

Shadow Graphs and Surface Reconstruction

Yizhou Yu and Johnny T. Chang

Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA
{yyz,jtchang}@uiuc.edu

Abstract. We present a method to solve shape-from-shadow using shadow graphs which give a new graph-based representation for shadow constraints. It can be shown that the shadow graph alone is enough to solve the shape-from-shadow problem from a dense set of images. Shadow graphs provide a simpler and more systematic approach to represent and integrate shadow constraints from multiple images. To recover shape from a sparse set of images, we propose a method for integrated shadow and shading constraints. Previous shape-from-shadow algorithms do not consider shading constraints while shape-from-shading usually assumes there is no shadow. Our method is based on collecting a set of images from a fixed viewpoint as a known light source changes its position. It first builds a shadow graph from shadow constraints from which an upper bound for each pixel can be derived if the height values of a small number of pixels are initialized properly. Finally, a constrained optimization procedure is designed to make the results from shape-from-shading consistent with the upper bounds derived from the shadow constraints. Our technique is demonstrated on both synthetic and real imagery.

Keywords. Surface Geometry, Shape-from-Shadow, Shadow Graph, Shading, Optimization

1 Introduction

In this paper, we consider the problem of shape-from-shadow and its integration with shape-from-shading. Shape-from-shadow tries to reconstruct a surface using multiple shadow images. It has a few advantages compared to other surface reconstruction techniques. For example, shadow constraints are insensitive to specular reflection and spatial variations of reflectance, and are able to impose long-range height constraints. The basic conclusion from previous work along this direction [20,13,9,4,14] says that with enough number of shadow images, the underlying surface can be recovered. However, the proposed algorithms for this problem are either complicated or heuristic. The major reason for this is that it was not clear how to effectively represent shadow constraints and integrate the information from multiple shadow images.

To clearly understand this problem, we introduce shadow graphs which can effectively represent and integrate shadow constraints from multiple images. We prove that the shadow graph alone is enough to solve the shape-from-shadow problem from a dense set of images. Simple operations on a shadow graph enable us to derive the structures of

the underlying surface. This approach is simpler and more systematic than the previous methods.

Usually most of the pixels in an image are not shadowed. However, shape-from-shadow neglects rest of the shading information in the input images. As we can see, shadow constraints are usually inequalities which are not as powerful as equalities. Consequently, it usually requires a dense set of input images to obtain good results. On the other hand, shape-from-shading [10] and photometric stereo [22] are effective approaches for a large class of surfaces including faces and sculptures. Both techniques use the pixelwise shading information to constrain surface normals, and do not allow shadows in the input images. They need an integration step to reconstruct a surface. This step tends to accumulate errors from pixel to pixel. Although theoretically they can uniquely recover the underlying surface, the final relative height values between distant points may not come out very accurately.

To take the advantages from both shape-from-shadow and shape-from-shading, we also develop a method of recovering shape from both shadow and shading constraints. A constrained optimization procedure is developed to make the results from shape-from-shading consistent with the upper bounds derived from shadow constraints.

1.1 Related Work

A few algorithms explicitly make use of shadow constraints [13,9,4,21]. Most of them belong to shape-from-shadow(darkness) algorithms. Some shape-from-shadow algorithms [13] use a shadowgram as an intermediate representation which is derived from a dense set of lighting directions. [9] assumes the underlying surface has a spline representation because shadows only provide a relatively sparse set of constraints. The number of unknown coefficients in the spline model is designed to scale with the number of shadow constraints. [4] introduces a shape-from-shadow algorithm using relaxation. A pair of upper-bound and lower-bound surfaces are constructed by updating the height values at pixels with violated shadow constraints. Like shape-from-shading, shape-from-shadow can recover unknown lighting directions as well [14].

The computation of shape-from-shading has been typically characterized as that of finding surface orientation from one single image followed by a step that converts the orientation information into height under integrability constraints. The surface is usually assumed to be Lambertian. [15] introduces an algorithm that allows direct computation of height from shading. Since the unknowns directly represent pixelwise height values, this approach can be more naturally integrated with other methods of recovering shape, such as stereo and shape-from-shadow. [5] presents provably convergent algorithms for this problem.

Photometric stereo [22] can usually obtain better results than shape-from-shading because of the use of multiple input images. This approach has been generalized to recover shape for metallic and hybrid surfaces with both diffuse and specular reflection [11,17]. The lighting direction for each image is usually assumed to be known. However, both surface shape and lighting directions can be recovered simultaneously from SVD decomposition up to a bas-relief transformation [2,1]. Shadowed pixels in each image can be masked out in the process with the hope that there are still enough images covering them [12,7].

[21] considers recovery of shape from shading under a uniform hemispherical light source. Partial shadowing is taken into account because only a part of the light source is visible from every surface point. Interreflections are also considered in the algorithm presented in [18].

2 Shadow Graphs

We consider recovering terrain-like height fields in this paper. For the convenience of discrete representation based on pixels, a height field is assumed to be a piecewise constant function with every pixel corresponding to a piece with constant height. Every piece of the height field is represented by the point corresponding to the center of the pixel. We also assume that the distance between the camera and the surface is large enough so that the orthographic projection model is accurate. Let us first check what




Fig. 1. 2D schematic of shadowed and nonshadowed regions on a terrain-like surface. L is the parallel lighting direction. x_0 is an occluder, x_1 is on the shadow boundary caused by x_0 , and x_2 is a non-shadowed point.

kind of constraints are available from images with shadows. Let $h(x)$ be a height field defined on a planar domain D with a finite area in the image plane and L be the lighting direction pointing downwards with a tilt angle $\theta (< 90^\circ)$. The normal orientation of this height field is denoted as $\mathbf{n}(x)$. The boundary curve of domain D is Ω . The projected vector of L in the domain D is L_p . Let \mathbf{x}_i and \mathbf{x}_j be two arbitrary 2D points in D . The line segment between them is denoted as a vector interval $[\mathbf{x}_i, \mathbf{x}_j]$ for convenience. Based on whether a point on the height field is in shadow or not under lighting direction L , there are two different sets of constraints (Fig. 1).

- If any point on the line segment $[\mathbf{x}_0, \mathbf{x}_1]$ is in shadow, the points at \mathbf{x}_0 and \mathbf{x}_1 are the delimiting points of this shadow segment, and the point at \mathbf{x}_0 is the occluding point generating this shadow segment, we have the following *shadow* constraints.

$$h(\mathbf{x}) \leq h(\mathbf{x}_0) - \frac{\|\mathbf{x} - \mathbf{x}_0\|_2}{\tan \theta}, \forall \mathbf{x} \in [\mathbf{x}_0, \mathbf{x}_1]; \quad (1)$$

$$h(\mathbf{x}_1) = h(\mathbf{x}_0) - \frac{\|\mathbf{x}_1 - \mathbf{x}_0\|_2}{\tan \theta}; \quad (2)$$

$$\mathbf{L} \cdot \mathbf{n}(x_0) = 0 \quad (3)$$

where the last equation means the lighting vector \mathbf{L} falls inside the tangential plane at \mathbf{x}_0 if the original continuous height field is locally differentiable at \mathbf{x}_0 .

– If the point at \mathbf{x}_2 is not in shadow, we have the following *antishadow* constraints.

$$h(\mathbf{x}) \leq h(\mathbf{x}_2) + \frac{\|\mathbf{x} - \mathbf{x}_2\|_2}{\tan \theta}, \forall \mathbf{x} \in [\mathbf{x}_b, \mathbf{x}_2] \quad (4)$$

where $\mathbf{x}_b \in \Omega$ and the line segment $[\mathbf{x}_b, \mathbf{x}_2]$ is in the same direction as \mathbf{L}_p .

Let us first focus on how to represent the inequality constraints (1) and (4) in a graph.

Definition 1. A shadow graph is a weighted directed graph $G = (V, E, W)$ where the set of nodes V is the set of points defined on domain D , an edge $e = (\mathbf{x}_i, \mathbf{x}_j) \in E$ indicates $h(\mathbf{x}_j)$ is dependent on $h(\mathbf{x}_i)$ and $h(\mathbf{x}_i) - h(\mathbf{x}_j) \geq W(e)$ where the edge weight $W(e)$ can be any real number.

A shadow graph can be induced from an image of the height field under an arbitrary lighting direction \mathbf{L} . Shadowed pixels can be detected from the image, and an occluder can be located for each continuous shadow segment with the knowledge of the lighting direction. For example, if $[\mathbf{x}_i, \mathbf{x}_j]$ is a shadow segment and the vector from \mathbf{x}_i to \mathbf{x}_j is in the direction of the projected lighting direction \mathbf{L}_p , the point at \mathbf{x}_i is the occluder of all the points in $[\mathbf{x}_i, \mathbf{x}_j]$. There should be an edge $(\mathbf{x}_i, \mathbf{x}_j)$ with weight $\frac{\|\mathbf{x}_j - \mathbf{x}_i\|_2}{\tan \theta}$ in the induced graph for all $\mathbf{x} \in (\mathbf{x}_i, \mathbf{x}_j]$. This graph basically encodes the shadow constraints available from the image. All the edge weights in this graph should be positive. However, this graph can have negative weights if the additional antishadow constraints in Eq. (4) are represented as well.

Suppose we have multiple images of the height field under a set of lighting directions $\{\mathbf{L}_k\}_{k=1}^m$. Each of the images has its own shadow graph. Finally, the edges from all of these individual graphs can be accumulated into one graph that is corresponding to all the images. Note that this graph does not have the specific lighting information, which is not particularly important because all the constraints essential to the height field are kept there.

Proposition 1. A shadow graph with positive weights is a directed acyclic graph.

Proof Suppose there is a circular path in the graph and a node v is on the path. Since all the arcs on this path have positive weights, it is easy to conclude that $h(v) > h(v)$ by starting from v , going through this path and back to v . A contradiction.

When dealing with real images with noise, shadow detection cannot be expected to be error free. Inaccurate shadow segmentations may result in cycles in the induced shadow graphs. Since cycles can lead to the above contradiction, we must convert a cyclic graph into an acyclic one by removing some of the edges in the graph. Since we would like to make the least amount of change to the graph, a sensible criterion for an optimal conversion is that the total accumulated weight for the removed edges should be minimized. However, graph conversion under this criterion is NP-hard [8]. To obtain

an efficient solution for this problem, we adopt the permutation-based approximation algorithm in [8] which tends to remove more edges than necessary. After applying this algorithm, for each of the removed edges, we still run a depth-first search to check whether the graph is still acyclic after the edge is inserted back into the graph. These two steps together lead to a polynomial time approximation that removes the least number of edges.

Definition 2. *The transitive closure of a shadow graph G is defined to be a new graph $G^c = (V, E^c, W^c)$ on the same set of nodes such that $(\mathbf{x}_i, \mathbf{x}_j) \in E^c$ as long as there is a path from \mathbf{x}_i to \mathbf{x}_j in G , and $W((\mathbf{x}_i, \mathbf{x}_j))$ is set to be the maximum accumulated weight among the paths from \mathbf{x}_i to \mathbf{x}_j .*

There are a set of nodes $V_t \subset V$ in G^c that do not have any incident edges with positive weights, which means they are not shadowed by any other points in any of the images. The highest point(s) of the height field surely belongs to this set because there is no other points which can occlude it(them) from the light sources. The absolute height values of the nodes in V_t are unrecoverable from shadow constraints. However, if we can recover their height values from other approaches such as stereo processing, the information embedded in G^c can be used for obtaining an upper bound of the height at any point in $V - V_t$. The set of edges in G^c connecting V_t and $V - V_t$ becomes the most important for this purpose. Suppose there is a node $v \in V - V_t$ and a set of associated edges $E_v \subset E^c$ such that if an edge $e = (v_t, v) \in E_v$, $v_t \in V_t$. The upper bound of the height at the point corresponding to node v can be obtained from

$$U(h(v)) = \min_{(v_t, v) \in E_v} (h(v_t) - W((v_t, v))). \quad (5)$$

Let us examine the asymptotic behavior of this upper bound when we increase the number of input images with lighting directions covering the whole lighting hemisphere. The set V_t will shrink and approach its limit which is the set of the highest points of the height field. Otherwise, assume there is a pair of nodes $v_1, v_2 \in V_t$ and $h(v_2) < h(v_1)$. We can always design a lighting direction from which the point corresponding to v_1 shadows the point corresponding to v_2 , which means $v_2 \notin V_t$, a contradiction. Since eventually V_t only has nodes at the same height, we do not need to seek their relative height through other reconstruction techniques. Our interest should be focused on the relative height of other points compared to the highest points whose height can always be set to zero.

Proposition 2. *Eq. (5) gives an upper bound for the height at any node in $V - V_t$ provided that the estimation of the height for the nodes in V_t is accurate. With an increasing number of input images with lighting directions covering the whole lighting hemisphere, Eq. (5) converges asymptotically to the correct relative height, with respect to the highest points in V_t , at any point in $V - V_t$.*

Proof The first part is obvious. The second part can be proved by induction. Since we only have a finite number of points according to our surface model, we can sort the points in decreasing order of their height. The highest points in the sorted list are assumed to be at height zero. Suppose the point at \mathbf{x}_m is the k -th element in the sorted list and the height

of its $k - 1$ preceding elements can be recovered to an arbitrary precision independently of the height of the rest of the elements in the list. Now we show that the height of the point at \mathbf{x}_m can also be recovered to an arbitrary precision independently of the height of its following elements in the list. Note that all the surface points are lit if we have a vertical lighting direction. If we increase the tilt angle of the light, the point at \mathbf{x}_m will certainly be shadowed since it is not one of the highest points. Given a certain density of the lighting direction, there exist two adjacent directions \mathbf{L}_r and \mathbf{L}_l such that this point at \mathbf{x}_m is non-shadowed when the light is at \mathbf{L}_r and becomes shadowed when the light moves to \mathbf{L}_l . An upper bound for this point can be obtained from \mathbf{L}_l and an occluder at \mathbf{x}_o whose height is recovered to an arbitrary precision. When we increase the density of the lighting direction, the difference between \mathbf{L}_r and \mathbf{L}_l becomes arbitrarily small and the upper bound for the point at \mathbf{x}_m also becomes arbitrarily close to its true height.

Shadowgrams introduced in [13,4] also have the capability to recover correct surface geometry. But they are more complicated than shadow graphs because they explicitly keep lighting directions in the representation.

It is clear that the antishadow constraints can be derived from the shadow constraints if we have a very dense set of images since the height field itself can be recovered from the shadow constraints alone according to the above Proposition. However, if we only have a sparse set of images, this is not necessarily true. Representing these antishadow constraints in a shadow graph usually can provide additional information. According to Eq. (4), antishadow constraints transform to additional edges with negative weights. Cycles can appear in the resulting graph. However, the accumulated weight of any cycle can not be positive according to the following Proposition.

Proposition 3. *The accumulated weight of a circular path in a shadow graph must be either zero or negative.*

Proof Suppose $\mathbf{x}_0, x_1, \dots, x_n \in V$ are consecutive nodes of a circular path, i.e. $(\mathbf{x}_i, x_{i+1}) \in E (i = 0, \dots, n - 1)$ and $(\mathbf{x}_n, x_0) \in E$. From the definition of a shadow graph, $h(\mathbf{x}_i) - h(x_{i+1}) \geq W((\mathbf{x}_i, x_{i+1})) (i = 0, \dots, n - 1)$ and $h(\mathbf{x}_n) - h(x_0) \geq W((\mathbf{x}_n, x_0))$. Therefore,

$$\begin{aligned} & \sum_{i=0}^{n-1} W((\mathbf{x}_i, x_{i+1})) + W((\mathbf{x}_n, x_0)) \\ & \leq \sum_{i=0}^{n-1} (h(\mathbf{x}_i) - h(x_{i+1})) + (h(\mathbf{x}_n) - h(x_0)) = 0. \end{aligned}$$

The transitive closure of a shadow graph G with cycles is still well-defined because negative cycles do not interfere with the objective to seek paths with maximum accumulated weights according to the definition. The resulting graph G^c can still be used for obtaining an upper bound of the height for any point in $V - V_t$. Since there may be negative edges pointing from nodes in $V - V_t$ to some nodes in V_t , these edges can be

used for obtaining a lower bound for some nodes in $V - V_t$. Since it is not guaranteed that there is an edge from each node in $V - V_t$ to some node in V_t given a sparse set of images, we can only obtain lower bounds for a subset of nodes in $V - V_t$. And these lower bounds may appear useful in combination with other surface reconstruction techniques.

3 Integrated Shadowing and Shading Constraints

Given a sparse set of images with known lighting directions, we would like to recover shape using both shadow and shading constraints. As we have seen, shadows impose explicit constraints over surface height values, but they are usually not sufficient if applied alone. On the other hand, shading information imposes constraints over normal orientation. We are going to explore two options for integrating shadow constraints with shading information.

3.1 Enforcing Shadowing Constraints with Penalty Terms

Since shape-from-shading is not the focus of this paper, we adopt the direct height from shading algorithm in [15] as the base for solving shading constraints. Since this technique computes a height field directly rather than through surface normals, it is relatively easy to incorporate shadow constraints and enforce surface upper/lower bounds from the previous section. The shape-from-shading problem is formulated to minimize the following cost function in [15]:

$$E_1 = \sum_{i,j} [\alpha(\rho R(p_{ij}, q_{ij}) - I(i, j))^2 + \lambda(u_{ij}^2 + v_{ij}^2)] \quad (6)$$

where ρ is the surface albedo, I is the observed image intensity, $p_{ij}, q_{ij}, u_{ij}, v_{ij}$ are the symmetric first and second finite differences of the surface height field $\{h_{ij}\}$, α and λ are two constant coefficients, and R is the Lambertian reflectance model. The first term in Eq. (6) corresponds to the photometric error term. And the second is a regularization term on the smoothness of the surface.

This formulation can be easily generalized to accommodate multiple input images and shadow masks as follows.

$$E_2 = \alpha \sum_k \sum_{i,j} l_{ij}^k (\rho_{ij} R^k(p_{ij}, q_{ij}) - I^k(i, j))^2 + \lambda \sum_{i,j} (u_{ij}^2 + v_{ij}^2) \quad (7)$$

where $I^k(i, j)$ represents the k -th input image with corresponding reflectance map R^k , l_{ij}^k is a binary shadow mask indicating whether pixel (i, j) in the k -th image is lit by the light source or not, and ρ_{ij} is the unknown pixelwise surface albedo. This treatment is similar to photometric stereo, but solves the height field directly instead. With multiple images, the regularization term becomes much less important, and can be set close to zero. However, it may still have some effects at pixels that are lit in less than three different images.

To further incorporate the constraints in Eq. (1) and (4) into the above formulation, we notice that the constraints have the same form which looks like

$$h_{ij} - h_{i'j'} \geq d_{ij i'j'} \quad (8)$$

To enforce this kind of inequalities in a gradient-based minimization method, a differentiable half-sided parabola is adopted as a penalty function.

$$S(i, j, i', j') = \begin{cases} (h_{ij} - h_{i'j'} - d_{ij i'j'})^2, & \text{if } h_{ij} - h_{i'j'} \\ & \leq d_{ij i'j'}; \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

The penalty functions for all the inequalities and equalities can be inserted as additional terms into Eq. (7). The new cost function for surface reconstruction is given as follows.

$$E_3 = \alpha \sum_k \sum_{i,j} l_{ij}^k (\rho_{ij} R^k(p_{ij}, q_{ij}) - I^k(i, j))^2 + \lambda \sum_{i,j} (u_{ij}^2 + v_{ij}^2) + \beta \sum_k (\sum_{m_k} S^k(i_{m_k}, j_{m_k}, i'_{m_k}, j'_{m_k}) + T_k) \quad (10)$$

where m_k is the index of the inequality constraints from the k -th image, $S^k(i_{m_k}, j_{m_k}, i'_{m_k}, j'_{m_k})$ represents the actual penalty terms contributed by the k -th image, and T_k represents the collection of penalty terms for the equality constraints associated with shadows, such as those in Eq. (2) and Eq. (3). In our experiments, we use iterative minimization algorithms and set $\alpha = 1 - \lambda$, $\beta = 1$. λ is initialized to 0.1 and divided by a constant factor after each iteration.

All the above three cost functions can be minimized by the standard conjugate gradient algorithm [19].

3.2 Enforcing Upper and Lower Bounds

In the above formulation, shadow constraints are enforced as soft constraints by using penalty terms in the original shape-from-shading algorithm. It is not guaranteed that all constraints are satisfied. Sometimes, it is more desirable to consider shadow constraints as hard constraints since they are less sensitive to specular reflection and albedo variations, and to consider shading constraints as soft ones since a little bit deviation in shading is not very noticeable. The upper and lower bounds discussed in Section 2 can serve this purpose and can be estimated as follows. (Note that the height of the nodes in the set V_t is unknown at the beginning, and they can be estimated from a solution of the height field from Section 3.1.)

1. Obtain an initial estimation of the height value for each point by minimizing Eq. (10);




Fig. 2. (a) Some parts of the height field recovered from minimization may exceed the upper bound; (b) We need to globally adjust the initial height field to maintain its original smoothness instead of simply clipping it against the upper bound.

2. Adjust the initial height values of the nodes in V_t to satisfy all the antishadow constraints among them as in the following convergent procedure;
 - a) fix the height of the highest point in V_t ;
 - b) loop through the rest of the points and check whether the considered point is in the shadow of some other point in V_t because of the violation of a antishadow constraint; if so, raise the considered point to the minimum height that can eliminate the violation.
3. Calculate the upper and lower bounds for nodes in $V - V_t$ from the transitive closure G^c .

To enforce the upper and lower bounds, our complete algorithm still takes an initial solution of the height field from minimizing Eq. (10). However, there are multiple possibilities to improve this initial solution:

1. For each point, if it is higher than its upper bound, push it down to the upper bound; if it is lower than its lower bound, raise it to the lower bound.
2. Use a constrained optimization algorithm such as sequential quadratic programming to enforce the upper and lower bounds.
3. Fix a subset of the adjusted points from the first step and minimize Eq. (10) with those fixed points as additional boundary conditions; alternate adjustment and minimization (with a few additional fixed points every iteration) until all the bounds are satisfied.

The first scheme chooses to satisfy all the hard constraints by using brute force and ignoring all the shading constraints, therefore tends to have unnatural discontinuities at those adjusted places. The second scheme chooses to apply some constrained optimization algorithm to automatically and iteratively adjust the heights so that the bounds are satisfied at the end. Unfortunately, constrained optimization algorithms such as sequential quadratic programming (SQP) are usually computationally expensive on high-dimensional data such as images. For example, the SQP software package [6] we tried took two hours to finish one iteration on 64x64 images on a Pentium III 800MHz workstation. The last scheme chooses to adapt unconstrained optimization algorithms so that they allow a part of the variables to be fixed. To achieve that, we can simply set the corresponding derivatives to be zero. We fix a few additional points within their bounds before unconstrained minimization takes place in every iteration, therefore can satisfy all the bounds in a finite number of iterations since we only try to recover height values at a finite number of points (pixels). An intuitive illustration is given in Fig. 2.

In practice, we chose the last scheme with some additional details. After initialization, the height values of the nodes in V_t , and the upper and lower bounds are fixed in all iterations. In every iteration, we subtract the upper bounds from the current estimation of the height field to obtain a difference field. Then the set of local maxima in the difference field are located. Those points corresponding to the local maxima are lowered to their corresponding upper bounds and fixed thereafter. The same procedure is repeated for lower bounds before the unconstrained minimization in Eq. (10) takes place once again with the newly fixed points as additional boundary conditions. We hope that the shading constraints solved during minimization can automatically adjust the neighborhoods of those fixed points so that there will be much less violated bounds in the next iteration. This

can also avoid having many unnatural discontinuities since the minimization procedure serves as a smoothing operator by considering all constraints simultaneously.

3.3 Experiments

We have tested the algorithm using integrated shadow and shading constraints on both synthetic and real imagery.




Fig. 3. (a) Input images for the pyramid scene. The tilt angle of the lighting directions in the top row is 45 degrees, the bottom row 60 degrees. (b) A synthetic image of the recovered height field illuminated from the same lighting direction as in the first input images; (c) A synthetic image of the recovered height field illuminated from a novel lighting direction.

Synthetic data. Eight synthetic images were generated as input for each of the three representative datasets we chose. Four of them were lit from a tilt angle of 45 degrees and the others were lit from 60 degrees to create images with significant amount of shadow. We also generated two images for each example from the recovered height field. The first image is lit from the same lighting direction as the first input image to verify both shadowed and non-shadowed regions. The second image is lit from a novel lighting direction which is different from the ones for the input images to show that the recovered height fields can be useful for creating images with correct appearance




Fig. 4. (a) Input images for the plaster material sample. It is lit from the same set of lighting directions as in Fig. 3. (b) A synthetic image of the recovered height field illuminated from the same lighting direction as in the first input images; (c) A synthetic image of the recovered height field illuminated from a novel lighting direction.

Table 1. Comparison of the three approaches on the three datasets: i) minimizing E_2 in Eq.(7), ii) minimizing E_3 in Eq.(10), iii) enforcing bounds as in Section 3.2. The top table shows the RMS errors of the recovered height fields using noise free input images, and the bottom one shows the RMS errors using images with 5% noise. All numbers are given in the unit of a pixel.

0% noise	E_2	E_3	E_3 +Bounds
Pyramids	3.6579	2.2424	2.1984
Plaster	1.9344	1.4548	1.4210
Face	4.4164	3.3335	3.3399

5% noise	E_2	E_3	E_3 +Bounds
Pyramids	3.7621	2.2675	2.2100
Plaster	1.9400	1.4089	1.3959
Face	4.4522	3.4298	3.4159

from novel lighting conditions. We also compared the recovered height fields with the ground truth to obtain error measurements which are shown in Table 1. In our examples, most points are lit from at least one lighting direction. The height field can be recovered from shape-from-shading or photometric stereo alone. However, the additional shadow constraints can definitely improve the accuracy of the results because shading-based techniques can introduce accumulated errors from pixel to pixel while shadow constraints are very good at enforcing long-range relative height constraints.

The first dataset is an artificial scene with four pyramids shown in Fig. 3(a). The pyramids have different height and orientation. The two synthetic images from the recovered




Fig. 5. (a) Input images for the face model. It is lit from the same set of lighting directions as in Fig. 3. (b) A synthetic image of the recovered height field illuminated from the same lighting direction as in the first input images; (c) A synthetic image of the recovered height field illuminated from a novel lighting direction.




Fig. 6. Comparison of the cross sections of four height fields: the ground truth is shown as 'original'; the one from minimizing Eq. (7) is shown as 'sfs'; the one from minimizing Eq. (10) is shown as 'sfs-sc'; and the one from enforcing bounds is shown as 'sfs-upper'. (a) Cross sections for the pyramid scene; (b) cross sections for the plaster sample.

height field are shown in Fig. 3(b)&(c). The second dataset is a previously recovered height field of a real plaster sample using the approach presented in [16]. This height

field serves as the ground truth to test the algorithm in this paper although we do not know the accuracy of this dataset. The input images are shown in Fig. 4(a) and the synthetic images from the height field recovered by the current algorithm are shown in Fig. 4(b)&(c). The third dataset is a face model shown in Fig. 5(a). And Fig. 5(b)&(c) give the images generated from the recovered face. In this example, the background plane is pushed down along the shadow boundaries in some of the input images to satisfy the shadow constraints. This is because shape-from-shading related techniques are better at estimating normal orientation than at estimating height values, and generated an inaccurate initial solution for our algorithm. A similar situation was also shown in the pyramid scene. Nevertheless, our algorithm still managed to enforce the shadow constraints and make the generated images look similar to the input ones.

Fig. 6 shows two comparisons of the cross sections. In each of the comparisons, there are four curves including the ground truth, the curve from minimizing Eq. (7), the curve from minimizing Eq. (10) and the curve from enforcing the upper bounds. The results from minimizing Eq. (7) are not as good as the other two versions because it does not consider shadow constraints.

Real data. We also did test on a real dataset. Three 128x128 images of a concrete sample from the CURET database [3] were used as the input to our final algorithm. They have various amount of shadow (Fig. 7(a)-(c)). We use 15 as the intensity threshold to detect shadowed pixels. The lighting directions of the input images are actually coplanar. Traditional photometric stereo would have problem to recover the height field. However, our algorithm successfully recovered it since it exploits shadow constraints and a regularization term. Minimizing E_3 in Eq. (10) took 5 minutes on a Pentium III 800MHz processor, and the iterative procedure for enforcing bounds took another half an hour. Synthetic images were generated from the recovered height field. The recovered dataset was illuminated from both original lighting directions (Fig. 7(d)-(f) of the input images and novel lighting directions (Fig. 7(g)-(h)).

4 Summary

We presented the concept of shadow graphs and proved that the shadow graph alone is enough to solve the shape-from-shadow problem from a dense set of images. We also developed a method of recovering shape from both shadow and shading constraints. A constrained optimization procedure has been developed to make the results from shape-from-shading consistent with the upper bounds derived from shadow constraints.

Future work includes more robust techniques that allow inaccurate shadow segmentation and the simultaneous recovery of shape and lighting directions from both shading and shadows.

Acknowledgment. This work was supported by National Science Foundation CAREER Award CCR-0132970.




Fig. 7. (a)-(c) Real images of a concrete sample; (d)-(f) synthetic images of the recovered height field illuminated from original lighting directions; (g)-(h) synthetic images of the recovered height field illuminated from two novel lighting directions.

References

1. P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible illumination conditions? *Int. Journal Comp. Vision*, 28(3):1–16, 1998.
2. P. Belhumeur, D. Kriegman, and A. Yuille. The bas-relief ambiguity. In *IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 1040–1046, 1997.
3. K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real-world surfaces. In *Proceedings of CVPR*, pages 151–157, 1997.
4. M. Daum and G. Dudek. On 3-d surface reconstruction using shape from shadows. In *IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 461–468, 1998.
5. P. Dupuis and J. Oliensis. Shape from shading: Provably convergent algorithms and uniqueness results. In *Computer Vision-ECCV 94*, pages 259–268, 1994.
6. Fsqp software. <http://gachinese.com/aemdesign/FSQPframe.htm>. Originally developed at the Institute for Systems Research, University of Maryland.

7. A. Georgiades, P. Belhumeur, and D. Kriegman. Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting. In *IEEE Workshop on Multi-View Modeling and Analysis of Visual Scenes*, pages 47–54, 1999.
8. R. Hassin and S. Rubinstein. Approximations for the maximum acyclic subgraph problem. *Information Processing Letters*, 51:133–140, 1994.
9. M. Hatzitheodorou. The derivation of 3-d surface shape from shadows. In *Proc. Image Understanding Workshop*, pages 1012–1020, 1989.
10. B.K.P. Horn and M.J. Brooks. The variational approach to shape from shading. *Computer Vision, Graphics & Image Processing*, 33:174–208, 1986.
11. K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Trans. Patt. Anal. Mach. Intel.*, 3(6):661–669, 1981.
12. D. Jacobs. Linear fitting with missing data: Applications to structure from motion and characterizing intensity images. In *IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 206–212, 1997.
13. J. Kender and E. Smith. Shape from darkness. In *Int. Conf. on Computer Vision*, pages 539–546, 1987.
14. D.J. Kriegman and P.N. Belhumeur. What shadows reveal about object structure. In *Computer Vision-ECCV 98*, 1998.
15. Y.G. Leclerc and A.F. Bobick. The direct computation of height from shading. In *Proc. of IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 552–558, 1991.
16. X. Liu, Y. Yu, and H.-Y. Shum. Synthesizing bidirectional texture functions for real-world surfaces. In *Proceedings of SIGGRAPH*, pages 97–106, 2001.
17. S.K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Trans. Robotics and Automation*, 6(4):418–431, 1990.
18. S.K. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. *International Journal of Computer Vision*, 6(3):2–11, 1991.
19. W.H. Press, B.P. Flannery, S.A. Teukolsky, and W.T. Vetterling. *Numerical Recipes in C*. Cambridge Univ. Press, New York, 1988.
20. S.A. Shafer. *Shadows and silhouettes in computer vision*. Kluwer Academic Publishers, 1985.
21. A. J. Stewart and M.S. Langer. Towards accurate recovery of shape from shading under diffuse lighting. *IEEE Patt. Anal. Mach. Intel.*, 19(9):1020–1025, 1997.
22. R.J. Woodham. Photometric method for determining surface orientation from multiple images. In B.K.P. Horn and M.J. Brooks, editors, *Shape from Shading*, pages 513–532. MIT Press, 1989.