Dense mirroring surface recovery from 1D homographies and sparse correspondences

Stas Rozenfeld, *Student Member, IEEE,* Ilan Shimshoni, *Member, IEEE,* and Michael Lindenbaum, *Member, IEEE*

***Abstract*— In this work we recover the 3D shape of mirrors, sunglasses, and stainless steel implements. A computer monitor displays several images of parallel stripes, each image at a different angle. Reflections of these stripes in a mirroring surface are captured by the camera. For every image point, the direction of the displayed stripes and their reflections in the image are related by a 1D homography matrix, computed with a robust version of the statistically accurate heteroscedastic approach. By focusing on a sparse set of image points for which monitor–image correspondence is computed, the depth and the local shape may be estimated from these homographies. The depth estimation relies on statistically correct minimization and provides accurate, reliable results. Even for the image points where the depth estimation process is inherently unstable, we are able to characterize this instability and develop an algorithm to detect and correct it. After correcting the instability, dense surface recovery of mirroring objects is performed using constrained interpolation, which does not simply interpolate the surface depth values but uses the locally computed 1D homographies to solve for the depth, the correspondence, and the local surface shape. The method was implemented and the shape of several objects was densely recovered at sub-millimeter accuracy.**

***Index Terms*— Mirroring objects, 3D Shape reconstruction, 1D homographies, stability**

I. INTRODUCTION

The problem of 3D shape reconstruction has been the focus of extensive study for many years. Many methods were developed for nonspecular objects such as structured light or stereo vision. Unfortunately, all those methods are inapplicable to both specular and to mirroring surfaces. As such surfaces are common in everyday life, this is a serious drawback, one that motivated a number of research efforts over the last few years.

The problem of specular surface reconstruction was first addressed by Blake and Brelstaff [3], who proposed a system that can recover the depth map and orientation of a specular surface, when enough Lambertian reference points exist on it. Several approaches focus on the reconstruction of surface curves. Zisserman, Giblin, and Blake [20] track the motion of specularities to obtain information on the surface. Oren and Nayar [10] show how to discriminate between Lambertian and specular (virtual) features. They track the specular features and, using known camera motion, recover the depth values of the surface at their traces (but not the whole surface). Known camera motion has also been used by Solem et al. in [17], where the full surface is reconstructed using a level-set based energy minimization technique that relies on surface smoothness. Park and Cho [11] proposed a system based on multiple reflections of a laser beam between the examined object and a reflecting sphere containing it. The 3D points are recovered one point at a time. Zheng and Murata [19] recover the whole shape of a rotating specular object by tracking the specularities created by a toroidal light source.

Several approaches focus on particular object classes. Ripsman and Jenkin [13] recover planar specular objects using a three camera system. Halstead et al. [5] recover roughly symmetric surfaces from the reflection of a light pattern containing concentric circles. A medical application for shape recovery of the human cornea was also demonstrated.

Other approaches attempt to recover a general object without moving it. In [2], Baba et al. proposed a laser scanning system which is able to scan specular surfaces as well as Lambertian ones. Unfortunately, the system works very slowly (it takes about 20 minutes to acquire 100 measurements). Bonfort and Sturm [4] introduce a local and effective process which recovers the shape of the surface using two cameras that observe the distorted reflections of images displayed on a monitor. A similar experimental setup was used by Knauer et al. in [6], where normals to the surface are recovered first and then used to recover the depth map. The calibration necessary for this technique is challenging. A method proposed by Tarini et al. [18] uses a single camera that observes the reflection of several images displayed on a monitor. It establishes dense correspondence between the image and the monitor using a color based process that relies on the uniformity of certain photometric properties. This correspondence provides local constraints on the depths and the surface normals. They are then integrated into a full surface estimate by global optimization, using smoothness assumptions. Recently, approaches have been proposed to estimate the shape of specular objects from optical flow generated by changes in the object's position with respect to the environment [1], [8] or from known camera motion where the object remains stationary [14].

The method we propose is closest to the approach suggested by Savarese, Chen, and Perona [16], and uses some of their results. Relying on a single pattern containing intersecting lines, they recover the surface depth as well as its higher order properties at a sparse set of points where at least three lines intersect and for which correspondence is available. Their process is fully local: using differential geometry to analyze the curve distortion leads to the construction of a matrix parameterized by the unknown depth. This matrix becomes degenerate for the correct depth.

In the proposed system a computer monitor displays several images of parallel stripes, each image at a different angle. Distorted reflections of these stripes in a mirroring surface are captured by the camera. For every image point, the directions of the displayed stripes and their reflections in the image are

S. Rozenfeld is with the Faculty of Industrial Engineering and Management, Technion, Israel Institute of Technology, Haifa, Israel, 32000. E-mail: rozefeld@tx.technion.ac.il

I. Shimshoni is with the Department of Management Information Systems, University of Haifa, Haifa, Israel, 31905. E-mail: ishimshoni@mis.hevra.haifa.ac.il

M. Lindenbaum is with the Computer Science Department, Technion, Israel Institute of Technology, Haifa, Israel, 32000. E-mail: mic@cs.technion.ac.il

related by a 1D homography, which differs for every image point. The direction of the stripes on the monitor are known and the direction of the reflected distorted stripes are estimated from the captured images. The homography is estimated using robust and statistically valid (heteroscedastic) methods (see [7] for related work), without any knowledge about correspondence between the monitor point and the camera points. Such correspondence is extracted only for a sparse set of points using the pseudo-random color pattern proposed by Morano [9]. Focusing on this sparse set and relying on the corresponding homography, we calculate the depth and first-order local shape by minimizing a statistically correct measure. Then dense surface recovery is performed using constrained interpolation, which does not simply interpolate the surface depth values, but rather uses the locally computed 1D homography to solve for the depth, the correspondence, and the local surface shape.

The cost function used for the proposed local shape recovery method as well as the one described in [16] may be insensitive to depth on a small part of the surface. This makes both methods inherently unstable there. We quantify the instability by measuring the derivative of the cost function and use it as a weight for a smoothing process. In addition, an algebraic constraint sufficient for the instability is formulated and a simple geometrical interpretation is given to this constraint for the case of planar and spherical surfaces.

The method was implemented and the shapes of a mirror, sunglasses, and a stainless steel ashtray were recovered. For the planar mirror, the estimated depths were consistent with a plane to sub-millimeter accuracy.

In short, this paper proposes a new method for shape recovery of mirror-like objects. The main innovations of the proposed method are:

1) A clear characterization of the information provided by the optical distortion about the surface depth and local shape using 1D homographies.
2) A statistically correct error measure that is used for local depth estimation. This criterion replaces the algebraic expression suggested in [16]. Besides being more statistically justified, this criterion is numerically stable and avoids the phantoms reported in [16]. Moreover, the new method is easier to minimize due to the cost function shape.
3) A method for dense depth estimation which uses the locally computed homographies and is therefore much more accurate than simple depth interpolation.
4) A characterization of the inherent limitations of depth recovery, a description of the loci of instability, and a practical method to circumvent this difficulty.
5) A robust and statistically correct technique that uses the heteroscedastic model [7].

This paper continues as follows. A basic mathematical analysis of the problem is presented in Section II. The detailed local recovery algorithm is described in detail in Section III. The instability problem is characterized and treated in Section IV. The second part of the algorithm, which provides the dense depth map, is described in Section V. A specific simpler algorithm for the special case of a planar mirror surface is provided in Section VI. Experimental results are presented in Section VII and some conclusions are suggested in Section VIII. Some of the more technical derivations are deferred to the appendices. A short version of this paper was presented in [15].

II. NOTATIONS AND MATHEMATICAL BACKGROUND

We start with notations and some previous results (from [16]) to be used in this paper. Consider a monitor that displays a known image and a camera that captures its reflection in a mirroring surface; see Figure 1 (left).

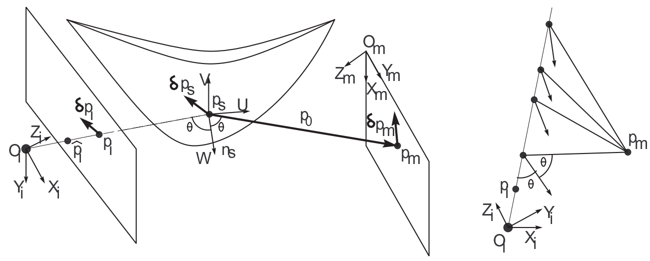

Fig. 1. Optical reflection geometry (left), and surface normal and depth correspondence (right).

Let $\boldsymbol{p_i}$ be an image point corresponding to the point $\boldsymbol{p_m}$ on the reference plane (monitor) and to its reflectance point $\boldsymbol{p_s}$ (on the surface to be reconstructed). Our goal is to estimate the depth s at this point, such that $s\hat{p}_i = \boldsymbol{p_s}$, where $\hat{\boldsymbol{p}}_{\boldsymbol{i}} = \boldsymbol{p_i}/\|\boldsymbol{p_i}\|$.

A basic rule of optics states that $\boldsymbol{n_s}$, the normal at $\boldsymbol{p_s}$, lies in the plane specified by $\boldsymbol{p_i}, \boldsymbol{p_m}$, and the camera's optical center $\boldsymbol{O_i}$. Moreover, $\boldsymbol{n_s}$ is the bisector of the angle $\angle \boldsymbol{O_i p_s p_m}$. Denote half of this angle as θ. Let $\boldsymbol{V}$ be the normal to this plane. Let $\{\boldsymbol{O_i}, [\boldsymbol{X_i}, \boldsymbol{Y_i}, \boldsymbol{Z_i}]\}$be the camera coordinate system specified relative to the camera's optical center, $\{\boldsymbol{p_m}, [\boldsymbol{X_m}, \boldsymbol{Y_m}, \boldsymbol{Z_m}]\}$the monitor coordinate system, and $\{\boldsymbol{p_s}, [\mathbf{U} = \mathbf{V} \times \boldsymbol{n_s}, \mathbf{V}, \mathbf{W} = \boldsymbol{n_s}]\}$ the (local) reflective surface coordinate system. The first coordinate system is chosen as our reference frame. The second coordinate system is estimated using an external calibration process that finds the coordinates' transformation from the camera to the monitor. The point $\boldsymbol{p_m}$ is specified by a correspondence process, which specifies $\boldsymbol{V}$ and a one-to-one relationship between s and $\boldsymbol{n_s}$, as illustrated in Figure 1 (right). For a known s, the third coordinate system is also fully specified. Let $\boldsymbol{p_0} = \boldsymbol{p_m} - \boldsymbol{p_s}$, be the vector from the reflecting point to the original monitor point. Most of our calculations are done in the surface coordinate system.

Consider a line in the monitor plane emanating from $\boldsymbol{p_m}$, $\boldsymbol{p_m}(t) = \boldsymbol{p_m} + t\boldsymbol{\delta p_m},\ t \in R$ where $\boldsymbol{\delta p_m} = \cos\alpha_m \boldsymbol{X_m} + \sin\alpha_m \boldsymbol{Y_m}$. The line $\boldsymbol{p_m}(t)$ induces, in a natural way, a curve $\boldsymbol{p_s}(t)$ on the mirroring surface and another curve $\boldsymbol{p_i}(t)$ in the image plane. Note that if the mirroring surface is a plane, then $\boldsymbol{p_s}(t)$ and $\boldsymbol{p_i}(t)$ are straight lines and not curves. Clearly, $\boldsymbol{p_s}(0) = \boldsymbol{p_s}$ and $\boldsymbol{p_i}(0) = \boldsymbol{p_i}$. Let $\boldsymbol{\delta p_s} = \dot{\boldsymbol{p_s}}(0)$ and $\boldsymbol{\delta p_i} = \dot{\boldsymbol{p_i}}(0)$ be the derivatives of $\boldsymbol{p_s}(t)$ at $\boldsymbol{p_s}$ and of $\boldsymbol{p_i}(t)$ at $\boldsymbol{p_i}$, respectively. Let α_i be the angle for which $\boldsymbol{\delta p_i} = \|\boldsymbol{\delta p_i}\|(\cos\alpha_i \boldsymbol{X_i} + \sin\alpha_i \boldsymbol{Y_i})$. The vector $\boldsymbol{\delta p_s}$ lies in the tangent plane of the surface at $\boldsymbol{p_s}$. Therefore, there is an angle α_s such that $\boldsymbol{\delta p_s} = \|\boldsymbol{\delta p_s}\|(\cos\alpha_s \boldsymbol{U} + \sin\alpha_s \boldsymbol{V})$. Finally, let a, b and c be the parameters describing the second order approximation $w = \frac{1}{2}au^2 + cuv + \frac{1}{2}bv^2$ of the mirroring surface close to $\boldsymbol{p_s}$ in its local coordinate frame.

The following expression, specifying a linear transformation between tangent vectors, follows from the results described in [16]:

$$\|\boldsymbol{\delta p_s}\| \begin{bmatrix} \cos\alpha_s \\ \sin\alpha_s \end{bmatrix} = \mathcal{A}\mathcal{B} \begin{bmatrix} \cos\alpha_m \\ \sin\alpha_m \end{bmatrix}, \tag{1}$$

where

$$\begin{aligned}
&\Delta = (J_u - 2a\cos\theta)(J_v - 2b\cos\theta) - (2c\cos\theta)^2,\\
&J_u = \cos^2\theta \tfrac{s+\|\mathbf{p}_0\|}{s\|\mathbf{p}_0\|},\\
&J_v = \tfrac{s+\|\mathbf{p}_0\|}{s\|\mathbf{p}_0\|},\\
&\mathcal{A} = \tfrac{1}{\Delta}\begin{bmatrix} J_v - 2b\cos\theta & 2c\cos\theta \\ 2c\cos\theta & J_u - 2a\cos\theta \end{bmatrix},\\
&\mathcal{B} = \tfrac{1}{\|\mathbf{p}_0\|}\begin{bmatrix} -\cos^2\theta & 0 & \cos\theta\sin\theta \\ 0 & -1 & 0 \end{bmatrix} \times\\
&\qquad [(\boldsymbol{X_m})^{\boldsymbol{s}}, (\boldsymbol{Y_m})^{\boldsymbol{s}}],
\end{aligned} \tag{2}$$

and $(\boldsymbol{X_m})^{\boldsymbol{s}}$ and $(\boldsymbol{Y_m})^{\boldsymbol{s}}$ are the column vectors $\mathbf{X}_m, \mathbf{Y}_m$ expressed in the surface coordinate system.

III. Local shape from 1D homographies

The linear transformation (1), denoted T_{ms}, clearly induces a 1D homography between the monitor angles and the corresponding surface angles. Denote this homography H_{ms}. Consider another 1D homography between the local surface tangent-plane angles and the image plane angles. This homography, H_{si}, can be computed from $\boldsymbol{U}$ and $\boldsymbol{V}$ as follows:

$$H_{si} \cong \begin{bmatrix} \boldsymbol{U}_x - (\boldsymbol{p_i})_x \boldsymbol{U}_z & \boldsymbol{V}_x - (\boldsymbol{p_i})_x \boldsymbol{V}_z \\ \boldsymbol{U}_y - (\boldsymbol{p_i})_y \boldsymbol{U}_z & \boldsymbol{V}_y - (\boldsymbol{p_i})_y \boldsymbol{V}_z . \end{bmatrix} \tag{3}$$

(see Appendix II). The combined homography

$$H_{mi} = H_{si} H_{ms} \cong H_{si}\mathcal{A}\mathcal{B} \tag{4}$$

relates the monitor angles and the image plane angles. The components of this homography (and thus the homography itself) can be computed from the depth s, the image-monitor correspondences between $\boldsymbol{p_i} = [(\boldsymbol{p_i})_x, (\boldsymbol{p_i})_y]^T$ and $\boldsymbol{p_m}$, and the local shape parameters a, b and c. The same homography H_{mi} may also be computed empirically, without knowing any of these parameters, but using only angle correspondences. Let H_{mi}^{emp} denote this estimated homography. Comparing the two versions of the same homography provides a constraint that can be used to estimate the depth.

To estimate H_{mi}^{emp} at some image point, we need to establish monitor–image angle correspondences at this point. A straightforward approach would rely on first establishing correspondence between monitor and image points. We take a different approach, which provides the homographies at any point in the image without location correspondences. We use a sequence of K displayed images, each associated with a single tangent direction α_m^k, common to all points in the displayed image; see Figure 2. Let $\{\alpha_i^k\}_{k=1}^K$ be the sequence of corresponding tangent angles associated with a particular image point. Now, using the sequence of angle correspondences $\{\alpha_m^k, \alpha_i^k\}_{k=1}^K$, H_{mi} can be estimated by combining a robust method with the heteroscedastic technique that we developed for 1D homography estimation; see Appendix III. This method provides results superior to standard linear regression techniques such as SVD. Note that H_{mi}^{emp} may be calculated independently for every point in the image.

A. Local depth recovery

We now describe how to estimate the depth s at a point by using the empirical homography estimate H_{mi}^{emp} and the analytical constraints 1,3, and 4 on H_{mi}.

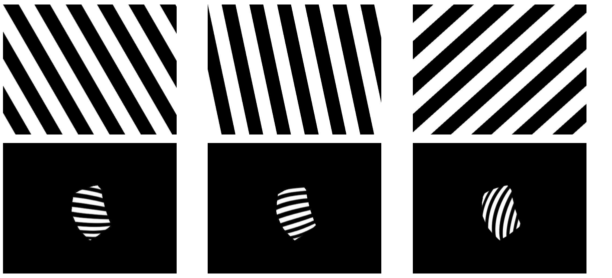

Fig. 2. Typical displayed images (top), and their corresponding captured images (bottom). Such pairs are used for the 1D homography estimation.

A straightforward approach would be to estimate s, and the local shape parameters a, b and c, in one optimization process. Instead, we will define a simpler optimization process where the score function depends on s only and relies on the symmetry of $\mathcal{A}$ (2).

Given a suggested value of s, the homography $H_{mi}(s)$ is constrained to be of the form

$$H_{mi}(s) \cong H_{si}(s)\mathcal{A}\mathcal{B}(s), \tag{5}$$

where $H_{si}(s)$ and $\mathcal{B}(s)$ are known because they are functions of s only. The matrix $\mathcal{A}$ associated with H_{mi}^{emp} may be estimated by

$$\widehat{\mathcal{A}}(s) = H_{si}^{-1}(s) H_{mi}^{emp} \mathcal{B}^{-1}(s) = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}. \tag{6}$$

$\mathcal{A}$ depends on the unknown parameters $a, b,$ and c and may be considered as an arbitrary symmetric matrix. Imposing this constraint yields

$$\widehat{\mathcal{A}}(s)_{sym} = \begin{bmatrix} a_{11} & (a_{12}+a_{21})/2 \\ (a_{12}+a_{21})/2 & a_{22} \end{bmatrix}.$$

$\widehat{\mathcal{A}}(s)_{sym}$ is closest to $\widehat{\mathcal{A}}(s)$ in the Frobenius norm sense. When the suggested value of s is close to the correct one, we expect that $\widehat{\mathcal{A}}(s)$ is close to $\mathcal{A}$ (up to a multiplicative constant), and that the symmetrization improves the accuracy of $\widehat{\mathcal{A}}(s)$.

Using $\widehat{\mathcal{A}}(s)$, we may now calculate the corresponding image plane angles tangent to the stripe:

$$\begin{bmatrix} \cos(\widehat{\alpha_i}^k(s)) \\ \sin(\widehat{\alpha_i}^k(s)) \end{bmatrix} \cong \widehat{H}_{mi}(s) \begin{bmatrix} \cos(\alpha_m^k) \\ \sin(\alpha_m^k) \end{bmatrix},$$

where $\widehat{H}_{mi}(s) = H_{si}(s)\widehat{\mathcal{A}}(s)_{sym}\mathcal{B}(s)$.

The penalty $f(s)$ is now specified as the sum of squared differences between the angles measured in the image and the depth dependent estimates:

$$f(s) = \sum_{k=1}^{K} (\widehat{\alpha_i}^k(s) - \alpha_i^k)^2. \tag{7}$$

Note that this function compares two estimates of the same meaningful quantities and should therefore perform well. If we make the common assumptions that the errors in the angle estimations are independent and identically normally distributed, then minimizing (7) is equivalent to solving a maximum likelihood problem. An alternative approach, proposed in [16], relies on similar constraints but does not use a geometrically meaningful interpretation (i.e., the effect of the 1D homography on the monitor angles). Instead these constraints lead to a set of homogeneous equations. The proposed penalty function relies on the rank deficiency of the coefficient matrix. In an earlier version of this work [15], we improved the method for estimating $\widehat{H}_{mi}(s)$. More importantly, we proposed an alternative criterion that measured the angle between H_{mi}^{emp} and an improved estimate of $\widehat{H}_{mi}(s)$, when they are regarded as vectors. This is an improvement over [16], as the penalty function measures a geometrically meaningful distance between $\widehat{H}_{mi}(s)$ and H_{mi}^{emp} but does not have a clear statistical meaning as the function suggested here. Two typical examples of the three penalty functions are shown in Figure 3. Note that unlike the rank deficiency criterion, the penalty functions we propose have a single minimum. A single basin of attraction covers the whole range, eliminating the phantom solutions phenomenon reported in [16].

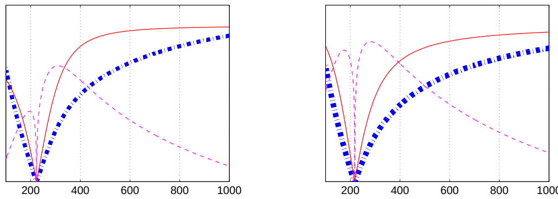

Fig. 3. The proposed penalty function (thick line) vs. the penalty function used in [16] (dashed line) and in [15] (solid line). The two plots correspond to two different image points.

B. *Local recovery of the second order surface parameters*

We now turn our attention to the estimation of the second order surface parameters a, b and c. This requires that we estimate the missing scale parameter to transform H_{ms} into the linear transformation T_{ms}. We do this with the scale-based method of [16]. The relationship between the length of a short vector on the monitor and its projection on the plane tangent to the surface is used to recover the missing scale parameter. The points on this plane are inferred by backprojecting the image points. After T_{ms} has been estimated, $\mathcal{A}$ can be recovered, and can then be used for the estimation of a, b and c.

IV. Inherent depth recovery instability and its correction

The surface reconstruction process is initialized by estimating the depth of a few initial points. These points, denoted $\{\boldsymbol{p_i^l}\}_{l=1}^{L}$, are those for which corresponding monitor points $\{\boldsymbol{p_m^l}\}_{l=1}^{L}$ are available. For each one of these points, the depth s is estimated using the procedure described in Section III-A.

A typical initial point reconstruction associated with a smooth surface is shown in Figure 4 (left). Clearly, the reconstruction is unreliable in the vicinity of some curve. This problem, which was already observed in [16], occurs in locations where the homography H_{mi} is highly insensitive to depth. In such cases, the inevitable inaccuracy in the image data (image point locations $\boldsymbol{p_i}$ and their corresponding angles α_i^k) is significantly amplified. See Figure 5 for two examples of the cost function (7) associated with a stable and an unstable point. We refer to this phenomenon as the *stability problem*.

Fortunately, it is easy to identify the unstable points by examining the absolute value of the numerical derivative of the cost function at the found minimum $\mathfrak{s}$:

$$\epsilon = \frac{1}{2}\left(\left|\frac{f(\mathfrak{s}+\Delta s)-f(\mathfrak{s})}{\Delta s}\right| + \left|\frac{f(\mathfrak{s})-f(\mathfrak{s}-\Delta s)}{\Delta s}\right|\right).$$

Figure 4 (right) shows that the inaccurately reconstructed points are exactly those which are associated with low *stability measure* values.

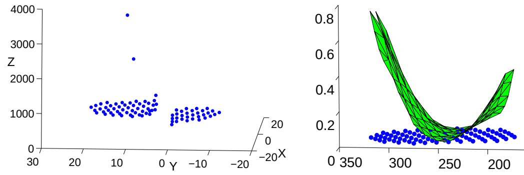

Fig. 4. Depth estimates at the initial point set obtained by the local optimization (left). The stability measure evaluated at the same points (right). Note that it is very low for the points where the calculation is inaccurate.

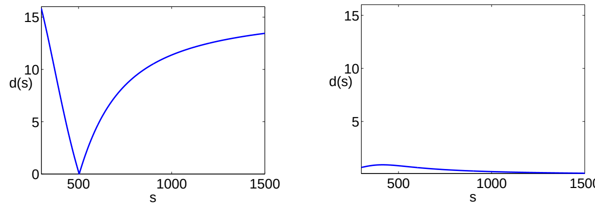

Fig. 5. The cost function associated with a stable point (left) and a problematic, unstable point (right).

A. *Algebraic and geometric conditions of instability.*

In general, the analysis of the loci of instability is quite complex and depends on the local surface geometry as well as the camera and monitor locations. We present here a general condition for instability and show that it leads to a simple geometric interpretation for loci of instability for the cases of planar and spherical surfaces.

Recall that in the absence of noise, $\widehat{\mathcal{A}}(s)$ will be symmetric for the correct value of s. Our depth estimation process relies on the symmetrization of $\widehat{\mathcal{A}}(s)$. Instability occurs when $\widehat{\mathcal{A}}(s)$ is symmetric for additional values of s. Technically, in the absence of noise, when $\widehat{\mathcal{A}}(s)$ is symmetric then $\widehat{\mathcal{A}}(s) = \widehat{\mathcal{A}}_{sym}(s)$, resulting in $\widehat{H}_{mi}(s) = H_{mi}^{emp} = H_{mi}$. This implies that the cost function $f(s) = 0$ for all values of s where $\widehat{\mathcal{A}}(s)$ is symmetric. Therefore, when $\widehat{\mathcal{A}}(s)$ is symmetric for all values of s, the depth cannot be estimated. The presence of noise does not help.

In Appendix I we consider the special case when the second order parameter $c = 0$. In this case we show that instability occurs when the local surface coordinate system axis $\boldsymbol{V}$ is parallel to the monitor. In general, the value of c depends both on the surface local geometry and its position relative to the camera. For planes and spheres, c is identically zero everywhere in all cases.

It is interesting to note that the condition on $\boldsymbol{V}$ is not general and does not hold when $c \neq 0$. To illustrate the condition's validity, we experiment with three synthetic surfaces: a plane, a sphere, and a cylinder. For each type of surface, Figure 6 shows a vertical triplet of images. The top image shows a captured image of stripes, with the red curve $C_{parallel}$ indicating points where $\boldsymbol{V}$ is parallel to the monitor plane. For each image we estimated the stability measure along three vertical lines. The image in the middle is a zoomed rotated version of the top image. The bottom image presents the stability estimates for the points on

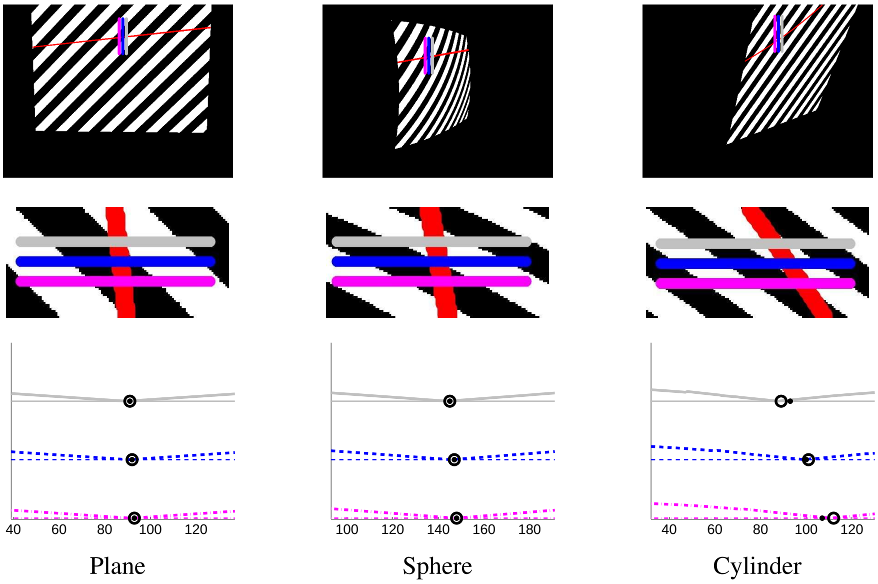

Fig. 6. A vertical triplet of images for a planar (left column), spherical (middle column), and cylindrical (right column) object. Top row: a typical captured striped image, with the (thin) red line $C_{parallel}$ marking all points where the local coordinate axis V is parallel to the monitor. Middle row: a zoomed rotated version of the same image. Bottom row: the empirical stability measure, calculated along each of the 3 colored lines.

the vertical lines as a function of the image's y coordinate. The minimal stability is marked by a black point. A circle indicates the point where the corresponding line intersects with $C_{parallel}$. As expected, the two coincide for the plane and the sphere, while for the cylinder they do not.

B. Correcting instability

We propose to correct the depth estimate of the unstable points using the accurate depth estimates of their stable neighbors. The neighborhoods associated with the initial set of points are specified via a Delaunay triangulation of these points (Figure 7). The neighborhood $\mathcal{N}_l$ of the point $\boldsymbol{p_i^l}$ is the set of direct neighbors of $\boldsymbol{p_i^l}$ in this triangulation as well as the point $\boldsymbol{p_i^l}$ itself.

Let s_l be a depth variable, associated with the initial point $\boldsymbol{p_i^l}$. Initialize s_l to the initially estimated depth. Let $s_{l\leftarrow l'}$ be a depth estimate at $\boldsymbol{p_i^l}$ obtained by a first order surface approximation relying on a neighbor $\boldsymbol{p_i^{l'}}$. The estimate is calculated from the depth $s_{l'}$ and the corresponding normal to the surface at that point. Clearly, $s_{l\leftarrow l} = s_l$. Let ϵ_l be the stability measure associated with the point $\boldsymbol{p_i^l}$. We update the depth estimation at every point, using the depth estimates associated with its neighbors and their stability, by minimizing the MSE. That is,

$$\{\hat{s}_l\} = \arg\min_{\{s_l\}} \sum_l \sum_{l'\in\mathcal{N}_l} \epsilon_{l'}^{\gamma}(s_l - s_{l\leftarrow l'})^2, \tag{8}$$

where γ is a constant. The optimal estimated depth satisfies

$$s_l = \frac{\sum_{l'\in\mathcal{N}_l} \epsilon_{l'}^{\gamma} s_{l\leftarrow l'}}{\sum_{l'\in\mathcal{N}_l} \epsilon_{l'}^{\gamma}}, \tag{9}$$

which is clearly the weighted average. This way, less stable neighbors contribute less to the depth estimate.

To minimize the MSE, we repeat this weighted smoothing iteratively. In each iteration all the initial points are traversed in descending stability measure order. The depth of each point is updated immediately after it is calculated, making the process similar to the Gauss-Seidel method. Note that updating the depth

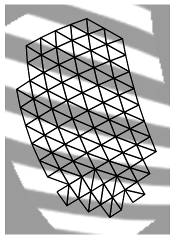

Fig. 7. Image triangulation

induces an updated surface normal as well. Initially we set $\gamma = 10$ and decrease its value after a few iterations until it reaches zero. Then the process continues with $\gamma = 0$ until convergence.

Initially, when γ gets high values, the depth at the highly unstable points is actually replaced by the interpolated depth, using their stable neighbors' depth and normals. Then, as γ decreases, the noise induced variation is reduced by a smoothing-like process. Note that this smoothing process is constrained by the correspondences between monitor points and image points, yielding a one-to-one correspondence between the depth and the surface normal. Therefore the process does not converge to a trivial planar surface solution. See Figure 8 for typical results.

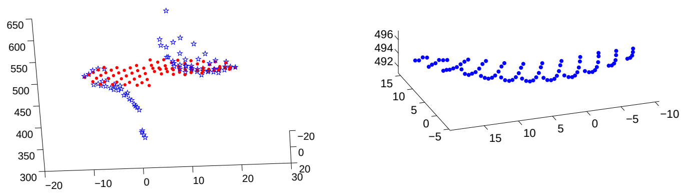

Fig. 8. The depth estimate before (blue stars) and after one smoothing iteration (red dots) (left). The final depth estimation of the initial points (right).

V. Increasing the density of the depth Map

We now augment the initial points with additional points, leading to a dense surface reconstruction. Additional points are added by an iterative process that does not change the existing points.

At the beginning of each iteration, we Delaunay triangulate the image plane using all the points recovered so far; see Figure 7. For each triangle we add one surface point that corresponds to its center of mass. For every point we now aim to calculate the depth, the surface normal, the corresponding point on the display, and the second order parameters a, b and c. All these parameters are required for the next iterations. The estimation of a, b and c for the initial set of points is described in Section III-B.

Depth - A (second order) estimate of the depth at the new point is obtained using the three vertices of its triangle. Every such vertex yields an estimate of the depth at the new point from its depth, its normal, and the second order shape parameters a, b and c, associated with it. The depth of the new point is estimated as the average of these three (Figure 9 (top)).

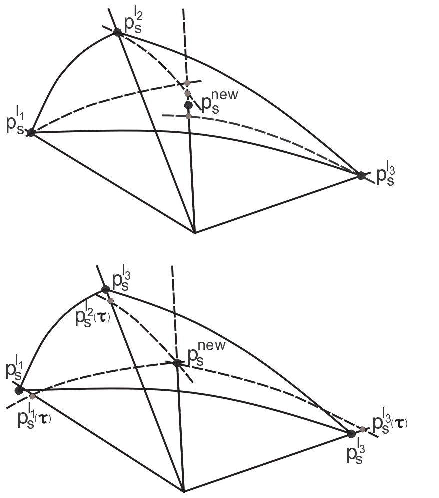

Fig. 9. Adding a single point within a triangle: estimating the depth of the new point (top) and testing depth consistencies as part of estimating the second order parameters (bottom).

Correspondence - The point $\boldsymbol{p_m}$, corresponding to the new image point $\boldsymbol{p_i}$, should satisfy the homography condition: the 1D homography, which depends on this corresponding point, should coincide with the empirical 1D homography available at every image point up to measurement errors. We found empirically that the angle between the two monitor-image homographies when treated as vectors in $\Re^4$ yields an error which is bounded by a value of $\lambda = 2^\circ$. As shown in Figure 10, points satisfying this constraint form a strip around a smooth curve. Let $\boldsymbol{p_m^{init}}$ be the center of mass associated with the monitor points corresponding to the vertices of the triangle. The monitor point corresponding to $\boldsymbol{p_i}$ is chosen:

$$\hat{\boldsymbol{p}}_{\boldsymbol{m}} = \underset{\boldsymbol{p_m} | d_{ang}(H_{mi}^{emp}, H_{mi}^{an}) < \lambda}{\arg\min} \|\boldsymbol{p_m^{init}} - \boldsymbol{p_m}\|^2. \quad (10)$$

Note that $\boldsymbol{p_m^{init}}$ is the natural choice, which would indeed be accurate if the surface was locally planar and the camera's projection model was affine.

The minimization is performed using the Nelder & Mead simplex derivative-free method [12]. Finally, the surface normal follows from the correspondence and the depth.

The second order surface parameters a, b and c - Given the correspondence and the depth, these parameters may be calculated up to an unknown scale parameter τ on the matrix $\mathcal{A}$; see (2). Given an hypothesized value for τ, the depths at the three triangle vertices may be estimated from the implied second order surface estimate and compared with the known depths at these points (estimated in earlier iterations). Thus, the sum of the three squared differences is used as the cost function for the 1D optimization process for finding τ (Figure 9, bottom). This method is much simpler than the one described in Section III-A for estimating the second order surface parameters for the initial points and does not require additional monitor-image correspondences. Note that this estimate is not a simple interpolation but relies on the empirically estimated homography (calculated at the new point) to set a powerful constraint on a, b and c. The correspondence estimate above is similarly constrained.

Now, after the depth, point correspondence, and second order surface parameters have been estimated at every new point, a new iteration which further increases the sampling density can begin. This process continues until the desired density is reached.

VI. The planar case

The reconstruction of planar specular surfaces deserves special attention. It turns out that this case is much simpler, has a more intuitive interpretation, and can be solved without numerical search. It is solved instead by solving several second degree equations.

For planar surfaces, the reflection of the planar monitor is a planar virtual object; see Figure 11, where the superscript r denotes the descriptors of this virtual object. The 1D homography between the virtual object and the image is identical to H_{mi}. This is because planar mirrors preserve angles and distances. In particular, $\alpha_m^k = {\alpha_m^k}^r$. The directions associated with the coordinate frame, $\{\boldsymbol{p_m^r}, [\boldsymbol{X_m^r}, \boldsymbol{Y_m^r}, \boldsymbol{Z_m^r}]\}$, of the reflected (virtual) monitor are specified by the three angles of a 3D rotation matrix. The 2×2 1D homography between the virtual object plane and the image plane is specified up to scale and thus provides three constraints on these parameters (3). Therefore this coordinate frame may be estimated directly from the empirical estimate of H_{mi}^{emp}. The solution from a single homography is non-unique; specifically, two candidate solutions satisfy the constraint. The false candidate is eliminated as described below.

For planar mirrors, the difference between the normal to the monitor and the normal to its reflection is the normal to the mirror,

$$\boldsymbol{W} = \frac{\boldsymbol{Z_m^r} - \boldsymbol{Z_m}}{\|\boldsymbol{Z_m^r} - \boldsymbol{Z_m}\|}.$$

Therefore the candidate normals to the mirror can be recovered at any image point using H_{mi}, with no need to establish point correspondence. Repeating this procedure for several image points yields corresponding pairs of candidate normals. The correct normals (one from each pair) form a tight cluster, which we detect. Averaging the normal estimates in this cluster yields a more accurate estimate of W. To recover the plane displacement it is sufficient to establish the correspondence $\{\boldsymbol{p_m}, \boldsymbol{p_i}\}$ at a single point.

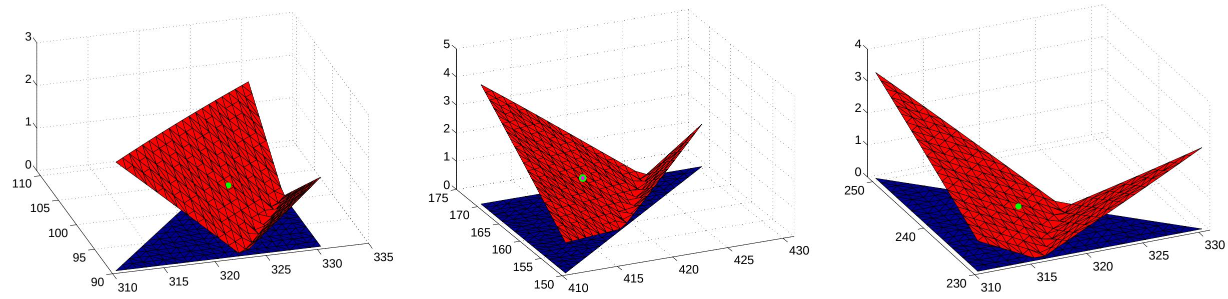

Fig. 10. Search for correspondence in the monitor plane (three typical examples). The blue triangle is specified by the vertex correspondences. The height of the red surface represents the value of the homography constraint (in degrees), namely $d_{ang}(H_{mi}^{emp}, H_{mi}^{an})$. The green point corresponds to the value of the constraint at the initial correspondence $\boldsymbol{p_m^{init}}$.

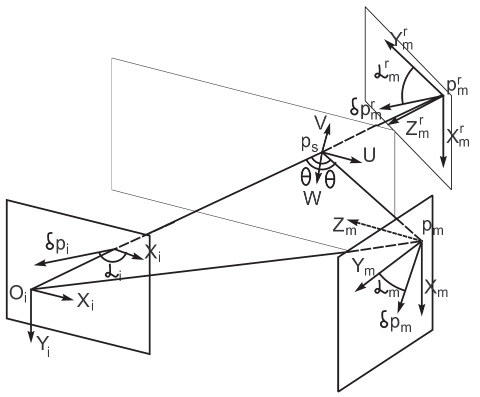

Fig. 11. Planar surface reconstruction. The reflection of the planar monitor caused by the planar mirror surface yields a reflected virtual plane denoted by the subscript r.

VII. Experimental results

A. General setting

We tested the proposed technique on three mirroring objects, each made from a different material. The objects were a planar mirror (glass), a pair of sunglasses (plastic), and a spherical ashtray (stainless steel).

The system consists of a monitor displaying illumination images, an object, reflecting these images, and a camera which acquires the illumination images after they have been distorted and reflected. The monitor-camera (3D) transformation is a bit tricky to find because the monitor is usually not directly visible to the camera. To overcome this problem we took an indirect approach: we replaced the object with a checkerboard and estimated its position relative to the camera. In addition, another camera captured both this board and the monitor in the same image, and we used this image to find the relative transformation between the two planes. Combining these estimates yields the monitor-camera transformation.

B. Illumination images

The illumination images displayed on the monitor were designed so that the algorithm can extract the following information:

1) A correspondence between the monitor and the camera for a sparse set of points, and the local monitor-image transformation scale at these points; see Section III-B.
2) A correspondence between the orientation at every image point and the orientation at its corresponding point on the monitor, $\{\alpha_i^k, \alpha_m^k\}_{k=1}^K$, which is required to estimate H_{mi}^{emp} at every image point.

1) Sparse correspondence: Following Morano *et al.* [9], the first correspondence is obtained by a single image containing a color pattern such as the one shown in Figure 12 (left). In this pattern, each 3×3 neighborhood is unique. Therefore, once we have identified the colors of the neighborhood of a captured patch, shown in Figure 12 (right), the corresponding neighborhood in the displayed image shown in Figure 12 (left) can be found. A second illumination image, containing a checkerboard with vertices at the colored patches' centers, is used to accurately find the positions of the corresponding points via corner detection, with sub-pixel accuracy.

The next illumination image is specified so that the scale characterizing the monitor-to-surface transformation T_{ms}, required for the recovery of the local second order estimate of the shape around the point, can be estimated. This image contains several points placed at known distances and directions around each of the initial points.

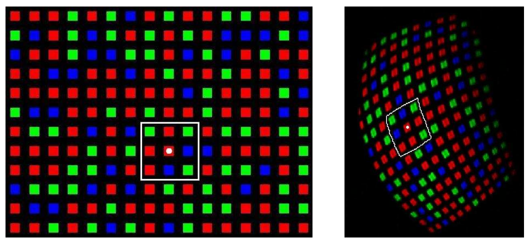

Fig. 12. Finding the correspondence of the initial points using a pseudo-random color template. Left, displayed image; right, captured image.

2) Orientation images: Achieving correspondence between all image orientations and the corresponding orientations at the monitor points generally requires dense correspondence of locations. To circumvent this requirement, we use only images with uniform orientation. That is, we display images of parallel black and white stripes, where the orientations are the same over the entire image. See Figure 2 (top). In the corresponding images (Figure 2 (bottom)), the corresponding local orientations are naturally nonuniform and are computed as the local tangent directions at the edges of the stripes. Localizing these edges is hard because the intensity of the stripes is highly nonuniform over the image, implying that every selected threshold would induce some bias in the edge's location and in the measured angle; see Figure 13a and b. Both the nonuniform orientation and the nonuniform intensities are due to the complex imaging process,

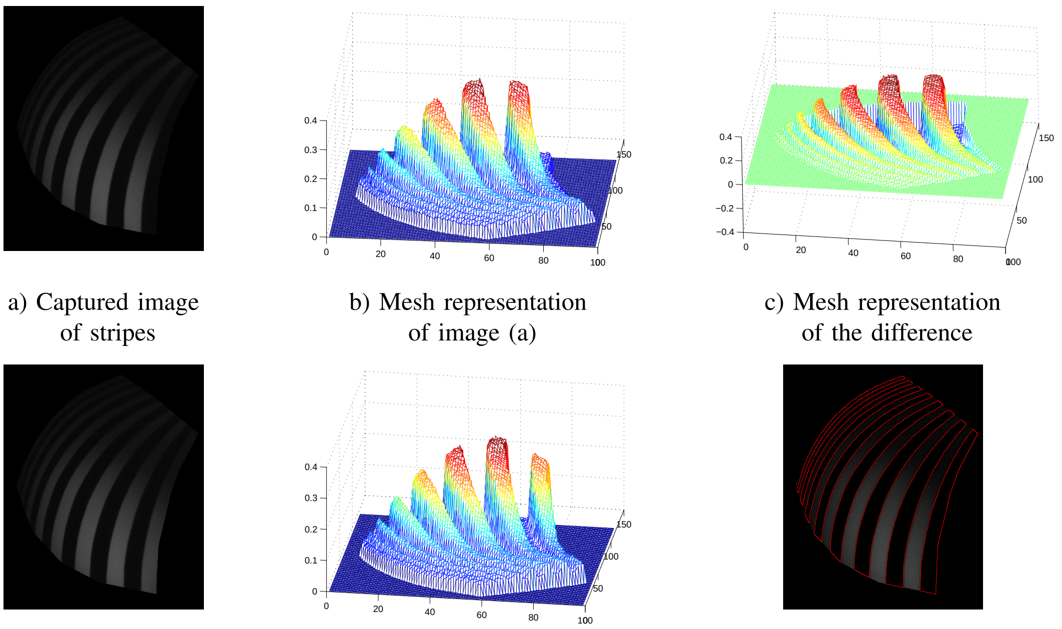

a) Captured image of stripes

b) Mesh representation of image (a)

c) Mesh representation of the difference

d) Captured negative image of stripes

e) Mesh representation of image (d)

f) Boundaries between the stripes found with zero threshold

Fig. 13. Using negative images to improve tangent direction recovery.

which includes, among others, the foreshortening by the curved mirror. To overcome the nonuniformity of the intensity, we display not only the striped image but also its inverted version. The zero crossing of the difference between the two reflected images serves now as the edge; see Figure 13. Moreover, to get a good estimate of the tangent direction all over the image, we display three shifted versions of each image. This results in angle estimates for all points on the curves. A multi-grid process is then used to provide an estimate of the tangent angle by interpolation at every image point. In most experiments we used 20 orientations which require $20 \times 2 \times 3 = 120$ images. As we show later, fewer orientations will often suffice. Given the tangent angles from several images as well as their corresponding (known) angles on the monitor, the homography may be estimated at any image point.

C. Results

The local and dense depths are estimated as described in sections IV-V. The results of the three experiments are shown in Figures 14-16. For each experiment, we show an image of the object, an intermediate reconstruction using only the sparse set of points, and two views of the final, dense reconstruction, obtained by three iterations of adding points.

We were able to evaluate the method's accuracy for two of the objects. For the planar mirror we fit a plane to the recovered points and compute the residuals. Figure 17 (left and middle) shows the residuals associated with the depths of the initial points, obtained by the local, homography based method, and after the constrained smoothing process. The third plot in Figure 17 shows a histogram of the residuals for a dense set of points obtained by the full algorithm. Clearly these residuals are very low and are bounded by $[-0.1, 0.05]$ millimeters. For planes, the second order parameters should be zero. The estimated values were indeed very low, with values smaller than $2 \cdot 10^{-4}$.

We perform a similar analysis for the spherical ashtray. First, we fit a sphere to the recovered sparse point set. The residuals are presented in the left-hand side of Figure 18 (blue), together with the stability measure (red, shifted down by 10). Note that many of these points are of low stability. Clearly, high residuals occur for these points. Following the constrained smoothing, the gross residuals are eliminated (Figure 18, middle). The residuals of the dense reconstructed points with respect to the sphere fitted to them is shown in Figure 18 (right). Note that the estimated depth is almost as good as the estimates obtained for the planar mirror. The second order parameters were estimated accurately as well. The radius of the ashtray (recovered from the parametric fit) is 44.64 mm. For an ideal spherical ashtray, the second order parameters should be $a = b = -1/r = -0.0224$ and $c = 0$ at every point. Figure 19 shows the high accuracy of the reconstructed second order parameters. To remove any doubt, we would like to emphasize that the process is local and does not assume any parametric shape.

The existence of the gross residuals demonstrated above can be explained by the stability problem of the cost function (7) and is discussed in Section IV. We would like to investigate the effect of the stability problem on the optimization of the cost functions from [16] and [15]. For that purpose we reconstruct the spherical ashtray again twice, once replacing the cost function (7) with the cost function from [15] and once replacing it with the cost function from [16]. Up to the cost function replacement, the two reconstruction processes are identical to the one described above. For both cases we build the residual graph of the local shape reconstruction results, similar to the one presented in Figure 18, left. Now we are only interested in gross residuals; therefore we count the number of locally reconstructed points with a residual larger then 10 millimeters. The results are 17, 11 and 16 for the cost function (7), the cost function from [15], and the cost function from [16] respectively. As can be seen, the numbers

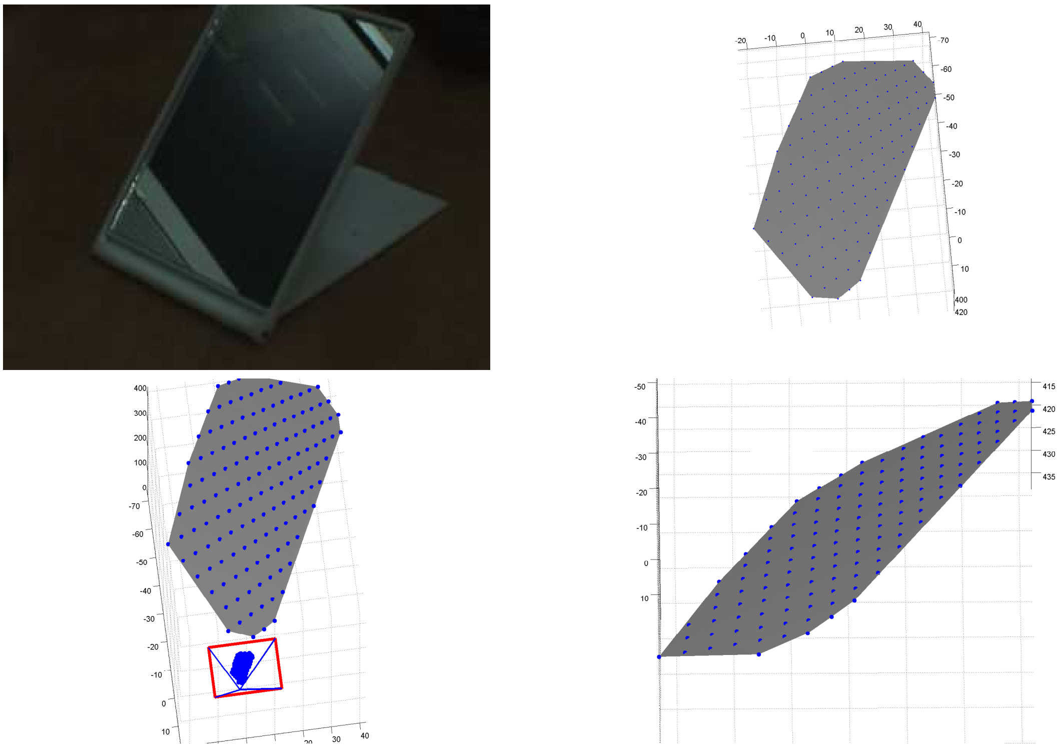

Fig. 14. Recovering the surface of a planar glass mirror. An image of the mirror (upper-left). A reconstruction using only the sparse (blue) point set (upper right). The bottom images are two views of the densely reconstructed surface. The camera location is included in one of them as well.

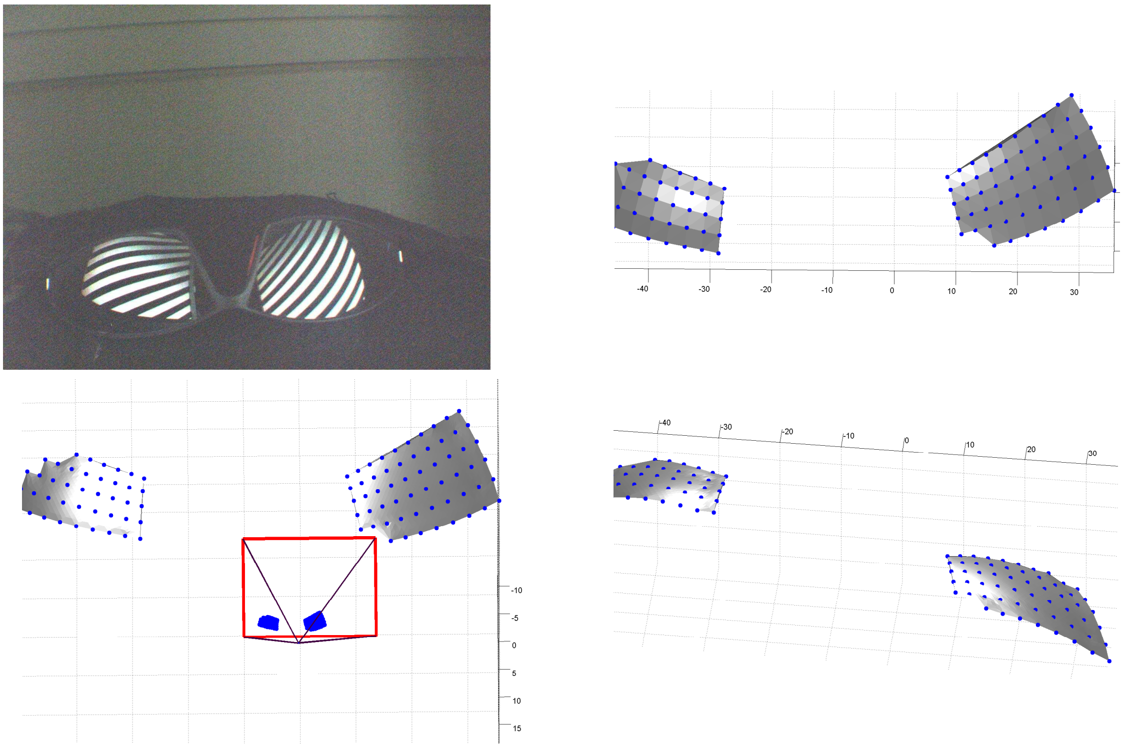

Fig. 15. Experimental results: recovering the surface of a pair of plastic sunglasses.

are similarly small, but (7) does not outperform its opponent. This should not bother us since the gross errors are eliminated at the smoothing stage, and it is much more important to have consistent convergence of the cost function optimization process, where the superiority of (7) was demonstrated. To emphasize this superiority, we would like to mention that while the correct depths were around 220 millimeters for this experiment, the initial solutions for (7) and [15] were chosen to be 500 millimeters for

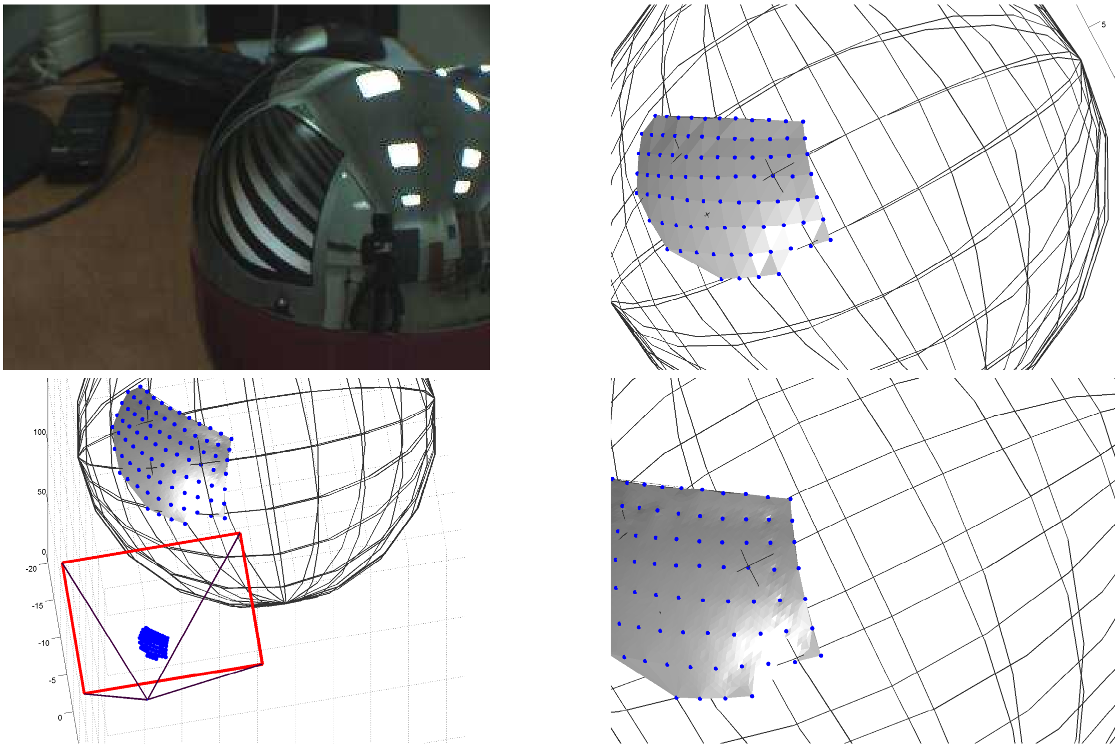

Fig. 16. Experimental results: recovering the surfaces of a stainless steel ashtray. The sphere drawn in the images is the one that best approximates the recovered surface.

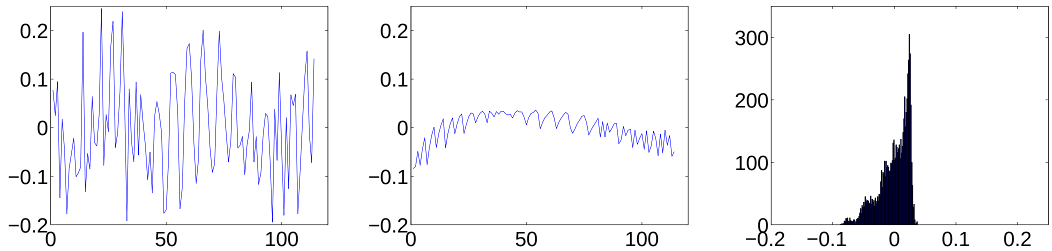

Fig. 17. Residuals in millimeters for planar surface reconstruction: (left) for the initial points after local depth estimation computation; (middle) after smoothing; (right) the residual histogram for the densely reconstructed surface.

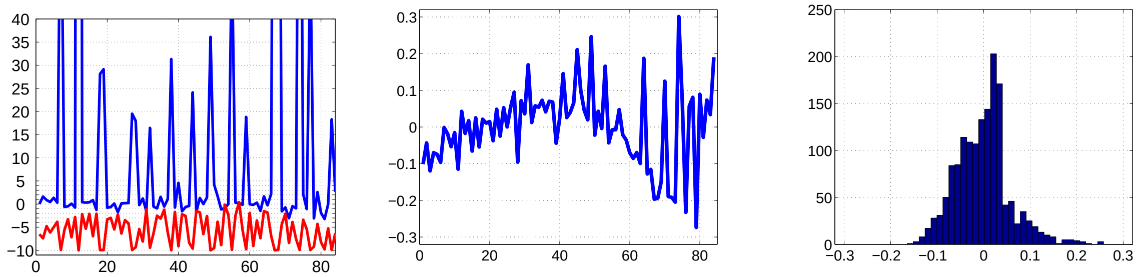

Fig. 18. Residuals in millimeters for ashtray approximated by a sphere: (left) for the initial points before smoothing, plotted together with the instability measure (red). Note the correlation (middle) for the initial points after smoothing. (Right), the residual histogram for the densely reconstructed surface.

all initial points and the optimization converged for all 84 of them. Moreover, to ensure convergence for all the initial points in the case of [16], we had to choose the initial solution to be 260 millimeters, which is much closer to the generally unknown correct solution.

In the next experiment we checked the dependence of the local depth estimation accuracy on the number of angle correspondences used to estimate the homography. For a given point, we plot the depth deviation from a reference depth specified to be the one estimated using a maximal number (20) of angle correspondences. This is repeated for 5 different points. See Figure 20. Clearly, $8-10$ angles could have been used without any significant decrease in accuracy.

In the following experiment we compared the general algorithm

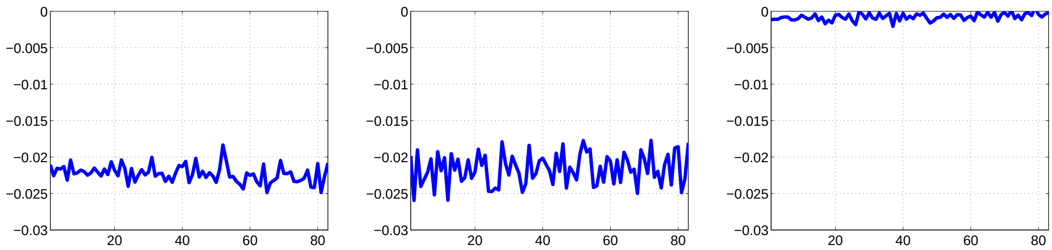

Fig. 19. Estimated values of the second order local parameters of the ashtray: (left) - a, (middle) - b, (right) - c.

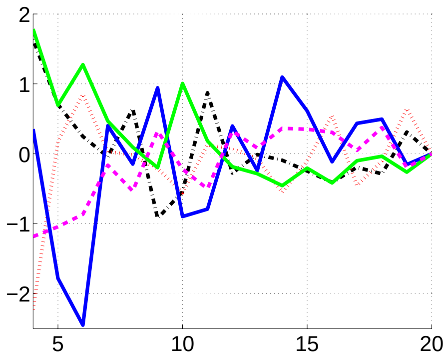

Fig. 20. The accuracy of the local depth estimation (in mm) as a function of the number of angle correspondences. The different curves correspond to 5 different surface points.

to the special algorithm designed to deal with planar objects (Section VI) by applying them to the planar mirror. The algorithms yield quite accurate results, although the more general algorithm is somewhat better. The planar object method is, however, faster. The algorithms were run on 114 points on the surface, recovering the surface normal at each point. In Figure 21 we plot the histogram of the angle differences between each pair of normals for both algorithms in degrees. Both algorithms estimate the normal with an average dispersion of less than 0.8°. No smoothing was done in this experiment.

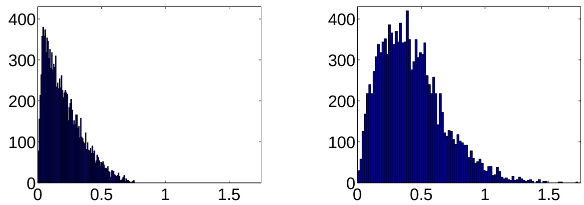

Fig. 21. Histogram of differences between the normals estimated at the initial points of the general algorithm: using the general algorithm (left); using the plane specific algorithm (right).

In our final experiment, we wanted to verify that the proposed constrained interpolation is indeed advantageous relative to the commonly used unconstrained one. We consider, again, a reconstruction of the spherical ashtray. Starting from a small initial set of 7 points, we use 5 iterations of our constrained interpolation to add 258 new points (Figure 22). Then, for comparison, we estimate depth at these points by two different unconstrained interpolation methods: linear and second order. For both types of unconstrained interpolation, we first perform a Delaunay triangulation of the initial points (Figure 22, right). In the linear interpolation, the Delaunay triangulation of the initial points also induces the triangulation of the surface and specifies the depths of the additional points (Figure 23, left). The second order unconstrained interpolation is somewhat similar to our constrained interpolation. Every new point lies within a Delaunay triangle (Figure 22, right). A second order approximation of the surface is available at the vertices of this triangle and is used to produce three independent depth estimates at the new point. These estimates are weighted proportionally to the distance to the opposite edge of the triangle (Figure 23, right), resulting in the combined depth estimate. Note that in contrast to the proposed method, in this unconstrained interpolation we cannot use additional image information to estimate the second order surface parameters at the new point.

To compare between the results of the interpolations, we fit a sphere to the reconstructed points and look at the dispersion of the points around this sphere. For the constrained interpolation, the average residual is 0.022 millimeters, while the linear and the second order unconstrained interpolations yield average residuals of 0.091 and 0.062 millimeters respectively. The standard deviations of the residuals in the three considered cases were 0.0226, 0.0713 and 0.0492 respectively. This demonstrates the superiority of the constrained interpolation.

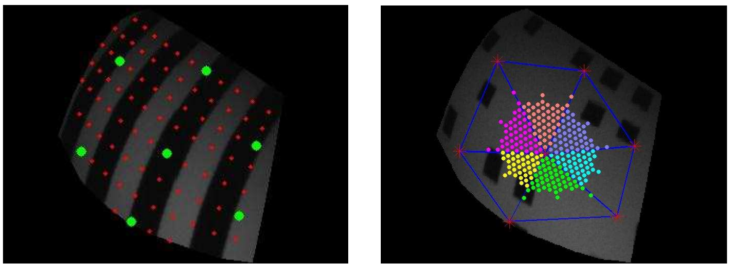

Fig. 22. Choosing points for comparing between the different interpolation methods. Left: a small set of 7 initial points (green dots). The red and the green dots were used in the shape reconstruction experiment of the ashtray. Right: the triangles are the results of the Delaunay triangulation. The dots represent the additional points. Points that belong to the same triangle have the same color.

VIII. Conclusions

This paper describes a new method for 3D shape recovery of mirroring surfaces. The proposed method builds on earlier approaches [16], [18] but differs from them in several important aspects. It shows that local surface reconstruction may be achieved by computing 1D homographies that specify the transformation of tangent angles. These homographies may be estimated by robust, statistically valid (heteroscedastic) methods, at any image point.

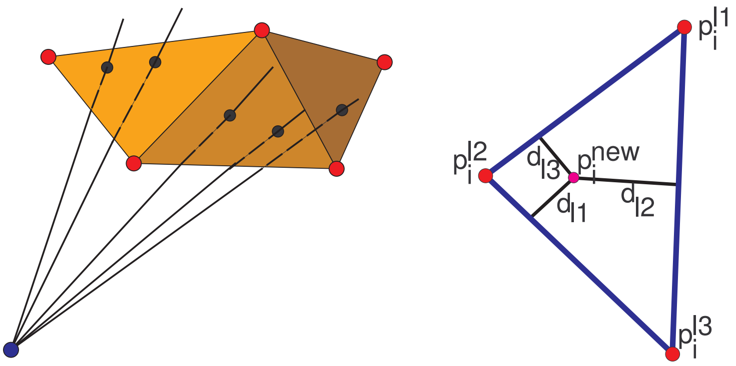

Fig. 23. The linear interpolation (left) and the weighting in the second order unconstrained interpolation (right). The blue dot in the left-hand image represents the camera's optical center; the red dots are the estimated set of initial points; the brown triangles form the triangulated surface; the black rays show the directions for which we want to estimate the depth; the black dots are the results of the linear interpolation. An image projection $\boldsymbol{p}_i^{\boldsymbol{new}}$ (right) belongs to a triangle formed by the image projections $\boldsymbol{p}_i^{\boldsymbol{l1}}$, $\boldsymbol{p}_i^{\boldsymbol{l1}}$ and $\boldsymbol{p}_i^{\boldsymbol{l3}}$ of three initial points. These points induce three second order estimates of the depth s_{l1}^{new}, s_{l2}^{new} and s_{l3}^{new} of the new point. For the point $\boldsymbol{p}_i^{\boldsymbol{lk}}$, d^{lk} is defined as the distance from $\boldsymbol{p}_i^{\boldsymbol{new}}$ to the edge opposite to $\boldsymbol{p}_i^{\boldsymbol{lk}}$. The weights of the three estimated depths $[s_{l1}^{new}, s_{l2}^{new}, s_{l3}^{new}]$ are proportional to $[d_{l1}, d_{l2}, d_{l3}]$.

Therefore, a dense surface may be recovered, requiring only a few initialization points where correspondence must be known.

The depth calculation at the initial points is similar to that described in [16], but uses a different cost function, which, because it is based on the homography and does not suffer from "ghost solutions," usually has a single minimum. Unlike [18], these calculations use only local information and do not require photometric calibration.

A special simple algorithm is also presented for planar surfaces. This algorithm is able to recover the plane normal accurately using homographies computed at many points, without any correspondence. A single point correspondence is needed to locate it.

The surface is recovered with high density using constrained interpolation, which does not simply interpolate the surface, but rather solves for the depth, the correspondence, and the local surface shape, simultaneously at each interpolated point, requiring consistency with the 1D homography. The superiority of the constrained interpolation was demonstrated empirically. This process is completely local and does not require a precalculated, dense correspondence, which may be difficult to obtain.

For a small fraction of the surface points, the ones which lie on a curve, the reconstruction is inherently unstable. We characterize these points both by an empirical measure and, for two special cases, by a simple geometric characterization. A smoothing process that relies on stability evaluation was able to overcome this problem.

The proposed method is essentially a structured light approach for recovering the shape of specular and mirror-like objects. Note that the method does not rely on photometric measurements. Therefore, no photometric calibration is needed. This characteristic, along with the dense, accurate surface recovery, make it suitable for real-world applications such as specular parts quality inspection.

References

[1] Y. Adato, Y. Vasilyev, O. Ben-Shahar, and T. Zickler. Toward a Theory of Shape from Specular Flow. In *Proc. Int. Conf. Comp. Vision*, pages 1–8, 2007.

[2] M. Baba, K. Ohtani, M. Imai, and T. Konishi. New laser rangefinder for three-dimensional shape measurement of specular objects. *Optical Engineering*, 40(1):53–60, 2001.

[3] A. Blake and G. Brelstaff. Geometry from specularities. In *Proc. Int. Conf. Comp. Vision*, pages 394–403, 1988.

[4] T. Bonfort and P. Sturm. Voxel carving for specular surfaces. In *Proc. Int. Conf. Comp. Vision*, pages 591–596, 2003.

[5] M. Halstead, B. Barsky, S. Klein, and R. Mandell. Reconstructing curved surfaces from specular reflection patterns using spline surfce fitting of normals. In *ACM Computer Graphics and Interactive Techniques*, pages 335–342, New Orleans, 1996.

[6] M. C. Knauer, J. Kaminski, and G. Häusler. Phase measuring deflectometry: a new approach to measure specular free-form surfaces. In *Optical Metrology in Production Engineering. Edited by Osten, Wolfgang; Takeda, Mitsuo. Proceedings of the SPIE, Volume 5457*, pages 366–376, 2004.

[7] Y. Leedan and P. Meer. Heteroscedastic regression in computer vision: Problems with bilinear constraint. *Int. J. of Comp. Vision*, 37(2):127–150, June 2000.

[8] J. Lellmann, J. Balzer, A. Rieder, and J. Beyerer. Shape from specular reflection and optical flow. *Int. J. of Comp. Vision*, 80(2):226–241, November 2008.

[9] R. Morano, C. Ozturk, R. Conn, S. Dubin, S. Zietz, and J. Nissanov. Structured light using pseudorandom codes. *IEEE Trans. Patt. Anal. Mach. Intell.*, 20(3):322–327, March 1998.

[10] M. Oren and S. Nayar. A theory of specular surface geometry. In *Proc. Int. Conf. Comp. Vision*, pages 740–747, 1995.

[11] W. Park and H. Cho. Measurement of the 3-dimensional shapes of specular objects using recursive triangulation. In *Pacific Conference on Manufacturing*, volume 1, pages 389–394, Seoul, Korea, 1996.

[12] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes in C*. Cambridge University Press, 1988.

[13] A. Ripsman and M. Jenkin. Local surface reconstruction of objects in space. *IEEE Int. Symposium on Computational Intelligence*, 2001.

[14] S. Roth and M. Black. Specular flow and the recovery of surface structure. In *Proc. IEEE Conf. Comp. Vision Patt. Recog.*, pages II: 1869–1876, 2006.

[15] S. Rozenfeld, I. Shimshoni, and M. Lindenbaum. Dense mirroring surface recovery from 1D homographies and sparse correspondences. In *Proc. IEEE Conf. Comp. Vision Patt. Recog.*, pages I:1–8, 2007.

[16] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. *Int. J. of Comp. Vision*, 64(1):31–67, August 2005.

[17] J. Solem, H. Aanaes, and A. Heyden. Variational surface interpolation from sparse point and normal data. *IEEE Trans. Patt. Anal. Mach. Intell.*, 29(1):181–184, January 2007.

[18] M. Tarini, H. P. A. Lensch, M. Goesele, and H.-P. Seidel. 3D acquisition of mirroring objects using striped patterns. *Graph. Models*, 67(4):233–259, 2005.

[19] J. Zheng and A. Murata. Acquiring a complete 3D model from specular motion under the illumination of circular-shaped light sources. *IEEE Trans. Patt. Anal. Mach. Intell.*, 22(8):913–920, 2000.

[20] A. Zisserman, P. J. Giblin, and A. Blake. The information available to a moving observer from specularities. *Image and Vision Computing*, 7(1):38–42, 1989.