ShadowCuts: Photometric Stereo with Shadows

Manmohan Chandraker[1] mkchandraker@cs.ucsd.edu
Sameer Agarwal[2] sagarwal@cs.washington.edu
David Kriegman[1] kriegman@cs.ucsd.edu

[1]Computer Science and Engineering
University of California, San Diego

[2]Computer Science and Engineering
University of Washington, Seattle

Abstract

We present an algorithm for performing Lambertian photometric stereo in the presence of shadows. The algorithm has three novel features. First, a fast graph cuts based method is used to estimate per pixel light source visibility. Second, it allows images to be acquired with multiple illuminants, and there can be fewer images than light sources. This leads to better surface coverage and improves the reconstruction accuracy by enhancing the signal to noise ratio and the condition number of the light source matrix. The ability to use fewer images than light sources means that the imaging effort grows sublinearly with the number of light sources. Finally, the recovered shadow maps are combined with shading information to perform constrained surface normal integration. This reduces the low frequency bias inherent to the normal integration process and ensures that the recovered surface is consistent with the shadowing configuration

The algorithm works with as few as four light sources and four images. We report results for light source visibility detection and high quality surface reconstructions for synthetic and real datasets.

1. Introduction

Shadowing is nearly unavoidable in any image. In fact under orthographic projection, the only light source direction illuminating a scene which can result in an entirely shadowless image is the one aligned with the viewing direction. Despite their ubiquitous presence, many computer vision algorithms either ignore or explicitly exclude data with shadowing.

In the context of surface reconstruction from photometric measurements, shadows are a double edged sword. While attached shadows can be understood from the local surface geometry relative to the light source, cast shadows are non-local and their causes are more difficult to identify. When not accounted for appropriately, they can lead to grossly distorted reconstructions. Yet, it is precisely the global nature of shadowing that can be an asset for a surface reconstruction algorithm since shadows constrain the shape of portions of the surface not in physical proximity.

Thus, the challenges are to first identify which pixels in an image correspond to surface patches that are in shadow with respect to each light source, and then to exploit the implied shadowing constraints to the fullest extent. In this paper, we consider shadowing within the context of photometric stereo and address both these issues.

A problem analogous to shadow detection is determining light source visibility at every pixel. In fact, the two problems are equivalent for a scene illuminated by a single point source. Traditional photometric stereo assumes that each image is acquired under a single point light [21]. Since at least three sources are needed to recover a surface normal, it is common to use additional light sources to ensure that most surface points are covered by at least three sources. However, this also leads to an increase in the number of acquired images. One way to reduce acquisition cost is to acquire images with multiple light sources activated simultaneously.

Using multiple light sources per image improves the reconstruction accuracy in more than one way. The signal strength will be enhanced at portions of the surface that are actually exposed to multiple sources, which has positive implications for boosting the signal to noise ratio (SNR) [19]. Moreover, light sources can be placed further from the viewing direction (while providing sufficient coverage of the surface), which leads to better conditioning of the light source matrix, and consequently better surface normal estimates.

A common heuristic for dealing with shadows is to simply compare an image pixel's value to a global threshold; if it falls below the threshold, the corresponding source is not used for estimating the surface normal at that pixel. However, this heuristic is clearly ill-suited even for single source images. For example, attached shadows are characterized by a gradual fall-off in intensity which may be indistinguishable from intensity change due to varying albedo.

Determining light source visibility when multiple sources illuminate each image is even more complicated. Is a low im-

1-4244-1180-7/07/$25.00 ©2007 IEEE

age intensity in a multi-source image due to only one source being visible? Or is it due to two sources placed obliquely with respect to the surface normal? Or is it simply an artifact of low surface albedo? These questions can be answered with the realization that, for most surfaces, visibility with respect to a light source is a piece-wise constant function. We exploit this fact to construct a fast and highly accurate shadow labeling algorithm based on graph cuts.

Once the shadowing and light source visibility issues have been resolved, it is trivial to compute the surface normal at every pixel. However, we can do better. Since we actually recover the shadow boundaries, we have the opportunity to ensure that our surface normal integration procedure produces a height field consistent with the observed shadowing behavior. Spatial constraints such as those lent by shadows can also overcome the low-frequency bias often observed in integrating a normal field. See [17] for instance. Our approach to surface normal integration with shadows can be succinctly considered as a quadratic program that can be solved using powerful interior point method based QP solvers [2].

In summary, our main contributions are the following:

- A highly accurate shadow labeling algorithm based on fast graph cuts.
- A photometric stereo algorithm that uses multiple light sources to enhance surface coverage and reconstruction accuracy, while requiring fewer images than sources.
- A normal integration algorithm that enforces the visibility constrained imposed by the estimated shadowmaps.

We begin with an overview of prior work in Section 2, outline our shadowing and visibility determination algorithm in Section 3 and discuss the problem of constrained normal integration in Section 4. We present experimental results on real and synthetic data in Section 5 and conclude with further discussions in Section 6.

2. Related Work

Ever since its introduction in the field of computer vision [21], calibrated photometric stereo has been a problem of significant interest because it can yield very high quality surface normal estimates for surfaces that approximate Lambertian behavior or where the reflectance maps are known. We point the reader to [14] for an overview of traditional photometric stereo. Photometric stereo under complex illumination has been demonstrated in [4]. The assumption of Lambertian BRDF in photometric stereo has been relaxed in a few cases, notably example-based photometric stereo [11].

Deriving shape information from only shadows also has a long history including a series of works by Kender and his colleagues starting with [15] as well as Daum and Dudek [9]. Shadow carving [18] is a visual hull carving approach to surface reconstruction from shadows, however, it works only with cast shadows and uses an *ad hoc* labeling of shadow boundaries. A theoretical analysis of shape recovery from shadows is discussed in [16] where it is shown that a surface can be reconstructed from shadowing due to a finite set of sources up to a four-parameter family of projective transformations, but no implementation was presented due to the difficulty of accurately detecting shadows. As expected in any algorithm that relies purely on shadows for reconstruction, a prohibitively large number of images are needed to yield reasonable results. In contrast, by combining shadow and shading constraints, we can produce comparable results with very few images.

Shadow detection or removal can be performed with a single image by incorporating additional assumptions on the lighting and acquisition set-up. For instance, [13] discuss a shadow removal algorithm based on assumptions such as Planckian lighting and narrow-band cameras. By instead using user-supplied hints in a Bayesian framework, smooth shadows can be extracted in [22].

Multiple light sources have been used in the past to increase surface coverage without any attempt to recover shadow boundaries, such as [8]. The advantages of multiplexed illumination are elucidated in [19], however, demultiplexing requires the assumption that as many images are acquired as light sources used (as does [8] above). A multiple source photometric stereo algorithm that uses color is presented in [3], however their method reduces to [8] for grayscale images of Lambertian scenes.

At the other end of the spectrum from single-source photometric stereo are frequency domain methods under complex illumination [4]. However, the theory of these methods applies only to convex smooth surfaces, and high frequency behavior such as cast shadows are a significant limitation. A dense photometric stereo method in [23] relies on around 50 light sources to cover the surface and eliminate shadow effects for recovering a discretized normal field through an MRF-like formulation.

The basic problem of surface normal integration can be considered equivalent to solving a Poisson equation [20]. Other approaches project the non-integrable gradient field onto a set of integrable slopes using some suitable choice of basis functions, such as Fourier [10]. By only imposing integrability, it can be shown that a surface can be recovered even in uncalibrated photometric stereo up to a three parameter generalized bas-relief ambiguity [5, 26]. A generalization of several of these classes of methods for integrating gradient fields is discussed in [1]. For integrating gradient fields arising in large images, a fast marching method can be used as an initialization to a more expensive conjugate gradient descent [12].

All the above methods of surface normal integration are

unconstrained. The problem of integrating surface normals with shadow constraints is discussed in [25]. However, the problem practically solved therein is a soft version, where the constraints are added in the objective function as a regularizer and may not be exactly satisfied. A heuristic attempt at shadow detection followed by constrained normal integration is presented in [24].

3. Detecting shadows and light source visibility

Since it forms the central workhorse of our algorithm, we will begin this section with a very short introduction to graph cuts as a method for combinatorial optimization. In Section 3.2, we will describe our approach to shadow labeling in $m \geq 4$ images, where each image is illuminated by exactly one point light source. In practice, a trivial selection procedure such as choosing the three brightest light sources can also yield good surface normal estimates. But such a selection procedure, of course, cannot assert with any certainty whether a particular pixel is shadowed or not. Neither can a global thresholding heuristic determine that information, for instance in the presence of attached shadows with texture variation, as discussed in Section 1. Our aim here is to use the single source setup to illustrate the design of our algorithm for determining light source visibility.

In section 3.3, we extend this algorithm to the much more difficult problem of determining light source visibility in multi-source photometric stereo with fewer number of images than light sources. Unlike the single source setup, there exist no reasonable heuristics for solving this problem.

3.1. Graph cuts review

Let $\mathcal{P}$ be a set of pixels and $\mathcal{L}$ be a set of discrete labels, then consider the problem of finding a labeling $f : \mathcal{P} \rightarrow \mathcal{L}$ such that a given energy function, $E(f)$ of the form

$$E(f) = \sum_{p \in \mathcal{P}} D_p(f_p) + \sum_{(p,q) \in \mathcal{N}} V_{p,q}(f_p, f_q), \qquad (1)$$

is minimized.

The first term in the above expression is called the data term which measures the disagreement between a given labeling and the observed data. The second term, called the smoothness term, imposes a penalty on variation of labeling within a neighborhood. Problems of this form occur frequently in physics, computer science and mathematics. One can show that finding the minimum energy solution to this problem can be interpreted as finding the maximum *a posteriori* labeling for a class of Markov Random Fields. In general, finding the global minimum of problems of this form is an NP-complete problem and one must resort to approximation algorithms.

The fast graph cuts based algorithm of Boykov et al. [7] is one such polynomial time approximation scheme. The algorithm operates by solving a max flow/min cut problems on a sequence of weighted graphs in which the set of pixels $\mathcal{P}$ is the vertex set the set $\mathcal{N}$ is the edgeset. The quality of approximation achieved depends on the properties of the function $V_{p,q}$.

3.2. Shadow Labeling

Let us begin by defining some notation. Let $\mathcal{P}$ be the set of pixels, $\mathbf{n}_i$ represent the surface normal at the i-th pixel and $\mathbf{s}_j$ stand for the j-th light source direction. If the pixel i is not shadowed from the light source j, the color of the pixel in the j^{th} image is given by

$$c_{ij} = \mathbf{n}_i^\top \mathbf{s}_j, \;\; i = 1, \cdots, N \;\; j = 1, \cdots, m \qquad (2)$$

Given 3 or more light sources it is then simple to invert the above equation as a least squares system to determine the normal vector on a per pixel basis. However, it is not always the case that every light source is visible to every pixel. In the hard shadowing case, we can write the image formation equation as

$$c_{ij} = \begin{cases} \mathbf{n}_i^\top \mathbf{s}_j & \text{if source } \mathbf{s}_j \text{ is visible to pixel } i \\ 0 & \text{otherwise} \end{cases} \qquad (3)$$

For each pixel i, define a $\{0, 1\}$-vector $\mathbf{w}_i$ such that $\mathbf{w}_{ij}$ is 1 if light source j is visible to pixel i, 0 otherwise. Then the image formation equation can be written as

$$c_{ij} = \mathbf{w}_{ij} \left(\mathbf{n}_i^\top \mathbf{s}_j \right) \qquad (4)$$

Since image measurements for shadowed pixels are not exactly zero, determining which pixels are shadowed is a hard combinatorial problem with no known solution.

If we consider visibility of a light source as a function on the object surface, while discontinuous, it is not completely non-smooth. We can exploit this fact to solve the problem of determining when certain light sources are shadowed or not (that is, recovering the vector $\mathbf{w}_i$) on a per pixel basis. Since the number of possible vectors $\mathbf{w}_i$ is finite ($O(2^m)$), we can consider this to be a pixel labeling problem, where each lighting configuration corresponds to a label. The smoothness prior is then encoded in terms of a penalty that demands that nearby pixels have similar lighting configurations. Thus, with $\mathbf{W}$ being the matrix with row vectors $\mathbf{w}_i$, we can formulate a discrete optimization problem:

$$\min_{\mathbf{W}} \sum_{r \in \mathcal{P}} \left[D_r(\mathbf{W}) + \lambda \sum_{(p,q) \in \mathcal{N}_r} D_{pq}(\mathbf{W}) \right] \qquad (5)$$

where $\mathcal{N}_r$ is some appropriate notion of a neighborhood.

Let $\mathbf{W}_r$ be the diagonal matrix formed using the entries of $\mathbf{w}_r$, the lighting indicator vector for the r-th pixel. Then,

the extent to which the image formation equation is satisfied at the r-th pixel is given by

$$D_r(\mathbf{W}) = \|\mathbf{W}_r(\mathbf{L}^\top \mathbf{n}_r) - \mathbf{c}_r\| \tag{6}$$

where $\mathbf{L}$ is a $3 \times m$ matrix whose columns are the light sources. $\mathbf{c}_r$ is a column vector with length equal to number of light sources. The cost can be divided by $\mathrm{Trace}(\mathbf{W}_\mathrm{r})$ to ensure that more light sources being visible does not amount to higher cost, but we did not find it necessary in practice.

The smoothness cost D_{pq} can be taken to be the Hamming distance between the vectors $\mathbf{w}_p$ and $\mathbf{w}_q$:

$$D_{pq}(\mathbf{w}_p, \mathbf{w}_q) = \|\mathbf{w}_p - \mathbf{w}_q\|_1. \tag{7}$$

While the Hamming distance is not a discontinuity preserving metric on the space of labels, we have found it to suffice for our application. The above functional is now an instance of the metric labelling problem as the Hamming distance is a metric on the set of configurations and can be solved using the algorithms proposed in [7]. Once the best lighting configuration has been determined, the normals are trivially determined.

Note that, in our algorithm, we do not need to simultaneously optimize over continuous normals and discrete shadow labels. Within each graph cut iteration, the putative shadow labeling assigns source visibility at every pixel, which determines the normal through the photometric stereo equation. If the shadow labeling is incorrect, residual error in (6) will be high. It is this error that graph cuts minimize by swapping shadow labels. Thus, normals do not form a direct part of our MRF formulation.

3.3. Multi-source Photometric Stereo

We will now extend the graph cut formulation of the previous section to recover visibility information and surface normals in the case of multi-source photometric stereo. Suppose $k \geq 4$ images are acquired using $m > k$ light sources, with some combinations of multiple sources illuminating each image. For concreteness, we assume that the pattern chosen was such that the light sources $j, \cdots, j+m-k$ are used to acquire the j-th image. The arguments that follow can be easily extended to include any pattern in which the light sources are combined.

It is not known *a priori* whether a pixel is occluded or visible with respect to any particular light source, however, an inherent assumption is that at least three sources are visible to every pixel (else, the normal becomes undefined). Depending upon the number of visible sources, there are $\sum_{l=3}^{m} \binom{m}{l}$ possible lighting configurations for every pixel. The label at pixel r can be represented by a $\{0,1\}$-vector, $\mathbf{w}_r$, of length m such that $\mathbf{w}_{rj}$ is 1 if the j-th source is visible to pixel r and 0 otherwise.

Let $\mathbf{W}_r$ be a $k \times m$ band-diagonal matrix, defined for every pixel r, such that

$$\mathbf{W}_r(l, l+j-1) = \begin{cases} 1 & \text{if source } j \text{ visible to pixel } r \text{ in image } l \\ 0 & \text{otherwise} \end{cases} \tag{8}$$

Now the data term corresponding to the energy function in (5) becomes

$$D_r(\mathbf{W}) = \|\mathbf{W}_r(\mathbf{L}^\top \mathbf{n}_r) - \mathbf{c}_r\| \tag{9}$$

while the smoothness cost is again the Hamming distance between the labels.

Note that the minimum number of images required by our algorithm is four, since with three images, whatever be the choice of labeling, a normal vector is uniquely determined. So, the residual which forms the data term will be identically zero for any random labeling and the graph cuts algorithm will not be guided by any image data.

An inherent ambiguity in our formulation is when the surface normal lies in the plane that bisects the angle between two light sources. In that case, the data term is not informative. Unless there happens to be a dominant plane in the scene which is fortuitously aligned with such a direction, the smoothness information from neighborhood pixels suffices to render this ambiguity difficult to observe in practice.

4. Combining shading and shadowing constraints

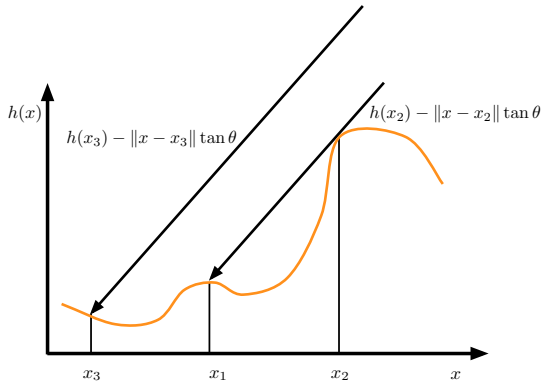

Figure 1. Shadow and anti shadow constraints

In this section we consider the problem of reconstructing a surface from an estimate of its normals, consistent with the estimated shadowmaps.

Let $H(x, y)$ be the surface we wish to reconstruct and let $P(x, y)$ and $Q(x, y)$ be estimates of $\partial_x H$ and $\partial_y H$, respectively. In the absence of shadows, one considers the

variational minimization problem:

$$\min_H \int_\Omega (H_x - P)^2 + (H_y - Q)^2 \, dxdy, \qquad \left.\frac{\partial H}{\partial n}\right|_{\partial\Omega} = 0 \tag{10}$$

Here, $\partial\Omega$ denotes the boundary of the domain Ω and n is the direction normal to $\partial\Omega$. The Euler-Lagrange equations for the above results in the following Poisson problem:

$$H_{xx} + H_{yy} = \partial_x P + \partial_y Q \qquad \left.\frac{\partial H}{\partial n}\right|_{\partial\Omega} = 0 \tag{11}$$

This is a standard problem in numerical linear algebra and a vast array of methods exist to solve it efficiently.

When shadows are present in the scene, the knowledge whether a pixel is shadowed or not imposes additional constraints on the height field [9]. Shadow graphs were developed as a representation of these constraints in [25]. In the following we briefly describe these constraints.

For simplicity of presentation, we consider a one dimensional image on the interval $[x_0, x_e]$. The generalization to two dimensional images requires applying the same logic by decomposing the image into a collection of one dimensional strips parallel to the projection of the light source direction on the image plane.

Without any loss of generality, we assume that the light source causing the shadowing is inclined at an angle less then $\pi/2$. We can now decompose the image into a set of intervals. Each interval is such that it only contains pixels which are shadowed and is maximal, that is, there is no other decomposition of the image possible in which another interval contains it as a proper subset. Let the line segment $[x_1, x_2]$ be such an interval where x_1 and x_2 are both shadow edges. Then, it is easy to see that all points $x \in [x_1, x_2]$ are shadowed by x_2 and obey the inequality

$$h(x) \leq h(x_2) - \|x - x_2\| \tan\theta \qquad \forall x \in [x_1, x_2] \tag{12}$$

These are the *shadowing* constraints. Now, if x_3 is not shadowed, then we know that there is no point between x_3 and the end of the image that occludes it from the light source, i.e. every point $x \in (x_3, x_e]$ has height less than $\|x_3 - x\| \tan\theta$, i.e

$$h(x) \leq h(x_3) + \|x - x_2\| \tan\theta \qquad \forall x \in [x_3, x_e] \tag{13}$$

These are the *anti-shadowing* constraints. Figure 1 illustrates the two sets of constraints. Imposing these constraints on the optimization problem in (10) results in a constrained optimization problem which can't be solved using a Poisson solver anymore, since we can't take variational derivatives.

Let $D_x \in \mathbb{R}^{m\times m}$ and $D_y \in \mathbb{R}^{n\times n}$ be matrices corresponding to the derivative operators on $\mathbb{R}^n$ and $\mathbb{R}^m$, and let $S_x = D_x \otimes I_n$ and $S_y = D_y \otimes I_m$, where $\otimes$ is the matrix Kronecker product and I_m is the m-dimensional identity operator. Finally let h, p and q be vectors obtained by concatenating the columns of H, P and Q, respectively. Then the quadratic programming problem we solve is

$$\begin{aligned} \min \quad & h^\top \left(S_x^\top S_x + S_y^\top S_y\right) h - 2(p^\top S_x h + q^\top S_y h) \\ \text{subject to} \quad & Ah \leq b \end{aligned} \tag{14}$$

where, the matrix A and vector b encodes the discrete shadow and anti-shadow constraints. The matrix A is extremely sparse with only two non-zero entries per row. Since $\left(S_x^\top S_x + S_y^\top S_y\right)$ is positive semidefinite, the above is a convex quadratic program which can be solved efficiently using modern solvers based on interior point methods [6]. We use the MOSEK solver [2].

5. Experimental Results

In this section, we report the performance of the our algorithm on a number of synthetic and real datasets. Our synthetic data is generated using the POVRay raytracer with 1% random noise added to image intensities.

We begin by considering a synthetic sphere with sharply varying albedo and illuminated by four light sources, one source per image. All shadows are attached. As is evident from the input images in Figure 2, it is very difficult to judge shadow boundaries using mere intensity thresholding for such textured surfaces. The shadow labeling recovered by ShadowCuts is nearly identical to ground truth.

Our next synthetic example consists of two convex hemispheres placed on a plane. For a complex scene, it makes sense to use more light sources for better coverage. Four images of the scene are obtained using six light sources, with three sources turned on at a time. Both attached and cast shadows are present and an inspection of the input images

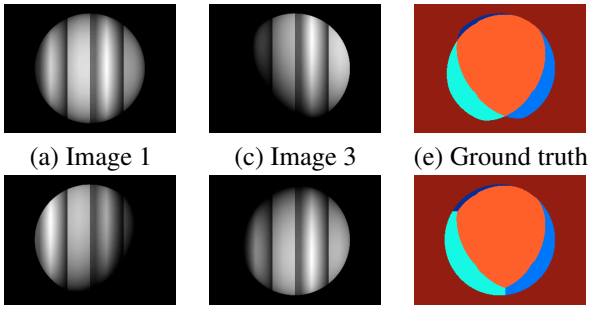

(a) Image 1 (c) Image 3 (e) Ground truth

(b)Image 2 (d) Image 4 (f) Estimated

Figure 2. Four source, four image photometric stereo with shadows. This figure illustrates the accurate recovery of attached shadows with a small number of images. Figures (a)–(d) are the source images used as input to the shadowcuts algorithm, (e) shows the ground truth shadowing configurating and (f) shows the labeling obtained using the shadowcuts algorithm. Notice that sharp changes in albedo do not affect the output of out algorithm.

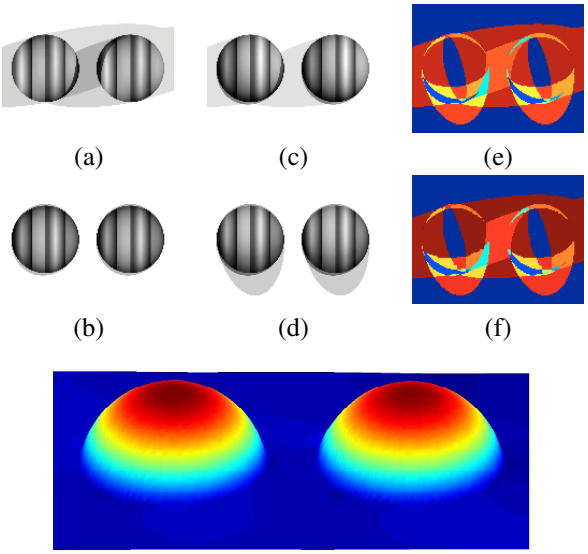

Figure 3. Six source, four image photometric stereo with shadows. The figure illustrates accurate shadow map recovery in the presence of varying albedo, cast as well as attached shadows. (a)-(d) Input images. (e) Ground truth shadow labeling. Each color stands for a different light source visibility configuration. (f) Labeling obtained using the ShadowCuts algorithm. Recovered surface has been false colored to indicate height above ground plane. Notice that sharp changes in albedo do not affect output of our algorithm.

shows there is no obvious way to decipher the underlying source visibilities. Our algorithm, on the other hand, correctly recovers the shadowing configuration almost everywhere. Figure 3 also shows the reconstruction obtained by integrating normals after imposing shadow constraints.

For comparison, we apply the algorithm of [8] to four single source images of the same scene. The resulting reconstruction (Figure 4) looks very distorted and the reason becomes clear if we compare the estimates of x and y gradients obtained by the two methods. While gradients obtained using ShadowCuts are symmetric with clean boundaries, those obtained using the method in [8] are severely biased by shadows. In particular, one can see the image of the cast shadow boundaries in the gradient estimates. This demonstrates our algorithm's advantage where a greater number of sources can be used for enhanced surface coverage without increasing the number of acquired images.

In the previous examples, the geometry of the scene was reasonably simple, even though the shadowing configurations were complicated. In the next synthetic example, we recover the shadowing configuration for a more intricate geometry, using four images from five sources. The geometry of the dragon model gives rise to a variety of cast and attached shadows of varying sizes. Again, the estimated shadow map matches the ground truth very closely (Figure 5).

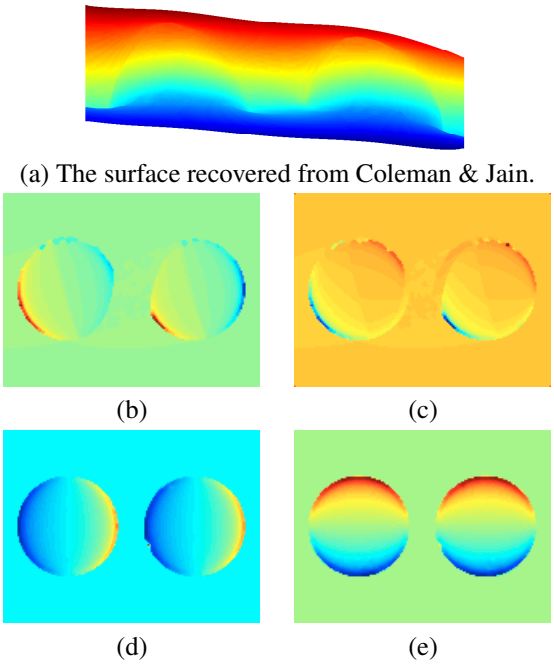

Figure 4. Comparison with Coleman & Jain's four source photometry method. (a) The reconstructed surface is badly distorted. (b)-(c) x and y gradients estimated using Coleman and Jain's method. Notice the significant bias in gradients caused by cast shadows. (d)-(e) Gradient estimates for the same geometry obtained using ShadowCuts.

We present reconstructions on some real objects where traditional algorithms would be hampered by the presence of shadows. Consider Figure 6, where four images of the object are obtained using five light sources, with two turned on at a time. Notice the faithful reconstruction in shadowed regions, such as the underside of the chin.

Another example of the reconstruction obtained by our algorithm for a real object is shown in Figure 7. Notice the cast shadow labeling on the left shoulder and attached shadow labelings on the sides of the skirt. Very accurate recovery of high frequency detail is consistent with expectations from a photometric stereo algorithm.

6. Discussions

We have demonstrated in this paper a novel and reliable algorithm for recovering cast and attached shadows in the presence of albedo variation and complex geometry. We have introduced a technique for multiple light source photometric stereo which enables better coverage of complex objects without increasing the number of images acquired. It has additional benefits of an improved SNR and better conditioning of the light source matrix.

The recovered shadow boundaries allow us to perform

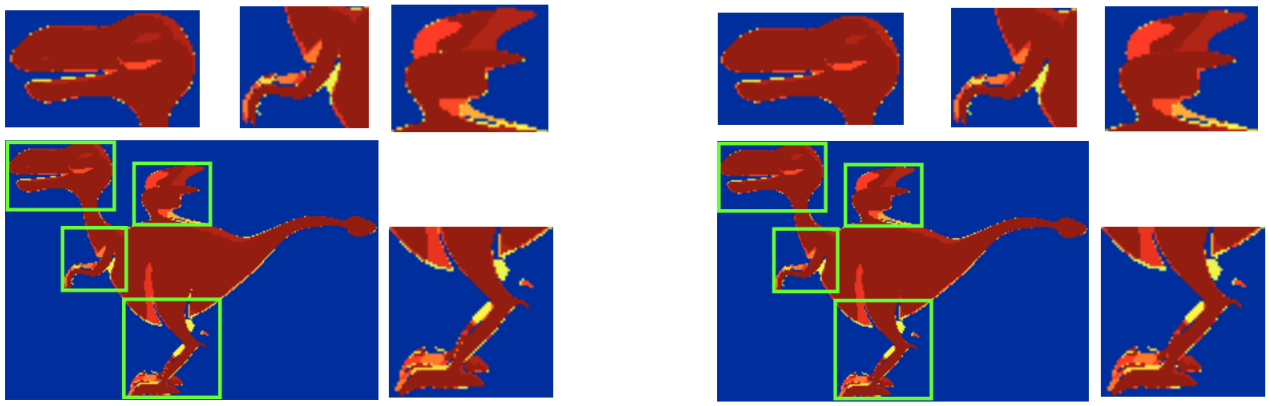

(e) (f)

Figure 5. Shadow recovery in the presence of complex geometry. (a) The ground truth shadow map. (b) Shadow map estimated using ShadowCuts algorithm. Notice the complexity of the shadowing pattern near (clockwise insets) the head, front legs, wings and hind legs.

normal integration with the constraint that the recovered surface respect these shadow boundaries. Applications for this might arise for objects with vicious geometries or for overcoming the low-frequency bias in normal integration. For regular objects considered in this paper, the shadow constraints do not result in an appreciably different surface. This is understandable as our normal estimates are excellent, so even unconstrained integration results in a surface that satisfies most of the shadow constraints.

An important avenue for future work is extension of this method to uncalibrated photometric stereo.

Acknowledgments

The authors would like to thank Satya Mallick, Fredrik Kahl and Michael Goesele for several helpful discussions. Manmohan Chandraker and David Kriegman were supported by NSF EIA-0303622 and IIS-0308185. Sameer Agarwal was funded by NSF EIA-0321235, University of Washington Animation Research Labs, Washington Research Foundation, Adobe and Microsoft.

References

[1] A. Agrawal, R. Raskar, and R. Chellappa. What is the Range of Surface Reconstructions from a Gradient Field? *ECCV*, 2006.

[2] E. Andersen, C. Roos, and T. Terlaky. On implementing a primal-dual interior-point method for conic quadratic optimization. *Math. Prog.*, 95(2):249–277, 2003.

[3] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *PAMI*, 25(10):1239–1252, October 2003.

[4] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. *Proc. of CVPR*, pages 374–381, 2001.

[5] P. Belhumeur, D. Kriegman, and A. Yuille. The Bas-Relief Ambiguity. *IJCV*, 35(1):33–44, 1999.

[6] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[7] Y. Boykov, O. Vexler, and R. Zabih. Efficient approximate energy minimization via graph cuts. *PAMI*, 20(12):1222–1239, 2001.

[8] E. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *CVGIP*, 18(4):309–328, April 1982.

[9] M. Daum and G. Dudek. On 3-D surface reconstruction using shape from shadows. In *CVPR*, pages 461–468, 1998.

[10] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *PAMI*, 10(4):439–451, 1988.

[11] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *PAMI*, 27(8):1254–1264, 2005.

[12] J. Ho, J. Lim, M. Yang, and D. Kriegman. Integrating Surface Normal Vectors Using Fast Marching Method. *Lecture Notes in Computer Science*, 3953:239, 2006.

[13] S. D. Hordley, G. D. Finlayson, and M. S. Drew. Removing shadows from images. In *ECCV*, pages 823–836, 2002.

[14] B. Horn and M. Brooks, editors. *Shape from Shading*. MIT Press, Cambridge, Mass., 1989.

[15] J. Kender and E. Smith. Shape from darkness. In *Proc. Int. Conf. on Computer Vision*, pages 539–546, London, 1987.

[16] D. Kriegman and P. Belhumeur. What shadows reveal about object structure. *JOSA*, 18(8):1804–1813, 2001.

[17] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi. Efficiently combining positions and normals for precise 3d geometry. *SIGGRAPH*, 2005.

[18] S. Savarese, H. Rushmeier, F. Bernardini, and P. Perona. Shadow carving. *ICCV 2001*, pages 190–197, 2001.

[19] Y. Y. Schechner, S. K. Nayar, and P. N. Belhumeur. A theory of multiplexed illumination. In *ICCV*, 2003.

(a) - (b) Sample input Images

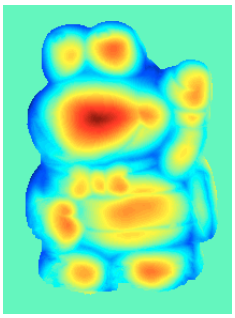

(c) Shadow map (d) Depth map

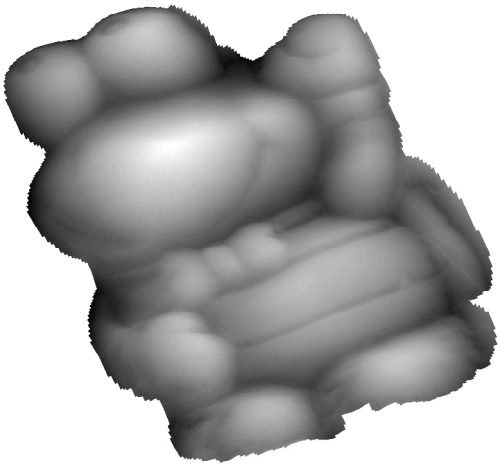

(e) Recovered surface

Figure 6. Shadow labeling and depth recovery for a real object. (a)-(b) Two of the four images obtained using five light sources. (c) Shadow map recovered by the ShadowCuts algorithm. (d) Depth map recovered using the estimated shadow map. (e) Geometry of the recovered surface. Notice the shadow labeling on the underside of the chin and faithful reconstruction in that region.

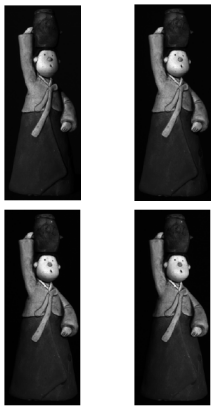

(a) - (d) Input Images

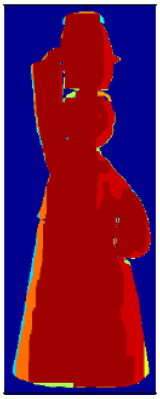

(e) Shadow map

(f) Reconstruction using ShadowCuts

Figure 7. Shadow detection and depth recovery for another real object. (a)-(d) Four images obtained using five light sources. (e) The recovered shadow map. (f) Reconstruction. Notice the recovery of high frequency detail.

[20] T. Simchony, R. Chellappa, and M. Shao. Direct analytical methods for solving poisson equations in computer vision problems. *PAMI*, 12(5):435–446, 1990.

[21] R. Woodham. Photometric stereo: A reflectance map technique for determining surface orientation from image intesity. In *SPIE*, volume 155, pages 136–143, 1978.

[22] T. Wu and C. Tang. A Bayesian Approach for Shadow Extraction from a Single Image. *ICCV*, 1, 2005.

[23] T.-P. Wu, K.-L. Tang, C.-K. Tang, and T.-T. Wong. Dense photometric stereo: A markov random field approach. *PAMI*, 28(11):1830–1846, 2006.

[24] Y. Yu and J. Chang. Shadow graphs and 3D texture reconstruction. *IJCV*, 62(1-2):35–60, 2005.

[25] Y. Yu and J. T. Chang. Shadow graphs and surface reconstruction. In *ECCV*, 2002.

[26] A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. *CVPR*, pages 158–164, 1997.