

A Differential Optical Range Camera

Hany Farid and Eero P. Simoncelli †

GRASP Laboratory

Department of Computer and Information Science

University of Pennsylvania, Philadelphia, PA 19104, USA

Abstract

We describe a novel formulation of the range recovery problem, based on computation of the differential variation in image intensities with respect to changes in camera position. The method uses a single stationary camera and a pair of calibrated optical attenuation masks to directly measure this differential quantity. The subsequent computation of the range image is simple and should be suitable for real-time implementation. We have constructed and tested a prototype range camera based on these principles.

Introduction

Visual images are formed via the projection of light from the three-dimensional world onto a two-dimensional sensor. In an idealized pinhole camera, all points lying on a ray passing through the pinhole will be imaged onto the same image position. Thus, information about the distance to objects in the scene (i.e., *range*) is lost. Range information can be recovered by measuring the change in appearance of the

world resulting from a change in viewing position. Traditionally, this is accomplished via simultaneous measurements with two cameras (binocular stereo), or via a sequence of measurements collected over time from a moving camera (structure from motion).

The recovery of range in these approaches frequently relies on an assumption of brightness constancy, which states that the brightness of the image of a point in the world is constant when viewed from different positions [4]. Consider the formulation of this assumption in one dimension (the extension to two dimensions is straightforward). Let $f(x; v)$ describe the intensity function measured through a pinhole camera system. The variable v corresponds to the pinhole position (along the direction perpendicular to the optical axis). The variable x parameterizes the position on the sensor. This configuration is illustrated in Figure 1. According to the assumption, the intensity function $f(x; v)$ is of the form:

$$f(x; v) = I\left(x - \frac{vd}{Z}\right), \quad (1)$$

where $I(x) = f(x; v)|_{v=0}$, d is the distance between the pinhole and the sensor and Z is the range (distance from the pinhole to a point in the world). Note that this assumption will typically be violated near occlusion boundaries, where points visible from one viewpoint are invisible from another.

Several complications arise in these approaches. The degree to which the bright-

This work was supported by NSF Grant STC SBR-89-20230; ARPA Grant N00014-92-J-1647; ARMY Grant DAAH04-96-1-0007; University Research Foundation Equipment Grant.

† Currently at Center for Neural Science and Courant Institute for Mathematical Sciences, New York University, New York, NY 10003.

Portions of this work have appeared in [10].




Figure 1: Geometry for a binocular stereo system with pinhole cameras. The variable V parameterizes the position of the camera pinholes. According to the brightness constancy constraint, the intensity of a point in the world, as recorded by the two pinhole cameras, should be the same (i.e. $f(x; 0) = f(x - \frac{vd}{Z}; v)$).

ness constancy assumption holds will, in general, decrease with increasing camera displacement. This is due to larger occluded image regions, and increased effects of the non-Lambertianity of surface reflectances. Violations of the brightness constancy assumption lead to difficulties in matching corresponding points in the images (the so-called “correspondence problem”). Furthermore, a two-camera stereo system (or a single moving camera) requires careful calibration of relative positions, orientations, and intrinsic parameters of the camera(s).

These problems are partially alleviated in techniques utilizing a single stationary camera. A number of these techniques are based on estimation of blur or relative blur from two or more images (e.g., [6, 9, 11, 12, 8]). Adelson and Wang [1] describes an unusual method in which a lenticular array is placed over the sensor, effectively allowing the camera to capture visual images from several viewpoints in a single exposure. Dowski and Cathey[2] and Jones and Lamb [5] each describe range imaging systems that use an optical attenuation mask in front of the lens. By observing local spectral

information in a single image, they are able to estimate range. Both techniques rely on power spectral assumptions about the scene.

In this paper, we propose a single-camera method which avoids some of the computational and technical difficulties of the single-camera approaches discussed above. In particular, we propose a “direct” differential method for range estimation which computes the image derivative with respect to viewing position using a single stationary camera and an optical attenuation mask. This approach avoids the correspondence problem, makes no spectral assumptions about the scene, is relatively straightforward to calibrate, and is computationally efficient. Based on these principles we have constructed and tested a prototype range camera. The construction of this camera, as well as some results are presented here.

Direct Viewpoint Derivatives

For the purpose of recovering range, we are interested in computing the change in the appearance of the world with respect to change in viewing position. It is thus natural to consider differential measurement techniques. Taking partial derivatives of the intensity function $f(x; v)$ (Equation (1)) with respect to the image and viewing positions, and evaluating at $v = 0$ gives:

$$\begin{aligned} I_x(x) &\equiv \left. \frac{\partial f(x;v)}{\partial x} \right|_{v=0} \\ &= I'(x), \end{aligned} \quad (2)$$

and

$$\begin{aligned} I_v(x) &\equiv \left. \frac{\partial f(x;v)}{\partial v} \right|_{v=0} \\ &= -\frac{d}{Z} I'(x), \end{aligned} \quad (3)$$

where $I'(\cdot)$ indicates the derivative of $I(\cdot)$ with respect to its argument. Combining these two expressions gives:

$$I_v(x) = -\frac{d}{Z} I_x(x). \quad (4)$$

Clearly, an estimate of the range, Z , can be computed using this equation. Note that in the case of differential binocular stereo (e.g., [7]), the derivative with respect to viewing position, I_v , is replaced by a difference, $I_{v_1} - I_{v_2}$. A similar relationship is used in computing structure from motion (for known camera motion), where I_v is typically replaced by differences of consecutive images. We now show a direct method for measurement of this derivative through the use of an optical attenuation mask.

Consider a world consisting of a single uniform intensity point light source and a standard lens-based imaging system with a variable-opacity optical attenuation mask, $M(u)$, placed directly in front of the lens (left side of Figure 2). The light striking the lens is attenuated by the value of the mask function at that particular spatial location.¹ With such a configuration, the image of the point source will be a scaled and dilated version of the mask function:

$$I(x) = \frac{1}{\alpha} M\left(\frac{x}{\alpha}\right). \quad (5)$$

The scale factor, α , is a monotonic function of the distance to the point source, Z , and may be derived from the imaging geometry:

$$\alpha = 1 - \frac{d}{f} + \frac{d}{Z}, \quad (6)$$

where d is the distance between lens and sensor, and f is the focal length of the lens.

In the system shown on the left side of Figure 2, the *effective* viewpoint may be altered by translating the mask, while leaving the lens and sensor stationary. The generalized intensity function, for a mask centered at position v is written as:

$$f(x; v) = \frac{1}{\alpha} M\left(\frac{x}{\alpha} - v\right), \quad (7)$$

assuming that the non-zero portion of the mask does not extend past the edge of the lens.

¹For our purposes, we assume that the values of such a mask function are real numbers in the range [0,1].

The *differential* change in the image (with respect to a change in the mask position) may be computed by taking the derivative of this equation with respect to the mask position, v , evaluated at $v = 0$:

$$\begin{aligned} I_v(x) &\equiv \left. \frac{\partial}{\partial v} f(x; v) \right|_{v=0} \\ &= -\frac{1}{\alpha} M'\left(\frac{x}{\alpha}\right), \end{aligned} \quad (8)$$

where $M'(\cdot)$ is the derivative of the mask function $M(\cdot)$ with respect to its argument. The derivative with respect to viewing position, $I_v(x)$, may thus be computed directly by imaging with the optical attenuation mask $M'(u)$!²

Finally, notice that the spatial derivative of the first image, $I(x)$, is closely related to the image $I_v(x)$:

$$\begin{aligned} I_x(x) &\equiv \left. \frac{\partial}{\partial x} f(x; v) \right|_{v=0} \\ &= \frac{1}{\alpha^2} M'\left(\frac{x}{\alpha}\right) \\ &= -\frac{1}{\alpha} I_v(x). \end{aligned} \quad (9)$$

From this relationship, the scaling parameter α may be computed as the ratio of the spatial derivative of the image formed through the mask $M(u)$, and the image formed through the derivative of that mask, $M'(u)$. This computation is illustrated in Figure 2. The distance to the point source can subsequently be computed from α using the monotonic relationship given in Equation (6). Note that the resulting equation for estimating range is identical to that of Equation (4) when $d = f$ (i.e., when the camera is focused at infinity).

Range Estimation

Equation (9) embodies the fundamental relationship used for the direct differential computation of range of a single point light source. A more realistic world consisting of a collection of many such uniform intensity point sources

²In practice, $M'(u)$ cannot be directly used as an attenuation mask, since it contains negative values. This issue is addressed later in the paper.




Figure 2: Direct differential range estimation for a single uniform intensity point source. Images of a point light source are formed using two different optical attenuation masks, $M(u)$ and its derivative, $M'(u)$. In each case, the image formed is a scaled and dilated copy of the mask function (by an amount α , monotonically related to the depth, Z). Computing the spatial (image) derivative of the image formed under mask $M(u)$ produces an image that is identical to the image formed under the derivative mask, $M'(u)$, except for a scale factor α . Thus, α may be estimated as the ratio of the two images. Range is computed from α using the relationship given in Equation (6).

imaged through an optical attenuation mask will produce an image consisting of a superposition of scaled and dilated versions of the masks. In particular, we can write an expression for the image by summing the images of the visible points, p , in the world:

$$f(x; v) = \int dx_p \frac{1}{\alpha_p} M\left(\frac{x-x_p}{\alpha_p} - v\right) L(x_p), \quad (10)$$

where the integral is performed over the variable x_p , the position in the sensor of a point p projected through the center of the lens. The intensity of the world point p is denoted as $L(x_p)$,

and α_p is monotonically related to the distance to p (as in Equation (6)). Note again that we must assume that each point produces a uniform light intensity across the mask.

Again, consider the derivatives of $f(x; v)$ with respect to viewing position, v , and image position, x :

$$\begin{aligned} \frac{\partial}{\partial v} f(x; v) &= \frac{\partial}{\partial v} \int dx_p \frac{1}{\alpha_p} M\left(\frac{x-x_p}{\alpha_p} - v\right) L(x_p) \\ &= - \int dx_p \frac{1}{\alpha_p} M'\left(\frac{x-x_p}{\alpha_p} - v\right) L(x_p) \end{aligned} \quad (11)$$

and

$$\begin{aligned} \frac{\partial}{\partial x} f(x; v) &= \frac{\partial}{\partial x} \int dx_p \frac{1}{\alpha_p} M\left(\frac{x-x_p}{\alpha_p} - v\right) L(x_p) \\ &= \int dx_p \frac{1}{\alpha_p^2} M'\left(\frac{x-x_p}{\alpha_p} - v\right) L(x_p), \end{aligned} \quad (12)$$

As in the previous section, the following two partial derivative images are defined:

$$\begin{aligned} I_v(x) &\equiv \frac{\partial}{\partial v} f(x; v) \Big|_{v=0} \\ &= - \int dx_p \frac{1}{\alpha_p} M'\left(\frac{x-x_p}{\alpha_p}\right) L(x_p), \end{aligned} \quad (13)$$

and

$$\begin{aligned} I_x(x) &\equiv \frac{\partial}{\partial x} f(x; v) \Big|_{v=0} \\ &= \int dx_p \frac{1}{\alpha_p^2} M'\left(\frac{x-x_p}{\alpha_p}\right) L(x_p). \end{aligned} \quad (14)$$

Equations (13) and (14) differ only in a multiplicative term of $\frac{1}{\alpha_p}$. Unfortunately, solving for α_p is nontrivial, since it is embedded in the integrand and depends on the integration variable. Consider, however, the special case where all points in the world lie on a frontal-parallel plane relative to the sensor.³ Under this condition, the scaling parameter α_p is the same for all points x_p and Equations (13) and (14) can be written as:

$$I_v(x) = -\frac{1}{\alpha} \int dx_p M'\left(\frac{x-x_p}{\alpha}\right) L(x_p) \quad (15)$$

$$I_x(x) = \frac{1}{\alpha^2} \int dx_p M'\left(\frac{x-x_p}{\alpha}\right) L(x_p). \quad (16)$$

³In actuality, this assumption need only be made locally.

The scaling parameter, α , a function of the distance to the points in the world (Equation (6)) can then be computed as the ratio:

$$I_v(x) = -\alpha I_x(x). \quad (17)$$

In order to deal with singularities (i.e., $I_x = 0$), a least-squares estimator can be used for α (as in [7]). Specifically, we minimize the following error function:

$$E(\alpha) = \sum_P (I_v(x) + \alpha I_x(x))^2, \quad (18)$$

where the summation is performed over a small patch in the image, P . Taking the derivative with respect to α , setting equal to zero and solving for α yields the minimal solution:

$$\alpha = -\frac{\sum_P I_v(x) I_x(x)}{\sum_P I_x(x)^2}. \quad (19)$$

The algorithm easily extends to a three-dimensional world: we need only consider two-dimensional masks $M(u, w)$, and the horizontal partial derivative $M_u(u, w) = \partial M(u, w) / \partial u$. For a more robust implementation, the vertical partial derivative mask $\partial M(u, w) / \partial w$ may also be included. The least-squares error function becomes:

$$E(\alpha) = \sum_P (I_u + \alpha I_x)^2 + (I_w + \alpha I_y)^2. \quad (20)$$

Solving for the minimizing α gives:

$$\alpha = -\frac{\sum_P (I_u I_x + I_w I_y)}{\sum_P (I_x^2 + I_y^2)}. \quad (21)$$

Optical Mask Design

Thus far, the only restriction placed on the optical masks, $M(u, w)$ and $M_u(u, w)$, is that the second be the derivative of the first. Illustrated in Figure 3 is an example of such a matched pair of masks based on a two-dimensional Gaussian. Typically, the function




Figure 3: Gaussian optical masks. Illustrated along the top row is a two-dimensional Gaussian mask, $M(u, w)$ (left), and its partial spatial derivative, $M_u(u, w)$ (right). Illustrated along the bottom row are a pair of non-negative masks, $M_1(u, w)$ and $M_2(u, w)$, computed from the Gaussian and its derivative masks using Equations (22) and (23).

$M_u(u, w)$ will have negative values (as in the case of the Gaussian) and thus is not feasible for use as an optical attenuation mask. Furthermore, a positive constant cannot simply be added to $M_u(u, w)$, since this will destroy the required derivative relationship between the two masks.

Due to the linearity of the imaging process, however, we can use masks that are linear combinations of the masks $M_u(u, w)$ and $M(u, w)$. In particular, a scalar multiple of $M(u, w)$ can be added to $M_u(u, w)$ in order to form a mask function that is entirely positive. The new mask, $M_1(u, w)$, shown in Figure 3, is given by:

$$M_1(u, w) = \beta M(u, w) + \gamma M_u(u, w), \quad (22)$$

where β, γ are scaling constants chosen to force the function $M_1(u, w)$ to fill the range $[0, 1]$. A second symmetrical mask can be formed by *subtracting* $M_u(u, w)$ from $M(u, w)$:

$$M_2(u, w) = \beta M(u, w) - \gamma M_u(u, w). \quad (23)$$

Note that $M_2(u, w)$ is equal to $M_1(u, w)$ rotated

180 degrees about its center, and that

$$M(u, w) = \frac{M_1(u, w) + M_2(u, w)}{2\beta} \quad (24)$$

$$M_u(u, w) = \frac{M_1(u, w) - M_2(u, w)}{2\gamma}. \quad (25)$$

Again, by linearity of the imaging process, the *images* that would have been obtained with the masks $M(u, w)$ and $M_u(u, w)$ can be recovered from images obtained with two masks $M_1(u, w)$ and $M_2(u, w)$. In particular, let $I_1(u, w)$ be the image obtained through the mask $M_1(u, w)$, and $I_2(u, w)$ the image obtained through the mask $M_2(u, w)$. Then:

$$I(u, w) = \frac{I_1(u, w) + I_2(u, w)}{2\beta} \quad (26)$$

$$I_v(u, w) = \frac{I_1(u, w) - I_2(u, w)}{2\gamma}, \quad (27)$$

where $I(u, w)$ and $I_v(u, w)$ are the desired quantities for estimating the range image using Equation (19).

Results

In addition to a series of simulations (results not shown here) we have constructed a prototype camera for validating the direct differential approach to range estimation. The camera consists of an optical attenuation mask (a liquid crystal spatial light modulator, LC SLM) sandwiched between a pair planar-convex lenses, and placed in front of an off-the-shelf SONY XC-77 CCD camera. The essential component of this system is the LC SLM purchased from CRL Smectic Technology (Middlesex, UK). This device is a fully programmable, fast-switching, twisted nematic liquid crystal display. This device has a display area of 28.48 mm (W) \times 20.16 mm (H); the spatial resolution is 640 \times 480, with 4 possible grayscale values. The LC-SLM was calibrated to correct for any non-linearities: light transmittance for each grayscale was measured using a photometer. This calibration information was incorporated into a standard stochastic error diffusion dithering algorithm (e.g., [3]) in order to




Figure 4: Illustrated are 1-D slices of the image of “point light source” taken through a pair of Gaussian-based masks (top). Also shown are 1-D slices of the images I and I_v (bottom: solid curve), and the their fit to a Gaussian and its derivative (dashed curve).

render the optical masks with reasonable accuracy. The CCD sensor was also calibrated and found to be nearly linear. The display is controlled through a PC VGA video interface. The LC SLM refreshes its display at 30 Hz; when synchronized with the frame grabber, the required images ($I(x, y)$ and $I_v(x, y)$) taken through the pair of masks may be acquired at 15 Hz.

In our first set of experiments with this range camera, a pinhole, at a distance of 20 cm, backlit by a desk lamp was imaged through a reduction tube. Illustrated in Figure 4 are 1-D slices of the images, I_1 and I_2 , taken through a pair of Gaussian-based masks. Also shown in this figure are 1-D slices of the images I and I_v computed by the taking the appropriate combinations of I_1 and I_2 (Equations (26) and (27)). Note that the resulting images are reasonably well fit to a Gaussian and its derivative (dashed line).

In our next set of experiments, a simple target consisting of a frontal-parallel planar textured pattern (random white noise) was imaged through a pair of Gaussian-based masks (see Figure 3). An estimate of depth was com-




Figure 5: Range images. Illustrated are the recovered range images for a frontal-parallel planar target at a distance of 11 and 17 cm from the camera. The computed range maps have a mean of 11.3 and 17.16 cm, with a standard deviation of 0.47 and 0.59 cm, respectively.

puted according to Equations (19) and (6). Illustrated in Figure 5 are the recovered range images for the target placed at a distance of 11cm and 17cm from the camera. For these targets, the mean estimate of range was 11.3 and 17.16 cm, with a standard deviation of 0.47 and 0.59 cm, respectively. In these results, the images were pre-processed to remove regions with low spatial derivative (this amounted to approximately 50% of the data), the resulting “holes” were filled with a simple bilinear interpolation scheme.

Discussion

An optical attenuation mask placed in front of a lens-based imaging system produces an image which is a superposition of scaled and dilated copies of the mask function. The derivative of this image is related by a scale factor to a second image created with the deriva-

tive of the first mask. The scale factor is monotonically related to range. This simple observation has led us to a direct differential method for estimating range from a single stationary camera. In particular, the derivative with respect to viewing position is computed directly: it is simply the image formed under the derivative mask. Based on these principles, we have constructed and tested a prototype range camera. We are able to acquire the necessary pair of images at 15 Hz through the use of a fast-switching LC SLM as an optical attenuation mask. Since the subsequent processing of the images is simple and fast, this technique should be amenable to a real-time implementation.

Two assumptions have been made in our solution to this problem. Both of these assumptions are made (although often not explicitly) in nearly every structure from stereo or motion algorithm. The first assumption is that the light emanating from each point in the scene is constant across the lens (i.e., the brightness constancy assumption). Note that this assumption will typically be violated at occlusion boundaries, because the light emanating from a partially occluded point will hit only a portion of the lens. One potential solution to this problem is to expand the function describing the light emanating from a point in a Taylor series. The coefficients of these terms may be estimated by collecting additional measurements (i.e., images) with higher-order derivative masks.

The second assumption is that of locally frontal-parallel surface orientation. This assumption was necessary in order to solve for α_p given the two image measurements described by Equations (13) and (14). Solving without this assumption is a nonlinear optimization problem (since the α_p appears inside the argument of $M(\cdot)$), which should be amenable to an iterative solution.

There are still several mask design issues that need to be resolved. First, our example of Gaussian-based masks was somewhat arbi-

trary. A pair of masks should be designed from a set of optimality constraints based on derivative accuracy, effective baseline, light transmittance, etc. Once an optimal function is determined, the construction of the actual optical masks must be calibrated to include nonlinearities in the printing process (e.g., halftoning), and the effects of the intrinsic point spread function of the camera. In particular, the image of a point light source recorded by the camera with mask $M_u(u, v)$ must be equal to the spatial derivative of the image recorded with mask $M(u, v)$. Finally, noise in the image measurements, $I_1(u, w)$ and $I_2(u, w)$, will be amplified by the computations in Equation (26) and Equation (27): small values of β or γ are thus undesirable.

As with most ranging techniques, accuracy behaves according to the rules of triangulation. In particular, errors will be proportional to the square of the range, and inversely proportional to both the focal length and baseline.⁴ We have verified these relationships via simple simulations. As in many other range-imaging systems, the accuracy may be improved with the use of structured illumination.

A counterintuitive aspect of our technique is that it relies on the defocus of the image. In particular, a perfectly focused image corresponds to $\alpha = 0$, leading to a singularity in Equation (9). In practice, this may be alleviated by focusing the camera at infinity (i.e., $d = f$), thus ensuring that points at distances within the operating range of the algorithm will be sufficiently blurred.

And finally, an interesting variant of the technique arises when considering a Gaussian mask, and its derivative with respect to σ :

$$\begin{aligned} G(u, w) &= \frac{1}{\sigma^2} e^{-(u^2+w^2)/2\sigma^2}, \\ G_\sigma(u, w) &= \frac{\partial}{\partial \sigma} G(u, w) \\ &= -\frac{2}{\sigma^3} e^{-(u^2+w^2)/2\sigma^2} \end{aligned} \quad (28)$$

⁴Effective baseline in our system depends on the mask function, and is proportional to the lens diameter.

$$+ \frac{(u^2+w^2)}{\sigma^5} e^{-(u^2+w^2)/2\sigma^2}. \quad (29)$$

Let $I(x, y)$ and $I_\sigma(x, y)$ be the images obtained through the masks $G(u, w)$ and $G_\sigma(u, w)$, respectively. Using the same techniques as in the previous section, it can be shown that these two images obey the following constraint:

$$\begin{aligned} I_\sigma(x, y) &= \alpha^2 \sigma [I_{xx}(x, y) + I_{yy}(x, y)], \\ &= \alpha^2 \sigma \nabla^2 I(x, y), \end{aligned} \quad (30)$$

where $I_{xx}(x, y)$ and $I_{yy}(x, y)$ correspond to the horizontal and vertical second partial derivatives of $I(x, y)$, and ∇^2 is the Laplace operator. As before, α is inversely proportional to range, and is given by Equation (6). This formulation provides a differential algorithm for range-from-defocus. Unlike previous formulations (e.g., [9]), this solution avoids the artifacts arising from the computation of local Fourier transforms.

References

- [1] E.H. Adelson and J.Y.A. Wang. Single lens stereo with a plenoptic camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):99–106, 1992.
- [2] E.R. Dowski and W.T. Cathey. Single-lens single-image incoherent passive-ranging systems. *Applied Optics*, 33(29):6762–6773, 1994.
- [3] R.W. Floyd and L. Steinberg. An adaptive algorithm for spatial grey scale. *Proceedings of the Society for Information Display*, 17(2):75–77, 1976.
- [4] B.K.P. Horn. *Robot Vision*. MIT Press, Cambridge, MA, 1986.
- [5] D.G. Jones and D.G. Lamb. Analyzing the visual echo: Passive 3-d imaging with a multiple aperture camera. Technical Report CIM-93-3, Department of Electrical Engineering, McGill University, 1993.

- [6] E. Krotkov. Focusing. *International Journal of Computer Vision*, 1:223–237, 1987.
- [7] B.D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 674–679, Vancouver, 1981.
- [8] S.K. Nayar, M. Watanabe, and M. Noguchi. Real-time focus range sensor. In *Proceedings of the International Conference on Computer Vision*, pages 995–1001, Cambridge, MA, 1995.
- [9] A.P. Pentland. A new sense for depth of field. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(4):523–531, 1987.
- [10] E.P. Simoncelli and H. Farid. Direct differential range estimation using optical masks. In *Proceedings of the European Conference on Computer Vision*, pages II: 82–93, Cambridge, UK, 1996.
- [11] M. Subbarao. Parallel depth recovery by changing camera parameters. In *Proceedings of the International Conference on Computer Vision*, pages 149–155, 1988.
- [12] Y. Xiong and S. Shafer. Depth from focusing and defocusing. In *Proc. of the DARPA Image Understanding Workshop*, pages 967–976, 1993.