. Specular flow and the recovery of surface structure. In *CVPR*, 2006.
- [15] S. Rozenfeld, I. Shimshoni, and M. Lindenbaum. Dense mirroring surface recovery from 1d homographies and sparse correspondences. In *CVPR*, 2007.
- [16] A. C. Sankaranarayanan, A. Veeraraghavan, O. Tuzel, and A. Agrawal. Specular surface reconstruction from sparse reflection correspondences. In *CVPR*, 2010.
- [17] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. *IJCV*, 2005.
- [18] M. Tarini, H. P. A. Lensch, M. Goesele, and H.-P. Seidel. 3d acquisition of mirroring objects using striped patterns. *Graphical Models*, 2005.
- [19] Y. Vasilyev, T. Zickler, S. J. Gortler, and O. Ben-Shahar. Shape from specular flow: Is one flow enough? In *CVPR*, 2011.
- [20] A. Zisserman, P. Giblin, and A. Blake. The information available to a moving observer from specularities. *IVC*, 1989.