

Reconstruction of Specular Surfaces using Polarization Imaging

Stefan Rahmann and Nikos Canterakis
Institute for Pattern Recognition and Image Processing
Computer Science Department, University of Freiburg
79110 Freiburg, Germany
{rahmann,canterakis}@informatik.uni-freiburg.de

Abstract

Traditional intensity imaging does not offer a general approach for the perception of textureless and specular reflecting surfaces. Intensity based methods for shape reconstruction of specular surfaces rely on virtual (i.e. mirrored) features moving over the surface under viewer motion. We present a novel method based on polarization imaging for shape recovery of specular surfaces. This method overcomes the limitations of the intensity based approach, because no virtual features are required. It recovers whole surface patches and not only single curves on the surface. The presented solution is general as it is independent of the illumination. The polarization image encodes the projection of the surface normals onto the image and therefore provides constraints on the surface geometry. Taking polarization images from multiple views produces enough constraints to infer the complete surface shape. The reconstruction problem is solved by an optimization scheme where the surface geometry is modelled by a set of hierarchical basis functions. The optimization algorithm proves to be well converging, accurate and noise resistant. The work is substantiated by experiments on synthetic and real data.

1. Introduction

The problem of shape recovery for the class of specular reflection objects has attracted a lot of researchers in the past. However, no general solution has been found so far. This paper presents an approach based on polarization imaging which overcomes previous restrictions.

An important contribution was the work by Oren and Nayar [8]; they tracked virtual features traveling along the surface under viewer motion. From the known motion and the trajectory of the virtual feature the corresponding surface profile was computed. The main drawbacks are twofold. First, this approach is highly illumination depending since a virtual feature corresponding to a highlight on the surface

must be observed. A more diffuse or dynamic illumination will corrupt the analysis. Secondly, the surface geometry can only be computed at profiles, not on the whole surface under inspection.

The aim of local shape analysis was pursued for example in [2] and [4] based on stereo and reflectance modelling respectively. Among the techniques using active illumination we mention the work [18] where objects are rotated on a turn table, and the work [9] where structured light was used. There also were a lot of attempts to model the reflection properties incorporating specularities for photometric stereo, e.g in [7], or stereo approaches as in [1].

An alternative approach is the reconstruction of objects using apparent contours which has been performed for example in [3]; however, a necessary condition is the convexity of the surface.

The first approach for shape analysis from multiple polarization images has been achieved by Wolff in [16] where the orientation of a plane was computed from two views. This previous work is generalized by our contribution. Under ideal lighting condition and a known refraction index of the object, shape recovery from a single polarization image was carried out in [11]. For the special case of a transparent planar surface a method determining the orientation from a single image is presented in [12]. The general problem of highlight analysis using polarization is treated in [6].

This paper presents a novel approach for the reconstruction of specular surfaces. It is designed to determine the depth of global surface regions. In this respect we claim, that our method is superior to the method proposed in [8], as no virtual features need to be present in the image and the depth of whole surface regions instead of only surface curves can be computed. The method is a passive one and hence do not requires special or active lighting conditions. It handles even modest concavities (as long as no inter-reflection occurs) and therefore covers a broader range of objects than approaches based on silhouettes, like for example [3]. We stress that the presented approach does not incorporate any additional intensity based information, such

as surface texture, surface edges or apparent contours. This will be done in the future and it is to be expected that the integration of intensity clues will increase the performance of the proposed method.

After presenting the relevant aspects of polarization imaging we state the shape from polarization problem. We summarize our previous work [10] where shape analysis from a single polarization image was carried out. It is shown, that a single polarization image provides not enough constraints on the surface geometry for a complete reconstruction. Then, the case of multiple polarization images is tackled. It is motivated that an adequate framework for shape recovery is a global optimization scheme. Using hierarchical basis functions for the approximation of the surface, guarantees good convergence in the optimization process. Experiments on synthetic and real world data show that a precise reconstruction is achieved.

2. Polarization Analysis

Unpolarized light becomes partially linear polarized upon reflection on both dielectrics and metals. Since common light sources emit unpolarized light, the analysis of partial linear polarization of the reflected light covers all materials. The cases of multiple interreflection, like in [15], or polarized light sources, are not treated in this context.

Polarization analysis is the determination of the complete state of polarization of the light. For the analysis of partial linear polarized light, the polarization image is a set of three images encoding the intensity (that is what a normal camera would see), the degree of polarization and the orientation of polarization (figure 5). The orientation of polarization is encoded in the so called phase image (the intensity encodes the angle of orientation).

The two basic assumptions for a geometric scene interpretation using polarization imaging are that first the object under investigation exhibits a smooth surface structure and secondly the light illuminating the scene is not polarized. A smooth surface has no micro-structure and as a consequence the geometry of specular reflection, which is the plane of reflection, is only defined by the ray of observation and the corresponding surface normal. Assuming the lighting to be unpolarized results in the fact that the phase image is invariant concerning the intensity of the illumination and a characteristic entity of the object's shape, which is clearly shown in figure 5(b).

The physics of the electromagnetic theory tells us, that upon specular or surface reflection, which is a reflection in a proper physical sense, unpolarized light becomes partial linear polarized with an orientation orthogonal to the plane of reflection. On the other hand, diffuse or body reflected light, which is, physically spoken, a refraction, is partial linear polarized parallel to the plane of reflection. There-

fore, the difference between a phase image produced by a specular and a diffuse reflecting object is just $\frac{\pi}{2}$. To get a unique phase image, independent of the type of reflection, we just add $\frac{\pi}{2}$ to phase values corresponding to specular reflection and keep the values corresponding to diffuse reflection. Then the phase values will correspond to the orientation of the surface normals projected onto the image plane. Notice that the phase images are defined modulo π ; in order to formulate addition and subtraction of phase values in a convenient way, phase values are defined over the interval $[-\frac{\pi}{2}, \frac{\pi}{2}[$.

Commonly the analysis of specular surfaces concentrates on highlights, that are surface regions mirroring strong light sources. But, it is understood that even specular surfaces can emit light resulting from diffuse or body reflection. The only object exhibiting only surface and no body reflection is the perfect mirror. Therefore, it is necessary to conceive a polarization based framework which covers both types of reflection. This can be performed based for example on the work in [17], where a method is proposed inferring the reflection type directly from the polarization image. Adding $\frac{\pi}{2}$ to phase values resulting from surface reflection and keeping phase values resulting from body reflection will lead to a phase image unequivocal in a geometric sense and invariant of the lighting conditions.

Even if our method can be applied to both types of reflection, surface reflection is the target application as the degree of polarization is in general much higher than for the case of body reflection. In most cases the degree of polarization is sufficient for reflection angles typically lying between 30 and 85 degrees. So, even if the phase image is illumination invariant the degree of polarization depends on the lighting conditions. Thus the crucial constraint for the employment of a shape from polarization approach is that the degree of polarization must be high enough for an accurate measurement of the orientation of polarization.

Finally, we want to note that the appropriate camera model is the scaled orthographic projection, because light has to project orthographically onto the polarizer filter for a correct polarization measurement.

3. The problem of shape from polarization

3.1. A single polarization image

To get into the shape from polarization problem we start with the analysis of a single image, summarizing the previous work in [10]. In the previous section it was stated that the phase image encodes the direction of the reflection plane. The reflection plane is spanned by the ray of observation and the surface normal of the observed surface point. Furthermore, the reflection plane intersects the image plane in the line defined by the image point and the orientation represented by the phase value ϕ . Hence, the reflection

plane can be spanned by the projecting ray and the vector \mathbf{v} generated by the phase value as $\mathbf{v} = (\cos(\phi), \sin(\phi))^T$. So, we can say that the phase value induces one constraint on the surface normal, as it has to lie within the reflection plane. Conversely, the phase value vector is the projection of the normal onto the image plane. Let \mathbf{X} be a world point projecting under orthographic projection onto the image point \mathbf{x} with $\mathbf{X}^T = (\mathbf{x}^T, Z)$. The corresponding surface normal \mathbf{N} can be parameterized as follows: $\mathbf{N} = (\mathbf{v}^T, \lambda)^T$, $\lambda \in \mathbb{R}$, where \mathbf{v} is the above defined phase value vector. So, the key observation is:

The phase image encodes the direction of the surface normals projecting onto the image plane.

We just mention, that this observation is not new, as it is used already in [16] and [11]. Instead of a point based interpretation of the phase values a global interpretation can be deduced as well. This leads to a complete new insight in the geometrical analysis of polarization imaging. Each phase image point encodes a direction $\mathbf{v}(\mathbf{x})$ which is the projection of the coresponding surface normal. Then, the image $\mathbf{v}_\perp(\mathbf{x})$ encodes a direction perpendicular to the surface normal. Going from one pixel to the next in the direction of $\mathbf{v}_\perp(\mathbf{x})$ we end up with a curve c . Mathematically more precisely, we can think of $\mathbf{v}_\perp(\mathbf{x})$ as a normalized vector field defined on the image domain by the phase image as $\mathbf{v}_\perp(\mathbf{x}) = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \mathbf{v}(\mathbf{x})$. Then, the curves c are the field lines of the vector field $\mathbf{v}_\perp(\mathbf{x})$. Field lines are the envelopes of the flow vectors or, in other words, flow vectors are the tangent vectors to the field lines. In the sequel we will call the field lines of $\mathbf{v}_\perp(\mathbf{x})$ level curves, which is due to the following fact proven in [10]:

Level curves are the (orthographic) projection of surface profiles parallel to the image plane (Iso-depth profiles).

The phenomenon is depicted in figure 1. From image measurements, i.e. the phase image, level curves can be computed. As outlined above, points on a single level curve are the projection of surface points on a profile, denoted C , all having constant depth. Surface points with constant depth are surface profiles, where the cutting plane is parallel to the image plane. Since the actual depth of the surface profiles is unknown the complete surface shape can be parameterized by the set of profile depth values. A very interesting aspect is that by the use of level curves the problem of reconstructing $N \times N$ depth values is reduced to the problem of finding the depth of only N level curves representing the complete surface.

3.2. The correspondence problem for multiple images

The classical approach for depth recovery is the triangulation of corresponding features seen in two or more images. Textureless and specular reflecting surfaces do not




Figure 1. Level curves are images of surface profiles

provide this type of features. As the phase image is the representation of a geometric feature, i.e. the projection of the surface normals onto the image plane, we have to investigate whether the correspondence problem can be solved on the basis of the phase image information. Then point-wise reconstruction via triangulation would be possible.

It is very straightforward to perceive that two polarization images do not provide enough constraints for a direct point-wise reconstruction algorithm: A given point in the first phase image provides one constraint on the surface normal: the normal has to lie in the reflection plane spanned by the ray of observation and the projection of the surface normal onto the image plane [10]. The corresponding point in the second image has to lie on the epipolar line and induces a second constraint on the surface normal. Both constraints uniquely determine the normal. Since the normal is not known any point on the epipolar line will generate a valid solution for the normal. Hence, a point-wise interpretation from only two polarization images is not sufficient for a solution of the correspondence problem and no reconstruction can be achieved.

Three images can offer a solution to the correspondence problem: The unknown depth parameterizes the corresponding points lying on the epipolar line in the second and the third image. The points in the first and second image determine the surface normals. Hence, for the correct depth the surface normal has to satisfy the constraint in the third image as well. In other words: the three reflection planes spanned by the three phase image points have to intersect in the same space line in order for them to correspond to the same surface point. Hence, three phase images provide a necessary set of constraints for finding correspondences. Therefore, we can state that:

In principle, three polarization views are sufficient for surface reconstruction.

However, local reconstruction using explicit correspondences does not seem to be the best choice for the following two reasons: first, it is expected to be susceptible to noise and second two redundant entities are computed which are the two parameters of the surface normal. In a global recon-

struction approach this would be redundant since the surface normals can be computed from the depth function. In the experimental section we show that a global scheme is adequate, because even on the basis of only two polarization images correct reconstruction is achieved.

4. Optimization approach for shape from polarization

Inspired by variational approaches used for solving the shape from shading problem, like for example in [5], we choose a similar way for the shape from polarization problem. Unlike the shape from shading problem, no additional constraints are required since the problem is overdetermined. We employ the idea of modelling the surface by hierarchical basis functions, which performed well in the case of the shape from shading problem [14] too. An implementation using an explicit formulation of the Jacobian results in a fast and well converging algorithm.

4.1. A functional for optimization

An object surface S observed by a camera C will produce a phase image $I = I(S, C)$. Let S be the unknown real surface and \hat{S} the reconstructed surface. The surface \hat{S} best approximates the real surface S if the squared differences between the phase images \hat{I} and I are minimal over all cameras i . Hence, the functional E to be minimized by the surface \hat{S} is:

$$E = \sum_i \int \left(\hat{I}(\hat{S}, C_i) - I_i \right)^2, \quad (1)$$

where the integral is evaluated over the image plane of the corresponding camera. Nevertheless, a formulation of the functional in world coordinates is more appropriate. In world coordinates the object surface S is modelled as a depth function $Z(X, Y)$ over some discrete lattice $(X, Y) \in D \subseteq \mathbb{R}^2$. The surface normals are defined as $\mathbf{N}(X, Y) = (-Z_X(X, Y), -Z_Y(X, Y), 1)^T$, with the partial derivatives $Z_X := \frac{\partial Z}{\partial X}$, $Z_Y := \frac{\partial Z}{\partial Y}$. World points $\mathbf{X} = (X, Y, Z)^T$ project onto image points $\mathbf{x}_i = (x_i, y_i)^T$ as $\mathbf{x}_i = k_i \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} R_i \mathbf{X} + \mathbf{t}_i$, where k_i is the scaling factor, R_i the rotation matrix between world and camera coordinate system and \mathbf{t}_i a vector incorporating camera translation and the principal point. The surface \hat{S} will generate a phase image \hat{I}_i as a function of the surface normals and the projection matrix P_i , more explicitly: $\hat{I}_i = \hat{I}(Z_X, Z_Y, P_i)$. The real phase image I_i in terms of world coordinates is $I_i = I_i(\mathbf{x}_i) = I(X, Y, Z, P_i)$. Now, the functional E can be defined over the definition domain D of the surface S :

$$E = \sum_i \int_D \left(\hat{I}_i(Z_X, Z_Y, P_i) - I_i(X, Y, Z, P_i) \right)^2 dX dY \quad (2)$$

4.2. Optimization based on hierarchical basis functions and a least squares algorithm

A standard strategy for minimizing the above functional would be to formulate the corresponding Euler-Lagrange equation. This would result in a partial differential equation in Z, Z_X, Z_Y . But, as the partial derivatives of Z depend on Z itself, we prefer a formulation where the functional E is function of only the depth Z . Then we have a classical least squares optimization problem to solve and standard techniques can be applied.

Using the functional as defined in equation (2) will result in an optimization algorithm which converges very slowly or even not at all depending on the initial solution. The reason for that can be seen in the local nature of the formulation: changing the depth at one point will change the error difference $\hat{I} - I$ at only five points given by the 4-connected neighborhood since the depth value itself and the first order partial derivatives occur in the formula 2. In a numerical context this is called the computational molecule or five-point star. In this scheme global errors have to be reduced through local interactions.

To overcome these problems, hierarchical basis functions are used and we follow the work in [13], [14], where the reader is referred for more details. The key idea is to set up a multiresolution pyramid where the entries at each level are the coefficients of the basis functions. Going up in this pyramid from fine to coarse corresponds to basis functions with a larger support. Assuming that the dimensions $N \times N$ of the definition domain are a power of two ($N = 2^L$), it follows that the pyramid has $L + 1$ levels. At the finest level $l = 1$ the basis functions have a support of two, which means that the corresponding coefficient changes the depth of the center point and the derivatives of the direct neighbors. The basis functions at general level l have a support of $2^l \times 2^l$ changing depth and derivatives likewise. The top most level $l = L + 1$ acts as a depth offset. To complete the scheme, a bilinear interpolation operator \mathcal{I}_l is selected which defines how each level is interpolated to the finest resolution. The operator \mathcal{I}_l acts on the coefficients Z_l of the basis function and we get the formula for the actual depth function Z :

$$Z = \sum_{l=1}^{L+1} \mathcal{I}_l(Z_l). \quad (3)$$

Rearranging the depth functions Z and Z_l in vector form the interpolation operators can be written in a matrix form with sparse structure. In [13] the multiresolution pyramid is only partially populated such that the total number of populated nodes in the complete pyramid is equal to the number of grid points at the finest resolution, which is N^2 . In contrast, we chose the pyramid to be completely populated resulting in $\frac{4}{3}(N^2 - 1)$ nodes or coefficients. This does not lead to an

ill-posed problem like in the shape from shading case, since the shape from polarization formulation provides $M * N^2$ data points, where M is the number of images. Now, the formulation has more degrees of freedom (over parameterization) than the original problem, but the complexity of the problem remains the same and the final algorithm still converges very well.

The optimization itself is implemented within MATLAB, which provides a least squares optimization procedure. MATLAB offers an algorithm specially tuned for large-scale problems, incorporating the sparsity structure of the problem and the explicit use of derivative information in the form of the Jacobian. The actual optimization procedure is of the type “trust-region reflective Newton”. Using the derivative information explicitly results in a more accurate and faster implementation. The Jacobian can be derived directly from the formulation of the functional E .

5. Experiments

5.1. Simulated data

In order to prove the convergence of our algorithm and to assess the accuracy of the reconstruction result we test our method on a synthetically generated object, like the one shown in figure 2. In a previous section we stated that, in principle, three polarization views are sufficient for surface reconstruction. Below we show that even two views can be sufficient for a complete reconstruction. To do so, two phase images of the synthetic surface were generated, see figure 3. These two images are used as input for our algorithm. As initial guess for the optimization routine we computed a plane, minimizing the error functional E , see figure 4(a). The evolution of the surface is shown in the images 4(a) to 4(d). After 20 iterations convergence is almost reached. The reconstruction result is identical to the original object, within an error ε of less than 0.001. The error ε is a relative error defined as $\varepsilon = |\frac{\hat{Z}-Z}{L}|$, where \hat{Z} is the reconstructed and Z the ground-truth depth; L is the maximal length of the reconstruction area.

Additionally, we tested the influence of the noise level on the resulting reconstruction as well. It turned out that the reconstruction error is proportional to the noise level. The root mean square error is always less than one fifth of the noise in the phase image. That means for a noise level of 0.05 the rms-error in the reconstruction is less than 0.01; The dependency on the number of images is negligible. Only in the case of two images the error increases slightly. This shows the strength of a global optimization process, which inherently smooths the reconstruction thus eliminating the influence of the noise.




Figure 2. The synthetically generated surface used for simulations. From this object phase images are generated from viewing directions orthogonal to the axis of rotation which is indicated by the line. The two viewing directions are represented by the arrows.



Figure 3. Left and right phase image of the synthetic object taken from the two viewing directions corresponding to the arrows in figure 2.

5.2. Real world experiment

In figure 5 a polarization image of a billiard ball is shown, comprising its three components: the intensity image, the phase image and the degree of polarization image. Since the ball is specular reflecting the surrounding illumination is mirrored on the surface, see figure 5(a). As the illumination does not produce any salient intensity features, the method proposed in [8] could not be applied in this situation. However, the problem can be solved by our polarization based method.

A sequence of five polarization images was taken from the ball by rotating it on a turntable. Then the third camera is chosen as reference coordinate system and a squared surface patch in the third image, as marked in figure 5(d), is chosen as reconstruction domain. The optimization procedure is initialized with a plane minimizing the error functional. The final reconstruction result can be seen in 6(a).

The marked square as shown in figure 5(d) is chosen as reconstruction domain, because the underlying surface patch projects in all five images onto regions, where the degree of polarization is higher than 0.1, like the depicted grey region in figure 5(d). Nevertheless we like to stress that the




Figure 4. The evolution of the surface during the optimization process: initial surface(a), after 2 (b), 5 (c) and 20 (d) iterations. The reconstruction result is identical to the original surface, like in figure 2.

proposed method works even if not the whole surface patch does produce enough polarization in all images. Parts of the surface patch, which produce enough polarization in only one image, can be reconstructed equally; this is based on the ideas presented in section 3.1.

As we know that the object is a sphere, we can assess the accuracy of the reconstruction by comparing it to ground truth data. From the intensity images we can compute the coordinates of the center of the sphere. In the case of a rotational symmetric object like the sphere the shape can be recovered only up to scale. This means that all the spheres with different radius, but centered at the same point, will produce the identical sequence of phase images. This is true, apart from the fact that smaller spheres will result in a smaller support of valid phase values in the phase image. Therefore, from the reconstruction result we have to determine the radius of the reference sphere, such that it best fits the reconstructed data. This can be done by using for example a nonlinear optimization method. Then, the difference between the reconstructed surface and the ground truth sphere is the reconstruction error, as shown in figure 6(b). The error ε is the unsigned error relative to the radius R of the sphere: $\varepsilon = \left| \frac{\hat{z} - z}{R} \right|$. It is remarkable that the overall root-mean-square error is only 0.0078 and the maximum error 0.026. These results agree well with the results obtained in the previous experiment based on simulated data.

Experiments based on only three phase images give similar results.




Figure 5. Polarization image of a billiard ball: intensity image (a), phase image (b) and degree of polarization (c). In (d) the grey region labels the points, where the degree of polarization is higher than 0.1. The area of reconstruction for the experiment in section 5.2 is the white square in (d).

6. Discussion and future work

We have presented a novel approach for the reconstruction of specular reflecting objects based on polarization imaging. The projection of the surface normals is directly provided by the polarization image. This geometric information is then employed for shape recovery. In a single image curves corresponding to surface profiles parallel to the image plane and with unknown depth can be computed. It is shown that in principle three polarization images are sufficient for reconstruction. Subsequently, the problem of shape recovery is stated in an optimization context. The object surface is modelled using hierarchical basis functions which guarantees good convergence properties and an accurate reconstruction. Results on simulated data prove that for a asymmetrical object a correct reconstruction can be




Figure 6. The reconstruction resulting from the analysis of a sequence of 5 phase images; one image out of this sequence is the one in figure 5(b). Figure (a) shows the reconstruction result for the labeled region in 5(d). Figure (b) depicts the reconstruction error, relative to the radius of the sphere. The maximum error is smaller than 0.03.

achieved even based on only two views. For a simple, symmetrical object like a sphere reconstruction based on three or more views can be achieved up to a scale factor. The algorithm performs well also on a real world experiment.

Since the theoretical proof for the fact that global reconstruction is possible from only two views is not yet given, further analysis has to be done on this issue. In the future the performance of the method will be enhanced by incorporating additional intensity based information like surface texture, surface boundaries or apparent contours. A more general surface model has to be used in order to describe arbitrary 3D objects.

Acknowledgments

This work was supported by the “Deutsche Forschungsgemeinschaft (DFG)”.

References

- [1] D. Bhat and S. Nayar. Stereo in the presence of specular reflection. In *Proc. of Intl. Conf. on Computer Vision (ICCV)*, pages 1086–1092, 1995.
- [2] A. Blake and G. Brelstaff. Geometry of specularities. In *Proc. of Intl. Conf. on Computer Vision (ICCV)*, pages 394–403, 1988.
- [3] R. Cipolla and A. Blake. Surface shape from the deformation of apparent contour. *International Journal of Computer Vision (IJCV)*, 9(2):83–112, 1992.
- [4] G. Healey and T. O. Binford. Local shape from specularity. *Computer Vision, Graphics and Image Processing*, 42:62–86, 1988.
- [5] B. Horn and M. Brooks, editors. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.
- [6] S. Nayar, X. Fang, and T. Boult. Separation of reflection components using color and polarization. *International Journal of Computer Vision (IJCV)*, 21(3):163–186, 1997.
- [7] S. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Transactions on Robotics and Automation*, 6(4):418–431, 1990.
- [8] M. Oren and S. Nayar. A theory of specular surface geometry. In *Proc. of Intl. Conf. on Computer Vision (ICCV)*, pages 740–747, 1995.
- [9] D. Perard and J. Beyerer. Three-dimensional measurement of specular free-form surfaces with a structured-lighting reflection technique. In *Conf. on Three-Dimensional Imaging and Laser-based Systems for Metrology and Inspection III*, volume 3204 of *SPIE Proceedings*, pages 74–80, 1997.
- [10] S. Rahmann. Polarization images: a geometric interpretation for shape analysis. In *Proc. of Intl. Conf. on Pattern Recognition (ICPR)*, volume 3, pages 542–546, 2000.
- [11] M. Saito, Y. Sato, K. Ikeuchi, and H. Kashiwagi. Measurement of surface orientations of transparent objects using polarization in highlight. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 381 – 386, 1999.
- [12] Y. Schechner, J. Shamir, and N. Kiryati. Polarization-based decorrelation of transparent layers: The inclination angle of an invisible surface. In *Proc. of Intl. Conf. on Computer Vision (ICCV)*, pages 814–819, 1999.
- [13] R. Szeliski. Fast surface interpolation using hierarchical basis functions. *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, 12(6):513–528, 1990.
- [14] R. Szeliski. Fast shape form shading. *CVGIP: Image Understanding*, 53(2):129–153, 1991.
- [15] A. Wallace, B. Liang, E. Trucco, and J. Clark. Improving depth image acquisition using polarized light. *International Journal of Computer Vision (IJCV)*, 32(2):87–109, 1999.
- [16] L. Wolff. Surface orientation from two camera stereo with polarizers. In *Optics, Illumination, Image Sensing for Machine Vision IV*, volume 1194 of *SPIE Proceedings*, pages 287–297, 1989.
- [17] L. Wolff. Scene understanding from propagation and consistency of polarization-based constraints. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1000–1005, 1994.
- [18] J. Zheng, Y. Fukagawa, and A. N. Shape and model from specular motion. In *Proc. of Intl. Conf. on Computer Vision (ICCV)*, pages 72–78, 1995.