

SHAPE FROM SHADINGS UNDER PERSPECTIVE PROJECTION AND TURNABLE MOTION

Miaomiao Liu and Kwan-Yee K. Wong

Department of Computer Science, The University of Hong Kong, Pokfulam, Hong Kong

Keywords: Two-frame-theory, Shape recovery, Turntable motion.

Abstract: Two-Frame-Theory is a recently proposed method for 3D shape recovery. It estimates shape by solving a first order quasi-linear partial differential equation through the method of characteristics. One major drawback of this method is that it assumes an orthographic camera which limits its application. This paper re-examines the basic idea of the Two-Frame-Theory under the assumption of a perspective camera, and derives a first order quasi-linear partial differential equation for shape recovery under turntable motion. Dynamic programming is used here to provide the Dirichlet boundary condition. The proposed method is tested against synthetic and real data. Experimental results show that perspective projection can be used in the framework of Two-Frame-Theory, and competitive results can be achieved.

1 INTRODUCTION

Shape recovery is a classical problem in computer vision. Many constructive methods have been proposed in the literature. They can generally be classified into two categories, namely multiple-view methods and single-view methods. Multiple-view methods such as structure from motion (Tomasi, 1992) mainly rely on finding point correspondences in different views, whereas single-view methods such as photometric stereo use shading information to recover the model.

Multiple-view methods can be further divided into point-based methods and silhouette-based methods. Point-based methods are the oldest technique for 3D reconstruction (Pollefeys et al., 2001). Once feature points across different views are matched, the shape of the object can be recovered. The major drawback of such methods is that they depend on finding point correspondences between views. This is the well-known correspondence problem which itself is a very tough task. Moreover, point-based methods do not work for featureless object. On the other hand, silhouette-based methods are a good choice for shape recovery of featureless object. Silhouettes are a prominent feature in an image, and they can be extracted reliably even when no knowledge about the surface is available. Silhouettes can provide rich information for both the shape and motion of an object (Wong and Cipolla, 2001; Liang and Wong, 2005).

Nonetheless, only sparse 3D points or a very coarse visual hull can be recovered if the number of images used for reconstruction is comparatively small. Photometric stereo, which is a single-view method, uses images taken from one fixed viewpoint under at least three different illumination conditions. No image correspondences are needed. If the albedo of the object and the lighting directions are known, the surface orientations of the object can be determined and the shape of the object can be recovered via integration (Woodham, 1980). However, most of the photometric stereo methods consider orthographic projection. Few works are related to perspective shape reconstruction (Tankus and Kiryati, 2005). If the albedo of the object is unknown, photometric stereo may not be feasible. Very few studies in the literature use both shading and motion cues under a general framework. In (Jin et al., 2008), the 3D reconstruction problem is formulated by combining the lighting and motion cues in a variational framework. No point correspondences is needed in the algorithm. However, the method in (Jin et al., 2008) is based on optimization and requires piecewise constant albedo to guarantee convergence to a local minimum. In (Zhang et al., 2003), Zhang et al. unified multi-view stereo, photometric stereo and structure from motion in one framework, and achieved good reconstruction results. Their method has a general setting of one fixed light source and one camera, but with the assumption of an orthographic camera model.

Similar to (Jin et al., 2008), the method in (Zhang et al., 2003) also greatly depends on optimization. A shape recovery method was proposed in (Moses and Shimshoni, 2006) by utilizing the shading and motion information in a framework under a general setting of perspective projection and large rotation angle. Nonetheless, it requires that one point correspondence should be known across the images and the object should have an uniform albedo.

A Two-Frame-Theory was proposed in (Basri and Frolova, 2008) which models the interaction of shape, motion and lighting by a first order quasi-linear partial differential equation. Two images are needed to derive the equation. If the camera and lighting are fully calibrated, the shape can be recovered by solving the first order quasi-linear partial differential equation with an appropriate Dirichlet boundary condition. This method does not require point correspondences across different views and the albedo of the object. However, it also has some limitations. For instance, it assumes an orthographic camera model which is a restrictive model. Furthermore, as stated in the paper, it is hard to use merely two orthographic images of an object to recover the angle of out-of-plane rotations.

This paper addresses the problem of 3D shape recovery under a fixed single light source and turntable motion. A multiple-view method that exploits both motion and shading cues will be developed. The fundamental theory of the Two-Frame-Theory will be re-examined under the more realistic perspective camera model. Turntable motion with small rotation angle is considered in this paper. With this assumption, it is easy to control the rotation angle compared to the setting in (Basri and Frolova, 2008). A new quasi-linear partial differential equation under turntable motion is derived, and a new Dirichlet boundary condition is obtained using dynamic programming. Competitive results are achieved for both synthetic and real data.

This paper is organized as follows: Section 2 describes the derivation of the first order quasi-linear partial differential equation. Section 3 describes how to obtain the Dirichlet boundary condition. Section 4 shows the experimental results for synthetic and real data. A brief conclusion is given in Section 5.

2 FIRST ORDER QUASI-LINEAR PDE

Turntable motion is considered in this paper. Similar to the setting in (Basri and Frolova, 2008), two images are used to derive the first order quasi-linear partial differential equation (PDE). Suppose that the object rotates around the Y-axis by a small angle. Let

$\mathbf{X} = (X, Y, Z)$ denotes the 3D coordinates of a point on the surface. The projection of \mathbf{X} in the first image is defined as $\lambda_i \bar{\mathbf{x}}_i = P_i \bar{\mathbf{X}}$, where P_i is the projection matrix, $\bar{\mathbf{X}}$ is the homogenous coordinates of \mathbf{X} and $\bar{\mathbf{x}}_i$ is the homogenous coordinates of the image point. Similarly, the projection of \mathbf{X} after the rotation is defined as $\lambda_j \bar{\mathbf{x}}_j = P_j \bar{\mathbf{X}}$, where P_j is the projection matrix after the rotation. The first image can be represented by $I(x_i, y_i)$, where (x_i, y_i) is the inhomogeneous coordinates of the image point. Similarly the second image can be represented by $J(x_j, y_j)$. Let the surface of the object be represented by $Z(X, Y)$. Suppose that the camera is set on the negative Z-axis. The unit normal of the surface is denoted by $\mathbf{n}(X, Y) = \frac{(Z_X, Z_Y, -1)^T}{\sqrt{Z_X^2 + Z_Y^2 + 1}}$, where $Z_X = \frac{\partial Z}{\partial X}$ and $Z_Y = \frac{\partial Z}{\partial Y}$. After rotating the object by a small angle θ around the Y-axis, the normal of the surface point becomes

$$\begin{aligned} \mathbf{n}_\theta(X, Y) &= \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix} \begin{pmatrix} Z_X \\ Z_Y \\ -1 \end{pmatrix} \\ &= \begin{pmatrix} Z_X \cos \theta - \sin \theta \\ Z_Y \\ -Z_X \sin \theta - \cos \theta \end{pmatrix}. \end{aligned} \quad (1)$$

Directional light is considered in this paper and it is expressed as a vector $\mathbf{l} = (l_1, l_2, l_3)^T$. Since the object is considered to have a lambertian surface, the intensities of the surface point in the two images are given by

$$I(x_i, y_i) = \rho \mathbf{l}^T \mathbf{n} = \frac{\rho(l_1 Z_X + l_2 Z_Y - l_3)}{\sqrt{Z_X^2 + Z_Y^2 + 1}}, \quad (2)$$

$$\begin{aligned} J(x_j, y_j) &= \rho \mathbf{l}^T \mathbf{n}_\theta \\ &= \frac{\rho((l_1 \cos \theta - l_3 \sin \theta) Z_X + l_2 Z_Y - l_1 \sin \theta - l_3 \cos \theta)}{\sqrt{Z_X^2 + Z_Y^2 + 1}}. \end{aligned} \quad (3)$$

where ρ is the albedo for the current point. If θ is very small, (3) can be approximated by

$$J(x_j, y_j) \approx \frac{\rho((l_1 - l_3 \theta) Z_X + l_2 Z_Y - l_1 \theta - l_3)}{\sqrt{Z_X^2 + Z_Y^2 + 1}}. \quad (4)$$

The albedo and the normal term (denominator) can be eliminated by subtracting (2) from (4), and dividing the result by (2). This gives

$$\begin{aligned} (l_1 Z_X + l_2 Z_Y - l_3)(J(x_j, y_j) - I(x_i, y_i)) \\ = (-l_1 \theta - l_3 Z_X \theta) I(x_i, y_i). \end{aligned} \quad (5)$$

Note that some points may become invisible after rotation and that the correspondences between image

points are unknown beforehand. If the object has a smooth surface, the intensity of the 3D point in the second image can be approximated by its intensity in the first image through the first-order 2D Taylor series expansion:

$$J(x_j, y_j) \approx J(x_i, y_i) + J_x(x_i, y_i)(x_j - x_i) + J_y(x_i, y_i)(y_j - y_i). \quad (6)$$

Therefore,

$$J(x_j, y_j) - I(x_i, y_i) \approx J(x_i, y_i) - I(x_i, y_i) + J_x(x_i, y_i)(x_j - x_i) + J_y(x_i, y_i)(y_j - y_i). \quad (7)$$

Substituting (7) into (5) gives

$$(l_1 V_{JI} + l_3 I(x_i, y_i) \theta) Z_X + l_2 V_{JI} Z_Y = -l_1 \theta I(x_i, y_i) \quad (8)$$

where

$$V_{JI} = J(x_i, y_i) - I(x_i, y_i) + J_x(x_i, y_i)(x_j - x_i) + J_y(x_i, y_i)(y_j - y_i).$$

Note that x_i , y_i , x_j , and y_j are functions of X , Y , and Z , and (8) can be written more succinctly as

$$a(X, Y, Z) Z_X + b(X, Y, Z) Z_Y = c(X, Y, Z) \quad (9)$$

where

$$\begin{aligned} a(X, Y, Z) &= l_1 V_{JI} + l_3 I(x_i, y_i) \theta, \\ b(X, Y, Z) &= l_2 V_{JI}, \\ c(X, Y, Z) &= -l_1 \theta I(x_i, y_i). \end{aligned} \quad (10)$$

(9) is a first-order partial differential equation in $Z(X, Y)$. Furthermore, it is a quasi-linear partial differential equation since it is linear in the derivatives of Z , and its coefficients, namely $a(X, Y, Z)$, $b(X, Y, Z)$, and $c(X, Y, Z)$, depend on Z . Therefore, the shape of the object can be recovered by solving this first order quasi-linear partial differential equation using the method of characteristics. The characteristic curves can be obtained by solving the following three ordinary differential equations:

$$\begin{aligned} \frac{dX(s)}{ds} &= a(X(s), Y(s), Z(s)), \\ \frac{dY(s)}{ds} &= b(X(s), Y(s), Z(s)), \\ \frac{dZ(s)}{ds} &= c(X(s), Y(s), Z(s)), \end{aligned} \quad (11)$$

where s is a parameter for the parameterization of the characteristic curves. (Basri and Frolova, 2008)

has given a detailed explanation of how the method of characteristics works. It is also noticed that the quasi-linear partial differential equation should have a unique solution. Otherwise, the recovered surface may not be unique. In the literature of quasi-linear partial differential equation, this is considered as the initial problem for quasi-linear first order equations. A theorem in (Zachmanoglou and Thoe, 1987) can guarantee that the solution is unique in the neighborhood of the initial boundary curve. However, the size of the neighborhood of the initial point is not constrained. It mainly depends on the differential equation and the initial curve. It is very important to find an appropriate Dirichlet boundary. In this paper, dynamic programming is used to derive the boundary curve.

3 BOUNDARY CONDITION

Under perspective projection, the Dirichlet Boundary condition cannot be obtained in the same way as in (Basri and Frolova, 2008). As noted in (Basri and Frolova, 2008), the intensities of the contour generator points are unaccessible. Visible points (X', Y', Z') nearest to the contour generator are a good choice for boundary condition. If the normal \mathbf{n}' of (X', Y', Z') is known, Z_X and Z_Y can be derived according to $\mathbf{n}(X, Y) = \frac{(Z_X, Z_Y, -1)^T}{\sqrt{Z_X^2 + Z_Y^2 + 1}}$. Note that (9) also holds for (X', Y', Z') . Therefore, (X', Y', Z') can be obtained by solving (9) with known Z_X and Z_Y values. The problem of obtaining (X', Y', Z') becomes the problem of how to obtain \mathbf{n}' .

Let (X, Y, Z) be a point on the contour generator (see Figure 1). Since (X, Y, Z) is a contour generator point, its normal \mathbf{n}_c must be orthogonal to its visual ray \mathbf{V} . Consider a curve C given by the intersection of the object surface with the plane π defined by \mathbf{V} and \mathbf{n}_c . In a close neighborhood of (X, Y, Z) , the angle between \mathbf{V} and the surface normal along C would change from just smaller than 90 degrees to just greater than 90 degrees. Now consider a visible point (X', Y', Z') on C close to (X, Y, Z) , its normal \mathbf{n}' should make an angle of just smaller than 90 degrees with \mathbf{V} . If (X', Y', Z') is very close to (X, Y, Z) , \mathbf{n}' would also be very close to \mathbf{n}_c . To simplify the estimation of \mathbf{n}' , it is assumed that \mathbf{n}' is coplanar with \mathbf{n}_c and lying on π . \mathbf{n}' can therefore be obtained by rotating \mathbf{n}_c around an axis given by the normal \mathbf{N} of π by an arbitrary angle γ . According to the lambertian law, \mathbf{n}' can be obtained by knowing the intensities of corresponding points across different views and the albedo. However, no prior knowledge of point cor-

correspondences and albedo ρ are available. Similar to the solution in (Basri and Frolova, 2008), ρ and γ can be obtained alternatively. The image coordinates of the visible point nearest to the contour generator can be obtained by searching along a line ℓ determined by the intersection of the image plane o and π . As for each γ in the range $0 < \gamma \leq \frac{\pi}{6}$, a corresponding albedo is computed and ρ is chosen as the mean of these computed values. γ is then computed by minimizing $(I - \rho \mathbf{I}^\top R(\gamma) \mathbf{n}_c)^2$, where $R(\gamma)$ is the rotation matrix defined by the rotation axis ℓ and rotation angle γ . \mathbf{n}' is finally determined as $R(\gamma) \mathbf{n}_c$ and (X', Y', Z') can be obtained by minimizing

$$E_{bon} = \|a(X', Y', Z')Z_X + b(X', Y', Z')Z_Y - c(X', Y', Z')\|^2. \quad (12)$$

Dynamic programming is used to obtain (X', Y', Z') . Two more constraints are applied in the framework of dynamic programming. One is called photometric consistency which is defined as

$$E_{pho}(p) = \sum_q (I_q(p) - \mathbf{I}^\top \mathbf{n}_q(p))^2, \quad (13)$$

where p is a 3D point on the contour generator, and $I_q(p)$ is a component of the normalized intensity measurement vector which is composed of the measurement of intensities in q neighboring views and $\mathbf{n}_q(p)$ is the normal of p in the q^{th} neighboring view.

The other constraint is called surface smoothness constraint which is defined as:

$$E_{con}(p, p') = \|pos(p) - pos(p')\|_2^2 \quad (14)$$

where $pos(p)$ denotes the 3D coordinates of point p and $pos(p')$ denotes the 3D coordinates of its neighbor point p' along the boundary curve.

The final Energy function is defined as

$$E = \mu E_{bon} + \eta E_{pho} + \omega E_{con} \quad (15)$$

where μ , η , ω are weighting parameters. By minimizing (15) for each visible point nearest to the contour generator, the nearest visible boundary curve can be obtained and used as Dirichlet Boundary condition for solving (9).

4 EXPERIMENTS

The proposed method is tested on synthetic models and real image sequence. The camera and light source are fixed and fully calibrated. Circular motion sequences are captured either through simulation or by rotating the object on a turntable for all experiments.

Since the derived first order quasi-linear partial differential equation is based on the assumption of small rotation angle, the image sequences are taken at a spacing of at most five degrees.

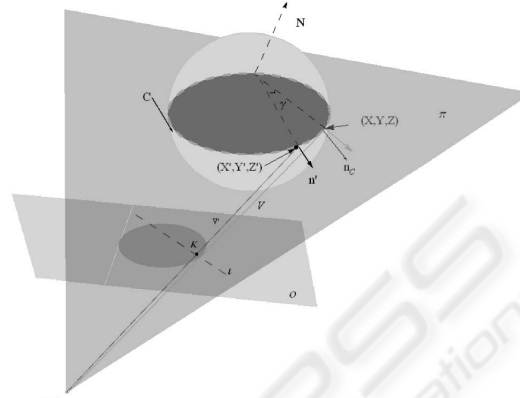


Figure 1: Computing boundary condition. The visible curve nearest to contour generator is a good choice for the boundary condition. The curve can be obtained by searching the points whose normal is coplanar with those of the respective contour generator points.




Figure 2: Shadow effect for shape recovery. Left column: original image for one view used for shape recovery and shadow appears near the left arm of the Venus model. Middle column: characteristic curves without using intensity difference threshold for the solution of partial differential equation. Right column: characteristic curves after using intensity difference threshold for the solution of partial differential equation.

It has been assumed that the synthetic model has a pure lambertian surface since the algorithm is derived according to the lambertian surface property. It is also quite important to get good boundary conditions for solving the partial differential equations.

As for synthetic experiment, it is easier to generate the image without specularities. However, it is hard to control the image sequence without shadows due to the lighting directions and the geometry of the object. From (9), (10), and (11), it can be noted that the solution of the first order quasi-linear partial differential equation mainly depends on the change of the intensities in the two images. If the projection of the point in the first view is in shadow, the estima-

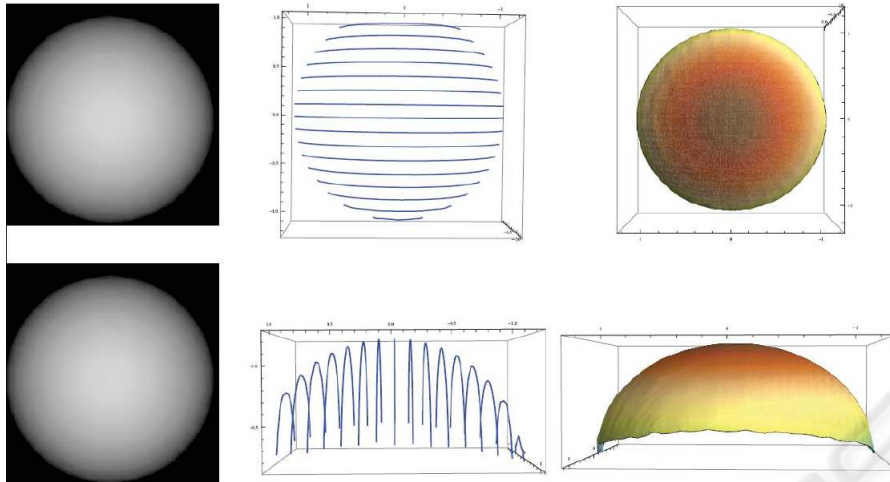


Figure 3: Reconstruction of the synthetic sphere model with uniform albedo. The rotation angle is 5 degrees. Light direction is $[0, 0, -1]^T$, namely spot light. Camera is set on the negative Z-axis. Left column: two original images. Middle column: characteristic curves for frontal view (top) and the characteristic surface observed in a different view (bottom). Right column: reconstructed surface for front view (top) and the reconstructed surface in a different view (bottom).




Figure 4: Shape recovery for the synthetic cat model. Left column: two images used for reconstruction. Middle column: characteristic curves for the view corresponding to the top image in the left column (top) and the characteristic curves observed in a different view (bottom). Right column: Recovered shape with shadings in two different views.

tion of the image intensity using Taylor expansion in the second view will also be inaccurate. The characteristic curves will go crazy (see Figure 2). In order to avoid the great error caused by shadow effect, the intensity difference for the estimation points should not be larger than a threshold δ_{thre} which is obtained through the experiments.

4.1 Experiment with Synthetic Model

The derived first order quasi-linear partial differential equation is applied to three models, namely the sphere model, cat model, and Venus model. $\delta_{thre} = 51$ is used for all the synthetic experiments to avoid the influence of shadow in the image.

The first simulation is implemented on the sphere model. Since the image of the sphere with uniform




Figure 5: Shape recovery for the Venus model. Left column: two images used for shape recovery. Middle column: characteristic curves for the view corresponding to the top image in the left column (top) and the characteristic curves observed in a different view (bottom). Right column: Recovered shape with shadings in two different views.

albedo does not change its appearance after its rotation around Y -axis (see Figure 3), only one image is used to recover the 3D shape. Any small angle can be used to the proposed equation. Nonetheless, if the sphere has non-uniform albedo, two images will be used for shape recovery. Figure 3 shows the original images, the characteristic curves observed in two different views, and the recovered 3D shape examined in two different views. The sphere model is assumed to have a constant albedo. The reconstruction error is tested by using $\|\hat{Z} - Z\|_2^2 / \|Z\|_2^2$, where \hat{Z} denotes the estimated depth for each surface point and $\|\cdot\|_2$ denotes the l_2 norm. The mean error is 4.63% for the sphere with constant albedo which is competitive to the result in (Basri and Frolova, 2008).

The second simulation is implemented on a cat model. The image sequence is captured under general lighting direction and with rotation angles at three degrees spacing. Seven images are used to get the boundary condition. Two neighboring images are used for solving the derived quasi-linear partial differential equation. The results are shown in Figure 4. It can be observed that the body of the cat can be recovered except a few errors appeared at the edge. The last simulation is implemented on a Venus model. The sequence is taken at a spacing of five degrees.

Similarly seven images are used for getting the boundary condition.

The front of the Venus model is recovered by using two neighboring view images for solving the derived quasi-linear partial differential equation. The result is shown in Figure 5.

After the recovery of characteristics curves, shapes for the Cat and Venus model with shadings are shown by using the existing points to mesh software VRmesh. The right columns of Figure 4, and Figure 5 show the results.

4.2 Experiment with Real Images

The real experiment is conducted on a ceramic mouse. The mouse toy is put on a turntable. The relative positions of the lighting and the camera is fixed. The image sequence is taken by a Cannon 450D camera with a 34 mm lens. The camera is calibrated using a chessboard pattern and a mirror sphere is used to calibrate the light. The image sequence is captured with the rotation angle at five degree spacing. It can be observed that there are specularities on the body of the mouse. The intensity threshold is used to eliminate the bad effect of the specularities since the proposed algorithm only works well on lambertian sur-




Figure 6: Shape recovery for a mouse toy. Left column: two images used for shape recovery. Middle column: characteristic curves for the view corresponding to the top image in the left column (top) and the characteristic curves observed in a different view (bottom). Right column: Recovered shape with shadings in two different views.

face. In order to avoid the shadow effect for solving the partial differential equation, $\delta_{thre} = 27$ is used which is smaller than the value used in the synthetic experiment since the captured images are comparatively darker. The original images and the characteristic curves are shown in Figure 6. Although the mouse toy has a complex topology and the images have shadows and specularities, good results can be obtained under this simple setting. Right column of Figure 6 shows the results.

5 CONCLUSIONS

This paper re-examines the fundamental ideas of the Two-Frame-Theory and derives a different form of first order quasi-linear partial differential equation for turntable motion. It extends the Two-Frame-Theory to perspective projection, and derives the Dirichlet boundary condition using dynamic programming. The shape of the object can be recovered by the method of characteristics. Turntable motion is considered in the paper as it is the most common setup for acquiring images around an object. Turntable motion also simplifies the analysis and avoids the difficulty of obtaining the rotation angle in (Basri and Frolova, 2008). The newly proposed partial differential equation makes the two frame method more useful for a more general setting. Although the proposed method is promising, it still has some limitations. For instance, the proposed algorithm cannot deal with object rotating with large angles. If the object rotates

with large angle, image intensities on the second image cannot be approximated by the two dimensional Taylor expansion. Some coarse-to-fine strategies can be used to derive new equations. In additions, the object is assumed to have a lambertian surface and the lighting is assumed to be a directional light source. The recovered surface can be used as an initialization for shape recovery method using optimization, which can finally get a full 3D model. Further improvement should be made before the method can be applied to an object under general lightings and general motion.

REFERENCES

- Basri, R. and Frolova, D. (2008). A two-frame theory of motion, lighting and shape. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 24-26 June 2008, Anchorage, Alaska, USA*, pages 1–7. IEEE Computer Society.
- Jin, H., Cremers, D., Wang, D., Prados, E., Yezzi, A. J., and Soatto, S. (2008). 3-d reconstruction of shaded objects from multiple images under unknown illumination. *International Journal of Computer Vision*, 76(3):245–256.
- Liang, C. and Wong, K.-Y. K. (2005). Complex 3d shape recovery using a dual-space approach. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 20-26 June 2005, San Diego, CA, USA*, pages 878–884.
- Moses, Y. and Shimshoni, I. (2006). 3d shape recovery of smooth surfaces: Dropping the fixed viewpoint assumption. In *Computer Vision - ACCV 2006*,

- 7th Asian Conference on Computer Vision, Hyderabad, India, January 13-16, 2006, Proceedings, Part I*, volume 3851 of *Lecture Notes in Computer Science*, pages 429–438. Springer.
- Pollefeys, M., Vergauwen, M., Verbiest, F., Cornelis, K., and Gool, L. V. (Jun. 2001). From image sequences to 3d models. In *In Proceedings 3rd international workshop on automatic extraction of man-made objects from aerial and space images*, pages 403–410.
- Tankus, A. and Kiryati, N. (2005). Photometric stereo under perspective projection. In *10th IEEE International Conference on Computer Vision, 17-20 October 2005, Beijing, China*, pages 611–616.
- Tomasi, C. (1992). Shape and motion from image streams under orthography: a factorization method. *International Journal of Computer Vision*, 9:137–154.
- Wong, K.-Y. K. and Cipolla, R. (2001). Structure and motion from silhouettes. In *8th International Conference on Computer Vision, July 7-14, 2001, Vancouver, British Columbia, Canada*, pages 217–222.
- Woodham, R. (Jan. 1980). Photometric method for determining surface orientation from multiple images. In *OptEng 19(1)*, pages 19(1):139–144.
- Zachmanoglou, E. C. and Thoe, D. W. (1987). *Introduction to partial differential equations with applications*. Dover Publications, New York:SUS.
- Zhang, L., Curless, B., Hertzmann, A., and Seitz, S. M. (2003). Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multi-view stereo. In *9th IEEE International Conference on Computer Vision, 14-17 October 2003, Nice, France*, pages 618–625.

