On 3-D Surface Reconstruction Using Shape from Shadows

M. Daum
mdaum@cim.mcgill.ca

G. Dudek
dudek@cim.mcgill.ca

McGill University
Centre For Intelligent Machines
School of Computer Science
Montreal, Quebec, Canada H3A 2A7

Keywords:
Shape From Darkness, 3D Reconstruction, Scene Recovery, Shadows

Abstract— ***In this paper we discuss new results on the Shape From Darkness problem: using the motion of cast shadows to recover scene structure. Our approach is based on collecting a set of images from a fixed viewpoint as a known light source moves "across the sky". Previously published solutions to this problem have performed the reconstruction only for cross sections of the scene.***

In this paper, we present a reconstruction algorithm and discuss the reconstruction of an entire 3-D scene under various light source trajectories. We also consider the constraints on reconstruction. We conclude with experimental results that illustrate the convergence properties of the solution process and its robustness properties.

I. Introduction

In this paper, we consider surface reconstruction from shadow information. That is, to use the shape and geometric properties of observed shadows to infer the shape of the surfaces casting the shadows as well as those that the shadows are cast upon. This problem is sometimes known as *shape-from-darkness* [11], [2].

The Shape From Darkness method allows one to construct a model of a scene using information on shadows cast within the scene under illumination from a moving light source. Consider, for example, the shadows cast by a mesas in the desert as the sun moves across the sky, or the shadows cast by a flare fired by search-and-rescue personnel working at night. By observing the shapes of shadows as the move across the scene, we can infer the shapes of the surfaces that cast the shadows. This method for shape reconstruction presents several advantages over other scene constructing algorithms in that it requires inexpensive instrumentation and allows for efficient computation due to the compact nature of shadow data. It also requires only weak assumptions about surface reflectance properties, as opposed to shape-from-shading's strong Lambertian assumptions. Furthermore, it has been shown that Shape From Darkness can be used to infer the shapes of surfaces in the scene which are not directly visible to the camera, providing information which other methods cannot. Pragmatically, the technique may be useful in contexts where traditional range sensors may not be suitable (eg. Martian Exploration).

Previously published formulations of Shape From Darkness have been limited in that they only allow the reconstruction of individual two dimensional cross sections of the scene. In past work, the depth map for each 2-dimensional section is constructed only relative to an arbitrary additive constant, obviating the inclusion of external data in order to combine the slices into the complete three dimensional scene. Implicit in this cross-sectional model also is the fact that the light source moves only through the plane containing both the cross section and the camera. This assumption limits the applicability of the method, even in seemingly natural applications such as terrain reconstruction from satellite imagery.

In practice, the recovery of 3-D structure from shadows entails the detection of shadows in the first place. Several interesting solutions to this sub-problem have been posed [15], [6], [1], [3], including some that are specifically suited to the shape-from-darkness context [12]. Due to lack of space, this paper will not discuss shadow detection explicitly and we assume it

This is a DRAFT COPY of a paper to appear in the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Los Angeles, CA, June, 1998. This work was supported by the Canadian Centres of Excellence IRIS Project IS-5 and a NSERC research grant to G. Dudek.

can be achieved by one of these existing methods (in practice, we use a combination of techniques).

The shape-from-shadows (shape-from-darkness) problem is also related in spirit to shape-from-shading methods as described by Horn and Ikeuchi [10], [9], or dealt with more recently by Dupuis and Oliensis [4], [5], among many others [14], [7], [13]. Shape from shading algorithms, however, are dominated by local constraints, while shape from darkness deals primarily with large scale global constraints. That is, it is the orientation of the surface normal, a differential (local) property, that determines the observed intensity under shading. In contrast, shape from darkness depends on the relationships of pixels that cast shadows on one another, potentially from widely separated locations.

Much of the prior ground breaking work on shape-from-darkness deals with two-dimensional instances of the problem, where the light source and the surface to be recovered lie in the same plane [11], [2]. Typical existing approaches to shape-from-darkness make three key assumptions regarding the problem [11], [8], [3], [12].

- The world, including the camera and light source is two-dimensional (i.e the light source and the surface to be recovered all lie in the same plane)
- The light source and camera geometry can both be modeled using orthographic projection (i.e. they are extremely distant)
- The surface to be reconstructed is a *terrain* described by a function $z(x)$ (i.e. a graph surface)

Our approach to the problem allows us to relax all three of the assumptions, although in this paper we will focus only on the first (a 3-D instead of a 2-D world).

In this paper, we present an approach to the reconstruction of 3-D scenes with limited constraints on the light source trajectory. In Section 2 we review the 2-D and 3-D shape from darkness problem, and define constraints that relate light source motion and scene geometry. In Section 2 we also consider basic equivalence classes of surfaces under shape from darkness. In Section 3 we describe our shape from darkness algorithm. In Section 4 we present experimental results and discuss the robustness of the solution with respect to errors in the estimated light source position, and in the shadow segmentation. Section 5 presents conclusions.

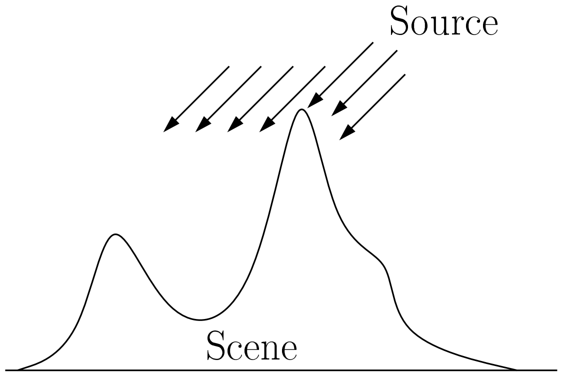

Fig. 1. Two Dimensional Model

II. Problem Definition

A. 2-D Problem

Previous work on Shape From Darkness has focused on the solution of the two dimensional version of the surface reconstruction problem. In this version of the problem a surface is defined as a function $z = f(x)$ where z and x are Euclidean spatial dimensions. If a surface $f(x)$ assigns a single value to each x in a given *range*, then this surface is *terrain-like.* If the surface is more complex, then it is *non-terrain-like.*

For non-terrains, it is convenient to define the concept of a *generated terrain* [12]. The generated terrain is defined as the terrain-like upper envelope of a scene. For scenes which are terrain-like, the generated terrain is identical its generator. The generated terrains of non-terrain-like scenes contain only those surfaces which face toward the camera, however. Thus the features of a non-terrain that are absent from the generated terrain are associated with hidden surfaces.

To reconstruct a surface, a light source must be moved through a trajectory of angles above the surface, as depicted in Figure 1(in practice, these can be arbitrary discrete sample locations.) A stationary camera records a series of images of the surface as the light source moves overhead. Both the light source and the camera are considered to be an "infinite" distance away from the scene. This has the effect of creating a camera with orthographic projection and a light source which casts rays which are parallel to one another. Thus, in the 2D formulation a single angular parameter θ suffices to describe the position of the light source. An important effect of this positioning is that every pixel in the image is guaranteed to be lit when the light source is directly overhead (in the "noon" position).

Shadow information can be described using an in-

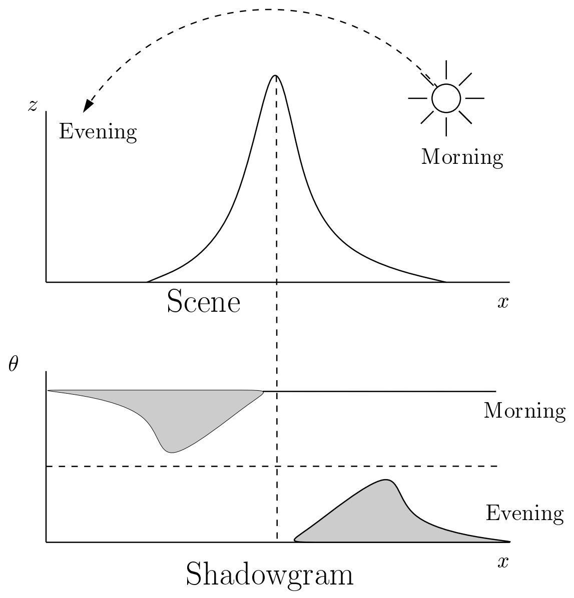

Fig. 2. Two Dimensional Shadowgram

termediate representation known as a *Shadowgram.* As shown in Figure 2, the shadowgram is a binary function $s(x, \theta)$ on the angle θ and a spatial dimension x. A white entry in the shadowgram indicates that image pixel x was lit when the light source was at angle θ, while a black entry indicates that it was shadowed. It was shown in [11] that the shadowgram *generated by a terrain-like surface* can be completely described by two curves: θ^+ and θ^-, representing the first light in the "morning" and the last in the "evening" respectively. If the light source travels from horizon to horizon, it is possible that one of these curves might disappear (i.e. the point is always lit from sunrise until noon). Both curves will not disappear simultaneously, however, unless the point in question is a global maximum of the scene (i.e. mountain tops are always lit). Because one has this guarantee, it is possible to reconstruct the surface by integrating θ^+ and θ^- [11].

When the surface is non terrain-like, the shadowgram possesses not only two curves, but also some white *holes* where one would expect darkness if the surface were a terrain. Here θ^+ and θ^- are not defined as first and last lighting curves but rather as the envelope of the shadowgram which lies closest to the noon position. It is shown in [12] that using these two curves to reconstruct the surface will in fact produce the surface's generated terrain. Furthermore, the holes in the shadowgram may then be used to carve pieces out of the generated terrain, allowing one to reconstruct some or all of the hidden surfaces in the scene.

B. Three Dimensional Problem

The three dimensional problem is a natural extension to its two dimensional analogue. Here the scene as presented to the camera consists of a two dimensional image rather than a single scan line. Thus the surface for reconstruction in the 3D problem is a function $z = f(x, y)$. As before, we say that the scene is terrain-like if this function is single valued over the ranges of x and y which are presented in the image.

Because the scene is not constrained to lie within a single vertical cross section, it is no longer necessary to force the light source to do the same. Instead, the light source direction is allowed to point freely, and must be described by two angles, ϕ and θ.

From this perspective one may view the original two dimensional problem as a special case of the larger 3D formulation. One in which the light source travels through a trajectory through a series of θ angles while keeping its ϕ fixed at noon. In fact, given this type of trajectory one can in fact reconstruct each scan-line of the image individually using the 2D algorithm. Because the process underlying the reconstruction is integration, however, these internally consistent scan-lines can not be combined into a whole as they are each reconstructed to within an unknown additive constant.

C. The Shadowgram in the Three Dimensional Model

Adding a dimension to both the image and the light source parameterization has a profound effect on the shadowgram. Unlike the original shadowgram, which is defined in the space $\mathcal{R} \times \mathbf{S}$ possessing two dimensions, the complete shadowgram lies in the space $\mathcal{R}^2 \times \mathbf{S}^2$ (i.e. $f : \mathcal{R}^2 \times \mathbf{S}^2 \rightarrow \{0, 1\}$). We define the *sandwich* and *tunnel* cross sections as specific three dimensional projections of the shadowgram.

The sandwich cross section refers to the projection of the shadowgram which includes both spatial dimensions and only one of the two angles ($\mathcal{R}^2 \times \mathbf{S}$). It is the cross section of the shadowgram obtained by fixing one of the angles at the noon position. As shown in Figure 3, the sandwich cross section is an extension of the 2-D shadowgram. The original θ^+ and θ^- curves are present, but are now entire surfaces. Likewise, the holes associated with a non-terrain surface now show up as three dimensional cavities.

The tunnel cross section is projected onto the space $\mathcal{R} \times \mathbf{S}^2$. It is equivalently defined as the cross section

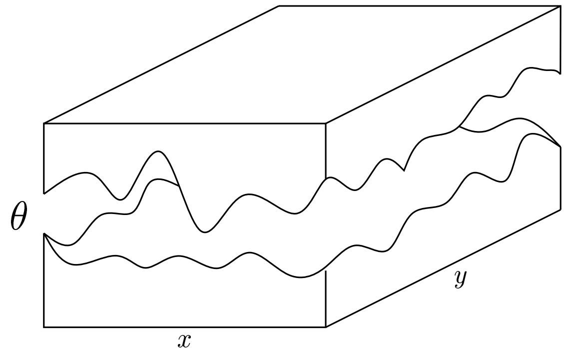

Fig. 3. Sandwich cross section

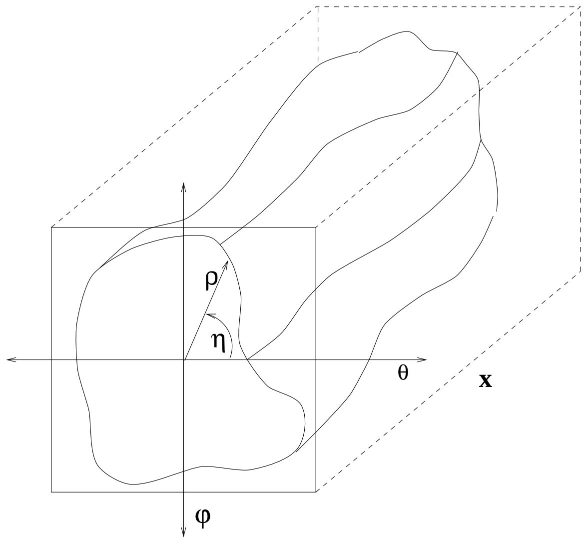

Fig. 4. Tunnel cross section

obtained by fixing one spatial coordinate of the shadowgram. As seen in Figure 4, the tunnel cross section resembles a snakelike open cavity.

The individual slices of the tunnel cross section are interesting in their own right. These slices present a complete picture of the shadowing of a single image pixel for all light source angles. Each slice possesses an inner bounding curve around the origin (the noon position). In addition, there may be holes around this curve if the shadowing surface is not a terrain. For the same reason that one of θ^+ or θ^- will always be present in the 2-D shadowgram, the inner curve is guaranteed to be a star-shaped polygon. In fact, this curve is a proper function from angle ϕ to radius $\rho(\rho = f(\phi))$. Furthermore, if we bound on the gradient of the original surface:

$$|\nabla f(x,y)| < n$$

then we can guarantee that the circle with radius $\tan^{-1}(n)$ will be a kernel of the polygon. In other words, the polygon will never pass inside of this circle. This is the strongest guarantee we can make.

D. Constraining Shadower and Exact Reconstruction

What do shadows tell us about 3D surface structure? Consider a shadowed point in an image representing a single light source position. If a ray is cast from this point in the direction of the light source (the point's *light seeking ray*), any surface point lying above this ray is a potential shadower of this point. Of these possible shadowing points, the point which lies highest above this ray (and furthest from the casting point, if this height is not unique) has special significance, and is called the point's *constraining shadower.*

If a point is a constraining shadower of another point, we have the guarantee that this point lies along a shadow boundary in the image in question. Furthermore, we know that this is a contact shadow boundary on the surface. As a result, we can identify the constraining shadower as the first shadow boundary point encountered in image space along the image projection of a point's light seeking ray. Note that this means that the identification of the constraining shadower can be performed prior to surface reconstruction.

This method of identification points to two interesting properties:

- All shadowed points in a connected neighborhood lying along the image projection of a light seeking ray share the same constraining shadower.
- The furthest of these points from the constraining shadower lies along a shadow boundary, and is in fact the unique cast shadow boundary generated by the constraining shadower.

From this we see that for a connected 2D region of shadowed points in a given image, each point on a "lightward" boundary of the region is a constraining shadower, and every other boundary point is a cast shadow boundary corresponding to a unique constraining shadower.

Given a two points representing a constraining shadower - cast shadow boundary pair for a certain light source direction, it is known that the relative height of the points is equal to the height travelled by a light seeking ray between the two points. If the distance were greater than this, then the cast shadow boundary would be moved away from the light. On the other hand if the distance were less than this the receiving pixel would not be shadowed. Thus the relative heights of all such boundary pairs are known exactly (Figure 5).

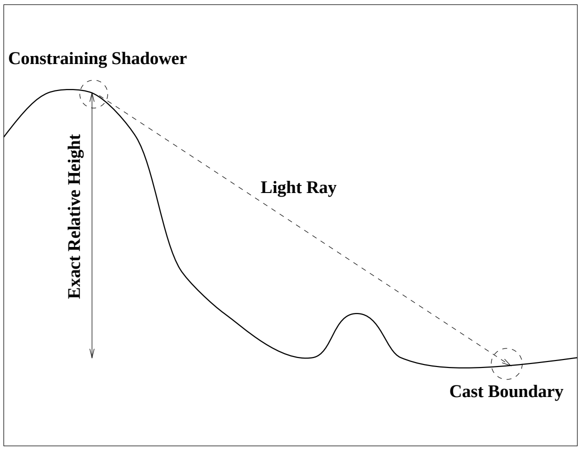

Fig. 5. Exact Relationship Between Heights of Shadow Pairs

When more than one light source direction is considered, it is possible for this exact relationship to apply to a large set of points, as constraining shadower's may have as many cast shadow boundaries as light source directions, and cast shadow boundaries may themselves become constraining shadowers. Consider a graph **G** containing a node corresponding to each image-space pixel, having an edge between two nodes only if the corresponding pixels are somewhere a constraining shadower - cast shadow boundary pair. Each fully connected subgraph of **G** represents a subset of surface pixels which is internally exactly reconstructed. If the entire graph is connected, then exact reconstruction is possible. However, if multiple sets exist, then the relative heights may not be exactly known, but can be bounded less strongly through application of weaker shadow constraints.

E. Shadow Equivalence Classes

It is clear that the action of shadowing is concerned only with the relative heights of points on the surface and not by the absolute height of the surface as a whole. This points to an important equivalence class of surfaces. Two surfaces $f(x,y)$ and $g(x,y)$ are said to be *lift equivalent* if they differ by an additive constant:

$$f(x,y) = g(x,y) + \mathbf{C} \qquad (1)$$

Surfaces which are lift equivalent are guaranteed to posses identical shadowgrams as a result of the use of parallel light rays. For this reason, we define the *anchoring operator* to remove lift equivalence in the following way:

$$\underline{f(x,y)} = f(x,y) - \inf\left(f(x,y)\right) \qquad (2)$$

As the name suggests, the anchoring operator has the effect of fixing the surface to the x-y plane. Hereafter, we will assume that all surfaces have been anchored.

In addition, there exists a less trivial equivalence class of *shadow equivalent* surfaces. Two surfaces are shadow equivalent if they generate the same shadowgram. If we consider the complete shadowgram, then all members of such a class will also be lift equivalent. However, if we restrict the set of light source directions to a single trajectory

$$\mathbf{t} = (\phi(t), \theta(t)) \; t = [0..1]$$

then it is possible that the shadowgram subsets generated will be identical for qualitatively different surfaces. Such surfaces are said to be *shadow equivalent under trajectory* $\mathbf{t}$.

III. Approach to Reconstruction

The reconstruction problem lends itself naturally to a solution through the iterative relaxation of constraints. Constraints based on the shadow information at each pixel are repeatedly applied to a working surface until each constraint is individually satisfied.

Due to the binary nature of shadow information, it is natural that two types of constraints exist. An *expect light* constraint exists for a given pixel and source direction if that pixel was lit by the source from the given pose. Likewise an *expect darkness* constraint exists if the given pixel was in shadow.

Consider the constraints in terms of the behaviour of a light seeking ray cast from a point on the surface. In the case of expect light, such a ray is expected to pass freely out of the scope of the image without intersecting the working surface. On the other hand, a ray expecting darkness must certainly intersect the surface in at least one place in order to shadow the pixel. It is assumed that all shadowers lie within the image.

During the reconstruction, we model the scene as a pair of bounding surfaces. In order for these surfaces to be initialized, it is necessary to make a base assumption about the maximum possible height spanned by structure present in the scene. This is not unreasonable, as it is possible to choose such a height based on the fact that given taller scene structure, very little of the shadow information would be present in the scene. Based on this assumption, the lower bound is initialized to a plane at $z = 0$ and the upper bound to a plane at $z = \mathbf{maxheight}$.

The two types of constraints (expect light and darkness) and the two surfaces (upper and lower bounds) yield four rules for extraction of shadow information:

- **Expect Light**
  - The **Upper Bound** of any pixel lying in the image projection of a light seeking ray cast from the lit point's upper bound is lowered to the level of the ray (if it was previously above). These points cannot be higher than this as they would then shadow the point, which is known to be lit, even if the upper bound was it's true location.
  - The lit point's **Lower Bound** is raised until a light seeking ray cast from said lower bound will be pass above (or just touch) all lower bounds along the ray. This bound cannot be lower as the point would then be shadowed even if the lightward points were all at their lower bounds.
- **Expect Darkness**
  - The **Upper Bound** of the shadowed pixel is lowered until the light seeking ray which it casts intersects the upper bound at some other point. If the bound were higher, then the point could not be shadowed even if the lightward points were located at the upper bound.
  - The **Lower Bound** of the shadowed point's *constraining shadower* is raised to the level of the light seeking ray cast from the lower bound of the shadowed point. The shadower must be at least this high in order to shadow the point at its lower bound.

From these rules one can see that the upper and lower bound surfaces are not directly coupled to one another. Points on the upper bound will only effect other upper bound points, while lower bounds only effect other lower bounds. Thus, the two surfaces are related only through the shadow information and may be computed separately.

It also follows that points on the upper bound are only lowered, while points on the lower bound are raised. This shows that at worst, the application of new shadow data will leave the bounds unchanged, and will never degrade the estimate. This also guarantees that the process will terminate. The distance between corresponding upper and lower bound points cannot decrease, making non-terminating cyclic oscillations impossible.

After the computation of the upper and lower bounds, it is necessary to merge the two in order to obtain a complete picture. Due to lift equivalence, it is possible to accomplish this by simply squeezing the surfaces together, adjusting the absolute heights so that contact is made only at the closest point (or set of points if non-unique).

Note that the absolute bounding of the scene is in fact an assumption of convenience. Because the upper and lower bound are reconstructed separately, it is possible to allow the reconstruction of scenes of arbitrary height. In this case unconstrained points on the lower bound would be represented as bottomless, and those on the upper bound as having no roof.

IV. Implementation and Results

The algorithm used to reconstruct the bounding surfaces corresponding to a particular set of shadow images involves the iterative application of the constraints outlined in the previous section. The work surfaces are initially set to be a flat plane. For each pixel in the work surfaces a constraint is enforced for each light source direction in the trajectory. The complexity of a single iteration is thus $\mathbf{O}(n \times m \times t \times r)$ for an image of width n and height m, a trajectory of t source directions, and an average ray length of r. If we assume a square image and an equal number of source directions then this reduces to $\mathbf{O}(n^4)$.

Each iteration involves two waves of reconstruction. In the first wave all of the expect darkness constraints are applied, followed by expect light in the second. During each wave, the constraints belonging to all pixels in a given image row are applied in parallel, with changes being written back between successive rows. The order in which rows are reconstructed (front to back or vice-versa) is chosen based on the type of constraints which are being applied, in order to minimize repetitive work. In order to avoid sensitivity to errors introduced by discretization, all constraints which lay in a neighborhood containing both shadow and light were discarded.

A. Experimental Data

A.1 Generated

Our implementation of shape from darkness accomplishes surface recovery on 64x64 pixel images with 64 shadow images in roughly 5 minutes on a sparc-20 workstation.

Figure (6) presents the input surfaces and the associated reconstructed upper bound for a surface of moderate complexity. The shadow information for this scene was extracted from a series of artificial images rendered from a CAD model of the scene. The upper bound was chosen for display as it typically

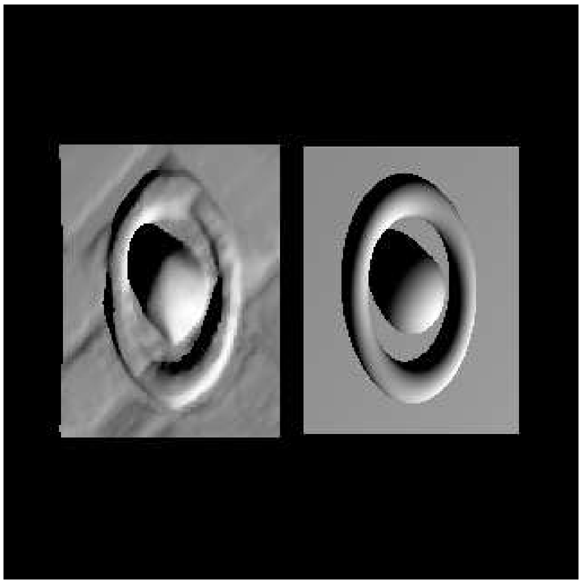

Fig. 6. Surface Reconstruction from Shadow Data for a Complex Terrain. Original surface is shown on the right

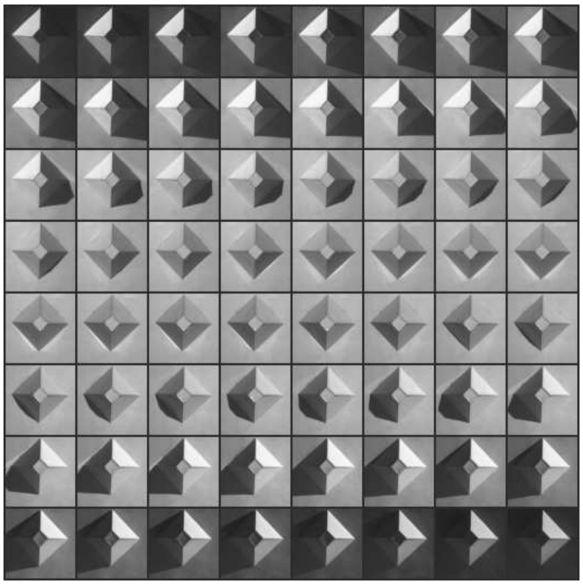

Fig. 7. Shadow Information Used in Pyramid Reconstruction

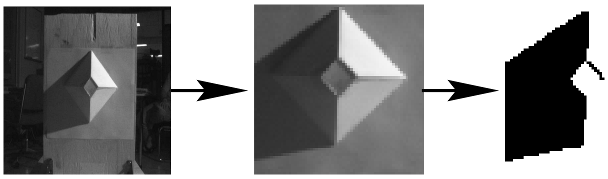

Fig. 8. Image Preparation for Reconstruction

resembles more closely the actual scene. Both the reconstruction and the CAD model were re-rendered for display. This method of display has a tendency to accentuate small errors in reconstruction as small changes in slope created large changes in Lambertian shading, making relatively flat areas appear mottled. The average error in this reconstruction was approximately 2.5% of the total scene height.

For portions of the surface that are not shadowed, or that are always in shadow, the absence of sufficient constraints on the surface geometry can sometimes lead to significant artifacts in the upper bound surface. These artifacts reflect the dearth of information about these points present in the shadow set. As expected, these effects are most often seen in the extreme fronts and backs of both images and the individual objects within them.

Observe that the surface is accurately reconstructed both on the front surface (facing the light source) as well as along the back surface (away from the light source). This is possible since information on surface geometry is obtained both based on shadows cast *by* a surface as well as by the behaviour of shadows that are cast *upon* a surface.

A.2 Real

As a demonstration of the applicability of this algorithm, reconstruction was performed on a simple scene containing a four-sided pyramid with a flat top. The camera and scene were both mounted on a platform, which was then rotated under constant lighting by a single source. The light-source "trajectory" generated was that of a cone of directions whose axis lay in the image plane. The base angle of the cone was approximately 70^o. Photographs of the scene were taken in 64 light source positions (figure 7). The resulting images were then cropped and thresholded, as depicted in figure (8).

The resulting reconstruction was performed in 33 iterations, taking 347 seconds on a sparc 20 workstation. A chart of the average difference between pixels in adjacent frames is presented in figure (9).

Figure (10) presents the reconstructed upper bound for the pyramid sequence. The main artifacts present are a shelf in the back generated by a lack of shadowing towards the rear of the image, and a rise in the front reflecting the difficulty in segmenting shadows on the base plane when light rays are highly oblique to its normal.

A.3 Robustness

Errors in Light Source Position

Two experiments were performed in order to determine the robustness of the algorithm. The first of

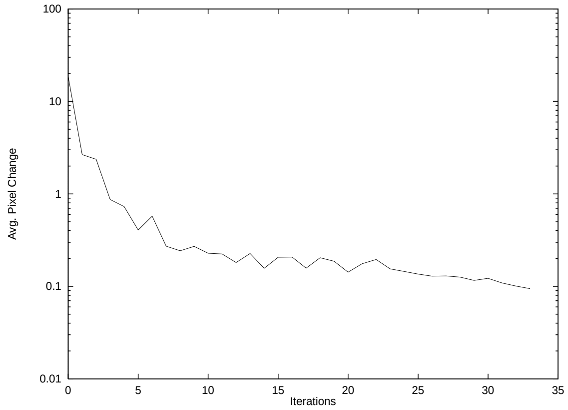

Fig. 9. Graph of the Convergence of The Pyramid Scene

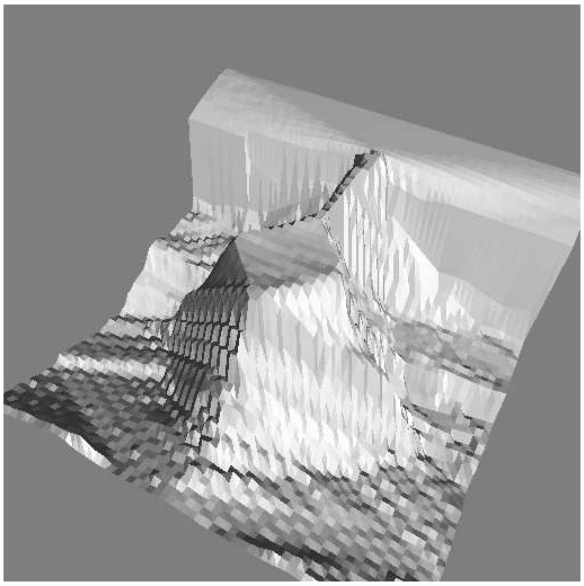

Fig. 10. Reconstruction of Pyramid from Photograph Sequence

these involved the reconstruction's sensitivity to errors in knowledge of the light source position. A jitter factor was applied to the assumed source positions during the reconstruction of a scene which had previously produced good results. The original, correct shadow data was used, but the algorithm was lied to about the corresponding source directions by a random amount bounded by a set number of degrees. It was found that the absolute difference between estimated angle and ground truth was the critical factor. Though minor artifacts appeared, convergence was obtained up to a maximum angular difference of approximately 6^o. Note that at this extreme value, position reordering occurred of shadow data up to 3 frames away along the trajectory.

Errors in Shadow Segmentation

To determine the algorithm's sensitivity to errors

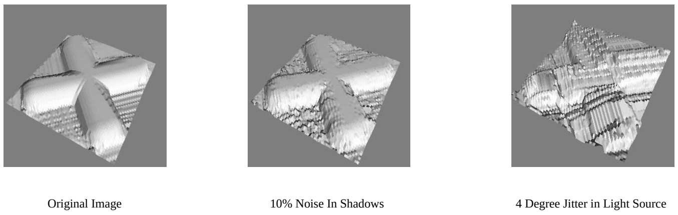

Fig. 11. Effects of Different Types of Noise on Reconstruction

in the segmentation, increasing levels of noise were added to the shadowgram of the previous scene. An increasing percentage of shadowgram constraints were randomly selected to be toggled from light to shadow or vice versa. It was found that all reconstructions converged up to a toggle percentage of 16%, and convergences were obtained for angles up to 28%. Among those which converged with hight amounts of shadow error (above 16%), the quality of the reconstructed surfaces varied greatly, indicating a sensitive dependence on the location of the inappropriate constraints.

Examples of these effects are described in figure (11).

V. Summary and Conclusion

In the paper we describe an algorithm for 3-D shape from darkness with limited constraints of the light source trajectory. Under a suitable trajectory, a good reconstruction of the original scene is possible, including "back faces" not directly observable to the camera. Our results indicate that the solution remains acceptable even in the face of various types of error in the input data.

Open problems for future work relate to the precise tradeoffs between light source trajectory and reconstruction quality. We are also investigating a coarse-to-fine resolution pyramid to accelerate the computation (with currently takes several minutes on a SGI R5000). How to exploit surface smoothness to produce more robust results, in the face of global constraints, is also under consideration.

References

[1] M.O. Ward C. Jiang. Shadow segmentation and classification in a constrained environment. *CVGIP: Image Understanding*, 59(2):213–225, 1994.

[2] K.A. Loparo D. Raviv, Y. Pao. Reconstruction of three-dimensional surfaces from two-dimensional binary images. *IEEE Transactions on Robotics and Automation*, 5(5):701–710, 1989.

[3] J. Kender D. Yang. Shape from shadows under error. In *Image Understanding Workshop 1993*, pages 1083–1090, Washington, D.C., August 1993.

[4] P. Dupuis and J. Oliensis. Direct method for reconstructing shape from shading. In *Proceedings, Conference on Computer Vision and Pattern Recognition*, pages 453–458, Champaign, Il., June 15-18 1992. Computer Society of the IEEE, IEEE Computer Society Press.

[5] P. Dupuis and J. Oliensis. Shape from shading: Provably convergent algorithms and uniqueness results. In *Computer Vision – ECCV 94*, volume 2, pages 259–268, Stockholm, Sweden, 1994. Springer-Verlag.

[6] D. Forsyth and A. Zisserman. Reflections on shading. 13:671–679, July 1991.

[7] Robert T. Frankot and Rama Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 10(4):439–451, July 1988.

[8] M. Hatzitheodorou and J.R. Kender. An optimal algorithm for the derivation of shape from shadows. In *Proceedings, Conference on Computer Vision and Pattern Recognition*, Ann Arbor,Michigan,USA, June 1988. Computer Society of the IEEE, IEEE Computer Society Press.

[9] Berthold Horn. *Robot Vision*. The MIT Press, Cambridge, Massachusetts, 1986.

[10] Katsushi Ikeuchi and Berthold K. P. Horn. Numerical shape from shading and occluding boundaries. In Michael Brady, editor, *Computer Vision*, pages 141–184. Elsevier Science Publishing Company, New York, NY, August 1981.

[11] E.M. Smith J.R. Kender. Shape from darkness: Deriving surface information from dynamic shadows. In *AIII*, pages 539–546, 1987.

[12] Michael Langer, Gregory Dudek, and Steven W. Zucker. Space occupancy using multiple shadowimages. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 390–396, Pittsburgh, PA, August 1995. IEEE Press.

[13] M.S. Langer and S. W. Zucker. Shape-from-shading on a cloudy day. *Journal of the Optical Society of America A*, 11(2):467–478, 1994.

[14] Alex P. Pentland. Local shading analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6(2):170–187, March 1984.

[15] S.A. Shafer. *Shadows and silhouettes in computer vision*. Kluwer Academic Publishers, 1985.