

Surface Determination by Photometric Ranging

Jeffrey B. Mulligan
NASA Ames Research Center
Jeffrey.B.Mulligan@nasa.gov

Xavier L. C. Brolly
NASA Ames Research Center
brollyx@eos.arc.nasa.gov

Abstract

We describe a method for recovering the three dimensional shape of a surface, using multiple images from a single camera, closely related to "light field reconstruction" described by Magda *et al.* [1]. A volume array of switchable illuminators is used to obtain a set of images that are used for the reconstruction. Unlike photometric stereo, where the illumination is assumed to be homogeneous, we exploit the inverse-square-law falloff of intensity with the distance between the illuminator and the surface. By placing illuminators at a variety of distances from the measurement volume, a direct estimate of surface range can be obtained, in addition to surface normal estimates obtained as in photometric stereo. We discuss solution methods and present simulation results for a few simple geometries. Aspects of illuminator design and calibration are also considered.

1. Introduction

This paper is concerned with the recovery of three-dimensional surface shape using images from a conventional camera. We introduce the term *photometric ranging* to describe the direct computation of depth or range from a set of photometric measurements. Our application is the measurement of the shape of the human face, but the methods we discuss are completely general. Here we present a method of photometric ranging based on a spatial array of compact diffuse illuminators. Our method, while developed independently, is equivalent in many respects to the light field reconstruction method described by Magda *et al.* [1].

Photometric ranging is a generalization of the well-established technique of *photometric stereo*[2, 3]. In photometric stereo, multiple images of a surface are obtained under different illumination conditions. The intensities obtained for different illumination conditions are used to compute an estimate of the local gradient vector (surface orientation). The gradient vectors are then integrated to recover the surface. It is normally assumed (for shape-from-shading as well as photometric stereo) that the illumination is uniform within the measurement volume, and that variations

in image intensity depend only on local surface albedo and orientation.

Photometric stereo can be contrasted with a class of approaches known as *structured light* [4, 5, 6]. In the structured light approach, a high-contrast pattern of spots or lines is projected onto the object to be measured. A camera views the object from a different direction, from which the projected pattern is distorted according to the surface shape. The intensities of the various illuminated pixels are discarded; only the locations of the illuminated points are used. Thus we see that photometric stereo and structured light are in some sense dual to one another: photometric stereo assumes homogeneous illumination and exploits the dependence of image intensity upon surface orientation, while structured light relies upon inhomogeneous illumination and ignores small variations in image intensity.

Our proposed technique falls somewhere in between these approaches. We create a complex three-dimensional light field (as in the structured light approach), and measure small differences in the gray level values produced by various illuminators (as in photometric stereo). We exploit the inverse-square-law dependence of illumination upon distance from a point source illuminator to obtain a direct estimate of surface range. The intuition behind the idea is as follows: imagine a small surface patch illuminated by a point source at distance r produces an image pixel with value v_1 ; now, imagine moving the illuminator away from the surface along the line joining the patch with the original position of the illuminator by a distance δ . For the patch in question, the angle of incidence of the illumination is unchanged, but the surface irradiance is decreased by a factor of $(\frac{r+\delta}{r})^2$. Therefore the image intensity measured under the new conditions will be: $v_2 = v_1(\frac{r}{r+\delta})^2$. Because the illuminator offset δ is known, we can solve for the original distance r in terms of the two intensities:

$$r = \frac{k\delta}{1-k}, \quad (1)$$

where $k = \sqrt{\frac{v_2}{v_1}}$.

In practice, this idealized case will not be realized. The illuminator displacement will only coincide with the illumination direction for one point on the surface; in general, mo-

tion of the illuminator will produce a change in illumination direction (as in photometric stereo) and a distance-related change in irradiance. Furthermore, rather than moving an illuminator, we wish to construct a stationary rig consisting of many fixed illuminators. In this case, it is impossible to have a pair of illuminators with the same incident direction, because the nearer illuminator would block the more distant. Thus, we must consider the combined effects of changes of both illuminator range and direction, although we will attempt to construct a geometry for which range effects dominate, so that equation 1 above is approximately correct.

Magda *et al.* [1] approached this problem by computing the effects of virtual illuminators, obtained by interpolating values obtained from a densely sampled grid of illuminator positions. In their setup, spatial sampling of illumination directions was accomplished by attaching an illuminator to the end of a robot arm, so the problem of shadows cast by near illuminators was not an issue. In their method, a different virtual illuminators must be computed for each possible range along a pixel’s line-of-sight; here we observe that, in the special case where the virtual illuminators can be made collinear with the camera nodal point, the illuminator axis will coincide with the line-of-sight, and equation (1) can be applied to obtain a direct estimate. The optimization performed by Magda *et al.* uses *all* the illuminators, and so is likely to be more immune to noise effects, but the direct estimate may provide a useful starting point for the optimization.

Many previous approaches to shape-from-shading and photometric stereo assume that the radiance of a surface patch is independent of the viewing direction. Surfaces for which this is true are known as Lambertian, and while there may be no perfectly Lambertian surfaces in nature, it is nevertheless a good approximation for most matte surfaces. The Lambertian assumption is violated for glossy surfaces which produce specular (mirror) reflections. Furthermore, materials with complex sub-surface scattering properties can depart from Lambertian behavior even in the diffuse component of their reflection. Most surfaces are *isotropic*, in the sense that there is no special orientation within the plane of the surface itself. Materials such as brushed metal are *anisotropic*; the reflection from such a surface cannot be determined without knowing the orientation of the surface microstructure relative to the illumination and viewing directions. All of these effects are captured by the bidirectional reflectance distribution function (BRDF). Because of the importance of the BRDF for both vision and graphics, there have been a number of recent studies concerning its measurement [7, 8] and description [9, 10]; a number of methods have also been proposed to recover surface shape for objects with unknown non-Lambertian reflectance properties [1, 11, 12]. While a non-Lambertian BRDF does

complicate the problem of photometric ranging, in principle we can recover not only the object shape but also information about the BRDF. This will be easier when we can assume that the BRDF is uniform over part or all of the surface.

The method of photometric sampling [13] estimates surface shape and the BRDF by combining many images obtained with an array of distant illuminators. Because the illuminator distance is assumed to be large with respect to the dimensions of the imaging volume, the irradiance and illumination direction for each illuminator is assumed to be uniform. The present approach may be seen as a generalization of their work to the near field of the illuminators.

2. Theoretical Framework

We consider the case of a surface \mathbf{S} lit sequentially by N lamps. Each pixel of the camera corresponds to a small patch of surface; let \mathbf{p}_i represent the three dimensional coordinates of the point at the center of the patch imaged by the i th pixel, with the total number of pixels represented by M . (Note that we use boldface symbols such as \mathbf{p}_i to represent vector quantities.) Similarly, let \mathbf{l}_j represent the coordinates of the j th lamp ($1 \leq j \leq N$). Finally, let v_{ij} represent the value of the i th pixel when the surface is illuminated by the j th lamp. The set of N images can be thought of as an $M \times N$ matrix \mathbf{V} with elements v_{ij} . Each row of \mathbf{V} , \mathbf{v}_i , is the set of values measured at the i th pixel, while each column is the image captured with the j th lamp.

Here we derive an expression for v_{ij} in terms of the coordinates of the surface and the lamps. We define r_{ij} to be the distance between the lamp at \mathbf{l}_j and the surface at \mathbf{p}_i , and \mathbf{d}_{ij} to be the unit direction vector from the lamp to the surface. As long as the dimensions of the lamp are small compared to r_{ij} , the irradiance falling on the patch will be proportional to $\frac{1}{r_{ij}^2}$. In general, the irradiance will also be a function of the direction \mathbf{d}_{ij} . We may seek to eliminate this dependence, by placing a wide-angle diffuser in front of the lamp, or we may seek to exploit it by using an anisotropic diffuser. We introduce the function $g(\mathbf{d})$ to represent the radiant flux per unit solid angle in the direction \mathbf{d} . Thus the irradiance falling on the i th patch from the j th lamp will be $\frac{g(\mathbf{d}_{ij})}{r_{ij}^2}$.

To predict the image intensity resulting from this situation, we need to specify a reflectance model. We define \mathbf{o} to be the location of the camera nodal point, and introduce the symbols $s_i = |\mathbf{o} - \mathbf{p}_i|$ to represent the range (distance from the camera) of the i th patch, and \mathbf{c}_i to represent the unit direction vector from the i th patch to the camera nodal point. In general, the image intensity will depend upon \mathbf{d}_{ij} , \mathbf{c}_i , and the surface normal vector \mathbf{n}_i . This dependence is captured by the BRDF. In this analysis we assume that the surface is isotropic, i.e. that it has no special axis,




Figure 1: Diagram showing a typical surface patch and the associated direction vectors; \mathbf{n}_i is the normal vector, \mathbf{d}_{ij} is a unit direction vector pointing back to the j th lamp, and \mathbf{c}_i is a direction vector pointing to the camera. The angle of incidence of the illumination θ is the angle between vectors \mathbf{n}_i and \mathbf{d}_{ij} , and the output scattering angle ϕ is the angle between \mathbf{n}_i and \mathbf{c}_i . The angle ω is the angle in the plane of the surface made by the projections of \mathbf{d}_{ij} and \mathbf{c}_i .

so that the image intensity resulting from a small patch is unchanged if the surface rotates in its plane. In this case, the BRDF can be expressed as a function of 3 angles: the angle of incidence θ formed by direction vectors \mathbf{d}_{ij} and \mathbf{n}_i , the slant ϕ of the surface relative to the camera defined as the angle between \mathbf{c}_i and \mathbf{n}_i , and the angle ω in the plane of the surface between the projections of \mathbf{d}_{ij} and \mathbf{c}_i (see figure 1). Thus,

$$v_{ij} = \frac{g(\mathbf{d}_{ij})}{r_{ij}^2} f(\theta, \phi, \omega). \quad (2)$$

Equation 2 describes the forward problem of image synthesis (and ignores the complications of self-shadowing and interreflections). But we are primarily interested in the inverse problem: how can we obtain the coordinates \mathbf{p}_i of the surface, given the image measurements v_{ij} ? Our goal is to compute the range s_i and surface orientation \mathbf{n}_i of each point on the surface, using only the local measurements \mathbf{v}_i . Let \hat{s}_i represent our estimate of the range of the i th patch; we wish to derive a function $h()$ which computes \hat{s}_i from the local measured intensities \mathbf{v}_i :

$$\hat{s}_i = h(\mathbf{v}_i). \quad (3)$$

Unfortunately, deriving the form of $h()$ analytically by inverting equation (2) has so far proven to be beyond our

mathematical abilities. Magda *et al.* [1] obtain \hat{s}_i using non-linear optimization; their method uses near and far arrays of illuminator positions; for each near illuminator, a virtual far illuminator is computed as a linear combination of the measured far illuminators, providing a pair of collinear illuminators for each value of \hat{s}_i considered during the optimization. Here we present an alternative non-iterative approach, which, while less exact, is considerably faster. We expect that a hybrid approach in which direct estimation is used to initialize an optimization or regularization may eventually prove to be effective.

Our approximate solution is obtained as follows: we tabulate a "training set" of simulated measurements for all possible states of patch i (where a "state" is comprised of a range, orientation, and reflectance). We then use multiple regression to solve for an approximation providing a least-squares estimate of the true values of the parameters. The simplest approximation is simple linear regression; for N illuminators, we solve for $N+1$ predictor coefficients. (Because each pixel corresponds to a different direction and set of positions in space, there will be a different set of coefficients for each pixel.) Obviously, the true relation between pixel values and range is not linear; one way to improve the approximation is to add higher-order terms of increasing degree. When N is large, however, the number of coefficients increases precipitously with polynomial order. An alternative approach to capturing the nonlinearity while keeping the number of coefficients to be estimated low is to compute simple nonlinear functions of the measurements (ratio, square root, etc.) and use these transformed variables in a linear regression. We might also see a benefit by transforming the output variables, such as solving for the z coordinate instead of the range (distance).

3. Simulation of Lambertian case

In this section, we test the method with simulated data for a simplified geometry. We consider the case in which the illuminators are perfectly diffused so that the irradiance does not depend on direction:

$$g(\mathbf{d}_{ij}) = 1. \quad (4)$$

Also, we assume that the surface is Lambertian, and can be described by an albedo a_i

$$f(\theta, \phi, \omega) = a_i \cos(\theta). \quad (5)$$

Under these assumptions,

$$v_{ij} = \frac{a_i \cos(\theta_j)}{r_{ij}^2}. \quad (6)$$

Here we demonstrate direct recovery using a simple two-dimensional simulation. Figures 2 and 3 show the geometry



Figure 2: Range recovery of planar surfaces; a simulated camera looks from left-to-right at a planar surface illuminated by four different lamps. Heavy lines show the reconstructed range. The range of depths of the test surfaces corresponds to the range used in the training set.

of our simulation. We simulate an illumination rig consisting of 4 perfectly diffused sources located at the corners of a rectangle, 9 inches to the left or right of the camera axis, with one pair located in the plane normal to the camera axis which contains the optical nodal point, and a second pair located 9 inches back.

Choosing the range of parameters over which to train is critical for approximation methods such as the one described here; the smaller the range, the fewer nonlinear terms will be required, and a better fit will be required. But generalization performance outside the training range will be degraded. We anticipate that a good balance between speed and accuracy may be obtained by having different sets of coefficients for different regions of the parameter space. Here we chose values roughly corresponding to our face-imaging setup. For range (distance from the camera nodal point), 15 values were used, from 27 to 38 inches. 15 angles for the angle of the normal vector (assumed to lie in the plane containing the illuminators) was uniformly sampled, in 15 steps ranging from -23 to 23 degrees. Finally, 15 values of albedo were sampled, ranging from 0.1 to 1 in linear steps. Thus, the total number of surface patches used for training was $15 \times 15 \times 15 = 3375$.

The training data for a single pixel consists of one simulated measurement for each illuminator and each set of surface parameters. The entire set of measurements (over all pixels) was scaled to have a maximum of 255. The values were then rounded to integers, simulating quantization by a frame-grabber. Training consists of solving for a set of coefficients to be applied to the measurements (and nonlinear functions of the measurements) which come closest to reproducing the original parameter values. This is done independently for each parameter, and each camera pixel. Various nonlinear transformations of the measurements were



Figure 3: Same as figure 2, but for oblique surfaces.

investigated; here we display results obtained using the 4 measurements and all 12 quotients. Including a constant term, a total of 17 parameters were estimated for each pixel. Figures 2 and 3 show typical reconstruction results for planar surfaces.

4. Challenges

The results shown in the previous section demonstrate that photometric ranging is a viable technique for direct range estimation, even using our approximate solution method. In this section we discuss a number of challenging problems that are present in the real world, and suggest approaches that we believe will yield effective solutions.

4.1. Nonlambertian Surfaces

Many surfaces of interest have a non-Lambertian BRDF. For example, human skin can have a strong specular component resulting from a film of oil on its surface. In the previous section, our calibration training set was obtained by simulating a Lambertian sample. Obviously the error will in general increase if these coefficients are used to measure a sample with a significantly different BRDF. Provided we have a sufficient number of measurements, we might be able to estimate both the object shape *and* the BRDF at each pixel. Such an approach depends critically on finding a low-dimensional representation of the space of possible BRDFs; Here we ask two questions: first, can we accurately determine the surface shape in the presence of a non-Lambertian BRDF? Second, can we estimate the form of the BRDF? For a human face, we are likely to have many pixels corresponding to surface patches with identical BRDFs, but in general we must be prepared to deal with spatial inhomogeneity.

The beauty of the method of Magda *et al.* is that by estimating the depth of each point using illuminators (real or virtual) which are collinear with the point in question, dependence of the measured intensity on the BRDF is elim-

inated. Their method relies on the ability to simulate virtual illuminators by linearly interpolating between the measurements obtained from nearby real illuminators. The validity of this approximation is dependent upon the real illuminators being spaced closely enough that the surface reflectance varies linearly between the two directions sampled by the illuminators.

How many parameters must be estimated to arrive at a sufficiently accurate description of a real BRDF? BRDFs and light fields are often represented as functions defined on the hemisphere. This representation allows elegant and efficient computation of reflectance: the output is the spherical convolution of the BRDF with the incident light field [14]. BRDFs are conveniently represented in terms of *spherical harmonics*[15], which are orthogonal basis functions defined on the sphere. Spherical harmonics are parametrized by a frequency, and can be used to compute a frequency spectrum. Just as in normal (flat) Fourier analysis, considerations of Nyquist sampling and aliasing apply. If we can assume that the BRDF is smooth, then we may be able to capture its form with a small number of low-frequency coefficients. Alternatively, a parametric description as a linear combination of diffuse and specular components [13] may reduce the number of parameters still further. Nevertheless, the concept of frequency bandwidth and critical sampling is of fundamental importance in determining good illuminator positions. In general, we will want to space the illuminators closely enough that we do not miss the specular lobe.

It is worth noting that, provided we are able to bootstrap a stable shape estimate, the BRDF can be measured at super-Nyquist frequencies using the techniques of *super-resolution* [16, 17], by collecting a sequence of images for different poses of the object. Super-Nyquist frequencies in the BRDF introduce aliasing into each measurement, but, by registering the objects between frames, high frequency information can be reconstructed. It should be emphasized that this involves more than simple averaging over time, which reduces noise and may so produce apparent sharpening, but does not increase the bandwidth of the estimate. While the Nyquist frequency is determined by the illuminator spacing, the highest frequency obtainable from super-resolution is determined by the angular extent of the illuminators themselves.

4.2. Shadows and Interreflections

We have assumed that the image intensity produced by a given patch of surface depend only on the range and orientation of the patch; this is false when the geometry of the object causes parts of the object to block the illumination to other parts (cast shadows), or when primary illumination from the lamps is augmented by secondary reflections from other parts of the object.

Interreflections between different parts of a concave sur-

face generally cause a surface patch to appear lighter than might be expected from its surface normal, and naive application of a local shape recovery algorithm results in an underestimate of the depth of the concavity [18]. Various methods have been proposed to address this situation [18, 3], which should in principle be applicable to photometric ranging.

Self-shadowing is in a sense the converse: here the amount of light reflected is *less* than predicted by the surface normal, because some of the illumination is blocked by other portions of the object. Here we outline a scheme to detect this situation. We consider a case in which the number of illuminators is large, so that the parameter recovery problem is overdetermined. In this case, the set of image samples \mathbf{v}_i from the i th pixel are constrained to lie on a manifold embedded in the full M dimensional space. What is the dimensionality of this manifold? Under the Lambertian assumption there are four degrees of freedom: one for albedo, another for range, and two more for surface orientation. If the patch lies in the shadow cast by the j th lamp, then v_{ij} will be smaller than it should be, and in most cases \mathbf{v}_i will lie off of the "valid" manifold. A shadow detector could be constructed by amassing valid samples from unshadowed surface patches, and invalid samples from patches for which one of the illuminators was blocked, and then generating a classifier which can separate these two groups. If this can be done reliably, detection of a shadow can trigger the application of a special set of coefficients to a reduced data vector (corresponding to the original data vector less the measurement from the blocked illuminator).

4.3. Target Motion

Our target application of the measurement of face shape poses special challenges, because of the fact that the face cannot be expected to remain completely stationary. We may be able to bootstrap the system by constructing a low-resolution model assuming no motion, and then refining the model as motion estimates are added. But when the number of illuminators is large, the possibility of target motion has important implications for how the illuminators are sampled. Multiplexed illumination [19] is an approach which, while offering significant advantages even for stationary targets, has particular utility when the target may move. In multiplexed illumination, multiple sources are energized in a particular temporal sequence, and the contributions of the individual illuminators are extracted by forming appropriate linear combinations of the stored images. A major advantage of this technique is that it enables accurate measurement of image components from weak or distant illuminators, even when the image from the isolated illuminator is below the camera's threshold. For a moving target, it has the additional benefit that the measurement is derived from *all* the frames in the interval, and this is true for all the illu-

minators. So while movement introduces temporal blurring of the parameter estimates, we do not have to worry about temporal asynchrony of the individual illuminator signals.

Schechner *et al.* [19] propose the use of *Hadamard codes* as the multiplexing signals. We plan to investigate the use of *binary M-sequences* [20] in this role as well. These patterns are often used as a surrogate for Gaussian white noise in nonlinear system analysis, as in the multifocal electroretinogram [21]. The first order Volterra-Wiener kernels correspond to the response to the individual illuminators. Because the physical superposition of light is perfectly linear, we do not expect to observe nonlinear interactions between the signals when everything is perfectly calibrated. Conversely, the higher-order kernels provide a means to calibrate the camera response and interactions between the illuminator control circuits, although perhaps not the simplest or most efficient.

For a given surface patch, some illuminators may be completely irrelevant, while a small subset may be maximally informative. It may prove useful to adaptively select the subset of illuminators based on where we are most uncertain about the surface geometry; this approach was taken by Lindsay and Blake [4], who used uncertainty to determine which areas of a structured light image to process in a real-time system.

4.4. Camera Calibration

Any surface reconstruction method which depends upon the precise values of image pixels requires that the camera(s) be radiometrically calibrated. Just as display monitors have a nonlinear "gamma" mapping from input voltages to output luminances, cameras are designed to have a complementary nonlinearity mapping incident light energy to output pixel values. We assume that this can be described accurately by a point nonlinearity, i.e. the value at a given pixel depends only on the light collected by that pixel, and is independent of the amount of light collected by other near or distant pixels. This condition is sometimes referred to as *spatial independence*.

A number of methods have been described which solve for the camera transfer function by collecting a series of images of the same scene at different exposure levels [22, 23, 24]. These methods are in principle applicable to our situation; the ability to control exposure level via the illuminators allows us to avoid complications arising from lens vignetting when an iris diaphragm is used to vary exposure.

While the sensor's transducer function is fixed, and can be measured as described above, the problem is complicated if the camera electronics include automatic gain control (AGC). The AGC circuit electronically boosts the signal when low light levels cause it to fall. For our application, manual gain control is preferred; it may be possible to use a

camera with automatic gain control, if multiplexed illumination [19] is used, if all of the frames can be made to have approximately the same total illumination, so that the gain is stable.

4.5. Illuminator Calibration

In our simulations above, we assumed isotropic point light sources of equal radiance. In practice, neither assumption can be taken for granted. There will be some variability in the radiant output of the individual light-emitting diodes (LEDs) used in the illuminator. To calibrate the rig, the relative radiance of each illuminator cell must be determined, either by direct measurement using a radiometer, or indirectly from a set of images.

Holographic diffusers are available in a variety of angular patterns, and it is not difficult to achieve good uniformity over a range of angle sufficient to cover a delimited volume. But is this always the best thing to do? A narrow fan of light produced by an anisotropic diffuser could be used to provide information concerning spatial location, instead of pure distance information - this can be thought of as a blurry version of the structured light approach in which stripes of light are projected. In this case, however, calibration will be crucial.

The brute force approach would be to place a known object, such as a ping-pong ball, at a series of known locations in the measurement volume. We could then calibrate the measured values, in much the same way as we determined our reconstruction weights from training data in section 2. One might imagine constructing a board with many balls attached to speed the process. We would like to avoid methods which require precise positioning of a test object, however. Perhaps an illuminator rig containing a mixture of isotropic and anisotropic radiation patterns could "self-calibrate" using a simple target such as a matte planar surface. The location of the plane could be determined using the values associated with the isotropic sources; given the pose of the planar target, the radiation pattern associated with an anisotropic source could be inferred from the measurements.

5. Implementation

We are in the process of constructing a versatile illumination rig with which we will explore the process described above. It will be installed in an existing imaging workstation developed for eye gaze tracking. In this setup, a stereo pair of wide-field cameras image the volume containing the subject's head; a third steerable narrow-field camera is also available. We are using square aluminum tubing to construct the illuminator rig. This tubing is easily machined, and is available with plastic connectors which allow easy reconfiguration. Figure 4 shows the ends of two such tubes,




Figure 4: Components used in our illuminator system. Illumination cells are created in 1 inch aluminum square tubing by drilling holes to accept LED clips, and milling diffuser windows on the opposite side. A small circuit board controls up to 64 led circuits. Cubes of foam are inserted inside the tubing between the cells to block and absorb stray light.

along with the components used to construct the light bar. Each illumination cell consists of 4 near-infrared LED's mounted to the rear surface of the tube; a square window cut into the front surface allows easy mounting of a piece of plastic diffusing film. We are using holographic diffusers (Physical Optics Corporation), which are available in a variety of angles. Small cubes of foam are inserted into the tube between cells to block and absorb stray light.

Our custom LED driver circuit can also be seen in figure 4. The circuit consists of a programmable microcontroller (Microchip PIC-16F876) which communicates with the host computer over an RS232 serial interface. The microcontroller uses the outputs of the sync separator (LM1881) to synchronize changes to the illuminator states with vertical blanking. Control of LED circuits is performed by a four monolithic LED driver chips (Texas Instruments TLC5921). Each chip controls 16 illuminator cells, giving the board a total capacity of 64 channels. The current to each cell is controlled by an on-board potentiometer (one per driver chip), and can range from 1-80 mA. Communication with the host computer is via an RS232 serial interface, which is used both for programming the microcontroller, and sending LED control commands to the firmware in real time. The circuit incorporates a video sync separator, driven by the camera signal, so that the LEDs can be synchronized with the cameras' electronic shutters. Our current firmware only changes LED states after the arrival of a sync pulse, so that each channel is either ON or OFF for the duration



Figure 5: View of our illumination rig from the point of view of a subject seated at the workstation. As of this writing, only two light bars have been populated with components (first and third from the top in this image). Note that the diffusers are not installed, with the exception of the second cell from the left in the top bar. The different ranges of the bars can be inferred by the different degrees to which each bar is lit by the camera flash.

of each frame; we plan to add a feature to the firmware allowing partial-frame exposures, allowing more than 1 bit of intensity control.

6. Summary and Conclusions

We have presented a local method for direct range estimation from a set of images. The method relies upon a three-dimensional array of compact illumination sources, and exploits the near-field dependence of irradiance upon distance. Preliminary results suggest that useful range estimates may be obtained without resort to other cues, while the incorporation of photometric range estimation into other shape estimation methods is bound to yield improvements.

Acknowledgments

This work was supported by the Airspace Operations Systems (AOS) project of NASA's Airspace Systems program, and by NASA's Aviation Security program.

References

- [1] S. Magda, T. Zickler, D. Kriegman, and P. Belhumeur, "Beyond lambert: reconstructing surfaces with arbi-

- trary brdfs,” in *International Conference on Computer Vision*, 2001, pp. 391–398.
- [2] R. Woodham, “Photometric method for determining surface orientation from multiple images,” *Optical Engineering*, vol. 19(1), pp. 139–144, 1980.
- [3] R. Woodham, “Gradient and curvature from the photometric-stereo method, including local confidence estimation,” *Journal of the Optical Society of America*, vol. 11(11), pp. 3050–3068, 1994.
- [4] P. Lindsey and A. Blake, “Real-time tracking of surfaces with structured light,” *Image and Vision Computing*, vol. 13(7), pp. 585–591, 1995.
- [5] E. Horn and N. Kiryati, “Toward optimal structured light patterns,” *Image and Vision Computing*, vol. 17, pp. 87–97, 1999.
- [6] A. Naftel and Z. Mao, “Recovering 3-D face models using a hybrid structured-light stereo matcher,” in *Proceedings 2nd IASTED International Conference on Visualization, Imaging and Image Processing*, September 2002, pp. 590–595.
- [7] W. Matusik and H. Pfister and M. Brand and L. McMillan, “Efficient isotropic BRDF measurement,” Tech. Rep., Mitsubishi Electric Research Laboratories, 2003.
- [8] K. J. Dana and J. Wang, “Device for convenient measurement of spatially varying bidirectional reflection,” *Journal of the Optical Society of America A*, pp. 1–12, 2004.
- [9] S. M. Rusinkiewicz, “A new change of variables for efficient brdf representation,” in *Proceedings Eurographics Workshop on Rendering*, 1998.
- [10] J. J. Koenderink and A. J. Van Doorn, “Phenomenological description of bidirectional surface reflection,” *Journal of the Optical Society of America A*, vol. 15(11), pp. 2903–2912, 1998.
- [11] T. Zickler, P. N. Belhumeur, and D. J. Kriegman, “Helmholtz stereopsis: exploiting reciprocity for surface reconstruction,” in *Proceedings European Conference on Computer Vision*, 2002.
- [12] J. Wang and K. J. Dana, “A novel approach for texture shape recovery,” in *Proceedings International Conference on Computer Vision*, 2003, pp. 1374–1380.
- [13] S. K. Nayar, K. Ikeuchi, and T. Kanade, “Determining the shape and reflectance of hybrid surfaces by photometric sampling,” *Journal of the Optical Society of America A*, vol. 15(11), pp. 2903–2912, 1998.
- [14] R. Green, “Spherical harmonic lighting: the gritty details,” in *Game Developers Conference*, 2003.
- [15] J. Kautz, P. Sloan, and J. Snyder, “Fast, arbitrary BRDF shading for low-frequency lighting using spherical harmonics,” in *Proceedings Thirteenth Eurographics Workshop on Rendering*, P. Debevec and S. Gibson, Ed., 2002.
- [16] N. Nguyen, P. Milanfar, and G. Golub, “A computationally efficient superresolution image reconstruction algorithm,” *IEEE Transactions on Image Processing*, vol. 10(4), pp. 573–583, 2001.
- [17] B. K. Gunturk, Y. Altunbasak, and R. Mersereau, “Gray-scale resolution enhancement,” in *IEEE Workshop on Multimedia Signal Processing*, October 2001, pp. 155–160.
- [18] S. K. Nayar, K. Ikeuchi, and T. Kanade, “Shape from interreflections,” *International Journal of Computer Vision*, vol. 6(3), pp. 173–195, 1991.
- [19] Y. Y. Schechner and S. K. Nayar and P. N. Belhumeur, “A theory of multiplexed illumination,” in *International Conference on Computer Vision*, 2003.
- [20] P. Z. Marmarelis and V. Z. Marmarelis, *Analysis of physiological systems: the white noise approach*, Plenum Press, 1978.
- [21] M. Bearse and E. Sutter, “Imaging localized retinal dysfunction with the multifocal electroretinogram,” *Journal of the Optical Society of America A*, vol. 13(3), pp. 634–640, 1996.
- [22] S. Mann and R. Picard, “Being ‘undigital’ with digital cameras: Extending dynamic range by combining differently exposed pictures,” in *Proceedings of 46th Annual Conference of the Society for Imaging Science and Technology*, May 1995, pp. 422–428.
- [23] P. E. Debevec and J. Malik, “Recovering high dynamic range radiance maps from photographs,” *Computer Graphics*, vol. 31, no. Annual Conference Series, pp. 369–378, 1997.
- [24] T. Mitsunaga and S. K. Nayar, “Radiometric self calibration,” in *Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, 1999.