

Shape Reconstruction of 3D Bilaterally Symmetric Surfaces*

ILAN SHIMSHONI

Dept. of Industrial Engineering and Management, The Technion, Haifa 32000, Israel

YAEL MOSES

Dept. of Computer Science, The Interdisciplinary Center, Herzliya 46150, Israel

MICHAEL LINDENBAUM

Dept. of Computer Science, The Technion, Haifa 32000, Israel

Received ??; Revised ??

Editors: ??

Abstract. The paper presents a new approach for shape recovery based on integrating geometric and photometric information. We consider 3D bilaterally symmetric objects, that is, objects which are symmetric with respect to a plane (e.g., faces), and their reconstruction from a single image. Both the viewpoint and the illumination are not necessarily frontal. Furthermore, no correspondence between symmetric points is required.

The basic idea is that an image taken from a general, non frontal viewpoint, under non-frontal illumination can be regarded as a pair of images. Each image of the pair is one half of the object, taken from different viewing positions and with different lighting directions. Thus, one-image-variants of geometric stereo and of photometric stereo can be used. Unlike the separate invocation of these approaches, which require point correspondence between the two images, we show that integrating the photometric and geometric information suffice to yield a dense correspondence between pairs of symmetric points, and as a result, a dense shape recovery of the object. Furthermore, the unknown lighting and viewing parameters, are also recovered in this process.

Unknown distant point light source, Lambertian surfaces, unknown constant albedo, and weak perspective projection are assumed. The method has been implemented and tested experimentally on simulated and real data.

Keywords: Shape Recovery, Shape-from-shading, Stereo, Bilaterally Symmetric Objects

1. Introduction

This paper presents a method for shape reconstruction of 3D bilaterally symmetric objects from a single image [9, 16, 26, 20, 27]. A bilaterally symmetric object is an object that is symmetric with respect to a plane

(as illustrated in Figure 1). Typical examples of this class of objects are faces, animals, human bodies and many man-made objects. We consider a natural setup where such objects are seen from a non-frontal view and with non-frontal illumination. Note that the resulting images are non-symmetric but, as we will show, are actually more informative than symmetric images.

*Ilan Shimshoni was supported in part by the Koret and Goldschmidt Foundations. A preliminary version of this paper has appeared in IEEE Int. Conf. on Image Analysis and Processing [22].




Fig. 1. An illustration of a bilaterally symmetric object viewed from an arbitrary viewing direction. For each point p^l there is a corresponding point p^r and a pair, n^l, n^r , of corresponding normals at those points.

Methods for general shape reconstruction usually rely on either geometric or photometric constraints. When a single image of a given object is available, geometric constraints are not sufficient by themselves for shape reconstruction of a general 3D object. However, when the object is known to belong to a given class of objects (hence providing some geometrical constraints) a single image may be sufficient [1]. Photometric methods such as shape-from-shading [14] can be used for shape reconstruction from a single image. However, these methods require the knowledge of various parameters (e.g., illumination), and some other assumptions. Here again, class constraints can often alleviate the need to be given these parameters and can substantially improve the reconstruction. See for example [10] where the shape of Straight Homogeneous Generalized Cylinders (SHGC's) was recovered by a special algorithm.

1.1. Our approach

Our approach can be classified as a class based 3D shape recovery algorithm that combines geometric and photometric constraints. In our case we deal with the class of bilaterally symmetric objects. The basic idea is that a single non-frontal image of a bilaterally symmetric object can be regarded as two images of another object taken from two different viewpoints and two different illumination directions. The object consists of one half of the bilaterally-symmetric object. It follows that the geometrical and photometric constraints

on an image of a bilaterally symmetric object is equivalent to the constraints of two images of a general 3D object. Therefore, one-image-variants of photometric-stereo [24, 18] and geometric-stereo [12, 5] can be used.

The straightforward application of the one-image variants of both geometric and photometric stereo methods requires however that the imaging parameters (viewpoint and illumination) as well as the correspondence between symmetric points, denoted *self-correspondence*, be given [9]. Finding dense correspondence is generally a non-trivial problem. Correspondences can usually be found only for a sparse set of feature points, which result only in a sparse depth map [12, 5]. Furthermore, as we describe below, the photometric or geometric constraints are insufficient by themselves for computing dense self-correspondence.

The essence of our approach for solving the self-correspondence and the global imaging parameters is to combine geometric and photometric constraints. To do so we specify the self-correspondence as a continuous function, and solve its derivatives by combining the two types of constraints. These constraints are local and are specified at every pixel. Our approach is to use them first to find the global (pose and illumination) parameters. Once these parameters are known, the local constraints enforce correct reconstruction locally. This stands in sharp contrast to traditional shape from shading solutions, which are neither local nor provide means for recovering the global parameters. Once one pair of corresponding points has been found using the local constraint, a combination of the photometric and geometric constraints are used to propagate the self-correspondence function over the entire image, yielding the 3D shape.

In this paper we present several algorithms, emphasizing different aspects of the theoretical results. Before we describe the general algorithm, we deal with the special case of recovering the shape of bilaterally-symmetric objects (a real human face) from a frontal view. Then we come to the the general case (general bilaterally-symmetric object, non-frontal view and non-frontal illumination), which we demonstrate in experiments done on both simulated data and on real images of a mannequin object. In a final experiment we demonstrate the automatic recovery of the global lighting and pose parameters from a real image of a mannequin.

1.2. Previous studies

Previous studies on symmetry have focused either on symmetric images [8, 25] or on skewed symmetric images [6, 3, 17, 19]. Note that in both cases, detecting the self-correspondence is a parametric problem, the solution of which depends only on a few parameters (characterizing the symmetry axis and possibly another direction in the image plane). Under general imaging conditions, images of 3D bilaterally-symmetric objects are not symmetric. Neither the image coordinates nor the grey levels associated with 3D symmetric objects, are symmetric. However, when self-correspondence is given then the 3D shape can be recovered [9, 16, 26, 20, 27]. Furthermore, the problem is non-parametric, and finding the self-correspondence is equivalent to recovering the 3D structure. Therefore, it cannot be based on the methods mentioned above.

Previous attempts to integrate photometric and geometric constraints were based on using each of the methods in different image regions. In some studies the results computed by say, a stereo algorithm, were transferred to a photometric shape recovery algorithm (e.g., [2, 4, 7, 15]). To the best of our knowledge, our scheme is the first to propose integration in the strong sense. The integration of the photometric and geometric constraints allow us not only to compute the self-correspondence function, but also to recover the shape of the object based on geometric constraints. This is a preferable method over photometric methods for shape reconstruction since it avoids the reconstruction based on the integration of surface normals, which cause errors to accumulate.

1.3. Paper outline

The paper continues as follows: in the next two sections (2 and 3) we derive the geometric and photometric constraints on images of a symmetric object. The special case of reconstruction from a frontal image is also considered in section 3. Then, in Section 4 we show how to integrate these constraints into one, stronger constraint on the correspondence function. In section 5 we turn to the reconstruction algorithm that uses the constraints efficiently and describe it in detail. All the presented algorithms have been implemented and results of testing them on simulated and real images are presented in Section 6. Finally, some limitations of our

algorithm and future research directions are discussed in Section 7.

2. Geometry

In this section we present the geometric setting, some notations, and the geometric constraints associated with bilaterally symmetric objects. In particular we present some known results using this notation. We first show that if self-correspondence is given, then a representation that is invariant to viewpoint can be computed. We then show that if the viewing parameter is given as well, then the 3D shape of the object can be directly recovered. Finally, we show that if the self-correspondence is unknown, then the 3D shape and the self-correspondence cannot be computed based only on geometrical constraints, even if the viewing direction is known. Then the photometric constraints (considered in the next section) must be added. The notations and the results are later used to build our integrated approach.

2.1. Notations and settings

Let $\mathbf{p}^r = (x, y, z)^T$, $\mathbf{p}^l = (-x, y, z)^T$ be a pair of symmetric points on a bilaterally symmetric object. We choose our coordinate system such that the *symmetry plane*, the plane relative to which the object is symmetric about, is the $y \times z$ plane. In this case, all lines that connect pairs of symmetric points are parallel to the x -axis.

The object points \mathbf{p}^l and \mathbf{p}^r are projected to the image points $\tilde{\mathbf{p}}^l$ and $\tilde{\mathbf{p}}^r$. We assume here weak perspective projection (i.e., orthographic projection of a rotated and scaled object). The projection of a point \mathbf{p} is $\tilde{\mathbf{p}} = sR\mathbf{p} + \mathbf{t}$ where s is a scale factor, R is two rows of a 3×3 rotation matrix, and \mathbf{t} is a 2D vector. Under this projection model, all image lines that connect pairs of symmetric points are parallel. Noting the similarity to the stereo setting, which will be discussed below, we refer to these set of image lines as *epipolar lines*. Estimating the direction of these lines is usually fairly easy and may be done by, say, finding a pair of corresponding points (e.g., the eyes of a person in the image (e.g. [23])). Therefore, we shall assume that this direction is known. Given the direction of the epipolar lines, finding the self-correspondence becomes a one-dimensional search problem.

Without loss of generality, it is sufficient in our case, to set $s = 1$ and $\mathbf{t} = 0$, and to reduce the viewing direction to a single parameter θ , which is the rotation of the object around the y -axis of the object. Thus R is reduced to:

$$R = \begin{pmatrix} \cos(\theta) & 0 & \sin(\theta) \\ 0 & 1 & 0 \end{pmatrix} \quad (1)$$

To show that a single parameter is indeed sufficient, let us first consider a general rotation matrix. Any 3D rotation can be described by three rotation angles: around the x , y and z axes (see Figure 2). When the epipolar direction is known the rotation in the image plane can be directly computed and compensated for. We can therefore assume that the rotation about the z -axis is 0. (Note that this is correct even though the z axis does not coincide with the optical axis.) The inherent ambiguity in the choice of the object coordinate system justifies the choice of zero rotation around the x axis, $s = 1$, and $\mathbf{t} = 0$. This ambiguity results in reconstruction of the object up to a scale factor, translation, and rotation around the x axis.




Fig. 2. (a) A bilaterally symmetric object (face); (b) (c) & (d) are its three rotations about the z , x and y axes respectively.

2.2. Geometric constraints

Under this choice of projection we can now compute the image locations: $\tilde{\mathbf{p}}^r = (\tilde{x}^r, \tilde{y}^r)^T$ and $\tilde{\mathbf{p}}^l = (\tilde{x}^l, \tilde{y}^l)^T$, of a pair of corresponding object points,

$\mathbf{p}^r = (x, y, z)^r$ and $\mathbf{p}^l = (-x, y, z)^l$. We obtain that

$$\begin{aligned} \tilde{y}^r &= y; \\ \tilde{x}^r &= x \cos(\theta) + z \sin(\theta); \\ \tilde{y}^l &= y; \\ \tilde{x}^l &= -x \cos(\theta) + z \sin(\theta). \end{aligned} \quad (2)$$

These equations can be inverted yielding the following expressions for the 3D coordinates of the object points.

$$\begin{aligned} x &= \frac{1}{2 \cos(\theta)} (\tilde{x}^r - \tilde{x}^l), \\ y &= \tilde{y}^r, \\ z &= \frac{1}{2 \sin(\theta)} (\tilde{x}^r + \tilde{x}^l). \end{aligned} \quad (3)$$

We call this shape reconstruction method *mono-geometric-stereo* because the two sides of the image are interpreted as a stereo pair of one half-object taken from two viewpoints. It is also clear now why we denoted the image lines parallel to the object's transversal axis by epipolar lines.

Finally, we claim that the location of the corresponding image points are not constrained by the fact that the object is symmetric, even if the viewing position θ is known. It is easy to see from Eq. 3 that given a right point in the image, any choice of a left point is consistent with an image of a bilaterally symmetric object. This ambiguity is equivalent to the inherent ambiguity in pairing points from two stereo images (see, e.g. [5]). Recall however, that even when θ is unknown it is possible to infer invariants to viewing direction directly from Eq. 3 [16, 20].

Note that under the Lambertian reflectance model, if the illumination is frontal (with respect to the object), then corresponding points are associated with the same grey level values, which may make correspondence finding algorithms used in traditional stereo processes, may be adapted to find the correspondence also here, from which the 3D shape is recoverable using Eq. 3. To demonstrate this case we present in Figure 3(a) a real image of a mannequin frontally illuminated with respect to the object. In Figure 3(b) Iso-intensity contours are overlaid over the image demonstrating the fact that corresponding points have the same intensity.



Fig. 3. Mono-geometric-stereo: (a) A real image of a mannequin with non-frontal viewing position and frontal illumination; (b) Iso-intensity contours are overlaid over the image demonstrating the equal intensity for corresponding points. For images like this, a variant of a stereo algorithm can be applied.

3. Photometry

In this section we derive the photometric constraint for bilaterally symmetric objects. In a rather artificial way, we ignore for now the evidence available from the location of these image points. This is done in order to isolate the information available from the two sources. We integrate the two sources of information in the next section. For the case when self correspondence is known (e.g. when the view is frontal) we describe the mono-photometric-stereo algorithm, which uses only photometric information, and is a specialized form of the proposed general algorithm. We also describe a specialized method for recovering the illumination parameters when these are not given. An interesting byproduct is an extension of this technique to the classic photometric-stereo task associated with unknown illuminations.

3.1. Notations and settings

The analysis is done in the context of a Lambertian surface and a distant light source. Different illumi-

nation and shading models lead to other photometric constraints, which are left for future research.

Let $\hat{\mathbf{n}} = (n_x, n_y, n_z)^T$ be the normal at a surface point of the object. Then, using the Lambertian reflectance model, the intensity of light reflected from the point is $I = \rho \mathbf{l} \cdot \hat{\mathbf{n}}$, where \mathbf{l} is the light source direction, $|\mathbf{l}|$ is the light source intensity, and ρ is the albedo. We assume that the albedo ρ , is fixed all over the surface. Equivalently, we may assume that $\rho = 1$ and let $|\mathbf{l}|$ be the product of the intensity and the albedo. We also assume that no shadows exist. We intend to relax these assumptions in future work.

3.2. Photometric constraints

We denote by $\hat{\mathbf{n}}^r$ and $\hat{\mathbf{n}}^l$ the normals of two symmetric points on the object. Denote the right side of the object surface by the function $(x, y, z(x, y))$. In this case the normals at two self-corresponding points are given by $\hat{\mathbf{n}}^r = \mathbf{n}^r / |\mathbf{n}^r|$ and $\hat{\mathbf{n}}^l = \mathbf{n}^l / |\mathbf{n}^l|$ where $\mathbf{n}^r = (z_x^r, z_y^r, -1)^T$ and $\mathbf{n}^l = (-z_x^r, z_y^r, -1)^T$. The intensity at the two points is

$$\begin{aligned} I^r &= \mathbf{l} \cdot \hat{\mathbf{n}}^r, \\ I^l &= \mathbf{l} \cdot \hat{\mathbf{n}}^l \end{aligned} \quad (4)$$

Taking the sum and the difference of Eq. 4 yields:

$$\begin{aligned} I_\Delta &= I^r - I^l = 2l_x n_x \\ I_\Sigma &= I^r + I^l = 2(l_y n_y + l_z n_z) \end{aligned} \quad (5)$$

Similarly to viewpoint invariants, when \mathbf{l} is unknown it is possible to infer invariants to illumination direction directly from Eq. 4 [16].

3.3. Mono-photometric-stereo

An immediate observation is that if the self-correspondence and \mathbf{l} are known, then the intensities reflected from a corresponding pair are sufficient for calculating the normal at these points by solving a quadratic equation up to a twofold ambiguity. First, n_x is calculated from the first part of Eq. 5. Then, the unit size of $\hat{\mathbf{n}}$ is called to provide an additional constraint, yielding a quadratic equation from which the two solutions for the other components of $\hat{\mathbf{n}}$ are recovered. We call this normal recovery method a *mono-photometric-stereo* algorithm. It is similar to photometric-stereo where the two sides of the image are interpreted as a stereo pair of one half-object taken with two different illuminations. The natural application for this algo-

rithm is when a frontal view of a bilaterally symmetric object illuminated by a non-frontal light source is considered. Then, the image is geometrically symmetric and finding the self-correspondence is easy.

3.4. Illumination recovery in the context of a frontal view

In photometric-stereo algorithms the lighting directions are supplied to the algorithm. Here however we will show that mono-photometric stereo reconstruction is possible even if the illumination \mathbf{l} is unknown. For finding the illumination we will use the first derivatives of the intensity. For a frontal view, the image and the object x and y coordinates are identical: $\tilde{x}^r = x$, $\tilde{y}^r = y$, $\tilde{x}^l = -x$, and $\tilde{y}^l = y$. Therefore, the derivatives of the image intensity are equal, up to a sign to the intensity derivatives with respect to the world coordinates x, y . Note that this is not true for non frontal

$$\begin{pmatrix} I^r_{\tilde{x}} \\ I^l_{\tilde{x}} \\ I^r_{\tilde{y}} \\ I^l_{\tilde{y}} \end{pmatrix} = \frac{1}{|\mathbf{n}|^3} \begin{pmatrix} (1 + z_y^2, -z_y z_x, z_x) \mathbf{l} & 0 & (-z_x z_y, 1 + z_x^2, z_y) \mathbf{l} \\ (-1 - z_y^2, -z_y z_x, z_x) \mathbf{l} & 0 & (z_x z_y, 1 + z_x^2, z_y) \mathbf{l} \\ 0 & (-z_x z_y, 1 + z_y^2, z_y) \mathbf{l} & (-1 - z_y^2, -z_y z_x, z_x) \mathbf{l} \\ 0 & (z_x z_y, 1 + z_y^2, z_y) \mathbf{l} & (1 + z_y^2, -z_y z_x, z_x) \mathbf{l} \end{pmatrix} \begin{pmatrix} z_{xx} \\ z_{yy} \\ z_{yx} \end{pmatrix} \quad (7)$$

Note that these equations are linear in z_{xx}, z_{xy} and z_{yy} . Their coefficients can be described as a tensor in the surface first derivatives multiplied by the illumination. For every symmetric pair, we have 6 equations in 5 local unknowns (2 from Eq. 5 and 4 from Eq. 7). In addition, we have also 3 global unknowns characterizing the illumination and the albedo. Enforcing the intensity constraints for (n) symmetric pairs, we get $6n$ equations in $5n + 3$ unknowns.

Solving this system of equations requires at least 3 corresponding pairs. The use of noisy intensity derivatives, however, is expected to yield an inaccurate solution. Therefore, we have chosen the following optimization based approach, which uses much more corresponding pairs, and thus has the potential to compensate for the noise. Indeed, our experiments with real data presented in Section 6.4 show that such a robust solution is achieved.

The solution searches, in a 3D space, for a feasible value of the global illumination parameters. For every candidate illumination, and for every correspond-

views due to the foreshortening effect. For non-frontal views the difference in the derivatives between two corresponding points depends on the viewpoint and the 3D shape as we discuss in section 4.3.

The frontal image intensity derivatives may be written in terms of the normals derivatives:

$$\begin{aligned} I^r_{\tilde{x}} &= I^r_x = \mathbf{l} \cdot \hat{\mathbf{n}}_x^r; \\ I^l_{\tilde{x}} &= -I^l_x = -\mathbf{l} \cdot \hat{\mathbf{n}}_x^l; \\ I^r_{\tilde{y}} &= I^r_y = \mathbf{l} \cdot \hat{\mathbf{n}}_y^r; \\ I^l_{\tilde{y}} &= I^l_y = \mathbf{l} \cdot \hat{\mathbf{n}}_y^l. \end{aligned} \quad (6)$$

where $I_e = \frac{\partial I}{\partial e}$.

Considering smooth integrable surfaces, where $z_{xy} = z_{yx}$, it follows that the normal derivatives, $\hat{\mathbf{n}}_x^r, \hat{\mathbf{n}}_x^l, \hat{\mathbf{n}}_y^r, \hat{\mathbf{n}}_y^l$ are functions of 5 local derivatives of the surface: $z_x, z_y, z_{xx}, z_{yy}, z_{xy} = z_{yx}$. Substituting into these equations the derivatives of the normals yields:

ing pair, we first use Eq. 5 to extract z_x, z_y , leading to a linear system 7 in z_{xx}, z_{xy} and z_{yy} .

Not for all candidate illuminations there is a solution for Eq. 5. These equations yield the following constraints on possible values of \mathbf{l} . Since $|\mathbf{n}| = 1$, it follows that: $l_x < I_\Delta/2$ and $|\mathbf{l}| \geq \max I^r, I^l$. The values l_y and l_z are further constrained by the second equation of Eq. 5. This is a quadratic equation in n_z (without loss of generality). In order for it to have a solution, its determinant must be non-negative. It follows that only pairs l_y and l_z which satisfy this constraint should be considered.

Once z_x and z_y have been found, the system of equations 7 is an over-determined set of equations, which is solvable (in a noiseless situation) only if the candidate illumination is correct. In the realistic noisy situation, the least squares residual is a measure of the candidate illumination correctness. Summing this correctness measure for all corresponding pairs, yields a total correctness measure for the candidate illumination. Standard optimization procedure (e.g. Levenberg-Marquardt) may use this score to find the optimal \mathbf{l} which gets the minimal score.

For non-frontal views, this technique has to be modified, as foreshortening affects the intensity derivatives at corresponding points differently. This more general situation is dealt with in the next section.

3.5. Illumination recovery for (traditional) photometric stereo

This method can be easily adapted for the classical photometric-stereo algorithm, when the illumination directions are unknown (without any points whose intensity is known [11]). In this case, we still have 5 parameters characterizing the local geometry at every point and 6 additional global parameters characterizing the two illuminations. Therefore, the information available from the intensities in $n \geq 6$ points suffices to determine both the illumination and the normals at these points.

4. Integrating Geometry and Photometry

We now come to the general reconstruction task, involving a non-frontal viewing direction and non-frontal illumination. Here we assume that the self-correspondence (correspondence between points in the left and the right part of the image) is not given and has to be recovered by the algorithm. In this section we present a method for integrating the geometric and photometric information in order to solve the self-correspondence problem. This method is also used for verifying self-correspondence locally, and for recovering the illumination vector and pose parameter. All these tools are used as components of the reconstruction algorithm, described in the next section.

4.1. The self-correspondence function

In this section we introduce the “self-correspondence” function. This function, which is initially unknown, can be constrained using the image data. To calculate this function, we start from a given pair of self-corresponding points, use the constraints to calculate dense correspondence in its neighborhood and further propagate the correspondence over the whole image, together with reconstructing the 3D shape. The correctness of the self-correspondence function can be verified locally which prevents the accumulation of errors.

In classical stereo algorithms the analogous notion to the self-correspondence function is the disparity function. The smoothness of the disparity function is often used to constrain the correspondence. In our case we will show how to use the photometric information to recover the derivatives of the self-correspondence function. This information is not available in the classical stereo setting since a single intensity value is available for each pair of corresponding points.

Given a point $\tilde{\mathbf{p}}^r = (\tilde{x}^r, \tilde{y}^r)^T$ in the image, its corresponding point is given by $\tilde{\mathbf{p}}^l = (\tilde{x}^l, \tilde{y}^l)^T$. Recall that $\tilde{y}^r = \tilde{y}^l$. Thus, it suffices to specify the correspondence by a scalar function C such that

$$C(\tilde{\mathbf{p}}^r) = C(\tilde{x}^r, \tilde{y}^r) = \tilde{x}^l. \quad (8)$$

The points $\tilde{\mathbf{p}}^r$ and $\tilde{\mathbf{p}}^l$ are functions of x , y , z and θ . Taking the derivatives of the function C with respect to the object coordinates x and y we obtain:

$$\begin{aligned} C_x &= C_{\tilde{x}^r} \tilde{x}_x^r + C_{\tilde{y}^r} \tilde{y}_x^r \\ C_y &= C_{\tilde{x}^r} \tilde{x}_y^r + C_{\tilde{y}^r} \tilde{y}_y^r \end{aligned} \quad (9)$$

The derivatives of the image coordinates \tilde{x}^r , \tilde{x}^l , \tilde{y}^r and \tilde{y}^l with respect to the object coordinates x and y can be computed from Eq. 2:

$$\begin{aligned} \tilde{x}_x^r &= \cos(\theta) + z_x \sin(\theta); \\ \tilde{x}_y^r &= z_y \sin(\theta); \\ \tilde{x}_x^l &= -\cos(\theta) + z_x \sin(\theta); \\ \tilde{x}_y^l &= z_y \sin(\theta); \\ \tilde{y}_x^r &= \tilde{y}_x^l = 0; \\ \tilde{y}_y^r &= \tilde{y}_y^l = 1. \end{aligned} \quad (10)$$

Inserting these derivatives into Eq. 9 yields the derivatives of the correspondence function with respect to the image coordinates as a function of the surface normals (z_x, z_y) :

$$C_{\tilde{x}^r} = C_x / \tilde{x}_x^r = \frac{\tilde{x}_x^l}{\tilde{x}_x^r} = \frac{-\cos(\theta) + z_x \sin(\theta)}{\cos(\theta) + z_x \sin(\theta)},$$

$$C_{\tilde{y}^r} = C_y - \tilde{x}_y^r C_{\tilde{x}^r} = \tilde{x}_y^r - \tilde{x}_y^r \frac{\tilde{x}_x^r}{\tilde{x}_x^r} = \frac{z_y \sin(2\theta)}{\cos(\theta) - z_x \sin(\theta)}. \quad (11)$$

Given a corresponding pair, the illumination and viewing direction, the normal at the point can be computed using Eq. 5. Then, using Eq. 11 the derivatives of the correspondence function at the image can also be computed. These derivatives are then used to estimate the corresponding position of a point close to $\tilde{\mathbf{p}}^r$ using the first order Taylor expansion of the correspondence function:

$$C(\tilde{x}^r + d_x, \tilde{y}^r + d_y) \approx C(\tilde{\mathbf{p}}^r) + d_x C_{\tilde{x}^r} + d_y C_{\tilde{y}^r}. \quad (12)$$

This procedure can be repeated again and again propagating the correspondence function to large regions of the image (as illustrated in Figure 4.1(a)).




Fig. 4. In this figure we draw the corresponding points in the left and the right parts of the image. (a) An illustration of the correspondence propagation process. Corresponding pairs are always on the same horizontal line. Note the difference between the propagation of the correspondence on the two sides. The location on the right side is the “independent variable” which was changed either horizontally or vertically. The corresponding locations on the left side are the same with respect to the y coordinate but depend on the self-correspondence function for their x coordinate. A horizontal step on the right corresponds to a horizontal step on the left. A vertical step on the right however, correspond to an oblique step on the left, because the self-correspondence function depends on both the x and y coordinates of the right point. (b) An illustration of the “circular tour” constraint. 8 different simple paths, specified so that every two of them end in the same (corner) point, are shown in the right image. The corresponding paths, should satisfy this constraint (every two paths end in the same point) as well, but they do not. Therefore either the starting points or the global parameters are incorrect.

4.2. Local self-correspondence verification by “circular tours”

In smooth parts of the image, the function

$$\tilde{\mathbf{p}}^l(\tilde{\mathbf{p}}^r) = (\tilde{x}^l, \tilde{y}^l) = (C(\tilde{\mathbf{p}}^r), \tilde{y}^l) = (C(\tilde{x}^r, \tilde{y}^r), \tilde{y}^r) \quad (13)$$

should be a continuous mapping from \mathbb{R}^2 to \mathbb{R}^2 . Consider some point pair $(\tilde{\mathbf{p}}_0^r, \tilde{\mathbf{p}}_0^l)$, which is believed to be a corresponding pair. A circular tour is a sequence of points on the image, starting from $\tilde{\mathbf{p}}_0^r$ and ending in it as well, $\tilde{\mathbf{p}}_0^r, \tilde{\mathbf{p}}_1^r, \dots, \tilde{\mathbf{p}}_n^r = \tilde{\mathbf{p}}_0^r$, where $(d_x)_j = \tilde{x}_j^r - \tilde{x}_{j-1}^r$ and $(d_y)_j = \tilde{y}_j^r - \tilde{y}_{j-1}^r$ are all small. The corresponding sequence of points, starting from $\tilde{\mathbf{p}}_0^l$, is specified by the following incremental corrections, $\tilde{x}_j^l = \tilde{x}_{j-1}^l + (d_x)_j C_{\tilde{x}} + (d_y)_j C_{\tilde{y}}$ and $\tilde{y}_j^l = \tilde{y}_{j-1}^l + (d_y)_j$. The final point $\tilde{\mathbf{p}}_n^l$ should be close to $\tilde{\mathbf{p}}_0^l$. Equivalently, all different paths starting at one point $\tilde{\mathbf{p}}_0^r$, and ending at another point $\tilde{\mathbf{p}}_n^r$, induce cor-

responding (different) paths which should end at the same point (see Figure 4.1(b)). The distance measured between the different end points is denoted the circular tour distance.

Note that while the circular tour constraint is testing the smoothness of the inferred surface, it also tests the correctness of the initial corresponding pair $(\tilde{\mathbf{p}}_0^r, \tilde{\mathbf{p}}_0^l)$, the global parameters, and the implied surface derivatives. We found that testing such closure is indeed a reliable indication for these value. In particular using this closure constraint is effective in eliminating accumulated errors within the correspondence propagation process.

4.3. Integrating Photometry and Geometry in a local differential constraint

The following differential version of the circular tour constraint, relates the intensity derivatives to the geometry, in a more local way. Unlike section 3.4, where a frontal view was considered, it is necessary here to deal with different foreshortening effects at the two points. Let $I^r(\tilde{\mathbf{p}}^r)$ and $I^l(\tilde{\mathbf{p}}^l)$ be the intensity at two corresponding points. In section 3.4 we have shown that the partial derivatives of I^r and I^l with respect to x and y can be expressed as functions of \mathbf{l} and the local shape derivatives $z_x, z_y, z_{xx}, z_{xy}, z_{yy}$. In that simpler case expressing the partial derivatives of I^r and I^l with respect to the image coordinates \tilde{x} and \tilde{y} was immediate because all the derivatives of object coordinates with respect to image coordinates were -1, 1 or 0. Here using the general derivatives (Eq. 10) yields the following relation:

$$\begin{aligned} I^r_{\tilde{x}^r} &= \mathbf{l} \cdot \hat{\mathbf{n}}_x^r / (\cos(\theta) + z_x \sin(\theta)); \\ I^l_{\tilde{x}^l} &= \mathbf{l} \cdot \hat{\mathbf{n}}_x^l / (-\cos(\theta) + z_x \sin(\theta)); \\ I^r_{\tilde{y}^r} &= \mathbf{l} \cdot \hat{\mathbf{n}}_y^r - z_y \sin(\theta) I^r_{\tilde{x}^r}; \\ I^l_{\tilde{y}^l} &= \mathbf{l} \cdot \hat{\mathbf{n}}_y^l - z_y \sin(\theta) I^l_{\tilde{x}^l}. \end{aligned} \quad (14)$$

These four equations, denoted the *integration constraint*, have the same structure as in the frontal case (Eq. 7). They are also linear in the three local second derivatives of the surface z_{xx}, z_{xy} and z_{yy} . Their coefficients are the illumination vector multiplied by a tensor whose components are functions of the first derivatives of the surface and the pose parameter.

When the global parameters are known a self-corresponding point can be found by eliminating from the equations all the local partial derivatives of z leaving us with a single integration constraint in the image intensity values and their first derivatives at the two points. These values can be searched in the image. Conversely when given four corresponding pairs we can use these equations to solve for the unknown global parameters.

These differential relations locally specify the combined geometric-photometric constraint. We found however that the (equivalent) circular tour constraint is more numerically stable and therefore used it in the algorithm.

5. The reconstruction algorithm

Basically, the algorithm consists of the following stages:

1. **Recover unknown illumination and pose:** Start here when illumination and pose are not known.
 - (a) **Choose candidates for illumination, albedo, pose (θ) and a few (1-3) corresponding pairs:** The choice is done as a part of an optimization process which is either continuous or just discriminates between a few options.
 - (b) **Verify hypothesis:** Using the circular tour constraint (or other constraints described in Section 6.4), test the hypothesized parameters. If the constraint is not satisfied return to step 1a.
2. **Estimate and verify dense self-correspondence:** Start here when illumination and pose are known. Starting from a correct match found in the previous stage, propagate the self-correspondence by evaluating the self correspondence function derivative. Occasionally, verify and possibly correct corresponding pairs using the circular tour constraint.
3. **Shape recovery:** Use Eq. 3 to recover the 3D shape from corresponding pairs.

The following notes describe some aspects of the algorithm:

1. **Searching for the global parameters:** The naive approach to implement the iterative process of stages 1a-1b, would be to perform a five parameter search for the illumination components

$\rho_x^l, \rho_y^l, \rho_z^l, \theta, \tilde{x}^l$. This search uses the circular tour distance which vanishes for the correct parameters. In general a small dimensional minimization process is not a difficult task, however, here we may reduce the search cost by exploiting the extrema of the image intensities, as described below. In our experiments we usually used these shortcuts for finding a starting point and then fine tuned the illumination direction using the common Levenberg-Marquardt optimization procedure.

2. **Shortcuts for accelerating the search:** First we use the common heuristic that the maximal intensity corresponds to a normal parallel to the illumination direction. Knowing the maximal intensity there reduces the search to a two dimensional manifold.
3. **Dense correspondence recovery:** Once the global parameters are known, the circular tour constraint, may be used to find more pairs of corresponding points. Starting from a corresponding pair of points, the algorithm calculates the surface normal and, using θ , solves for $C_{\tilde{x}^r}$. Now the algorithm starts a propagation step: a new point, close to the right point is chosen, and an estimated position of its corresponding left point is determined using Eq. 12. Usually, the position is accurate enough, but occasionally the circular tour constraint may be called upon to fine tune it.

Special algorithms, which are not described here, were developed for cylindrical (bilaterally symmetric) objects. Those are characterized by a surface $z = f(x)$, implying that neither the object nor the intensity change in the y direction. In this case, stronger constraints are available, enabling a faster, less complex, algorithm. Knowing two corresponding point pairs, for example, allows to solve for the illumination direction, and then, to find the dense correspondence using a local, grey level based, criterion. Then the pose parameter θ is easily found as well. This algorithm is described in [21].

6. Experiments

In this section we present the results of running the algorithm on simulated and real data.

6.1. Results on simulated data

We now test our algorithm on the image shown in Figure 5(d), synthesized from the object shown in Figure 5(a). In order to speed up the minimization procedure, we recovered the global pose and lighting parameters by using shortcuts to obtain initial estimates for them. The minimization process then uses these initial estimates to fine-tune these global parameters.

6.1.1. Recovering the light source We estimate ρ as the maximal intensity in the image and then use the boundary points of the object in the image to estimate \mathbf{l} . There are two types of boundary points: occlud-

ing boundary points where $\mathbf{n} \cdot \mathbf{v} = 0$, where \mathbf{v} is the viewing direction and shading boundary points where $\mathbf{n} \cdot \mathbf{l} = 0$. They can be distinguished by the fact that the intensity close to the shading boundary points reduces gradually to zero, where as at occluding contour the intensity can drop abruptly to zero. We are only interested in occluding boundary points (Figure 5(c)). At these points the normal is known and lies in the image plane. We extrapolate the intensity at the boundary point from neighboring points and obtain at each point a linear equation in the components of lighting direction l_x and l_y . Using a least-squares procedure we recover l_x and l_y and obtain l_z from the unity constraint.




Fig. 5. (a) The ground truth of the 3D structure of the object studied. (b) The results of the shape recovery algorithm. (c) A frontal view and frontal illumination image of a symmetric object. This image is not used in the reconstruction. It is brought here only to help visualizing the object better. (d) An image of the object taken with $\theta = 0.2$ rad. (e) The occluding contour of the object.

6.1.2. Recovering the Pose Parameter We recover θ by performing a two parameter numerical search for corresponding points for several (10) points chosen arbitrarily. For each point and θ value, the corresponding point, which minimizes the circular tour measure was found. Figure 6(a) is a typical plot of the measure as a function of θ for one of these points. Then, the value of θ and the corresponding minimal measure value,

were collected from all these points (see Figure 6(b)). Note that most of the θ 's are close to the correct value of 0.2rad and have low scores. The θ we chose was the value found by the mean shift mode operator (the most dense point of the distribution of θ 's) which was $\theta = 0.204$ rad. It is important to note that this value of θ was achieved with the approximate value of \mathbf{l} found in the previous step.

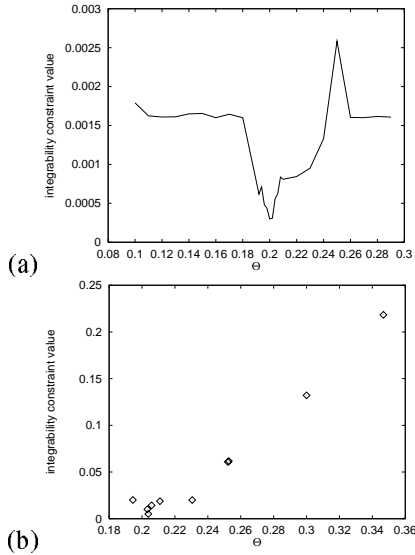


Fig. 6. (a) The results of searching for θ using 10 initial matches where from each point a sequence of twenty pairs of corresponding points are propagated. The y coordinate is the value of the aggregated integration constraint. If for a certain value of θ the value of the integration constraint exceeded 0.016 the computation terminated with a failure. (b) The candidate θ 's suggested by 10 arbitrary points. These are verified on 10 other sequences of length 10. Note that even in the initial stage many results were close to $\theta = 0.2$ rad and the best result was achieved for $\theta = 0.204$ rad.

Once the global parameters are estimated a global optimization procedure is run to fine-tune their values minimizing the value of the integration constraint.

6.1.3. Running of the algorithm In the final step the 3D shape is recovered using a combination of the propagation procedure and the integration constraint to correct correspondence errors accumulated during the procedure. Once the correspondences are found the 3D structure is recovered using Eq. 3 without needing to perform integration on the recovered normals [13]. Figure 5(a) shows the ground-truth of the 3D structure and Figure 5(b) shows the recovered shape.

6.2. Real image experiment: mono-photometric-stereo

Next we tested the mono-photometric-stereo algorithm on a frontal image of a face with non-frontal and ambient illumination (Figure 7(a)). The algorithm was given a rough estimate of the lighting information and the location of the symmetry line in the image. The mono-photometric-stereo algorithm was used to recover the

normals. Points for which there was no solution were left blank. This was mainly due to changes in albedo and due to shadows. As this is a frontal view Eq. 3 could not be used as in the previous example. Instead a simple normal integration procedure was employed to recover the shape. The results are displayed by overlaying a constant height contour map over the original image (Figure 7(b)). As expected, the results are only qualitatively correct and could be improved when a more exact reflectance model specially developed for human skin is used.




Fig. 7. (a) A frontal image of a face with non-frontal illumination. (b) An image which overlays a contour map over the original image. White points in the image represent points for which no normals could be recovered. This is usually due to albedo changes.

6.3. Real image experiments: mannequin images

In another experiment (Figure 8) we took several images of a mannequin painted with matte white paint to produce a close approximation to a Lambertian surface. The mannequin is rotated yielding images with different values of θ and l . In this section we present shape reconstruction results in which the global parameters were given to the algorithm and in the next section we present an experiment which computes them automatically.

For each image we estimated the albedo as the maximum intensity and gave the algorithm the lighting direction and pose parameter. The algorithm was provided with the point on the tip of the nose, which corresponds to itself. We assume that such a point or some other pair of corresponds points can be specified using face feature detection processes (e.g. eye, nose or mouth detection algorithms). Other methods of automatically finding a small number of corresponding pairs can also be used (e.g. using points above or below the corners of the eyes). We applied the propagation procedure to this “pair” of points recovering the normals and the 3D coordinates in object coordinates at each point. We used a slightly enhanced version of

the general algorithm, where the propagation procedure was run on the left and right sides of the image independently. This way, we obtained a corresponding point to every pixel center in the image (and not only to the pixels centers on the right side). This had also the advantage that incorrectly chosen global parameters result in a large discontinuity between the two sides of the recovered shape and can be detected automatically. Two views of the recovered shape for each of the experiments is results is shown in Figure 8. As we can see, the results are mixed, and some reconstructions are better than others. As expected, there are problems where shadows are dominant. We associated the inac-

curacy of the results mainly to these shadows and to the approximate model which does not take into account non-Lambertian effects and inter-reflections.

To test the quality of the recovered results we compared two horizontal cross-sections of the recovered shape in the four results, one on the forehead and the other on the nose beneath the eyes. The cross-sections for the forehead are very close in the four surfaces. This due to the fact that for convex surfaces with no inter-reflections and shadows the algorithm works well. The cross-sections for the nose are less close. This is due to errors in the propagation due to shadows in the eye regions. Even so the results are qualitatively similar.

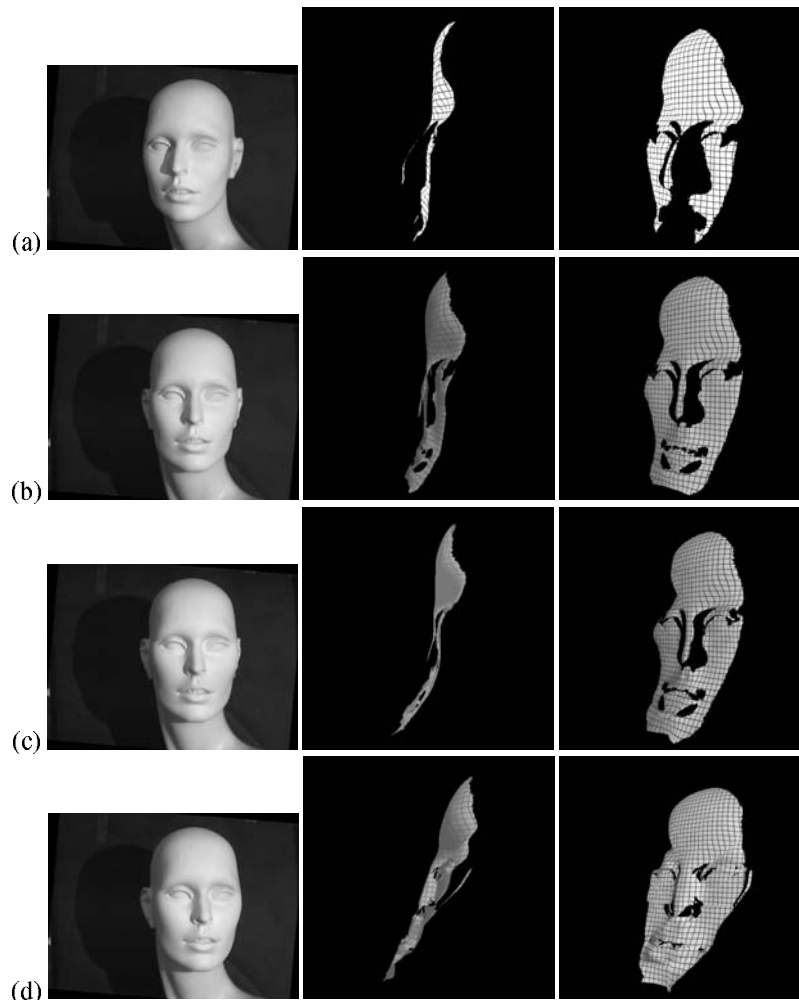


Fig. 8. Several images of a Mannequin, a profile view and another view of the recovered shapes. (a) $\theta = -0.165$ (b) $\theta = -0.005$ (c) $\theta = 0.10$ (d) $\theta = 0.22$

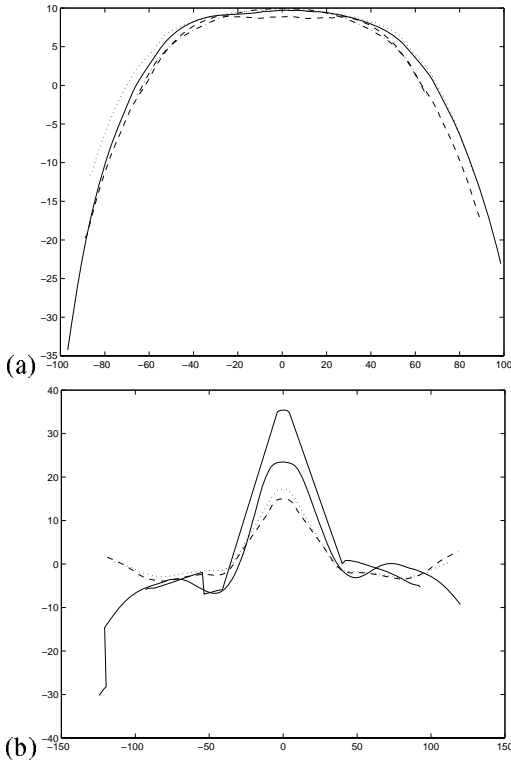


Fig. 9. Two horizontal cross-sections of the recovered shapes in Fig 8; (a) Forehead cross-section; (b) Nose cross-section. The line-styles of the cross-sections of the four images in Fig 8(a-d) are: solid, dotted, dashed and dash-dot, respectively.

6.4. Recovering Global Parameters

In this section we focus on the recovery of the lighting direction \mathbf{l} and the pose parameter θ . We examine a measure characterizing the quality of reconstruction. This measure and the reconstruction itself, naturally depend on the hypothesized global parameters. We demonstrate that this measure gets its minimum only when the hypothesized global parameters are close to their correct values. Therefore, evaluating this measure within the algorithm may be used for estimating the global parameters reliably.

These experiments were done using a new image of the mannequin (see Figure 10 for the image and two views of the reconstruction).

The quality measure is composed of two components. The first component is the number of pixels appearing in the reconstructed surface. A pixel is not considered a part of the reconstruction if \mathbf{l} and the intensity pair I_l, I_r do not satisfy the constraints described in Section 3.3. This can result from incorrect correspon-

dence or an incorrect \mathbf{l} . This component eliminates most of the \mathbf{l}, θ values which are not close to the correct values.




Fig. 10. An image and two views of the recovered shape. This image was used in recovering θ and \mathbf{l} .

The second component relies on additional corresponding pairs (which are different than those used to initialize the propagation). If the reconstruction is correct, then the true value of $C(\tilde{\mathbf{p}})$ should agree with its propagated value, resulting from the algorithm. Thus, the second component is the distance between the calculated and the actual location of the left point, corresponding to $\tilde{\mathbf{p}}$. (Getting a few corresponding pairs is usually easy. We used, for example, the forehead points above the center of the eye.)

In the experiment, we scanned 1701 values of global parameters, corresponding to a 3D grid, where two coordinates being the deviations of \mathbf{l} about the ground truth direction, and the third coordinate being θ . The angle between the correct illumination vector \mathbf{l} and the scanned values was up to 30° and θ ran between -0.5 and 0.5 . The correct value of θ was -0.2 . Three additional corresponding pairs were used for the second component. Values which were very far from the correct values were easily discarded, usually by failing the second test (985 from 1701).




Fig. 11. Two dimensional projections of the 3D space (θ, l_1, l_2) . Better results are depicted by lighter gray values. (a) Projection to the (l_1, l_2) space of the first criterion. (b) Projection to the (l_1, θ) space of the first criterion. (c) Projection to the (l_1, θ) space of the second criterion for one additional point pair. (d) Projection to the (l_1, l_2) space of the second criterion for the sum of the distances of three point pairs. (e) Projection to the (l_1, θ) space of the second criterion for the sum of the distances of three point pairs.

We considered several combinations of these two components, and got good results with all of them. To present the results we have projected the 3D space (θ, l_1, l_2) onto 2D subspaces and show the values obtained in Figure 11. In these figures the measure was scaled so that higher values represent higher quality. The projection was done by taking the maximal measure over all values of the hidden parameter. In Figure 11. The first two images show the projection of the first criterion. Good results are obtained close to the correct values but also for some extreme values of l_1 and l_2 . The second criterion (for one point) is displayed in Figure 11(c). Better results are obtained when the sum of distances associated with three point pairs are added (Figure 11(d-e)). Note that this measure gets its maximal value close to the correct values of l and θ , implying that automatic recovery of them is possible.

6.5. Degenerate Cases

In the previous sections we described three variants of the algorithm for three different cases. The case dealing with a frontal view of the object (mono-photometric-stereo), the case dealing with frontal illumination with respect to the object (mono-geometric-stereo), and the general algorithm for non-frontal illumination and a non-frontal image. Specialized algorithms were proposed for the special cases. We also tested how does the general algorithm perform on images close to these singular configurations. We have found experimentally, that for images which are close to the frontal view (e.g Figure 8(b)), the algorithm still provides reasonable results. On the other hand, when the configuration gets close to the other singular configuration (shown in Figure 3), the reconstruction quality deteriorates much more severely.

The reason for this difference in performance around the two singular configurations is as follows. The general algorithm uses the photometric information to propagate the correspondences. Then for each corresponding pair the 3D shape is recovered using the geometric information. Thus, around the first singular configuration the photometric component of the algorithm works correctly yielding correct correspondences. There is some inaccuracy in the recovered 3D shape due to the fact that we are near a singular point of the geometric constraints. However, these errors do not accumulate to the neighboring pixels. Around the second singular configuration, the first photometric constraint which is the difference between the two corresponding intensities is very close to zero, yielding a very unstable solution to this equation, resulting in a very inaccurate surface normal which cause errors in the self-correspondence function to accumulate. This results in a very inaccurate correspondence function which yields a very inaccurate recovered shape. Such a bad result should be detectable by comparing the given image with the Lambertian image implied by the reconstructed model. Here the mono-geometric-stereo algorithm should be used.

7. Summary and discussion

In this paper we have presented a shape recovery method for bilaterally symmetric objects. The major contributions of the paper are a strong, integrated, local constraint which combines photometric and geometric

information, and its application for extracting dense correspondence, for the implied correspondence-less 3D shape reconstruction algorithm, and for recovering unknown lighting and viewing parameters.

We focused on the theoretical aspects of integrating photometric and geometric constraints. We assume a Lambertian reflectance function which is well defined and an easy model to analyze. The theoretical results were supported by empirical experiments. The experiments on images of real objects were limited by the constant albedo assumption, the Lambertian reflectance function, and the assumption that neither occlusions nor shadows are present. We believe that we can overcome these limitation in future studies. Regarding occlusion and shadows: these are usually characterized by steep edges, which we can try to detect, and eliminate from the path taken by the propagation process. We also intend to extend our method to other reflectance functions that should be studied separately. Another direction is to extend our approach to non-constant albedo surfaces such that the albedo function is restricted (otherwise any object can project to any image). A relatively simple extension of our method to non-constant albedo surfaces is to use the local properties of our method to handle surfaces with patch-wise constant albedo. In future work we intend to study this and other non-constant albedo models.

Acknowledgements

We would like to thank Dalit Ramon for donating the mannequin used in the experiments. Yael Moses would like to thank Mike Brady for useful discussion in early stages of this study.

References

1. R. Basri and Y. Moses. When is it possible to identify 3D objects from single images using class constraints? *Int. J. of Comp. Vision*, 33(2):1–22, September 1999.
2. A. Blake, A. Zisserman, and G. Knowles. Surface description from stereo and shading. *Image Vision Computation*, 3(4):183–191, 1985.
3. T. Cham and R. Cipolla. Symmetry detection through local skewed symmetries. *Image and Vision Computing Journal*, 1995.
4. J.E. Cryer, P.S. Tsai, and M. Shah. Integration of shape from shading and stereo. *PR*, 28(7):1033–1043, July 1995.
5. O.D. Faugeras. *Three-Dimensional Computer Vision*. MIT Press, Boston MA, 1993.
6. S.A. Freiberg. Finding axes of skewed symmetry. *Computer Vision Graphics and Image Processing*, 34:138–155, 1986.
7. P. Fua and Y. Leclerc. Object-centered surface reconstruction: Combining multi-image stereo and shading. In *Proc. DARPA Image Understanding Workshop*, pages 1097–1120, 1993.
8. Y. Gofman and N. Kiryati. Detecting symmetry in grey level images: The global optimization approach. In *ICPR*, page A94.2, 1996.
9. G. G. Gordon. Shape from symmetry. In *Proc. of SPIE Intelligent Robots And Computer Vision VIII*, volume 1192, November 1989.
10. A.D. Gross and T.E. Boult. Recovery of SHGCs from a single intensity view. *IEEE Trans. Patt. Anal. Mach. Intell.*, 18(2):161–180, February 1996.
11. H. Hayakawa. Photometric stereo under a light-source with arbitrary motion. *J. Opt. Soc. Am. A*, 11(11):3079–3089, November 1994.
12. B.K.P. Horn. *Computer Vision*. MIT Press, Cambridge, Mass., 1986.
13. B.K.P. Horn and M.J. Brooks. The variational approach to shape from shading. *Comp. Vis. Graph. Im. Proc.*, 33:174–203, 1986.
14. B.K.P. Horn and M.J. Brooks. *Shape from Shading*. MIT Press, 1989.
15. D.R. Hougen and N. Ahuja. Estimation of the light source distribution and its use in integrated shape recovery from stereo and shading. In *Proc. Int. Conf. Comp. Vision*, pages 148–155, 1993.
16. Y. Moses and S. Ullman. Generalization to novel views: Universal, class-based, and model-based processing. *Int. J. of Comp. Vision*, 29(3):233–253, September 1998.
17. D.P. Mukherjee, A. Zisserman, and J.M. Brady. Shape from symmetry - detecting and exploiting symmetry in affine images. Technical Report Report No. OUEL 1988/93, University of Oxford, Department of Engineering Science, 1993.
18. R. Onn and A.M. Bruckstein. Integrability disambiguates surface recovery in two-image photometric stereo. *IJCV*, 5(1):105–113, August 1990.
19. J. Ponce. On characterizing ribbons and finding skewed symmetries. *Comp. Vis. Graph. Im. Proc.*, 52:328–340, 1990.
20. C.A. Rothwell. Object recognition through invariant indexing. In *Oxford University Press*, 1995.
21. I. Shimshoni, Y. Moses, and M. Lindenbaum. Auto photometric/geometric stereo from a single image or class based reconstruction of bilaterally symmetric objects. Technical Report CIS-9710-1997, The Technion, 1997.
22. I. Shimshoni, Y. Moses, and M. Lindenbaum. Shape reconstruction of 3D bilaterally symmetric surfaces. In *Proc. Int. Conf. Image Ana. and Proc.*, pages 76–81, 1999.
23. B. Takacs and H. Wechsler. Detection of faces and facial landmarks using iconic filter banks. *PR*, 30(10):1623–1636, October 1997.
24. R. J. Woodham. Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proc. SPIE*, 155:136–143, 1978.
25. H. Zabrodsky. *Computational aspects of pattern characterization - continuous symmetry*. PhD thesis, Hebrew University, Jerusalem, 1993.
26. Z. Zhang and H.T. Tsui. 3D reconstruction from a single view of an object and its image in a plane mirror. In *Proc. International Conf. Patt. Recog.*, page SAP1, 1998.
27. A. Zisserman, D. Forsyth, J. Mundy, C. Rothwell, J. Liu, and N. Pillow. 3D object recognition using invariance. *AI*, 78(1-2):239–288, October 1995.