

Out of the Dark: Using Shadows to Reconstruct 3D Surfaces ^{*}

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. *Shape From Darkness refers to using the shadows cast by a scene to reconstruct the structure of the scene. A collection of images associated with different light source positions is used. Previously published solutions to this problem have performed the reconstruction only for cross sections of the scene.*

We propose a variant of Shape From Darkness which is capable of reconstructing the entire 3-D scene. In addition, this algorithm can be applied to a broader class of light source trajectories, including trajectories which mimic the motion of the sun during the day.

We present a formal statement of the 3-D problem and some of its characteristics, and an algorithm for recovering a surface from shadows. Experimental results are presented and discussed for both real data and synthetic data with associated ground truth.

1 Introduction

The Shape From Darkness method allows one to construct a model of a scene using information on cast shadows under illumination from a moving light source [9, 1, 10]. By observing the shapes of shadows as they move across the scene, we can infer the shapes of the surfaces that cast them. This method 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 reflectance (eg. Lambertian) assumptions. Furthermore, it has been shown that Shape From Darkness can be used to infer the shapes of surfaces in the scene even if they are not directly visible to the camera. Pragmatically, the technique may be useful in contexts where traditional range sensors may not be suitable (eg. Martian Exploration).

^{*} This work was supported by the Canadian Centres of Excellence IRIS Project IS-5

Shape from darkness is superficially related to the shape-from-shading problem. Shape from shading has been examined in a variety of contexts, using methods including relaxation labelling and regularization [7, 8, 12, 5, 3, 4, 11]. Shape from shading also appears to be a psychophysically relevant process [14]. Note, however, that shape from shading is a fundamentally local phenomenon (assuming the standard assumptions including that of a single distant light source and an absence of mutual illumination are made).

Much of the prior work on shape-from-darkness deals with two-dimensional instances of the problem, where the light source and the surface to be recovered all lie in the same plane [9, 1].

Typical existing approaches to shape-from-darkness make three critical assumptions regarding the problem [9, 6, 2, 10].

- 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).

We present a method by which scene reconstruction can be performed given a light source moving through an arbitrary set of three dimensional positions. During the reconstruction, the scene is modelled as a region lying between two bounding surfaces. As more shadow information is integrated into the estimate, the bounding surfaces move closer together until an exact reconstruction is achieved.

2 Problem Definition

2.1 Two Dimensional 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 $z = f(x)$. 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*.

To reconstruct a surface, a light source must be moved through a trajectory of angles above the surface (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 intermediate representation known as a *Shadowgram*. [9] As shown in Figure 1, the shadowgram is a binary

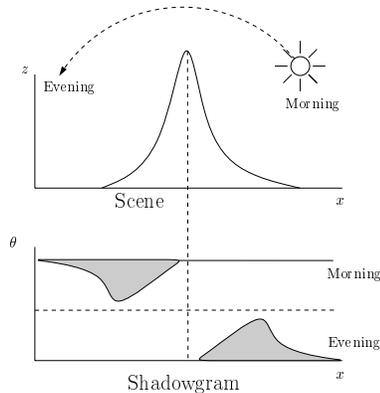


Fig. 1. Two Dimensional Shadowgram

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 [9] that the shadowgram of 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. It is possible to reconstruct the surface by integrating θ^+ and θ^- [9].

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 [10] that using these two curves to reconstruct the surface will in fact produce a terrain-like superset of the surface. Furthermore, the holes in the shadowgram may then be used to carve pieces out of this terrain, allowing one to reconstruct some or all of the hidden surfaces in the scene.

2.2 Three Dimensional Problem

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 no

longer lies within a single vertical cross section, the light source is allowed to point freely, and must be described by two angles, ϕ and θ .

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.

2.3 Constraining Shadower and Exact Reconstruction

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. 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.

3 Approach to Reconstruction

The reconstruction problem lends itself naturally to a solution through the iterative relaxation of constraints. It is natural that two types of constraints exist: *expect light* for pixels and source directions resulting in light, and *expect darkness* for pixels 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 an upper and a lower bounding surface which are incrementally brought together. 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 shadow the point.

- 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.

– **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.
- 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 ever 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. The distance between corresponding upper and lower bound points cannot increase, guaranteeing termination.

The global character of the discretized shape from darkness problem suggests that it is an ill-posed problem in the sense of Tikhonov [13]. That is, small changes in one part of the surface or shadowgram can have very large repercussions for the solution. As a result, it is necessary when solving the discrete version of the problem to apply a stringent confidence-based threshold on the shadow information. Only constraints which pass this test are allowed to contribute to the reconstruction. In practice, it is possible to base such a filter on the immediate neighborhood in shadow-space of the constraint in question, avoiding the complicating effects of non-local shadowing relationships.

4 Implementation and Results

The algorithm used to reconstruct the bounding surfaces involves the iterative application of constraints . The work surfaces are initially set to be flat. 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. Each image row's constraints are applied in parallel, with changes being written back between successive rows. Rows are processed either front-to-back or vice-versa in order to minimize repetitive work. To avoid sensitivity to errors introduced by discretization, all constraints in a neighborhood containing both shadow and light are discarded.

4.1 Experimental Data

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 (2) presents the input surfaces and the associated reconstructed upper

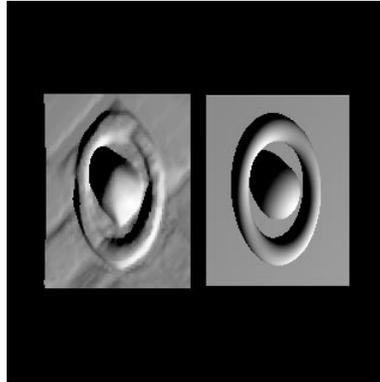


Fig. 2. Surface Reconstruction from Shadow Data for a Complex Terrain. Original surface is shown on the right

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 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.

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° . Photographs of the scene were taken in 64 light source positions (figure 3). The resulting images were then cropped and thresholded, as depicted in figure (4). The resulting reconstruction

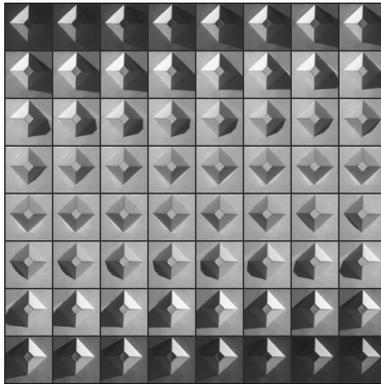


Fig. 3. Shadow Information Used in Pyramid Reconstruction

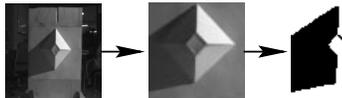


Fig. 4. Image Preparation for Reconstruction

was performed in 33 iterations, taking 347 seconds on a sparc 20 workstation.

5 Summary and Conclusion

Initial results show that the 3D Shape From Darkness algorithm is capable of estimating the structure of a scene using only the shadows within the scene. Furthermore, this reconstruction can be performed in a manner which uses far less data than the theoretically complete set, and which generates a full, three dimensional description of the surface.

This algorithm is well suited to terrain reconstruction from images taken of the earth from a geosynchronous satellite throughout a single day. It can also

be applied to the problem of environment mapping by allowing for two explorer robots, one of which holds a camera while the other moves with a light source. Questions for future and ongoing research are those of the stability of uniqueness properties of reconstructed surface and its envelope. Our results show that the robustness of the solution is acceptable.

The use of shadow data provides for a robust reconstruction which uses extremely compact data due to the boolean nature of shadow information, which makes this method a computationally efficient tool for terrain reconstruction.

The use of shadow data permits a new form of three-dimensional surface reconstruction that exploits the rich geometric information contained in cast shadows. This appears to permit computational surface reconstruction in new contexts. Several interesting problems remain to be resolved, including uniqueness of the computed solutions and smoothness constraints.

References

1. 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.
2. J. Kender D. Yang. Shape from shadows under error. In *Image Understanding Workshop 1993*, pages 1083–1090, Washington, D.C., August 1993.
3. P. Dupuis and J. Oliensis. Direct method for reconstructing shape from shading. pages 453–458.
4. P. Dupuis and J. Oliensis. Shape from shading: Provably convergent algorithms and uniqueness results. volume 2, pages 259–268.
5. 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.
6. M. Hatzitheodorou and J.R. Kender. An optimal algorithm for the derivation of shape from shadows.
7. Berthold Horn. *Robot Vision*. The MIT Press, Cambridge, Massachusetts, 1986.
8. 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.
9. E.M. Smith J.R. Kender. Shape from darkness: Deriving surface information from dynamic shadows. In *AIII*, pages 539–546, 1987.
10. 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.
11. 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.
12. Alex P. Pentland. Local shading analysis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6(2):170–187, March 1984.
13. Andrei Nikolaevich Tikhonov and Vasilii Iakovlevich Arsenin. *Solutions of ill-posed problems [Metody resheniia nekorrektnykh zadach]*. Halsted Press, New York, 1977.
14. James T. Todd and Ennio Mingolla. Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance*, 9(4):583–595, 1983.