

Resolving the Generalized Bas-Relief Ambiguity by Entropy Minimization

Neil G. Aldrin

Satya P. Mallick

David J. Kriegman

nalldrin@cs.ucsd.edu spmallick@vision.ucsd.edu kriegman@cs.ucsd.edu

University of California, San Diego

9500 Gilman Dr. Dept. #0404, La Jolla, CA 92093

Abstract

It is well known in the photometric stereo literature that uncalibrated photometric stereo, where light source strength and direction are unknown, can recover the surface geometry of a Lambertian object up to a 3-parameter linear transform known as the generalized bas relief (GBR) ambiguity. Many techniques have been proposed for resolving the GBR ambiguity, typically by exploiting prior knowledge of the light sources, the object geometry, or non-Lambertian effects such as specularities. A less celebrated consequence of the GBR transformation is that the albedo at each surface point is transformed along with the geometry. Thus, it should be possible to resolve the GBR ambiguity by exploiting priors on the albedo distribution. To the best of our knowledge, the only time the albedo distribution has been used to resolve the GBR is in the case of uniform albedo. We propose a new prior on the albedo distribution : that the entropy of the distribution should be low. This prior is justified by the fact that many objects in the real-world are composed of a small finite set of albedo values.

1. Introduction

A fundamental problem in machine vision is to recover the surface geometry of an object given a set of input images taken of the object. In photometric stereo, first proposed by Woodham [14], the illumination is varied in each image while the camera and object remain fixed. When the object is Lambertian and the illumination is a distant point light source, the surface normal and albedo at each position on the object's surface can be recovered from as few as three images if the light source positions are known (assuming non-Lambertian effects like shadows and inter-reflections are negligible). If the light source positions are unknown, then the surface normals and albedo map can be recovered up to a 3×3 linear transformation using singular value decomposition. Moreover, it has been shown that only a 3-parameter subset of these transformations, known as the

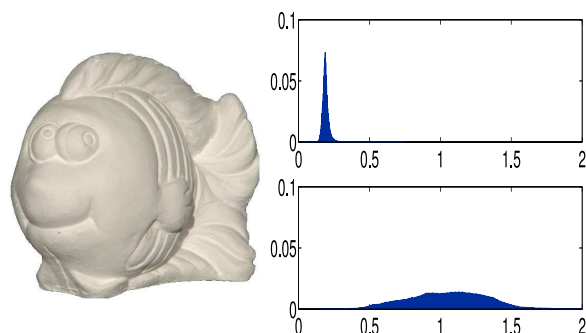


Figure 1. An illustration of the effects of a GBR transformation on the distribution of surface albedos. The top row shows the true albedo distribution which is very peaked while the GBR transformed ($\mu = 3, \nu = 1.5, \lambda = 0.5$) albedo distribution shown in the bottom row is significantly spread out.

Generalized Bas-Relief (GBR) ambiguity, preserve surface integrability [1, 10]. Thus, given three or more images of a Lambertian scene acquired under light sources of unknown direction and strength, the surface can be reconstructed up to a GBR transformation by enforcing surface integrability (using, for example, the algorithm of Yuille and Snow [15]). To resolve the GBR ambiguity, additional constraints must be imposed.

Prior works on resolving the GBR ambiguity involve assumptions about (a) the light sources [15], (b) the distribution of albedos on the surface [1, 8], (c) the surface reflectance [1, 8, 4], and (d) the geometry of the surface [3, 7]. Yuille and Snow [15] assume knowledge of light source intensities, while several algorithms [1, 8] assume constant albedo over the entire surface to resolve the parameters of the GBR. Drbohlav and Sara [4] assume non-Lambertian reflectance, and recover the parameters of the GBR by assuming that the surface normal corresponding to a specular highlight bisects the viewing direction and the light source direction. Georgiades [6] shows that the GBR can be resolved by assuming that the reflectance of the surface is well described by the Torrance-Sparrow reflectance model. Georgiades *et al.* [7] use priors on surface geometry, while

Chandraker *et al.* [3] assume that the surface geometry leads to interreflections which can then be utilized to resolve the GBR.

In this paper, we propose a novel technique for resolving the GBR ambiguity based on minimization of the entropy of the recovered albedos. The hypothesis is based on the observation that in general the GBR transformation smears the distribution of albedos. For example, if the albedo is constant over the entire surface, the distribution of the albedos is a delta function with zero entropy. After a GBR transformation, the observed albedo distribution is a function of the surface geometry and the GBR parameters, and is therefore no longer constant across the surface. This smearing of the albedo values in turn increases the entropy of the albedo distribution (see Figure 1).

This paper makes the following contributions,

- Presents an intuitive conjecture that the GBR can be resolved by minimizing the entropy of the GBR transformed albedos.
- Proves the above conjecture under certain assumptions about the distribution of normals and albedos.
- Empirically shows that the above conjecture has excellent performance on synthetic as well as real images.
- Identifies degenerate configurations of the surface normals and the albedos for which the conjecture fails to resolve the GBR.

2. Background

2.1. Photometric Stereo

The classical photometric stereo algorithm assumes a Lambertian surface illuminated by distant point light sources and viewed by an orthographic camera. Consider acquiring M images, each taken with a different lighting configuration, with N pixels in each image. Ignoring shadows and interreflections, the image intensity at the i th pixel in the j th image is given by,

$$I_{ij} = \rho_i \mathbf{n}_i^\top \mathbf{s}_j \quad (1)$$

where ρ_i and \mathbf{n}_i are the albedo and surface normal at the i th pixel position and \mathbf{s}_j encodes the magnitude and direction of the j th directional light source. To simplify analysis, we find it convenient to express Equation 1 in matrix form. Stacking the pixels in each image row-wise and the images column-wise, we get

$$\mathbf{I} = \mathbf{B}\mathbf{S} \quad (2)$$

where,

- $\mathbf{I} \in \mathbb{R}^{N \times M}$ stores the pixels from all input images. Each column contains the set of pixels from a single image, and each row corresponds to a different lighting condition.
- $\mathbf{B} \in \mathbb{R}^{N \times 3}$ encodes the surface normal/albedo at each pixel. The i th row corresponds to the product of the albedo with the surface normal vector at the i th pixel; i.e., $\mathbf{B}_{i,:} = \rho_i \mathbf{n}_i^\top$.
- $\mathbf{S} \in \mathbb{R}^{3 \times M}$ encodes the light source direction and intensity for each image. The j th column contains the j th light source vector \mathbf{s}_j .

If \mathbf{S} is known, \mathbf{B} can be obtained by solving the over-constrained linear system in Equation 2. On the other hand, if \mathbf{S} is unknown, \mathbf{B} and \mathbf{S} can be recovered only up to a GBR transformation as mentioned in Section 1.

2.2. The Generalized Bas-Relief Ambiguity

Consider a Lambertian surface defined by a height map $z = f(x, y)$ and albedo map $\rho(x, y)$ with surface normal $\mathbf{n} = \frac{(-\frac{\partial z}{\partial x}, -\frac{\partial z}{\partial y}, 1)}{\sqrt{(\frac{\partial z}{\partial x})^2 + (\frac{\partial z}{\partial y})^2 + 1}}$. Then a GBR transformation has the following effect on the albedo and surface normal,

$$\hat{\rho} = \rho \|\mathbf{n}^\top \mathbf{G}^{-1}\| \quad \hat{\mathbf{n}}^\top = \frac{\mathbf{n}^\top \mathbf{G}^{-1}}{\|\mathbf{n}^\top \mathbf{G}^{-1}\|} \quad (3)$$

where $\hat{\rho}$, $\hat{\mathbf{n}}$, and \mathbf{G} are the transformed albedo, transformed surface normal, and GBR transformation matrix respectively. A GBR transformation depends on three parameters μ , ν , and λ and has the following form [1],

$$\mathbf{G} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{pmatrix}; \quad \mathbf{G}^{-1} = \frac{1}{\lambda} \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ -\mu & -\nu & 1 \end{pmatrix}. \quad (4)$$

In terms of Equation 2, a GBR yields transformed \mathbf{B} and \mathbf{S} matrices,

$$\mathbf{I} = \mathbf{B}\mathbf{G}^{-1}\mathbf{G}\mathbf{S} = \hat{\mathbf{B}}\hat{\mathbf{S}}. \quad (5)$$

3. Entropy Minimization

Differential entropy is a natural measure of the "peakiness" of a probability density function f and is defined as

$$H(f) = - \int_S f(x) \log(f(x)) dx \quad (6)$$

where S is the support of f (ie, $f(x \notin S) = 0, \forall x$). Entropy minimization has been previously used in vision algorithms to estimate parameters that result in a peaky distribution of some observed quantity. Finlayson *et al.* [5] use an intuitive argument to justify the use of entropy minimization to estimate the direction of projection in log-chromaticity space

for obtaining an intrinsic image. Similarly, entropy has been used to define a prior in unsupervised clustering when the number of clusters are not known [11].

In the following subsection, in addition to providing intuitive arguments in support of minimizing the entropy for resolution of the GBR parameters, we prove that under certain assumptions about the distribution of the albedos and the surface normals, the minimum entropy solution is locally optimal.

3.1. An Entropy Based Cost Function

A large percentage of man-made (for example, toys) and natural objects (for example, fruits and vegetables) are composed primarily of a small set of dominant albedo values¹. The probability density function (pdf) of the albedos of such objects will be very close to the sum of a set of delta functions. Even the objects that don't consist of a small set of dominant albedos (for example, human skin) will typically have pdf's that are peaked (*i.e.*, it is very uncommon for an object to have a truly uniform set of albedos).

Equation 3 suggests that the distribution of GBR transformed albedos depends on the the distribution of the true albedos as well as the distribution of $\gamma = \|\mathbf{n}^\top \mathbf{G}^{-1}\|$. In the absence of a GBR transformation (*i.e.* $G = I_{3 \times 3}$), the pdf of γ is a delta function centered at one (*i.e.* $\delta(\gamma - 1)$). The GBR has the effect of smearing this delta function, and thereby increasing the entropy of the distribution of the GBR transformed albedos. To motivate this, Figure 2 shows the effects of a few GBR transformations on both the surface geometry and the albedo distribution of a synthetic surface. For this surface, GBR transformations clearly smooth out the distribution of albedos.

3.2. The k-Albedo Configuration

We now consider a special case when the minimum entropy solution is a local optimum. Consider an object consisting of k different albedo values ρ_i where $i = 1 \dots k$, and $\rho_{i+1} > \rho_i \forall i$. Let α_i denote the fraction of all pixels that have albedo ρ_i , and n_i denote the set of normals at these pixels. Note that $\sum_{i=1}^k \alpha_i = 1$. We assume that the corresponding sets of surface normals \mathbf{n}_i are identically distributed on a Gauss sphere for all $i = 1 \dots k$. This implies that after GBR transformation, the density of the term $\gamma_i = \|\mathbf{n}_i^\top \mathbf{G}^{-1}\|$ will be identical to the density of $\gamma_j = \|\mathbf{n}_j^\top \mathbf{G}^{-1}\|$ for all $i, j = 1 \dots k$. The density of the true albedos can be written as

$$f_\rho(x) = \sum_{i=1}^k \alpha_i \delta(x - \rho_i) \quad (7)$$

¹In fact, in a recent work [9], a lower hue count in images was used as a measure of good quality.

and the entropy can be written as

$$H(f_\rho) = - \sum_{i=1}^k \alpha_i \log \alpha_i \quad (8)$$

In the presence of a GBR transformation, the density of $\gamma_i = \|\mathbf{n}_i^\top \mathbf{G}^{-1}\|$ is no longer a delta function. For simplicity of derivation, we assume that f_γ represents the distribution of any γ_i $i = 1 \dots k$ after the GBR transformation (recall that each γ_i is identically distributed). We also assume that f_γ is a small perturbation of the delta function, and has finite support. In other words,

$$f_\gamma(x) = 0, \quad 1 + \Delta < x < 1 - \Delta \quad (9)$$

and,

$$\int_{1-\Delta}^{1+\Delta} f_\gamma(x) dx = 1 \quad (10)$$

The GBR transformed albedo $\hat{\rho}$ can be treated as a product of two independent random variables ρ and γ . Therefore the probability distribution of $\hat{\rho}$ can be written as [13],

$$\begin{aligned} f_{\hat{\rho}}(y) &= \int_0^\infty f_\rho(x) f_\gamma\left(\frac{y}{x}\right) \frac{1}{x} dx \\ &= \int_0^\infty \sum_{i=1}^k \alpha_i \delta(x - \rho_i) f_\gamma\left(\frac{y}{x}\right) \frac{1}{x} dx \\ &= \sum_{i=1}^k \frac{\alpha_i}{\rho_i} f_\gamma\left(\frac{y}{\rho_i}\right) \end{aligned} \quad (11)$$

where, f_γ is the same for all values of i owing to identical distribution of n_i and n_j for all $i, j = 1 \dots k$ over the Gauss sphere.

The entropy of the distribution of the GBR transformed albedos can be written as

$$\begin{aligned} H(f_{\hat{\rho}}) &= - \int_0^\infty \sum_{i=1}^k \frac{\alpha_i}{\rho_i} f_\gamma\left(\frac{y}{\rho_i}\right) \log \left(\sum_{j=1}^k \frac{\alpha_j}{\rho_j} f_\gamma\left(\frac{y}{\rho_j}\right) \right) dy \end{aligned} \quad (12)$$

If $\Delta < \frac{\rho_{i+1} - \rho_i}{2(\rho_{i+1} + \rho_i)}$ holds for $i = 1 \dots k - 1$, then the distributions of transformed albedos corresponding to two different albedos will not overlap². Under this condition, Equa-

²This follows since $(1 - \Delta)\rho_{i+1} > (1 + \Delta)\rho_i$.

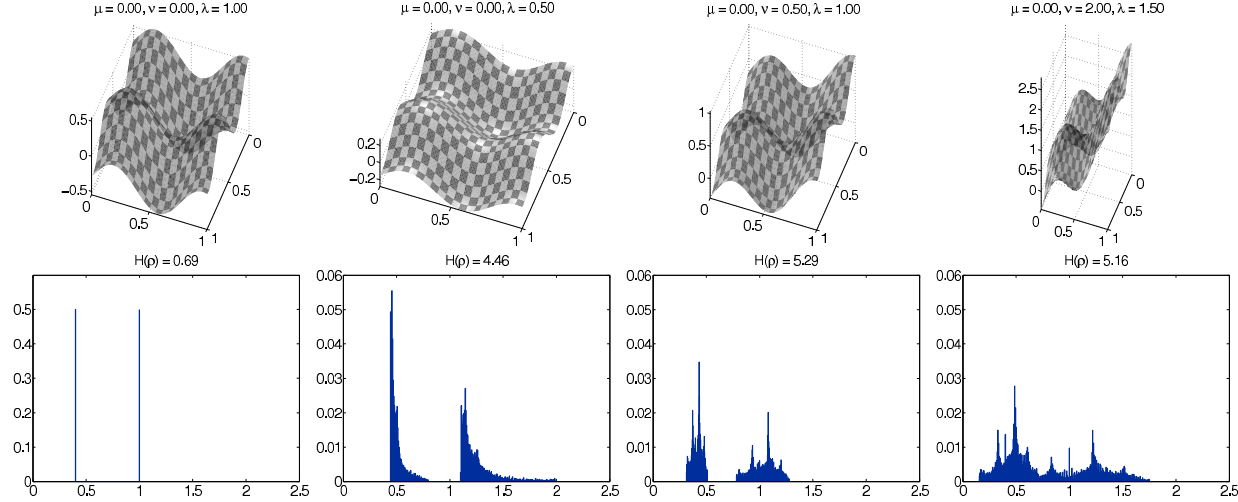


Figure 2. (Top row) A synthetically generated surface, $f(x, y)$, with two albedos transformed by a GBR, $\tilde{f}(x, y) = (x, y, f(x, y))\mathbf{G}^\top$ and $\tilde{\rho} = \|\tilde{\mathbf{n}}^\top \mathbf{G}^{-1}\|$. (Bottom row) The histogram of GBR transformed albedo values.

tion 12 can be re-written as,

$$\begin{aligned}
& H(f_{\tilde{\rho}}) \\
&= - \sum_{i=1}^k \frac{\alpha_i}{\rho_i} \int_{(1-\Delta)\rho_i}^{(1+\Delta)\rho_i} f_\gamma\left(\frac{y}{\rho_i}\right) \log \frac{\alpha_i}{\rho_i} f_\gamma\left(\frac{y}{\rho_i}\right) dy \\
&= - \sum_{i=1}^k \alpha_i \int_{(1-\Delta)\rho_i}^{(1+\Delta)\rho_i} f_\gamma\left(\frac{y}{\rho_i}\right) \log f_\gamma\left(\frac{y}{\rho_i}\right) d\frac{y}{\rho_i} \\
&\quad - \sum_{i=1}^k \alpha_i \log \alpha_i + \sum_{i=1}^k \alpha_i \log \rho_i \\
&= H(f_\gamma) + H(f_\rho) + \sum_{i=1}^k \alpha_i \log \rho_i \tag{13}
\end{aligned}$$

Therefore the entropy of the distribution of the transformed albedos is greater than the entropy of the distribution of the correct albedos when $H(f_\gamma) + \sum_{i=1}^k \alpha_i \log \rho_i > 0$.

3.3. Degenerate Configurations

In the previous section we showed that minimum entropy solution corresponds to the correct solution under certain assumptions. Empirical evidence suggests that the minimum entropy hypothesis works for a wide variety of configurations. However, we have theoretically identified degenerate cases in which the distribution of surface normals and the surface albedo conspire such that the entropy of the distribution of the GBR transformed albedos is lower than or equal to the entropy of true albedo distribution.

Consider the image of a polyhedron with k visible faces, and let \mathbf{n}_i be the normal associated with the i^{th} face. The transformed albedo of the i^{th} face is

$$\hat{\rho}_i = \rho_i \|\mathbf{n}_i^\top \mathbf{G}^{-1}\| \tag{14}$$

where, ρ_i is the albedo associated with the i^{th} face, and \mathbf{G} is the GBR transform. Given a GBR \mathbf{G} and a collection of normals \mathbf{n}_i $i = 1 \dots k$, the i^{th} face of the polyhedron can be painted with albedo $\rho_i = 1/\|\mathbf{n}_i^\top \mathbf{G}^{-1}\|$ so that all GBR transformed albedos have value equal to unity, and their distribution is a delta function and the entropy is minimum. This is clearly a degenerate configuration.

Consider a second case in which the surface is a plane, and it contains k distinct albedos. It is easy to see that under this condition, the distribution of $\hat{\rho}$ is a sum of k delta functions, regardless of what GBR transform is chosen. Therefore, this configuration is also degenerate, and entropy does not provide information about the GBR.

4. Experimental Validation

To empirically validate that the minimum entropy solution indeed resolves the GBR, we formulate an optimization algorithm that solves for the parameters of the GBR (μ , ν , and λ) given an uncalibrated set of surface normals and albedos. Our experiments consist of the following high-level steps,

- Photograph or render multiple images of a Lambertian object under different lighting conditions.
- Recover the surface and light sources $\mathbf{I} = \hat{\mathbf{B}}\hat{\mathbf{S}}$ up to a GBR transformation using the algorithm of Yuille and Snow [15].
- Find the GBR parameters that minimize the entropy of the albedo distribution and apply the GBR to recover the true surface $\rho_i \mathbf{n}_i^\top = \mathbf{B}_i = \hat{\mathbf{B}}_i \mathbf{G}$.
- Validate the solution by comparing to calibrated photometric stereo results.

The following subsections describe our methodology in more detail.

4.1. Approximating the Differential Entropy

In Section 3 we assumed knowledge of the underlying probability density function of $\hat{\rho}$. In practice, we have measurements of this quantity at each pixel location and seek to approximate the underlying distribution using these samples. While there are many powerful methods for approximating continuous distributions from samples (e.g., Parzen windows, the mean-shift algorithm, kernel based estimators, etc.), for the purpose of computing entropy we find it sufficient to use a histogram approximation.

Histograms are beneficial because (1) they are simple to generate from sampled data and (2) the entropy of a histogram can be computed as a discrete summation instead of an integral. These properties are particularly important in our case since we need to evaluate the entropy of the albedo distribution under many different GBR transformations. Consider an m -bin histogram $\{a_i\}$, $i = 1 \dots m$ generated from n i.i.d. samples with underlying distribution f . Then the simplest estimator of the entropy $H(f)$ is,

$$\hat{H}_{MLE}(f) = - \sum_{i=1}^m \frac{a_i}{n} \log \frac{a_i}{n} \quad (15)$$

which is the maximum likelihood estimator of $H(f)$ given its histogram. While $\hat{H}_{MLE}(f)$ can exhibit significant bias³, its variance is typically low and thus should be well suited in a minimization framework.

A downside to using histograms is that the number of bins and the support region must be chosen appropriately. For the purpose of comparing albedo distributions we fix the number of bins to some constant (256 bins were used for most of the experimental results shown in Section 4.4). Moreover, since we don't want to overly favor distributions with narrow support, we set the support to the range of the albedo samples.

4.2. Optimization

We now turn to the problem of finding the GBR parameters μ , ν , and λ that minimize the entropy of the albedo distribution. Let $\{\hat{\mathbf{b}}_i = \hat{\rho}_i \hat{\mathbf{n}}_i\}$, $i = 1 \dots N$ be a set of GBR transformed surface normals scaled by the GBR transformed albedo. Then, under a *proposed* set of GBR parameters encoded in matrix $\hat{\mathbf{G}}$, the obtained set of albedos is $\{\hat{\rho}_i = \|\hat{\mathbf{b}}_i \hat{\mathbf{G}}\|\}$ from which the entropy $H(f_{\hat{\rho}}) \simeq \hat{H}_{MLE}(\text{hist}(\hat{\rho}))$ can be estimated as described in Section 4.1. The first thing to note is that the entropy estimate is neither convex, nor differentiable with respect to the GBR parameters. Given these facts, it seems finding the minima

will be very difficult. Luckily, the number of parameters is low and the error surface is empirically smooth. Moreover, by making weak assumptions on the surface geometry or lighting configuration we can bound the parameters. Empirically, we find that the GBR parameters induced by the uncalibrated photometric stereo algorithm very rarely exceed an absolute value of 5 – even with varied light source intensity. Also, because of the concave/convex ambiguity, we can restrict λ to be positive. Based on these observations, we restrict our search space throughout our experiments to,

$$\begin{aligned} -5 &\leq \mu \leq 5 \\ -5 &\leq \nu \leq 5 \\ 0 &\leq \lambda \leq 5. \end{aligned}$$

Given these bounds on the parameters, along with our observation of smoothness in the error surface, we solve for the GBR parameters using discrete search with coarse-to-fine refinement. The algorithm works by iteratively sampling the error function at uniform positions within the parameter space. At each iteration, the bounds are tightened and centered around the best solution found in the previous iteration. This process continues until the sampling interval falls below some threshold. While other strategies such as simulated annealing could also have been used, we found that our method worked well enough to find an optimum in most cases⁴.

4.3. Image Acquisition and Pre-Processing

Prior to using any input images, we perform a number of pre-processing steps. First of all, whenever possible we remove ambient lighting by subtracting each image from an ambient image. Next, we generate a mask image to isolate the object from the background. This is either done manually or automatically using thresholding. Once this is done, we randomly scale each image by a scalar between 0.5 and 1.5. This has the effect of randomizing the light source magnitude, which is important because many datasets are collected using uniform magnitude which can resolve the GBR. By randomizing the intensities we can be certain that our algorithm is not taking advantage of uniform light source magnitude.

After this step, we compute a visibility matrix, \mathbf{V} , which is used to reduce the effect of outlier pixels in the input images. \mathbf{V} is the same size as the image matrix \mathbf{I} and takes value 1 if the corresponding pixel in \mathbf{I} is an inlier or value 0 if the corresponding pixel is an outlier. We compute the visibility matrix by excluding pixels that don't lie close to the linear subspace predicted by the Lambertian model (this can be done using SVD).

³See [12] for a bias corrected estimator.

⁴Depending on the “coarseness” of the search grid, as well as the number of samples, the optimization takes anywhere from a few minutes to a few hours.

Next we perform uncalibrated photometric stereo on the set of input images, discounting pixels masked by V^5 , which yields a set of GBR transformed surface normals and albedos. Prior to searching for the GBR parameters as described in Section 4.2, we aggressively filter any surface points that could be outliers. This includes edge points (in any image) as well as points that do not lie close to the predicted Lambertian subspace. Finally, we search over the GBR parameter space to find GBR that minimizes the entropy of the albedo distribution.

4.4. Results

We show results on three datasets. Figure 4 shows the reconstruction of a synthetically generated Stanford bunny. There are a total of 6 input images in this dataset, each rendered in Povray with a different light source position. As is clearly seen in the figure, the output of the uncalibrated photometric stereo algorithm results in an extremely flat distribution of albedos. More importantly, our optimization routine clearly finds the correct GBR parameters.

Our second dataset consists of 15 images from the YaleB face database [7]. Looking at figure 5, two things are noteworthy: first, our minimum entropy solution appears qualitatively better than the reconstruction obtained from calibrated photometric stereo. For example, the calibrated photometric stereo result is clearly skewed in the x direction, while the minimum entropy solution is not. We conjecture that slight errors in the light source vectors are biasing the reconstructed surface in this case. There are also many shadow and non-Lambertian effects in these images that may be affecting the calibrated case. These results empirically show that the albedo entropy can resolve the GBR for images consisting of k albedos as well as more general albedo distributions such as human faces.

Finally, Figure 6 shows the output of our algorithm on a real surface with two albedos.

5. Conclusion

This paper presents an algorithm to resolve the Generalized Bas-Relief ambiguity based on a novel prior on the observed albedos. In addition to providing an intuitive reasoning in favor of entropy minimization for estimation of GBR parameters, we show that entropy minimization provides provably correct results under certain assumptions about the surface normal and albedo distribution. Finally, we validate our results on real and synthetic data.

Acknowledgements

This work was partially funded under NSF grants IIS-0308185 and EIA-0303622.

⁵To handle missing values, we utilize the SVD algorithm derived by Brand [2].




Figure 3. Input images. (Top row) Three of the 6 synthetic input images used to generate the results in Figure 4. (Middle row) Three of the 15 images (obtained from the Yale face database) used to generate the results in Figure 5. (Bottom row) Three of the 5 input images used to generate the results in Figure 6.

References

- [1] P. Belhumeur, D. Kriegman, and A. Yuille. The Bas-Relief Ambiguity. *IJCV*, 35(1):33–44, 1999.
- [2] M. Brand. Incremental singular value decomposition of uncertain data with missing values. *ECCV '02*, pages 707–720, 2002.
- [3] M. Chandraker, F. Kahl, and D. Kriegman. Reflections on the Generalized Bas-Relief Ambiguity. *CVPR '05*, 1, 2005.
- [4] O. Drbohlav and R. Sara. Specularities reduce ambiguity of uncalibrated photometric stereo. *ECCV '02*, pages 46–60.
- [5] G. Finlayson, M. Drew, and C. Lu. Intrinsic images by entropy minimization. In *ECCV '04*, 2004.
- [6] A. Georghiades. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. *ICCV '03*, pages 816–823, 2003.
- [7] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *PAMI*, 23(6):643–660, 2001.
- [8] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *J. Optical Soc. Am. A*, 11(11):3079–3089, 1994.
- [9] Y. Ke, X. Tang, and F. Jing. The design of high-level features for photo quality assessment. In *CVPR '06*, pages 419–426, 2006.

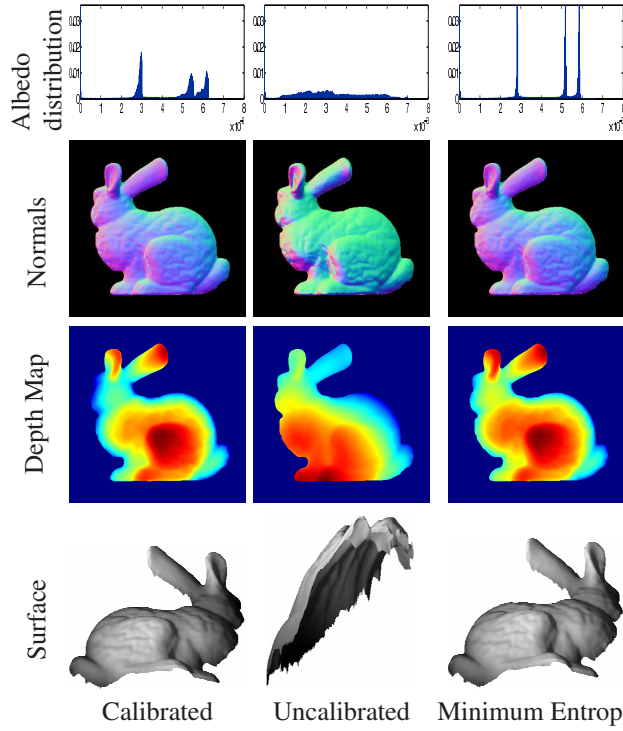


Figure 4. Results from synthetic images. (Column 1) Calibrated photometric stereo. (Column 2) Uncalibrated photometric stereo. (Column 3) Minimum entropy solution. (Row 1) Distribution of albedos. (Row 2) Color encoded surfaces normals. (Row 3) Color encoded depth map. (Row 4) Reconstructed surfaces from the same viewpoint.

- [10] D. Kriegman and P. Belhumeur. What shadows reveal about object structure. In *ECCV '98*, volume 2, page 399, 1998.
- [11] G. Palubinskas, X. Descombes, and F. Kruggel. An unsupervised clustering method using the entropy minimization. In *ICPR*, volume 2, pages 1816–1818, 1998.
- [12] L. Paninski. Estimation of entropy and mutual information. *Neural Comput.*, 15(6):1191–1253, 2003.
- [13] V. K. Rohatgi. "An Introduction to Probability Theory and Mathematical Statistics". John Wiley & Sons Inc., New York, 1976.
- [14] R. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, January 1980.
- [15] A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. *CVPR '97*, pages 158–164, 1997.

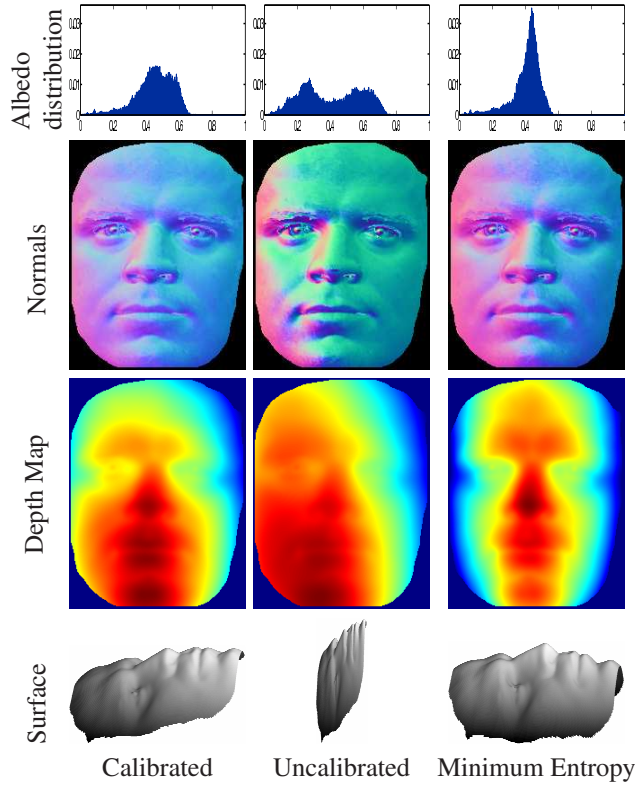


Figure 5. Results from a human face (from the Yale face database).

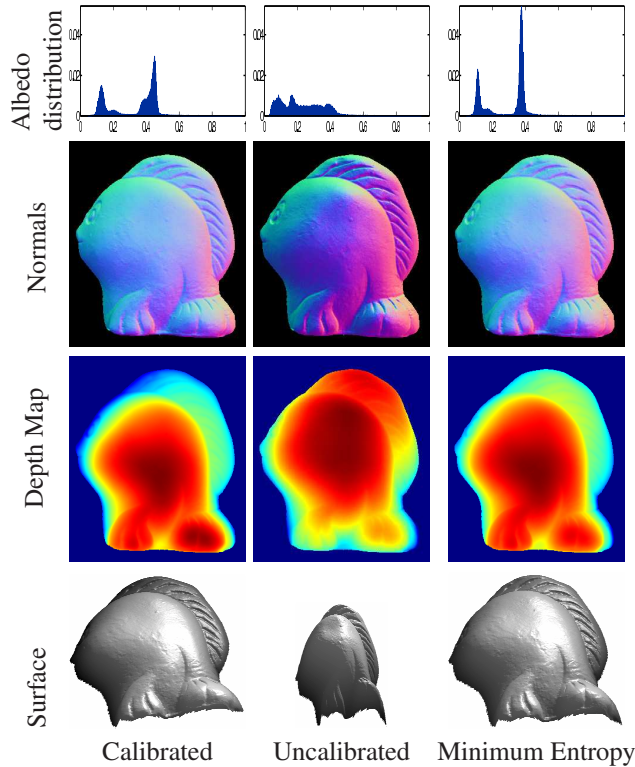


Figure 6. Results from (non-synthetic) images of a fish.