Shape Reconstruction from Multiple Images of the Ocean Surface

Howard Schultz

Department of Computer Science

University of Massachusetts, Amherst

May 20, 1994

Abstract

In many remote sensing applications the detailed shape of a water surface must be determined instantaneously, and under natural lighting conditions. Because of their nondestructive measurement capabilities, stereo image analysis methods seem to be ideally suited for this application. Most of these techniques, however, assume surface reflectance properties which are incompatible with the reflectance properties of water, especially for centimeter and smaller wave lengths. To analyze the fine scale structure of the ocean surface, a new technique (known as Specular Surface Stereo) is proposed that makes use of the unique optical properties of water. In this method, the surface shape is recovered by solving the inverse problem of finding a surface shape (defined by an elevation and gradient map) that results in a set of synthetic (computer generated) images that closely match a set of observed images. To accomplish this, the algorithm requires an image of the illumination source, three images of the water surface, a few initial elevation estimates, and an image synthesis model that predicts the appearance of an image of the surface. The performance of the Specular Surface Stereo technique was tested by processing simulated data. Based on these numerical results, it is shown that the Specular Surface Stereo technique can potentially recover both surface elevations and gradients over a wide range of wavenumbers.

Index Terms: Optical remote sensing, Stereo vision, small scale wave measurements, nondestructive measurements, shape from shading

1 Introduction

Microwave scatter from the ocean is highly sensitive to the small scale structure of the surface, which in turn is modulated by several geophysical processes. Microwave measurements, therefore, are indirectly influenced by environmental parameters such as atmospheric temperature and pressure, wind speed and direction, sea state and ocean currents. This sensitivity to geophysical processes forms the basis for microwave remote sensing. Interpretation of microwave signatures, therefore, requires an understanding of electromagnetic scattering from rough surfaces, as well as the hydrodynamic processes that influence the ocean surface. Consequently, many remote sensing research programs have focused on EM scattering theories that predict microwave scatter from the ocean surface, and hydrodynamic theories (air-sea, wave-wave and wave-current) that predict the response of the ocean surface to geophysical forces. Implementation and verification of many of these theories requires detailed shape information. For example, Chen (1990 and 1992b) and Nickolaev (1992) , have shown the importance of higher order surface statistics by deriving a relationship between the upwind/downwind asymmetry in the radar cross section of the surface and non-Gaussian terms in the distribution of surface elevations and slopes. Many of these higher order descriptions of the ocean surface require joint elevation–slope maps of the surface over a dense grid of at least 512×512 samples.

Because of their nondestructive capabilities, optical remote sensing techniques are ideally suited for this application. Several investigators have developed active optical systems specifically designed to recover shape information from a smooth water surface. Keller and Gotwols (1983) , and Jähne and Riener (1990) retrieved two-dimensional spectra using a single camera located above the surface and a one-dimensional graduated light source located below the surface. The instantaneous surface slopes in the direction of the light gradient are then inferred from the observed image grey levels and temporal resolution is achieved by analyzing an image sequence. The advantage of these light gradient methods are that the data analysis is simple (and can be performed in real-time) and the techniques can be readily extended to two spatial dimensions. Although these methods have be very successful in the laboratory, interference from ambient light, and the need for a shallow submerged light source limit the number of operational scenarios in the field. In another approach, Wass and Jähne (1992) made combined slope-height measurements using a Reflective Stereo Slope Gauge (RSSG). This instrument consists of a pair of point light sources and video cameras housed in a single package suspended over the surface. Linear polarization filters are placed in front of the lights and cameras so that only specular reflections from the left light are imaged by the right camera, and only specular reflections from the right light are imaged by the left camera. The apparent displacement of specular reflections between the left and right images (parallax) are proportional to the elevation of the surface facet where the reflection occurred. By computing the parallax for all specular reflections the position and slope of specular points on the surface are found. Because a bright light source is used and the instrument package does not cross the air-sea interface, the instrument is very rugged and can be operated under any lighting conditions, day or night. Although this type of joint height-slope observations adds significantly to the understanding of the dynamics of the ocean surface, it is not possible to reconstruct the shape or retrieve high order statistical description of the ocean surface from the data.

In another approach Bock and Hara (1992) , Martinsen and Bock(1992) and Lee et al. (1992) employed a scanning laser slope gauge in which a scanning laser beam is located below the surface, and a receiver is located above the surface. Then by recording the instantaneous position of the transmitted and received beams, a two-dimensional slope map of the surface is reconstructed. These instruments have been successfully deployed in the laboratory and field, and are capable of recovering significant statistical information about dynamics of the ocean surface. Although these devices are capable of good temporal resolution, frame rates in excess 30 Hz, they have limited spatial resolution generally less than 100×100 samples. Furthermore, as in the light gradient method, these systems cross the air-sea interface and may disturb the surface.

To design a system capable of recovering instantaneous joint elevation-slope maps over a 512×512 grid, while not disturbing the air-surface interface, a variety of passive image analysis techniques were investigated. Only methods that could be adapted to analyze a water surface where considered. In the millimeter to centimeter length scales, a water surface displays specular reflectance properties. That is, at the air-sea boundary, light rays are reflected back into the atmosphere or refracted into the water volume, and multiple reflections and scattering at the surface can be neglected. Consequently, the water surface acts like a partially silvered mirror the reflects the sky brightness into the camera. Most image analysis techniques designed to recover shape information from images (feature matching, photometric stereo, and shape-from-shading) assume Lambertian surface reflectance properties. For a Lambertian type surfaces the apparent brightness of a surface point is strongly dependent on the position of the light source relative to the surface, and weakly dependent on the position of the observer. Feature matching algorithms take advantage of this reflectance property by assuming that small surface features can be identified in images taken from different positions [Horn, 1986] [Marr and Poggio, 1976]. Similarly, photometric stereo[Horn, 1986] [Woodham, 1978] and shape-from-shading algorithms [Hartt and Carlotto, 1989] [Horn and Brooks, 1989] assume that the pixel brightness depends only on the position of the light source relative to the surface. Several researchers have extended the photometric stereo technique to specular surfaces [Blake and Brelstaff, 1988] [Coleman and Jain, 1982] [Ikeuchi, 1981] [Tagare and deFigueiredo, 1989]. These methods, however, require a single camera viewing a stationary surface illuminated by a sequence of extended light sources. Shape-from-specularity techniques [Healey and Benford, 1987] have successfully recovered shape information from specular surfaces by analyzing the reflections produced by a small number of point light sources [Blake and Brelstaff, 1988] [Blake, 1985] [Sanderson et al., 1988] [Zisserman et al., 1989]. These methods cannot be extended to natural lighting conditions because it is not possible to isolate individual specular reflection. Optical flow techniques have been employed, recently, to recover the instantaneous shape of a water surface by analyzing the apparent motion of a submerged pattern [Murase, 1992].

This technique is limited, however, to gradually undulating surfaces where the maximum slope does not exceed 5° . Because of these restrictions, none of the above methods can be readily adapted to recovering the shape of a moving specular surface (such as water waves) under natural lighting conditions.

Thomas et al. (1991) and Hartt et al. (1989) employed a multiple image shape-from-shading technique to recover the shape of a surface with diffuse reflectance properties. The technique finds a surface elevation map $z(x, y)$ that minimizes a cost function $U(z)$,

$$U\left(z(x,y)\right) = \int\int \left[\sum \left(\tilde{I}_k(x,y) - I_k(x,y)\right)^2 + \lambda\left(\frac{\partial z(x,y)^2}{\partial x \partial x} + 2\frac{\partial z(x,y)^2}{\partial x \partial y} + \frac{\partial z(x,y)^2}{\partial y \partial y}\right)\right] dxdy \tag{1}$$

where $z(x, y)$ is the elevation map, $\tilde{I}_k(x, y)$ and $I_k(x, y)$ are the observed and predicted (or synthetic) pixel intensities in the kth image of the surface and λ is a constant. The first term $\left(\tilde{I}_k - I_k\right)^2$ is small when the predicted pixel grey levels are close to the observed one, and the second term, $\lambda\left(\frac{\partial z^2}{\partial x \partial x} + 2\frac{\partial z^2}{\partial x \partial y} + \frac{\partial z^2}{\partial y \partial y}\right)$, is small when the surface is smooth. The second term acts as a penalty function that limits the roughness of the surface. The purpose of this term is to ensure that the retrieved surface is smooth, resolve grey level ambiguities (cases where more than one elevation results in the same set of pixel intensities), and to fill in regions where either the surface is occluded or the image synthesis model cannot be used. Because of the general form of this algorithm, which includes a mechanism for assimilating images taken from different viewing positions, the specular surface stereo algorithm is based on the multiple image shape-from-shading approach. This technique is adapted to analyze specular surface by incorporating a specular surface image synthesis model instead of a diffuse surface image synthesis model [Schultz, 1992] [Schultz, 1994]. Furthermore, to reduce the number of hardware dependent operations, units of irradiance are used instead of pixel intensities. To further simplify the algorithm development λ is set to zero. In addition to removing any smoothness constraints, neglecting the second term in Equation 1 reduces the retrieval process to a point-by-point scheme. Thus, the cost function is minimized by minimizing the difference term at each node in the retrieval grid.

2 The Imaging Environment

Before constructing an algorithm, the viewing geometry and image synthesis model must be specified. For the purpose of describing the specular surface stereo method, an idealized camera, surface, and illumination models are employed.

2.1 Viewing Geometry

The viewing geometry shown in Figure 1 depicts an ideal camera with a focal length f, viewing a surface described by an elevation map $z(x, y)$ that has continuous derivatives $p(x, y)$ and $q(x, y)$ with respect to x and y. The camera focal point is located at P, and O is a point on the surface. The observation vector $\hat{r}$ is a unit vector that points upward from O to P, the surface normal vector $\hat{s}$ is perpendicular to the surface at O, and the incidence vector $\hat{t}$ is the reflection of $\hat{r}$ about $\hat{s}$. From Snell's law, $\hat{t}$ is related to $\hat{r}$ and $\hat{s}$ by,

$$\hat{t} = -\hat{r} + 2\hat{s}\,(\hat{s} \cdot \hat{r}) = -\hat{r} + 2\hat{s}\cos(\theta) \tag{2}$$

where the observation angle θ (the angle between $\hat{r}$ and $\hat{s}$) is equal to the incidence angle (the angle between $\hat{s}$ and $\hat{t}$). The extension of the observation vector $\hat{r}$ intersects the image plane of camera k at pixel ij with image space coordinates $u_{ijk} = (x_{ijk}, y_{ijk}, -f_k)$. The object space coordinates of O are related to the image space coordinates u_{ijk} by the projective transformation [Slama, 1980],

$$O = P_k + \lambda_{ijk} M_k u_{ijk} \tag{3}$$

where P_k is the position of the focal point of camera k, M_k is the rotation of the image plane coordinate system of camera k relative to the object space coordinate system, and λ_{ijk} is a scale factor for pixel ijk. Since the incidence angle is completely specified by the coordinates of the pixel ijk that views O, and the orientation (p, q) at O, the incidence angle can be written as $\theta_{ijk}(p, q)$.

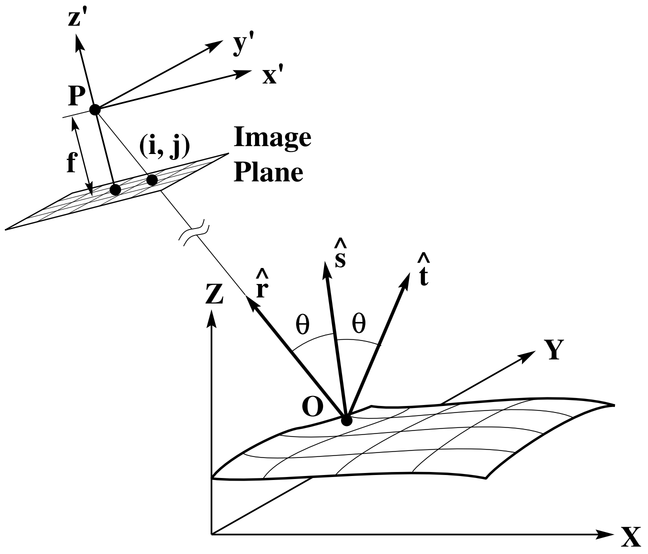

Figure 1: Viewing geometry: The focal point of an ideal camera with focal length f is located at P; the observation point on the surface is located at O; $\hat{r}$ is the unit vector that points from O to P; $\hat{s}$ is the surface normal unit vector; $\hat{t}$ is the reflection of $\hat{r}$ about $\hat{s}$

2.2 Image Synthesis

Once the coordinates of the pixel that views a point on the surface is determined, the irradiance at that pixel must be determined. That is, we need an image synthesis model $E_{ijk}(p,q)$ that predicts the irradiance at the pixel ijk as a function of the surface orientation (p,q). The specular surface stereo technique requires only that the image synthesis model predict accurately the irradiance at each pixel, and that the pixel irradiance depends on (p,q). To illustrate the technique, an idealized image synthesis model will be employed that imposes the following conditions on the surface and illumination:

1. The radius of curvature of the surface patch within the pixel field-of-view is large compared to the dimensions of the patch. Thus, the surface is modeled by a mosaic of small, partially silvered, flat mirrors.

2. The illumination originates from a distant, extended source, such that the illumination radiance $L(p^t, q^t)$ depends only on the orientation (p^t, q^t) of the incidence vector $\hat{t}$.

3. The light incident on the surface is unpolarized.

4. The reflecting material is a homogeneous dielectric.

5. No explicit restrictions are placed on the surface slopes, however, to ensure that most of the incidence vectors come from the sky, the maximum surface slope should not exceed $\pm 45°$.

When light is incident on the surface of a smooth dielectric media some of the light is reflected from the surface and some is refracted through the surface. The Fresnel reflection coefficient for a dielectric media gives the fraction of radiant energy reflected at the surface as a function of the polarization relative to the plane of incidence, the real part of the index of refraction n, and the incidence angle θ [Born and Wolf, 1980] [Wolff and Boult, 1991]. The reflection coefficients for the parallel $R_{||}(\theta)$ and perpendicularly $R_{\perp}(\theta)$ polarized components of the incident light are given by,

$$R_{||}(\theta) = \frac{\sin^2(\theta - \theta_1)}{\sin^2(\theta + \theta_1)}, \qquad R_{\perp}(\theta) = \frac{\tan^2(\theta - \theta_1)}{\tan^2(\theta + \theta_1)} \tag{4}$$

where

$$\theta_1 = \sin^{-1}\left(\frac{\sin(\theta)}{n}\right)$$

is the angle of refraction. For unpolarized light, the parallel and perpendicularly polarized components have equal intensities. The reflection coefficient $R(\theta)$ for unpolarized incident light, therefore, is given by

$$R(\theta) = \frac{1}{2}R_{||} + \frac{1}{2}R_{\perp}. \tag{5}$$

Since θ depends on the observation vector and the surface normal, the reflection coefficient can be expresses as a function of the pixel coordinates and surface gradient $R_{ijk}(p,q)$. Likewise, the orientation of the incidence vector can be expressed in terms of (ijk) and (p,q), therefore the radiance directed along $\hat{t}$ can be written as $L_{ijk}(p,q)$. Provided that the radius of curvature at the surface point is large compared to the surface area within the pixel field-of-view (see condition 1 above), the radiance reflected into the field-of-view of the pixel ijk is the product of the radiance $L_{ijk}(p,q)$ and the reflection coefficient $R_{ijk}(p,q)$. The predicted irradiance $E_{ijk}(p,q)$ is the product of the radiance directed along the observation vector and a constant C_{ijk} [Horn, 1986] which is a function of the optical properties of the camera, such as the focal length and aperture. The image synthesis model, therefore can be written as,

$$E_{ijk}(p,q) = C_{ijk} L_{ijk}(p,q) R_{ijk}(p,q). \tag{6}$$

The calibration constants C_{ijk} are measured once for each imaging device, the radiance distribution $L_{ijk}(p,q)$ must be measured every time the illumination source changes, and $R_{ijk}(p,q)$ is a known function of the electrical properties of the water.

3 The Specular Surface Stereo (S^3) Algorithm

Starting with at least one known elevation, an iterative process is used to propagate the surface shape. Each iteration is comprised of two steps—surface orientation and surface propagation. For a point on the surface with a known elevation z, the surface orientation algorithm computes the orientation of the surface normal (p,q). Next, the surface propagation algorithm computes the elevations at surrounding points. The process is repeated until all surface elevations and gradients within the retrieval grid are known. In Sections 3.1 and 3.2, the surface orientation and surface propagation procedures are discussed in detail. In Section 3.3, a method is presented for implementing the specular surface stereo technique on a two-dimensional grid. In Section 4, the technique is illustrated with numerical simulations.

3.1 Surface Orientation

The surface orientation procedure finds the orientation (p, q) of a point on the surface given the elevation z of the point, N images of the surface $\tilde{E}_{ijk}, k = 1, \cdot\cdot, N$ (which have been converted to units of irradiance), and a radiance map of the illumination $L(p^t, q^t)$. The first step is to identify the pixels that view the surface point, by evaluating the inverse of Equation 3 for each camera. Next, a value of (p, q) is found such that the predicted irradiance at these pixels $E_{ijk}(p, q), k = 1, \cdot\cdot, N$ match the observed values $\tilde{E}_{ijk}, k = 1, \cdot\cdot, N$. Although several methods exist for obtaining a match, a weighted least squares objective function of the form

$$J = \sum_{k=1}^{N} w_{ijk} \left(\tilde{E}_{ijk} - E_{ijk}(p, q) \right)^2 \tag{7}$$

is used to illustrate the procedure. The objective function is further simplified by setting all weights to unity, $w_{ijk} = 1$ for all ijk. In general, however, w_{ijk} may depend on variables such as the signal-to-noise ratio. A simple linear search is employed to minimize J with respect to (p, q). Although several methods exist for performing this minimization, a simple linear search was employed to avoid false solutions that may arise from local minima in (p, q) space, and to facilitate implementing the algorithm on a massively parallel computer. The search was done by dividing the solution space into equally spaced bins, evaluating J at each bin, and selecting the value of (p, q) at the center of the bin where J is a minimum.

If a point on the surface is not visible in one or more images, or the image synthesis model does not apply (for example, if the point lies on a rough spot), or the orientation of the incidence vector points outside the domain of $L(p^t, q^t)$, then the surface orientation algorithm will not predict correctly the surface orientation. Nodes were failures occur are identified by comparing the minimum residual of J to a rejection threshold. If the minimum residual is above the rejection threshold, the point is marked as missing and excluded from further analysis.

3.2 Surface Propagation

If at least one surface elevation is known anywhere on the surface, it is possible to recover the surface shape by propagating a solution. Consider a surface point with horizontal coordinates (x_0, y_0) and a known elevation and gradient (z_0, p_0, q_0). Denoting the horizontal coordinates of a nearby point by $x_1 = x_0 + \delta x$ and $y_1 = y_0 + \delta y$, the elevation z_1 is found by evaluating the truncated Taylor series,

$$z_1 = z_0 + p_0 \delta x + q_0 \delta y. \tag{8}$$

Initial seed elevations may be obtained in a variety of ways. If, for example, features with Lambertian reflectance properties, such as bubbles or foam, are on the surface, conventional stereo matching techniques can be used to calculate the elevations of these non-specular surface features. If the surface is predominately smooth, electro-optical devices such as the Reflective Stereo Slope Gauge [Wass and Jähne, 1992] can be used to obtain seed elevations and slopes. Furthermore, under some circumstances seed orientations may be obtained. On a clear day, for example, the orientation of any pixel where a sun glint is visible is known to within $\pm 0.25°$ (half the angular extent of the sun).

3.3 Implementation

On a two-dimensional grid, the specular surface stereo technique is implemented as a dilation filter. That is, the surface grows around any surface point where the elevation is known. At each iteration, a 3×3 window is passed over the retrieval grid, and the action taken at the window center depends on the state of the nodes that lie within the window. Defining a complete node as any node where (z, p, q) are known, the possible actions taken at iteration m and node ij are:

1. If $(z, p, q)_{ij}^{m-1}$ are known, then set $(z, p, q)_{ij}^{m} = (z, p, q)_{ij}^{m-1}$.

2. If z_{ij}^{m-1} is known and $(p, q)_{ij}^{m-1}$ are missing, then set $z_{ij}^{m} = z_{ij}^{m-1}$ and use the surface orientation algorithm to compute $(p, q)_{ij}^{m}$.

3. If z_{ij}^{m-1} is missing and $(p, q)_{ij}^{m-1}$ are known, and there are no complete surrounding nodes, then set $(z, p, q)_{ij}^{m} = (z, p, q)_{ij}^{m-1}$.

4. If z_{ij}^{m-1} is missing and $(p, q)_{ij}^{m-1}$ are known, and there are at least n_c complete surrounding nodes, then for each complete node use the surface propagation algorithm to estimate the elevation at the window center, set z_{ij}^{m} equal to the average elevation estimate, and set $(p, q)_{ij}^{m} = (p, q)_{ij}^{m-1}$.

5. If $(z, p, q)_{ij}^{m-1}$ are missing, and there are no complete surrounding nodes, then set $(z, p, q)_{ij}^{m} = (z, p, q)_{ij}^{m-1}$.

6. If $(z, p, q)_{ij}^{m-1}$ are missing, and there are at least n_c complete surrounding nodes, then for each complete node use the surface propagation algorithm to estimate the elevation

at the window center, set z_{ij}^m equal to the average elevation estimate, and use the surface orientation algorithm to compute $(p,q)_{ij}^m$.

Initially n_c is set to two, thus, at least two complete nodes are required to estimate the elevation at the window center. By requiring a minimum of two complete nodes instead of one, numerical errors grow more slowly as the surfaces propagates away from the seed point. If no new elevation are computed after an iteration, then n_c is reduced to one, and a new iteration is started. If, again, no new elevations are computed with $n_c = 1$, then the procedure has converged. Otherwise, if at least one new elevation was computed, n_c is reset to two, and a new iteration is started. The procedure converges when no new elevations are computed, or the maximum number of iterations is exceeded.

Since only two out of the eight surrounding nodes must be complete to estimate the elevation at the window center, solutions can grow around regions of bad or missing data. If, for example, a small region on the surface is occluded in some of the images, no elevations or gradients will be recovered in the occluded region. Along the edges of the occluded region, however, where at least two complete nodes are found, the algorithm will successfully compute (z,p,q).

4 Numerical Simulations

The $\mathbf{S^3}$ technique was tested by processing numerical simulations. For each simulation, the sky radiance map, surface elevation map, and camera positions were specified, and a set of three synthetic images of the surface generated. Then from the synthetic images, the sky radiance map and a few seed elevations, the surface elevation and gradient maps were retrieved using the $\mathbf{S^3}$ technique and compared to the original ones. The same sky radiance map and camera placement were used for all simulations. The focal length of each lens was 170mm, the focal points were located on the vertices of a $3 \times 3 \times 3$ meter equilateral triangle elevated 10m over the mean surface, and the cameras were oriented so that the optic axes passed through the center of the retrieval grid. The surface elevation maps were defined on a 512×512 grid in object space that covered a horizontal area of 1m $\times$ 1m. The radiance map $L(p, q)$ was defined on a 512×512 grid in gradient space where $-4 \leq p \leq 4$ and $-4 \leq q \leq 4$. The computer generated images $\tilde{E}_{ijk}$, k=1,2,3 were defined on 512×512 pixel grid in the image coordinate system of each camera and covered an area of 17.4mm $\times$ 17.4mm. The camera placement is shown in Figure 2, and the sky radiance map along with a sample simulated elevation map and three synthesized images are shown in Figure 3.

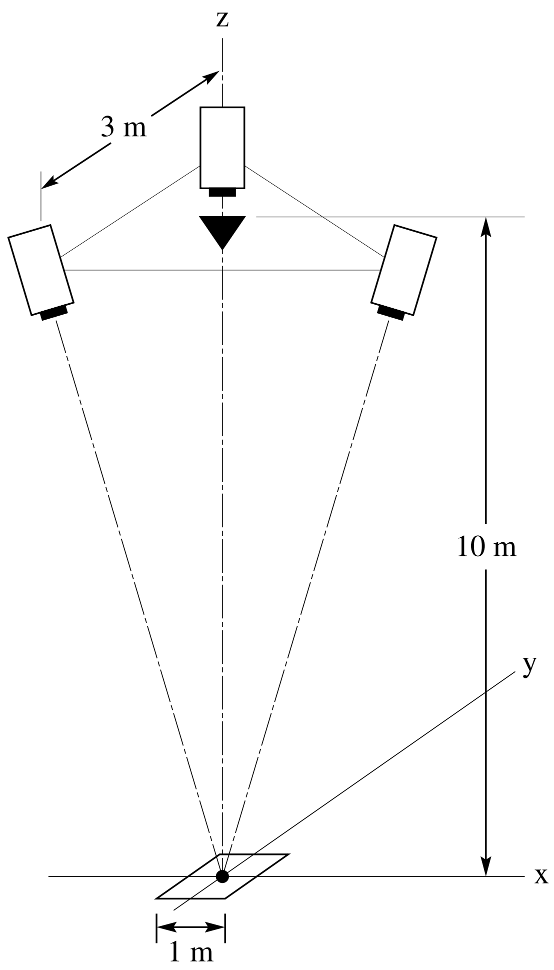

Figure 2: Camera placement for the numerical simulations

The simulated elevation maps were produced by generating a two-dimensional sequence of Fourier coefficients in wavenumber space (k_x, k_y) with a known amplitude spectrum and random phases, taking an inverse Fourier transform, and storing the results in a two-dimensional array $z(x, y)$. To study the relationship between spectral shape and retrieval accuracy three spectral forms were used. Form 1 had no dependency on the wind direction,

$$\begin{aligned} |H| &= 0.0007Uk^{-4}\,, \qquad \text{for} \;\; k \leq 1.0cm^{-1} \\ |H| &= 0, \qquad\qquad\quad\;\; \text{for} \;\; k > 1.0cm^{-1} \end{aligned} \tag{9}$$

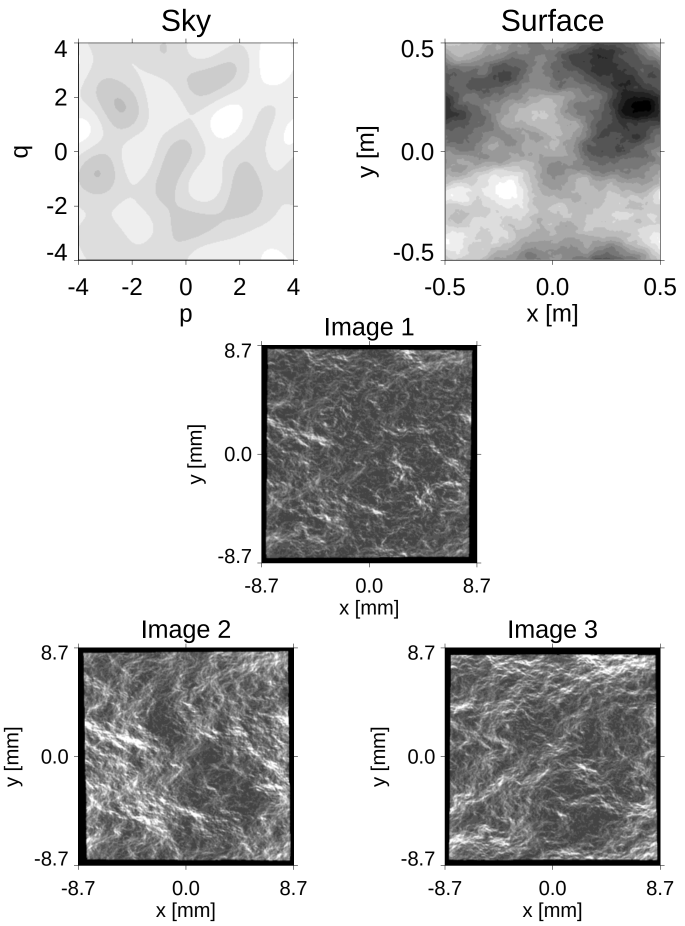

Figure 3: The sky radiance map (used for all simulations), a sample surface elevation map, and a set of three synthesized images. For the sky radiance map, light pixels represent higher radiance values and dark pixels represent lower radiance values. For the surface elevation map, light pixels represent higher elevations and dark pixels represent lower elevations

form 2 had a moderate dependency on the wind direction

$$
\begin{aligned}
|H| &= 0.0007 U k^{-4}\left[0.5 + 0.5\,\cos^2(\phi - \phi_0)\right], && \text{for} \quad k \le 1.0 cm^{-1} \\
|H| &= 0, && \text{for} \quad k > 1.0 cm^{-1}
\end{aligned}
\tag{10}
$$

and form 3 had a strong dependency on the wind direction

$$
\begin{aligned}
|H| &= 0.0007Uk^{-4}\cos^2(\phi-\phi_0), \qquad && \text{for} \quad k \le 1.0cm^{-1} \\
|H| &= 0, && \text{for} \quad k > 1.0cm^{-1}
\end{aligned}
\tag{11}
$$

where ϕ_0 is the wind direction, $k = \sqrt{k_x^2 + k_y^2}$ is the wavenumber, and $\phi = \tan^{-1}\left[\frac{k_y}{k_x}\right]$ is the wave direction. In all cases, the wind speed U was set to 10 m/sec and the wind direction ϕ_0 was set to 0°. The shapes of the elevation and slope spectra for each form are shown in Figure 4, and sample elevation maps generated using spectral forms 1, 2 and 3 are shown in Figure 5.

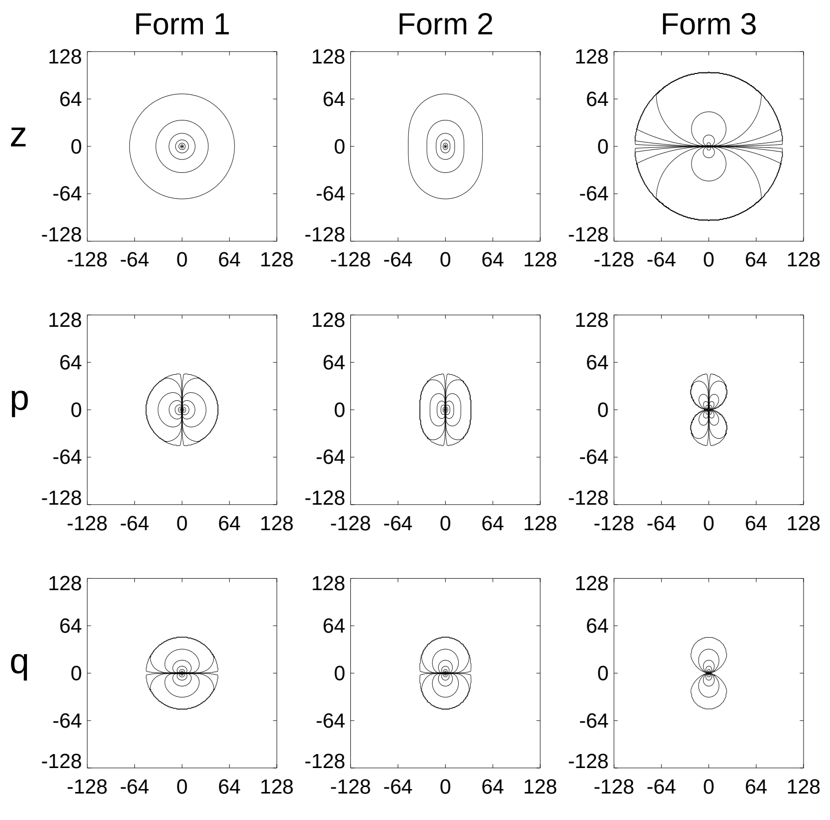

Figure 4: Contour maps of the elevation and slope wavenumber spectra for forms 1, 2 and 3

To achieve statistically reliability twenty realizations of each spectral form were produced and analyzed. For each simulated surface, three synthetic images were generated and 16

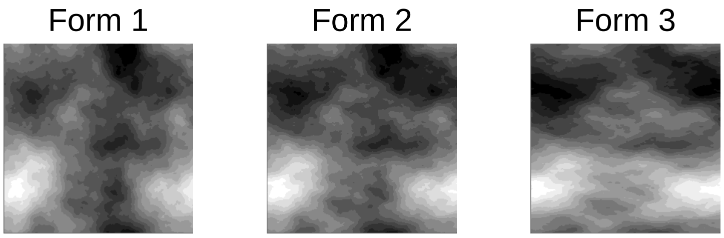

Figure 5: Surface realizations of spectral forms 1, 2 and 3. Higher elevations are depicted by lighter shades and lower elevations are depicted by darker shades.

equally spaced seed elevations were selected. The surface elevation and gradient maps were then retrieved using the $\mathbf{S^3}$ technique. The residual threshold, used by the surface orientation algorithm, was set so that approximately 85 percent of the retrieved points had a residual below the threshold. In other words, after convergence approximately 15 percent of the nodes in the retrieval grid were classified as missing. Afterwards missing points in the elevation and gradient maps were filled in using a median replacement filter, which replaced missing data with the median of the surrounding non-missing data. To illustrate this procedure Figure 6 shows an example of an input (simulated) and retrieved elevation and gradient maps, as well as the difference between the input and retrieved maps.

The accuracy of the retrieval was tested by computing the coherency spectrum between the known (simulated) and recovered elevation and slope maps. The coherency spectrum $\Gamma_z(k_x, k_y)$ for the elevation maps is defined by

$$\Gamma_z(k_x, k_y) = \frac{\left|\sum_{m=0}^{N-1} Z_m \tilde{Z}_m^*\right|}{\sqrt{\left|\sum_{m=0}^{N-1} \tilde{Z}_m \tilde{Z}_m^*\right| \left|\sum_{m=0}^{N-1} Z_m Z_m^*\right|}} \tag{12}$$

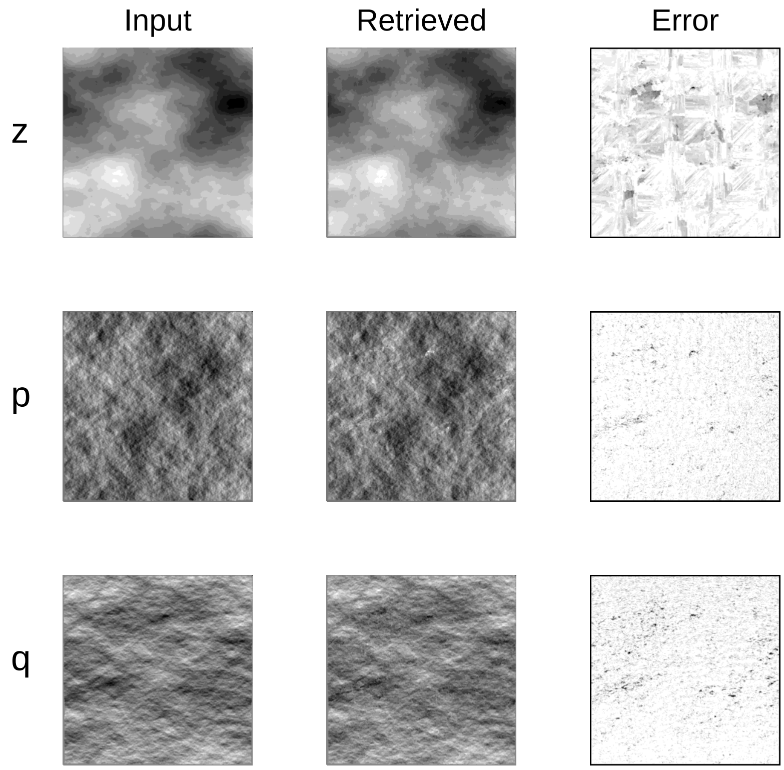

Figure 6: Sample input and retrieved (z, p, q) maps as well as the difference between the input and retrieved maps. For the error maps low errors are shown in lighter shades and higher errors are shown in dark shades.

where N is the number of simulations, $z_m(x, y)$ is the m'th simulated elevation map, $\tilde{z}_m(x, y)$ is the m'th retrieved elevation map, $w(x, y)$ is a spectral window, and Z_m and $\tilde{Z}_m$ are the discrete Fourier transforms of $z_m(x, y) \cdot w(x, y)$ and $\tilde{z}_m(x, y) \cdot w(x, y)$ respectively. Similar expression exist for the slope coherency spectra $\Gamma_p(k_x, k_y)$ and $\Gamma_q(k_x, k_y)$. A coherency spectrum may be interpreted as the correlation coefficient for each wavenumber, between the input spectrum (derived from the simulated surface) and the output spectrum (derived from the retrieved surface). A significance test may be applied to the coefficients $\Gamma_z(k_x, k_y)$

to determine which spectral components are correlated. In the above examples, $\Gamma_z(k_x, k_y)$ was computed using 20 degrees of freedom, which has a minimum significance value of 0.28 [Jenkins and Watts, 1968] [Julian, 1975]. Thus, when $\Gamma_z(k_x, k_y) \geq 0.28$ the input and output spectral amplitudes at (k_x, k_y) are correlated.

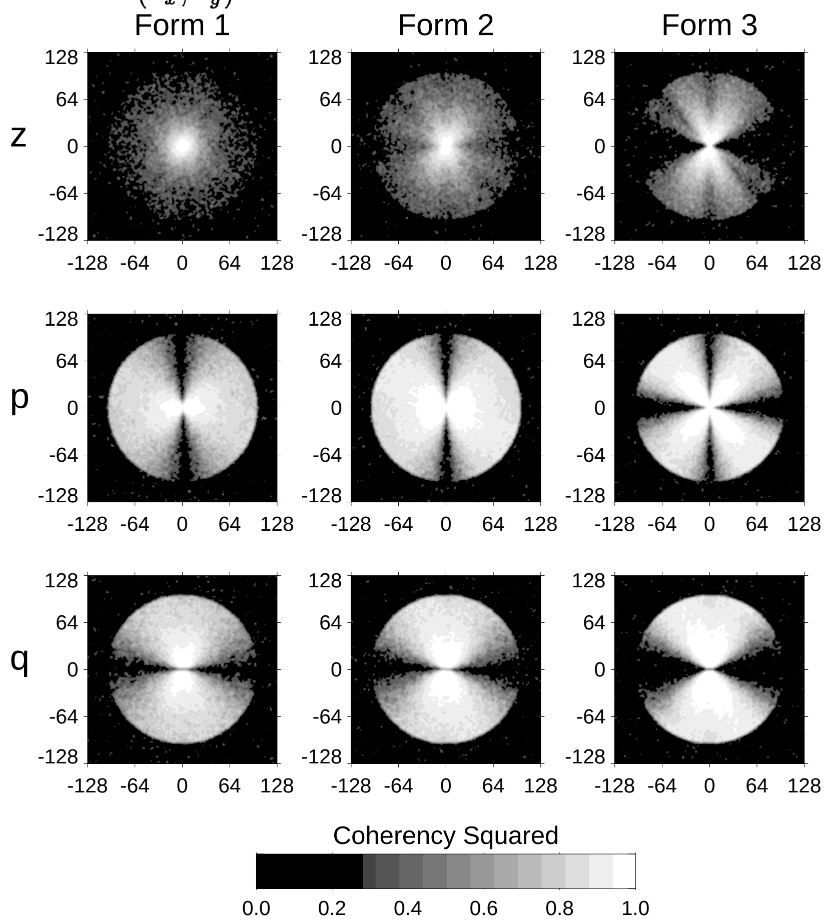

Figure 7: Coherency spectra between the simulated and retrieved elevation and slope maps for spectral forms 1, 2 and 3.

The coherency spectra for spectral forms 1, 2 and 3 are shown in Figure 7. When the variance, or spectral amplitude, associated with a particular wavenumber is small, the variability of the spectral amplitude become large, which results in a poor correlation (low coherency) between the input and retrieved elevation and slope maps. Thus, when the spectral amplitudes are small it becomes harder to reject the null hypothesis—to determine if two spectral amplitudes are correlated. Inspection of Figures 4 and 7 clearly shows this relationship between spectral shape and coherency. It also is apparent from Figure 7 that the $\mathbf{S^3}$ technique recovers slopes better than elevations. This is not unexpected because the appearance of an image of a water surface is primarily a function of the surface slopes.

5 Discussion and Conclusions

The specular surface stereo technique is specifically designed to retrieve the shape of a water surface under natural lighting conditions. The technique consists of two parts—a surface orientation procedure, that computes the gradient at a point on the surface by minimizing the difference between observed and synthesized irradiance values, and a surface propagation procedure. The surface orientation procedure incorporates an image synthesis model, that predicts the irradiance at each image pixel as a function of the source radiance, reflectance properties of the surface, and the optical properties of the imaging device. There are four general requirements for using the specular surface stereo method, (1) the existence of an image synthesis model where the brightness of the surface depends on the orientation of the surface facets, (2) a radiance map of the illumination source, (3) a sufficient number of images of surface so that the surface orientation procedure usually has a unique solution, and (4) a few seed elevation estimates.

To test the feasibility the specular surface stereo technique, an idealized surface, radiance map and camera were specified, and three images of the surface were synthesized. Then from the radiance map, synthesized images and sixteen seed elevation estimates, the surface elevations and gradients were retrieved. Under these controlled circumstances, the results demonstrated the feasibility of the method. When applied to real world situations, additional factors must be included in the image synthesis and illumination models. In a field environment, for example, sun glitter, sky polarization, rough surface features, finite radius of curvature, and the absorption and multiple scattering of light within the water volume must be included in the models.

6 Acknowledgements

This research was sponsored by a grant (N00014-89-J-3229) from the Office of Naval Research.

References

[Blake, 1985] Blake, A. (1985). Specular stereo. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI)*, 2:973–976.

[Blake and Brelstaff, 1988] Blake, A. and Brelstaff, G. (1988). Geometry from specularities. *International Conference on Computer Vision (ICCV)*, pages 394–403.

[Bock and Hara, 1992] Bock, E. J. and Hara, T. (1992). Optical measurements of ripples using a scanning laser slope gauge, part ii: Data analysis and interpretation from a laboratory wave tank. In Estep, L., editor, *Proc. SPIE. Optics of the Air-Sea Interface: Theory and Measurements*, pages 272–282, San Diego, CA.

[Born and Wolf, 1980] Born, M. and Wolf, E. (1980). *Principles of Optics, Sixth Ed.* Pergamon Press.

[Chen and Fung, 1992] Chen, K. S. and Fung, A. K. (1992). Bispectrum model for sea surface scattering. In *IEEE International Geoscience and Remote Sensing Symposium*, Houston, Texas.

[Chen et al., 1992] Chen, K. S., Fung, A. K., and Weissman, D. E. (1992). A backscattering model for ocean surface. *IEEE Trans. on Geoscience and Remote Sensing*, 8(4):811–817.

[Coleman and Jain, 1982] Coleman, E. N. and Jain, R. (1982). Obtaining 3-dimensional shape of textured and specular surface using four-source photometry. *Computer Graphics and Image Processing*, 18:309–328.

[Hartt and Carlotto, 1989] Hartt, E. and Carlotto, M. (1989). A method for shape-from-shading using multiple images acquired under different viewing and lighting conditions. In *Proc. IEEE Computer Vision and Pattern Recognition*, pages 53–60, San Diego, CA.

[Healey and Benford, 1987] Healey, G. and Benford, T. O. (1987). Local shape from specularity. In *Proc. Image Understanding Workshop*, pages 874–887, Los Angeles, CA.

[Horn, 1986] Horn, B. K. P. (1986). *Robot Vision.* MIT Press, Cambridge MA.

[Horn and Brooks, 1989] Horn, B. K. P. and Brooks, M. J., editors (1989). *Shape from Shading.* MIT Press, Cambridge MA.

[Ikeuchi, 1981] Ikeuchi, K. (1981). Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-3(6):661–669.

[Jähne and Riemer, 1990] Jähne, B. and Riemer, K. S. (1990). Two-dimensional wave number spectra of small-scale water surface waves. *Journal of Geophysical Research*, 95(C7):11531–11546.

[Jenkins and Watts, 1968] Jenkins, G. M. and Watts, D. G. (1968). *Spectral Analysis and its Applications*. Holden-Day, Oakland, CA.

[Julian, 1975] Julian, P. R. (1975). Comments on the determination of significance levels of the coherency statistic. *Journal of the Atmospheric Sciences*, 32:836–837.

[Keller and Gotwols, 1983] Keller, W. and Gotwols, B. L. (1983). Two-dimensional optical measurement of wave slope. *Applied Optics*, 22:3476–3478.

[Lee et al., 1992] Lee, P. H. Y., J. D. Barter, K. L. B., Hindman, C. L., Lake, B. M., Rungaldier, H., Schatzman, J. C., Shelton, J. C., Wagner, R. N., Williams, A. B., Yee, R., and Yuenn, H. C. (1992). Recent advances in ocean surface characterization by a scanning laser slope gauge. In Estep, L., editor, *Proc. SPIE. Optics of the Air-Sea Interface: Theory and Measurements*, pages 234–244, San Diego, CA.

[Marr and Poggio, 1976] Marr, D. and Poggio, T. (1976). Cooperative computation of stereo disparity. *Science*, 194:283–287.

[Martinsen and Bock, 1992] Martinsen, R. J. and Bock, E. J. (1992). Optical measurements of ripples using a scanning laser slope gauge, part i: Instrumentation and preliminary results. In Estep, L., editor, *Proc. SPIE. Optics of the Air-Sea Interface: Theory and Measurements*, pages 258–271, San Diego, CA.

[Murase, 1992] Murase, H. (1992). Surface shape reconstruction of a nonrigid transparent object using refraction and motion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-14(10):1045–1052.

[Nickolaev et al., 1992] Nickolaev, N. I., Yordanov, O. I., and Michalev, M. A. (1992). Non-gaussian effects in the two-scale model from rough surface scattering. *Journal of Geophysical Research*, 97(C10):15,617–15,624.

[Sanderson et al., 1988] Sanderson, A. C., Weiss, L. E., and Nayar, S. K. (1988). Structured highlight inspection of specular surfaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-10:44–55.

[Schultz, 1992] Schultz, H. (1992). Specular surface stereo: A new method for retrieving the shape of a water surface. In Estep, L., editor, *Proc. SPIE. Optics of the Air-Sea Interface: Theory and Measurements*, pages 283–294, San Diego, CA.

[Schultz, 1994] Schultz, H. (1994). Retrieving shape information from multiple images of a specular surface. *IEEE Trans. on Pattern Analysis and Machine Intelligence.*, 16(2):195–201.

[Slama, 1980] Slama, C. C., editor (1980). *Manual of Photogrammetry, fourth Edition.* American Society of Photogrammetry, Falls Church, VA.

[Tagare and deFigueiredo, 1989] Tagare, H. D. and deFigueiredo, R. J. P. (1989). A theory of photometric stereo for a general class of reflectance maps. In *Proc. IEEE Computer Vision and Pattern Recognition*, pages 38–45, San Diego, CA.

[Thomas et al., 1991] Thomas, J., Kober, W., and Leberl, F. (1991). Multiple image sar shape-from-shading. *Photogrammetric Engineering and Remote Sensing*, 57(1):51–59.

[Wass and Jähne, 1992] Wass, S. and Jähne, B. (1992). Combined slope-height measurements of short wind waves: First results from field and laboratory measurements. In Estep, L., editor, *Proc. SPIE Optics of the Air-Sea Interface: Theory and Measurement*, volume 1749, pages 295–306, San Diego, CA.

[Wolff and Boult, 1991] Wolff, L. B. and Boult, T. E. (1991). Contrasting object features using a polarization reflection model. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-13:635–657.

[Woodham, 1978] Woodham, R. J. (1978). Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proc. SPIE 22nd Annual Technical Symp.*, volume 155, pages 136–143, San Diego, CA.

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

Fig. 1. Viewing geometry. The focal point of an idea camera with focal length f is located at P; the observation point on the surface is located at O; $\hat{r}$ is the unit vector that points from O to P; $\hat{s}$ is the surface normal unit vector; $\hat{t}$ is the reflection of $\hat{r}$ about $\hat{s}$.

Fig. 2. Camera placement for the numerical simulations.

Fig. 3. The sky radiance map (used for all simulations), a sample surface elevation map, and a set of three synthesized images. For the sky radiance map, light pixels represent higher radiance values and dark pixels represent lower radiance values. For the surface elevation map, light pixels represent higher elevations and dark pixels represent lower elevations.

Fig. 4. Contour maps of the elevation and slope wavenumber spectra for forms 1, 2 and 3.

Fig. 5. Surface realizations of spectral forms 1, 2 and 3. Higher elevations are depicted by lighter shades and lower elevations are depicted by darker shades.

Fig. 6. Sample input and retrieved (z, p, q) maps as well as the difference between the input and retrieved maps. For the error maps low errors are shown in lighter shades and higher errors are shown in dark shades.

Fig. 7. Coherency spectra between the simulated and retrieved elevation and slope maps for spectral forms 1, 2 and 3.