

What shadows reveal about object structure

David J. Kriegman

*Beckman Institute, Department of Computer Science, University of Illinois, Urbana-Champaign,
Urbana, Illinois 61801*

Peter N. Belhumeur

*Center for Computational Vision and Control, Department of Electrical Engineering and Computer Science,
Yale University, New Haven, Connecticut 06520-8267*

Received May 5, 2000; revised manuscript received January 2, 2001; accepted January 3, 2001

In a scene observed from a fixed viewpoint, the set of shadow boundaries in an image changes as a point light source (nearby or at infinity) assumes different locations. We show that for any finite set of point light sources illuminating an object viewed under either orthographic or perspective projection, there is an equivalence class of object shapes having the same set of shadows. Members of this equivalence class differ by a four-parameter family of projective transformations, and the shadows of a transformed object are identical when the same transformation is applied to the light source locations. Under orthographic projection, this family is the generalized bas-relief (GBR) transformation, and we show that the GBR transformation is the only family of transformations of an object's shape for which the complete set of imaged shadows is identical. Finally, we show that given multiple images under differing and unknown light source directions, it is possible to reconstruct both an object's surface and the light source locations up to this family of transformations from the shadows alone. © 2001 Optical Society of America

OCIS codes: 330.7310, 150.2950, 150.0150.

1. INTRODUCTION

The shadows produced by an object depend on both the object's shape and the position of the light sources by which it is illuminated. If the light source positions remain fixed but the shape of the object changes, then shadows change as well. If, however, the object is flattened in the direction that the observer is gazing (as is done in the case of classical relief sculpture), then for any flattening of the surface relief, there is a corresponding change in the light source direction such that the shadows appear the same. This is not restricted to classical reliefs but, as we will show below, applies equally to a greater set of projective transformations.

More specifically, when an object is viewed from a fixed viewpoint, there is a four-parameter family of projective transformations of the object's structure and the light source locations such that the images of the shadows remain the same. It follows, then, that when light source positions are unknown, one cannot determine the Euclidean structure of an object from its shadows alone. Yet in all past work on reconstruction from shadows, it is explicitly assumed that the directions or the locations of the light sources are *known*. An implication of these results is that two objects differing by these transformations cannot be recognized solely from their shadow lines.

In early work, Waltz considered labelings of shadow edges in line drawing interpretation.¹ Subsequently, Shafer and Kanade showed how geometric constraints on surface orientation could be obtained from labeled line drawings by using shadow and surface outlines under orthographic projection.² The entry-exit method was developed to segment and label shadow boundaries by using

information about the projection onto the image plane of the light source direction.³ Kender and his colleagues have undertaken a series of studies pertaining to metric reconstruction of surfaces from the shadows in multiple images of an object in fixed pose when the light source direction is known,⁴⁻⁶ and this problem was recently reexamined by Daum and Dudek.⁷ Shadows have also been used in the interpretation of aerial images, particularly to locate and reconstruct buildings when the sun direction is known.⁸⁻¹¹

Here we consider shadows on unknown objects produced by light sources whose directions are also unknown. In Section 2, we show that seen from a fixed viewpoint under perspective projection, two surfaces produce the same shadows if they differ by a particular projective transformation—which we call the generalized perspective bas-relief (GPBR) transformation. See Fig. 1 for an example of this transformation. This result holds for any number of proximal or distant point light sources. Furthermore, under conditions where perspective can be approximated by orthographic projection, this transformation is the generalized bas-relief (GBR) transformation.¹² As will be shown in Section 3, the GBR transformation is unique in that any two smooth surfaces that produce the same shadows must differ by a GBR.

In Section 4, we propose an algorithm for reconstructing, from the attached shadow boundaries, the structure of an object up to a GBR transformation. The algorithm assumes that the object is viewed orthographically and that it is illuminated by a set of point light sources at infinity. We do not propose this algorithm with the belief that its present form has great applicability, but rather




Fig. 1. Illustration of the effect of applying a generalized perspective bas-relief (GPBR) transformation to a scene composed of a teapot resting on a supporting plane. The first image shows the original teapot. The second image shows the teapot after having undergone a GPBR transformation $(a_1, a_2, a_3, a_4) = (0.05, 0.05, 0.05, 1)$ with respect to the viewpoint used to generate the first image. Note that the attached and cast shadows as well as the occluding contour are identical in the first two images. The third image shows the original teapot from a second viewpoint. The fourth image reveals the nature of the GPBR transformation, showing the transformed teapot from the second viewpoint.

we give it to demonstrate that under ideal conditions, information from shadows alone is sufficient to determine the structure of the object up to a GBR transformation. Some of the results presented in this paper originally appeared in Ref. 13.

2. SHADOWING AMBIGUITY

Let S be the set of all point light sources illuminating one object and S' be the set of all light sources illuminating a second object. We say that these two objects are shadow equivalent if there exists a nonempty, open subset $\mathcal{L} \subset S$ such that for every light source $\mathbf{s} \in \mathcal{L}$ illuminating one object, there exists a light source $\mathbf{s}' \in S'$ illuminating the second object such that the shadowing in both images is identical. Let us further define two objects as being strongly shadow equivalent if for every light source illuminating one object, there exists a source illuminating the second object such that shadowing is identical, i.e., shadow equivalent with $\mathcal{L} = S$. In this section, we will show that two objects are shadow equivalent if they differ by a particular set of projective transformations.

Consider a camera-centered coordinate system whose origin is at the focal point, whose x and y axes span the image plane, and whose z axis points in the direction of the optical axis. Let a smooth surface f be defined with respect to this coordinate system and lie in the half-space $z > 0$. Since the surface is smooth, the surface normal $\mathbf{n}(\mathbf{p})$ is defined at all points $\mathbf{p} \in f$.

We model illumination as a finite set of point light sources, located nearby or at infinity. Note that this is a restriction of the lighting model presented by Langer and Zucker,¹⁴ which permits light sources whose intensity is a function of direction as well as extended sources such as line or area sources. In this paper, we will represent surfaces, light sources, and the camera center as lying in either a two- or a three-dimensional real projective space ($\mathbb{RP}^2$ or $\mathbb{RP}^3$). (For a concise treatment of real projective spaces, see Ref. 15.) This allows a unified treatment of both point light sources that are nearby (proximal) or distant (at infinity) and camera models that use perspective or orthographic projection.

When a point light source is proximal, its coordinates can be expressed as $\mathbf{s} = (s_x, s_y, s_z)^T$. Note that throughout this paper, Euclidean coordinates are written in the roman alphabet, while projective coordinates are denoted

by the corresponding greek character (e.g., the homogeneous coordinates of $\mathbf{s}$ and $\mathbf{p}$ are, respectively, given by σ and π). When a point light source is at infinity, all light rays are parallel, and so one is concerned with the direction of the light source. The direction can be represented as a unit vector in $\mathbb{R}^3$ or as a point on an illumination sphere $\mathbf{s} \in S^2$. In projective coordinates, the fourth homogeneous coordinate of a point at infinity is zero, and so the light source can be expressed as $\sigma = (s_x, s_y, s_z, 0)^T$. (Note that when the light source at infinity is represented in projective coordinates, the antipodal points from S^2 are equated.)

For a single point source $\sigma \in \mathbb{RP}^3$, let us define the set of light rays emanating from σ as the lines in $\mathbb{RP}^3$ passing through σ . For any other point $\pi \in \mathbb{RP}^3$ with $\pi \neq \sigma$, there is a single light ray passing through π . Naturally, it is the intersection of the light rays with the surface f that determines the shadows. Following Refs. 16 and 17, we differentiate between two types of shadows: attached shadows and cast shadows. We also consider both local and global conditions for these to occur and arrive at the following four definitions (see Figs. 2 and 3).

Definition 1: Local attached shadow. A surface point $\mathbf{p}$ is called a local attached shadow boundary point if the ray from light source $\mathbf{s}$ lies in the tangent plane to the surface $\mathbf{p}$. Algebraically, this condition can be expressed as $\mathbf{n}(\mathbf{p}) \cdot (\mathbf{p} - \mathbf{s}) = 0$ for a nearby light source (here $\mathbf{p}$ and $\mathbf{s}$ denote Euclidean coordinates) and as $\mathbf{n}(\mathbf{p}) \cdot \mathbf{s} = 0$ for a distant light source (here $\mathbf{s}$ denotes the direction of the light source).

Definition 2: Global attached shadow. A surface point $\mathbf{p}$ is called a global attached shadow boundary point if it satisfies both the local condition stated above and the light ray does not intersect the surface between $\mathbf{p}$ and $\mathbf{s}$, i.e., the light source is not occluded at $\mathbf{p}$.

Definition 3: Local cast shadow. A surface point $\mathbf{p}$ is called a local cast shadow boundary point if the light ray from $\mathbf{s}$ through $\mathbf{p}$ grazes the surface at some other point $\mathbf{q}$ (i.e., $\mathbf{q}$ is a local attached shadow boundary point for $\mathbf{s}$).

Definition 4: Global cast shadow. A surface point $\mathbf{p}$ is called a global cast shadow boundary point if it satisfies both the local condition stated above and the light ray between $\mathbf{p}$ and $\mathbf{s}$ does not intersect some other point on the surface that is not a global attached shadow boundary point for $\mathbf{s}$.

Before proceeding, we consider only surfaces for which

the set of points satisfying the local attached shadow condition form a curve for every $\mathbf{s}$. This restriction is imposed to avoid surfaces such as planes and cylinders that may have shadow boundaries (as defined above) that are whole regions of the surface (e.g., when a light source at infinity is in the direction of a cylindrical ruling).

Now consider applying an arbitrary projective transformation $\alpha: \mathbb{RP}^3 \rightarrow \mathbb{RP}^3$ to both the surface and the light source. Under this transformation, let $\mathbf{p}' = \alpha(\mathbf{p})$ and $\mathbf{s}' = \alpha(\mathbf{s})$.

Lemma 1. A point $\mathbf{p}$ on a smooth surface is a local attached shadow boundary point for point light source $\mathbf{s}$ iff $\mathbf{p}'$ on a transformed surface is a local attached shadow boundary point for point light source $\mathbf{s}'$.

Proof. At a local attached shadow boundary point $\mathbf{p}$, the line defined by $\mathbf{p} \in \mathbb{RP}^3$ and light source $\mathbf{s} \in \mathbb{RP}^3$ lies in the tangent plane at $\mathbf{p}$. Since the order of contact (e.g., tangency) of a curve and a surface is preserved under projective transformations, the line defined by $\mathbf{p}'$ and $\mathbf{s}'$ lies in the tangent plane at $\mathbf{p}'$. ■

Lemma 2. A point $\mathbf{p}$ on a smooth surface is a local cast




Fig. 2. In this 2-D illustration of the GPBR transformation, the lower shadow is an attached shadow, while the upper one is composed of both attached and cast components. A GPBR transformation has been applied to the left surface, yielding the right one. Note that under a GPBR, all surface points and the light source are transformed along the optical rays through the center of projection. Transforming the light source from $\mathbf{s}$ to $\mathbf{s}'$ preserves the shadows.




Fig. 3. Under orthographic projection, the image points that lie in shadow for a surface under light source $\mathbf{s}$ are identical to those in shadow for a transformed surface under light source $\mathbf{s}'$. In this 2-D illustration, the lower shadow is an attached shadow, while the upper one is composed of both attached and cast components. A generalized bas-relief (GBR) transformation with both flattening and an additive plane has been applied to the left surface, yielding the right one.

shadow boundary point for point light source $\mathbf{s}$ iff $\mathbf{p}'$ on a transformed surface is a local cast shadow boundary point for point light source $\mathbf{s}'$.

Proof. For a local cast shadow boundary point $\mathbf{p} \in \mathbb{RP}^3$ and a light source $\mathbf{s} \in \mathbb{RP}^3$, there exists another point $\mathbf{q} \in \mathbb{RP}^3$ on the line defined by $\mathbf{p}$ and $\mathbf{s}$ such that $\mathbf{q}$ lies on an attached shadow. Since collinearity is preserved under projective transformations, $\mathbf{p}'$, $\mathbf{q}'$, and $\mathbf{s}'$ are collinear. Hence, from Lemma 1, $\mathbf{q}'$ is also an attached shadow point. ■

Taken together, Lemmas 1 and 2 indicate that under an arbitrary projective transformation of a surface and a light source, the set of local shadow boundaries is a projective transformation of the local shadow boundaries of the original surface and light source. However, these two lemmas do not imply that the two surfaces are shadow equivalent, since the transformed points may project to different image points or the global conditions in Definitions 2 and 4 may not hold.

A. Perspective Projection: Generalized Perspective Bas-Relief Transformation

We will further restrict the set of projective transformations. Modeling the camera as a function $\gamma: \mathbb{RP}^3 \rightarrow \mathbb{RP}^2$ (see Ref. 18), we require that for any point $\mathbf{p}$ on the surface, we have $\gamma(\mathbf{p}) = \gamma(\alpha(\mathbf{p}))$, where α is a projective transformation—that is, $\mathbf{p}$ and $\alpha(\mathbf{p})$ must project to the same image point. We will consider two specific camera models in turn: perspective projection γ_p and orthographic projection γ_o .

Without loss of generality, consider a pinhole perspective camera with unit focal length located at the origin of the coordinate system and with the optical axis pointed in the direction of the z axis. If we let the homogeneous coordinates of an image point be given by $\omega \in \mathbb{RP}^2$, then pinhole perspective projection of a point $\mathbf{p}$ whose homogeneous coordinates are π is given by $\omega = \Gamma_p \pi$, where

$$\Gamma_p = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}. \quad (1)$$

For $\gamma_p(\mathbf{p}) = \gamma_p(\alpha(\mathbf{p}))$ to be true for any point $\mathbf{p}$, the transformation must move $\mathbf{p}$ along the optical ray between the camera center and $\mathbf{p}$. This can be accomplished by the projective transformation $\alpha: \pi \rightarrow A\pi$, where

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ a_1 & a_2 & a_3 & a_4 \end{bmatrix}. \quad (2)$$

We call this transformation the generalized perspective bas-relief (GPBR) transformation. Note that GPBR is a three-dimensional (3-D) relation, which is a homology, where the vertex of the homology is the origin, the vertex lies on the plane of fixed points, and all lines through the vertex are fixed.¹⁹

In Euclidean coordinates, the transformed surface and light source are given by

$$\mathbf{p}' = \frac{1}{\mathbf{a} \cdot \mathbf{p} + a_4} \mathbf{p}, \quad \mathbf{s}' = \frac{1}{\mathbf{a} \cdot \mathbf{s} + a_4} \mathbf{s}, \quad (3)$$

where $\mathbf{a} = (a_1, a_2, a_3)^T$. Note that Eqs. (3) define a four-parameter family of surfaces and light source positions. Figure 2 shows a two-dimensional (2-D) example of a GPBR being applied to a planar curve and a single light source. The effect is to move points on the surface and the light sources along lines through the camera center in a manner that preserves shadows. The sign of $\mathbf{a} \cdot \mathbf{p} + a_4$ plays a critical role: If it is positive, all points on f move inward or outward from the camera center, remaining in the half-space $z > 0$. On the other hand, if the sign is negative for some points on f , these points will move through the camera center to points with $z < 0$, i.e., they will not be visible to the camera. The equation $\mathbf{a} \cdot \mathbf{p} + a_4 = 0$ defines a plane that divides $\mathbb{R}^3$ into these two cases; all points on this plane map to the plane at infinity. A similar effect on the transformed light source location is determined by the sign of $\mathbf{a} \cdot \mathbf{s} + a_4$.

Proposition 2.1. The image of the shadow boundaries for a surface f and a light source $\mathbf{s}$ is identical to the image of the shadow boundaries for a surface f' and a light source $\mathbf{s}'$ transformed by a GPBR if $\mathbf{a} \cdot \mathbf{s} + a_4 > 0$ and $\mathbf{a} \cdot \mathbf{p} + a_4 > 0$ for all $\mathbf{p} \in f$.

Proof. Since a GPBR is a projective transformation, Lemmas 1 and 2 show that the local attached and cast shadow boundaries on the transformed surface f' from light source $\mathbf{s}'$ are a GPBR transformation of the local shadow boundaries on f from light source $\mathbf{s}$. For any point $\mathbf{p}$ on the surface and any GPBR transformation A , we have $\Gamma_p \pi = \Gamma_{pA} \pi$, and so the images of the local shadow boundaries are identical.

To show that the global condition for an attached shadow in Definition 2 is also satisfied, we note that projective transformations preserve collinearity; therefore the only intersections of the line defined by $\mathbf{s}'$ and $\mathbf{p}'$ with f' are transformations of the intersections of the line defined by $\mathbf{s}$ and $\mathbf{p}$ with f . Within each light ray (a projective line), the points are subjected to a projective transformation; in general, the order of the transformed intersection points on the line may be a combination of a cyclic permutation and a reversal of the order of the original points. However, the restriction that $\mathbf{a} \cdot \mathbf{p} + a_4 > 0$ for all $\mathbf{p} \in f$ and that $\mathbf{a} \cdot \mathbf{s} + a_4 > 0$ has the effect of preserving the order of points between $\mathbf{p}$ and $\mathbf{s}$ on the original line and between $\mathbf{p}'$ and $\mathbf{s}'$ on the transformed line. ■

It should be noted that for any $\mathbf{a}$ and a_4 , there exists a light source $\mathbf{s}$ such that $\mathbf{a} \cdot \mathbf{s} + a_4 < 0$. When f is illuminated by such a source, the transformed source passes through the camera center, and the global shadowing conditions in Definitions 2 and 4 may not be satisfied. Hence two objects differing by a GPBR are not strongly shadow equivalent. On the other hand, for any bounded set of light sources and bounded object f , there exists a set of $a_1, \dots, a_4$ such that $\mathbf{a} \cdot \mathbf{s} + a_4 > 0$ and $\mathbf{a} \cdot \mathbf{p} + a_4 > 0$. Hence there exists a set of objects that are shadow equivalent.

Since the shadow boundaries of multiple light sources are the union of the shadow boundaries from the individual light sources, this also holds for multiple light sources. It should also be noted that the occluding con-

tours (silhouette) of f and f' are identical, since the camera center is a fixed point under a GPBR and the occluding contour is the same as the attached shadow boundary produced by a light source located at the camera center.

Figure 1 shows an example of the GPBR transformation being applied to a scene containing a teapot resting on a support plane. The images were generated by using the VORT ray-tracing package; the scene contained a single proximal point light source, the surfaces were modeled as Lambertian, and a perspective camera model was used. When the light source is transformed with the surface, the shadows are the same for both the original and transformed scenes. Even the shading is similar in both images, so much so that it is nearly impossible to distinguish the two surfaces. However, from another viewpoint, the effect of the GPBR transformation on the object's shape is apparent.

This result complements past work on structure from motion in which the aim of structure recovery is a weaker, non-Euclidean representation, such as affine,^{20–23} projective,²⁴ or ordinal.²⁵

B. Orthographic Projection: Generalized Bas-Relief Transformation

When a camera is distant and can be modeled by an orthographic projection, the visual rays are all parallel to the direction of the optical axis. In $\mathbb{RP}^3$, these rays intersect at the camera center, which is a point at infinity. Without loss of generality, we consider the viewing direction to be in the direction of the z axis and the x and y axes to span the image plane. Again, letting the homogeneous coordinates of an image point be given by $\omega \in \mathbb{RP}^2$, we can express orthographic projection of $\pi \in \mathbb{RP}^3$ as $\omega = \Gamma_o \pi$, where

$$\Gamma_o = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (4)$$

Now let us consider another set of projective transformations $g: \mathbb{RP}^3 \rightarrow \mathbb{RP}^3$. For $\gamma_o(\mathbf{p}) = \gamma_o(g(\mathbf{p}))$ to be true for any point $\mathbf{p}$, the transformation g must move $\mathbf{p}$ along the viewing direction. This can be accomplished by the projective transformation $g: \pi \rightarrow G\pi$, where

$$G = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ g_1 & g_2 & g_3 & g_4 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (5)$$

with $g_3 > 0$. The mapping g is an affine transformation that was introduced in Ref. 12 and is called the generalized bas-relief (GBR) transformation. Again, this transformation is an elation, but in this case the vertex is at infinity and the plane at infinity is fixed. Consider the effect of applying a GBR to a surface parameterized as the graph of a depth function, $(x, y, f(x, y))$. This yields a transformed surface

$$\begin{pmatrix} x' \\ y' \\ z' \end{pmatrix} = \begin{pmatrix} x \\ y \\ g_1x + g_2y + g_3f(x, y) + g_4 \end{pmatrix}. \quad (6)$$

See Fig. 3 for an example. The parameter g_3 has the effect of scaling the relief of the surface, g_1 and g_2 characterize an additive plane, and g_4 provides a depth offset. As described in Ref. 12, when $g_1 = g_2 = 0$ and $0 < g_3 < 1$, the resulting transformation is simply a compression of the surface relief, as in relief sculpture.

Proposition 2.2. The image of the shadow boundaries for a surface f and a light source $\mathbf{s}$ are identical to the image of the shadow boundaries for a surface f' and a light source $\mathbf{s}'$ transformed by any GBR.

Proof. The proof follows that of Proposition 2.1. ■

It should be noted that Proposition 2.2 applies to both nearby light sources and those at infinity. However, in contrast to the GPBR transformation, nearby light sources do not move to infinity nor do light sources at infinity become nearby light sources, since a GBR is an affine transformation that fixes the plane at infinity. Since Proposition 2.2 holds for *any* light source, all objects differing by a GBR transformation are strongly shadow equivalent.

An implication of Propositions 2.1 and 2.2 is that when an object is observed from a fixed viewpoint (whether perspective or orthographic projection), one can at best reconstruct its surface up to a four-parameter family of transformations (GPBR or GBR) from shadow or occluding contour information, irrespective of the number of images and the number of light sources. Under the same conditions, it is impossible to distinguish (recognize) two objects that differ by these transformations from shadows or silhouettes.

3. UNIQUENESS OF THE GENERALIZED BAS-RELIEF TRANSFORMATION

In Subsection 2.B, we have shown that objects differing by a GBR transformation are strongly shadow equivalent, but it is conceivable that there could exist transformations, projective or otherwise, producing families of objects that are also strongly shadow equivalent. Here we prove that under orthographic projection, the GBR transformation is in fact unique in that there is no other transformation of an object's surface that preserves the set of shadows produced by illuminating the object with all possible point sources at infinity. Our proof is a generalization of the one in Ref. 26 that considered only convex objects. Our only restrictions are that the surface is regular,²⁷ that the points satisfying Definition 1 are curves (not regions), and that no portion of the surface is occluded by another object.

Let the vector $\sigma = (s_x, s_y, s_z)^T$ denote in homogeneous coordinates a point light source at infinity, where all light sources producing the same set of global attached shadow boundary points are equated, i.e., $(s_x, s_y, s_z)^T \equiv (ks_x, ks_y, ks_z)^T \forall k \in \mathbb{R}, k \neq 0$. Hence the space $\mathcal{S}$ of light source directions is equivalent to the real projective plane ($\mathbb{RP}^2$), with the line at infinity given by coordinates of the form $(s_x, s_y, 0)$. Note that in Section 2, we represented light sources as points in $\mathbb{RP}^3$; here we restrict ourselves only to distant light sources lying on the plane at infinity of $\mathbb{RP}^3$ (a real projective plane).

Let $\mathbf{n} = (n_x, n_y, n_z)^T$ denote the direction of a surface normal. Again, the magnitude and the sign are unimportant, so we have $(n_x, n_y, n_z)^T \equiv (kn_x, kn_y, kn_z)^T \forall k \in \mathbb{R}, k \neq 0$. Thus the space $\mathcal{N}$ of surface normals is, likewise, equivalent to $\mathbb{RP}^2$. Note that under the equation $\mathbf{n} \cdot \mathbf{s} = 0$, the surface normals are the dual of the light sources. Each point in the $\mathbb{RP}^2$ of light sources has a corresponding line in the $\mathbb{RP}^2$ of surface normals, and vice versa.

Let us now consider the image contours determined by the image points (x, y) that correspond to global attached shadow points for some $\mathbf{s}$. These image contours are the global attached shadow boundaries orthographically projected onto the image plane. For lack of a better name, we will refer to them as the imaged attached shadow boundaries. Likewise, we define the imaged cast shadow boundaries as the image contours determined by the image points (x, y) that correspond to the projection of the global cast shadow points for some $\mathbf{s}$.

For each light source direction, there is a unique imaged attached shadow boundary. Note that this imaged attached shadow boundary is not necessarily a single image contour—as would be the case for a convex object—but may be composed of a collection of image contours. The contours may not be smooth or even closed; see Ref. 28 for a discussion of the generic structure of shadows. Thus we have a bijection mapping the light source direction (points in $\mathbb{RP}^2$) to the imaged attached shadow boundaries. Likewise, for each great circle of light source directions, there is a unique collection of imaged attached shadow boundaries. Thus we have a bijection mapping the great circles of light source directions (lines in $\mathbb{RP}^2$) to collections of imaged attached shadow boundaries. Any two great circles of light source directions intersect at one and only one point in $\mathbb{RP}^2$, as antipodal points of the illumination sphere are equated. Thus any two collections of imaged attached shadow boundaries (corresponding to two great circles of light source directions) have one and only one imaged attached shadow boundary in common.

It follows, then, that we have a projective isomorphism between the real projective plane $\mathcal{S}$ of light source directions and an abstract projective plane $\mathbb{P}^2$ of imaged attached shadow boundaries, where a “point” in the abstract projective plane is a single attached shadow boundary and a “line” in the abstract projective plane is the collection of imaged attached shadow boundaries arising from a great circle of light source directions.

Now let us say that we are given two objects whose visible surfaces are described by respective functions $f(x, y)$ and $f'(x, y)$. If the set of imaged attached shadow boundaries for all light source directions is the same for both objects (i.e., if the objects are strongly shadow equivalent), then the question arises: How are the two surfaces $f(x, y)$ and $f'(x, y)$ related?

Proposition 3.1. Consider two objects satisfying the following: The surfaces are regular,²⁷ the points on the surfaces determined by Definition 1 are curves (not regions), and no portion of the object is occluded by another object. If the visible surfaces of the two objects f and f' are strongly shadow equivalent, then the visible surfaces are related by a GBR transformation.

Proof. As discussed above and illustrated in Fig. 4, we




Fig. 4. Relation of different spaces in the proof of Proposition 3.1.

can construct a projective isomorphism between the abstract projective plane P^2 of imaged attached shadow boundaries and the real projective plane $S \equiv RP^2$ of light source directions illuminating surface $f(x, y)$. Consider the space N of surface normals of f , which, as mentioned above, is the dual space S of the light sources. Each point in the RP^2 of light sources S has a corresponding line in the RP^2 of surface normals N and vice versa. The above mappings are not chosen arbitrarily but are pinned down by the intersection points of the imaged attached shadow boundaries. Each “line” in the abstract projective plane of imaged attached shadow boundaries is formed by the collection of boundaries arising from a line (great circle) of light source directions. The imaged attached shadow boundaries forming a “line” in P^2 intersect at one or more points (x, y) in the image, and these intersection points have a common and unique surface normal. Note that because the surface is regular and is not occluded, the visible surface has normals covering half the Gauss sphere.

In the same manner, we can construct a projective isomorphism between P^2 and the real projective plane of light source directions $S' \equiv RP^2$ illuminating the surface $f'(x, y)$. We can again define N' to be the space of surface normals of f' , which is similarly the dual space of the light sources S' .

Under these two mappings, we have a projective isomorphism between S and S' , which in turn is a projective transformation (collineation).²⁹ Because N and N' are the duals of S and S' , respectively, the surface normals of $f(x, y)$ are also related to the surface normals of $f'(x, y)$ by a projective transformation, i.e., $\mathbf{n}'(x, y) = P\mathbf{n}(x, y)$, where P is a 3×3 invertible matrix.

The transformation P is further restricted in that the surface normals along the occluding contour of f and f' are equivalent, since the occluding contour is after all identical to the attached shadow for a light source aligned with the viewing direction. This implies that the transformation P pointwise fixes the line at infinity of surface normals, and so P must be of the form

$$P = \begin{bmatrix} 1 & 0 & p_1 \\ 0 & 1 & p_2 \\ 0 & 0 & p_3 \end{bmatrix},$$

where $p_3 \neq 0$.

The effect of applying P to the surface normals is the same as that of applying G in Eq. (5) to the surface if $p_1 = -g_1/g_3$, $p_2 = -g_2/g_3$, and $p_3 = 1/g_3$. That is, P has the form of the GBR transformation. Note that the shadows are independent of the translation g_4 along the line of sight under orthographic projection.

The above discussion establishes that the surface normals of $f(x, y)$ are related to the surface normals of $f'(x, y)$ by $\mathbf{n}'(x, y)$, provided that there are at least two image attached shadow boundaries passing through a point (x, y) . Yet because of the possible nonconvexity of the surface and the resulting cast shadows, some points may not lie on the projection of multiple attached shadow boundaries. Furthermore, even though the surface is regular, the visible portion of the surface may be only piecewise continuous, as in Fig. 5. Nevertheless, we can still relate the surfaces $f(x, y)$ and $f'(x, y)$ at these points by considering the imaged cast shadow boundaries.

Consider a point $\mathbf{p}$ that does not fall on more than one imaged attached shadow boundary (analogous to $\mathbf{p}$ in Fig. 5). In three dimensions, it is simple to show that there are multiple light source directions for which $\mathbf{p}$ lies on a cast shadow boundary, and so there is a set Q of corresponding attached shadow points (e.g., $\mathbf{q}$ in Fig. 5) corresponding to each light source direction. The point $\mathbf{p}$ must




Fig. 5. Illustration of some of the ideas in the proof of Proposition 3.1. In two dimensions, consider the shadows of a non-convex object observed under orthographic projection with viewing direction $\mathbf{v}$. Over all light source directions, only points on the thickened arcs can fall on a visible global attached shadow boundary. The points between the thickened arcs and the dotted horizontal lines are visible to the observer, but there is no light source direction for which these points satisfy the global attached shadow conditions. Nonetheless, there exist light source directions for which each of these points can lie on a cast shadow boundary (e.g., $\mathbf{p}$ lies on a cast shadow boundary for light source direction $\mathbf{s}$). Since $\mathbf{p}$ is a cast shadow, it lies in the tangent plane to an attached shadow point $\mathbf{q}$; when the portion of the surface about $\mathbf{q}$ is transformed by a GBR transformation, $\mathbf{p}$ must also be transformed by the same GBR transformation. Similarly, points such as $\mathbf{p}'$ and $\mathbf{q}'$ are, respectively, cast and attached shadow points for light source $\mathbf{s}'$. Note that $\mathbf{p}'$ can also be an attached shadow boundary point. Because $\mathbf{p}'$ lies in the tangent plane to $\mathbf{q}'$, the same GBR that is applied to $\mathbf{q}'$ and its tangent plane must be applied to $\mathbf{p}'$. Hence the two thickened arcs containing $\mathbf{p}'$ and $\mathbf{q}'$ must be transformed by the same GBR if they are to be strongly shadow equivalent.

lie in the tangent plane of each global attached shadow boundary point $\mathbf{q} \in \mathcal{Q}$, and so $\mathbf{p}$ must satisfy $(\mathbf{p} - \mathbf{q}) \cdot \mathbf{n}(\mathbf{q}) = 0$ for all $\mathbf{q} \in \mathcal{Q}$. When all $\mathbf{q} \in \mathcal{Q}$ and $\mathbf{n}(\mathbf{q})$ are transformed by some GBR, $\mathbf{p}$ must also be transformed by the same GBR transformation. In the same way, cast shadow boundaries also establish that points of the visible surface on opposite sides of a discontinuity (e.g., $\mathbf{p}'$ and $\mathbf{q}'$ in Fig. 5) must be related by the same GBR. ■

4. RECONSTRUCTION FROM ATTACHED SHADOWS

In Section 3, we showed that under orthographic projection with distant light sources, the only transformation of a surface that preserves the set of imaged shadow contours is the GBR transformation. However, Proposition 3.1 does not provide a prescription for actually reconstructing a surface up to a GBR. In this section, we consider the problem of reconstruction from the attached shadow boundaries measured in n images of a surface, each illuminated by a single distant light source. We will show that it is possible to estimate the n light source directions and the surface normals at a finite number of points, all up to a GBR. In general, we expect to reconstruct the surface normals at $O(n^2)$ points. From the reconstructed normals, an approximation to the underlying surface can be computed for a fixed GBR. Alternatively, existing shape-from-shadow methods can be used to reconstruct the surface from the estimated light source directions (for a fixed GBR) and from the measured attached and cast shadow boundaries.⁴⁻⁷

First, consider the occluding contour (silhouette) of a surface, which will be denoted by C_0 . This contour is equivalent to the attached shadow produced by a light source whose direction is the viewing direction. Define a coordinate system with $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ spanning the image plane and with $\hat{\mathbf{z}}$ pointing in the viewing direction. For all points $\mathbf{p}$ on the occluding contour, the viewing direction lies in the tangent plane [i.e., $\mathbf{n}(\mathbf{p}) \cdot \hat{\mathbf{z}} = 0$], and the surface normal $\mathbf{n}(\mathbf{p})$ is parallel to the image normal. Hence if the normal to the image contour is (n_x, n_y) , the surface normal is $\mathbf{n} = (n_x, n_y, 0)^T$. In $\mathbb{RP}^2$, the surface normals to all points on the occluding contour correspond to the line at infinity.

Now consider the attached shadow boundary C_1 produced by a light source whose direction is $\mathbf{s}_1$ [see Fig. 6(a)]. For all points $\mathbf{p} \in C_1$, $\mathbf{s}_1$ lies in the tangent plane, i.e., $\mathbf{s}_1 \cdot \mathbf{n}(\mathbf{p}) = 0$. Where C_1 intersects the occluding contour, the normal $\mathbf{n}_1$ can be directly determined from the measured contour as described above. It should be noted that while C_1 and the occluding contour intersect transversally on the surface, their images generically share a common tangent and form the crescent moon image singularity.²⁸ Note that by measuring $\mathbf{n}_1$ along the occluding contour, we obtain a constraint on the light source direction: $\mathbf{s}_1 \cdot \mathbf{n}_1 = 0$. This restricts the light source to a line in $\mathbb{RP}^2$ or to a great circle on the illumination sphere S^2 . The source $\mathbf{s}_1$ can be expressed parametrically in the camera coordinate system as

$$\mathbf{s}_1(\theta_1) = \cos \theta_1 \mathbf{n}_1 + \sin \theta_1 \hat{\mathbf{z}},$$

where θ_1 is unknown. From the shadows in a single image, it is not possible to further constrain $\mathbf{s}_1$, nor does it




Fig. 6. Reconstruction up to a GBR from attached shadows. For a single object in fixed pose, these figures show superimposed attached shadow contours C_i for light source direction $\mathbf{s}_i$. The surface normal where C_i intersects the occluding contour is denoted by $\mathbf{n}_i$. The normal at the intersection of C_i and C_j is denoted by $\mathbf{n}_{i,j}$. (a) The three contours intersect at three points in the image. (b) The three contours meet at a common point, implying that $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_3$ lie on a great circle of the illumination sphere. (c) Eight attached shadow boundaries, of which four intersect at $\mathbf{p}_{1,2}$ and four intersect at $\mathbf{p}_{1,3}$; the direction of the light sources $\mathbf{s}_1, \dots, \mathbf{s}_8$ and the surface normals at the intersection points can be determined up to a GBR. (d) Structure of the illumination sphere S^2 for the light source directions generating the attached shadow boundaries in (c).

seem possible to obtain any further information about points on C_1 . Hence we need to consider the shadows in additional images, and we shall see that only two more unknowns need to be introduced (θ_2 and θ_3) and that these unknowns amount to the GBR parameters.

Now consider a second attached shadow boundary C_2 formed by a second light source direction $\mathbf{s}_2$. Again, the measurement of $\mathbf{n}_2$ (where C_2 intersects C_0) determines a projective line in $\mathbb{RP}^2$ (or a great circle on S^2) that the light source $\mathbf{s}_2$ must lie on. In general, C_1 and C_2 will intersect at one or more visible surface points. If the object is convex and the Gauss map is bijective, then they intersect at only one point $\mathbf{p}_{1,2}$. For a nonconvex surface, C_1 and C_2 may intersect more than once. However, in all cases, the direction of the surface normal $\mathbf{n}_{1,2}$ at the intersections is

$$\mathbf{n}_{1,2} = \mathbf{s}_1(\theta_1) \times \mathbf{s}_2(\theta_2). \quad (7)$$

Thus from the attached shadows in two images, we directly measure $\mathbf{n}_1$ and $\mathbf{n}_2$ and obtain estimates for $\mathbf{n}_{1,2}$, $\mathbf{s}_1$, and $\mathbf{s}_2$ as functions of θ_1 and θ_2 .

Consider a third image illuminated by $\mathbf{s}_3$, in which the attached shadow boundary C_3 does *not* pass through $\mathbf{p}_{1,2}$ [Fig. 6(a)]. Again, we can estimate a projective line (great circle on S^2) containing $\mathbf{s}_3$. We also obtain the surface normal at two additional points, the intersections of C_3 with C_1 and C_2 . From the attached shadow boundaries for a convex surface measured in n images (if no three contours intersect at a common point), the surface normal can be determined at $n(n-1)$ points as a function of n unknowns θ_i , $i = 1, \dots, n$.

However, the number of unknowns can be reduced when three contours intersect at a common point. Consider Fig. 6(b), where contour C_4 intersects C_1 and C_2 at $\mathbf{p}_{1,2}$. In this case, we can infer from the images that $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_4$ all lie in the tangent plane to $\mathbf{p}_{1,2}$. In $\mathbb{RP}^2$, this means that $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_4$ all lie on the same projective line. Since $\mathbf{n}_4$ can be measured, $\mathbf{s}_4$ can be expressed as a function of θ_1 and θ_2 , i.e.,

$$\mathbf{s}_4(\theta_1, \theta_2) = \mathbf{n}_4 \times [\mathbf{s}_1(\theta_1) \times \mathbf{s}_2(\theta_2)].$$

Thus a set of attached shadow boundaries [C_1 , C_2 , and C_4 in Fig. 6(b)] passing through a common point ($\mathbf{p}_{1,2}$) is generated by light sources [$\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_4$ in Fig. 6(d)] located on a great circle of S^2 . The light source directions can be determined up to 2 degrees of freedom θ_1 and θ_2 . Now if in addition a second set of light sources lies along another projective line [the great circle in Fig. 6(d) containing $\mathbf{s}_1$, $\mathbf{s}_3$, $\mathbf{s}_6$, and $\mathbf{s}_7$], the corresponding shadow contours [C_1 , C_3 , C_6 , and C_7 in Fig. 6(c)] intersect at another point on the surface ($\mathbf{p}_{1,3}$). Again, we can express the location of light sources ($\mathbf{s}_6$ and $\mathbf{s}_7$) on this great circle as functions of the locations of two other sources ($\mathbf{s}_1$ and $\mathbf{s}_3$):

$$\mathbf{s}_i(\theta_1, \theta_3) = \mathbf{n}_i \times [\mathbf{s}_1(\theta_1) \times \mathbf{s}_3(\theta_3)].$$

Since $\mathbf{s}_1$ lies at the intersection of both projective lines, we can estimate the direction of any light source located on either line up to just 3 degrees of freedom θ_1 , θ_2 , and θ_3 . Furthermore, the direction of any other light source [$\mathbf{s}_8$ in Fig. 6(d)] can be determined if it lies on a projective line defined by two light sources whose directions are known up to θ_1 , θ_2 , and θ_3 . From the estimated light

source directions, the surface normal can be determined by using Eq. (7) at all points where the shadow boundaries intersect. As mentioned above, there are $O(n^2)$ such points—observe the number of intersections in Fig. 6(c). It is easy to verify algebraically that the 3 degrees of freedom θ_1 , θ_2 , and θ_3 correspond to the degrees of freedom g_1 , g_2 , and g_3 in a GBR. The translation g_4 of the surface along the line of sight cannot be determined under orthographic projection.

5. DISCUSSION

We have defined notions of shadow equivalence for an object, showing that two objects differing by a four-parameter family of projective transformations [generalized perspective bas-relief (GPBR)] are shadow equivalent under perspective projection. Furthermore, under orthographic projection, two objects differing by a generalized bas-relief (GBR) transformation are strongly shadow equivalent; i.e., for any light source illuminating an object, there exists a light source illuminating a transformed object such that the shadows are identical. We have proven that a GBR is the only transformation having this property. While we have shown that the occluding contour is also preserved under a GPBR and a GBR, it should be noted that image intensity discontinuities (step edges) arising from surface normal discontinuities or albedo discontinuities are also preserved under these transformations, since these points move along the line of sight and are viewpoint and (generically) illumination independent. Consequently, edge-based recognition algorithms should not be able to distinguish objects differing by these transformations, nor should edge-based reconstruction algorithms be able to perform Euclidean reconstruction without additional information.

In earlier work, where we concentrated on light sources at infinity,¹² we showed that for any set of point light sources, the shading and the shadowing on an object with Lambertian reflectance are identical to the shading and the shadowing on any GBR transformation of the object, i.e., the illumination cones are identical.³⁰ This implies that a surface can be reconstructed up to a GBR only when the light sources are unknown or that, as Fan and Wolff³¹ showed, one can recover the Hessian of a surface from unknown multiple light sources and integrability. These results are consistent with the effectiveness of well-crafted relief sculptures in conveying a greater sense of the depth than is present; see Ref. 12 for connections of a GBR to relief sculpture. It is clear that shading is not preserved for a GPBR or for a GBR when the light sources are proximal; the image intensity falls off by the reciprocal of the squared distance between the surface and the light source, and distance is not preserved under these transformations. Nonetheless, for a range of transformations and for some sets of light sources, it is expected that the intensity may vary only slightly.

Furthermore, we have shown that it is possible to reconstruct a surface up to a GBR from the shadow boundaries in a set of images; obviously, such a reconstruction does not depend on the material characteristics (e.g., Lambertian model) of the surface. To implement a reconstruction algorithm based on the ideas in Section 4 re-

quires detection of cast and attached shadow boundaries. While detection methods have been presented,^{32–34} it is unclear how effective these techniques would be in practice. In particular, attached shadows are particularly difficult to detect and localize, since for a Lambertian surface with constant albedo, there is a discontinuity in the intensity gradient or the shading flow field but not in the intensity itself. On the other hand, there is a step edge at a cast shadow boundary, and so extensions of the method described in Section 4 that use information about cast shadows to constrain the light source direction may lead to practical implementations.

Finally, Helmholtz³⁵ has shown that for binocular stereopsis, when the baseline/vergence of the two cameras varies but the image pair is the same, there is a corresponding transformation of the scene. (See also related work by Koenderink and van Doorn.³⁶) When the baseline/vergence is unknown, this leads to a shape ambiguity that is a subgroup of a GPBR. It would be interesting to consider the relations between this stereo ambiguity, the shadow ambiguities considered here, and the motion and shading ambiguities considered in Ref. 12.

ACKNOWLEDGMENTS

D. J. Kriegman was supported by a National Science Foundation Young Investigator Award (IRI-9257990) and by Army Research Office grant DAAG55-98-1-0168. P. N. Belhumeur was supported by a Presidential Early Career Award for Scientists and Engineers, National Science Foundation Career Award IRI-9703134, and Army Research Office grant DAAH04-95-1-0494. Many thanks to David Mumford for leading us to the proof of Proposition 3.1, Alan Yuille for many discussions about the GBR, Jean Ponce for pointing out the stereo ambiguity, and Andrew Zisserman and Ian Reid for enlightening us on elations.

Address correspondence to David J. Kriegman at the location on the title page or by e-mail, kriegman@uiuc.edu.

REFERENCES

1. D.L. Waltz, "Understanding line drawings of scenes with shadows," in *The Psychology of Computer Vision*, P. H. Winston, ed. (McGraw-Hill, New York, 1975), pp. 19–91.
2. S. Shafer and T. Kanade, "Using shadows in finding surface orientation," *Comput. Vision Graph. Image Process.* **22**, 145–176 (1983).
3. L. N. Hambrick, M. H. Loew, and R. L. Carroll, Jr., "The entry-exit method of shadow boundary segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-9**, 597–607 (1987).
4. M. Hatzitheodorou, "The derivation of 3-D surface shape from shadows," in *Proceedings of the Image Understanding Workshop* (Morgan Kaufmann, Los Altos, Calif., 1989), pp. 1012–1020.
5. J. R. Kender and E. M. Smith, "Shape from darkness," in *International Conference on Computer Vision* (Morgan Kaufmann, Los Altos, Calif., 1987), pp. 539–546.
6. D. Yang and J. R. Kender, "Shape from shadows under error," in *Proceedings of the Image Understanding Workshop* (Morgan Kaufmann, Los Altos, Calif., 1993), pp. 1083–1090.
7. M. Daum and G. Dudek, "On 3-D surface reconstruction using shape from shadows," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society Press, Los Alamitos, Calif., 1998), pp. 461–468.
8. F. Cheng and K. H. Thiel, "Delimiting the building heights in a city from the shadow in a panchromatic spot-image. 1 test of 42 buildings," *Int. J. Remote Sens.* **16**, 409–415 (1995).
9. A. Huertas and R. Nevatia, "Detection of buildings in aerial images using shape and shadows," in *Proceedings of the International Joint Conference on Artificial Intelligence* (Morgan Kaufmann, San Francisco, Calif., 1983), pp. 1099–1103.
10. R. B. Irvin and D. M. McKeown, "Methods for exploiting the relationship between buildings and their shadows in aerial imagery," *IEEE Trans. Syst. Man Cybern.* **19**, 1564–1575 (1989).
11. G. G. Medioni, "Obtaining 3-D from shadows in aerial images," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society Press, Los Alamitos, Calif., 1983), pp. 73–76.
12. P. Belhumeur, D. Kriegman, and A. Yuille, "The bas-relief ambiguity," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society Press, Los Alamitos, Calif., 1997), pp. 1040–1046.
13. D. Kriegman and P. Belhumeur, "What shadows reveal about object structure," in *Proceedings of the European Conference on Computer Vision* (Springer, Berlin, 1998), pp. 399–414.
14. M. S. Langer and S. W. Zucker, "What is a light source?" in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society Press, Los Alamitos, Calif., 1997), pp. 172–178.
15. J. Mundy and A. Zisserman, *Geometric Invariance in Computer Vision* (MIT Press, Cambridge, Mass., 1992).
16. M. Baxandall, *Shadows and Enlightenment* (Yale U. Press, New Haven, Conn., 1995).
17. A. Shashua, "Geometry and photometry in 3D visual recognition," Ph.D. thesis (Massachusetts Institute of Technology, Cambridge, Mass., 1992).
18. O. Faugeras, *Three Dimensional Computer Vision* (MIT Press, Cambridge, Mass., 1993).
19. A. Criminisi, I. Reid, and A. Zisserman, "Duality, rigidity and planar parallax," in *Proceedings of the European Conference on Computer Vision* (Springer, Berlin, 1998), Vol. II, pp. 846–861.
20. J. J. Koenderink and A. J. Van Doorn, "Affine structure from motion," *J. Opt. Soc. Am. A* **8**, 377–385 (1991).
21. R. Rosenholtz and J. J. Koenderink, "Affine structure and photometry," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society Press, Los Alamitos, Calif., 1996), pp. 18–20.
22. L. S. Shapiro, A. Zisserman, and M. Brady, "3D motion recovery via affine epipolar geometry," *Int. J. Comput. Vis.* **16**, 147–182 (1995).
23. S. Ullman and R. Basri, "Recognition by a linear combination of models," *IEEE Trans. Pattern Anal. Mach. Intell.* **13**, 992–1006 (1991).
24. O. Faugeras, "Stratification of 3-D vision: projective, affine, and metric representations," *J. Opt. Soc. Am. A* **12**, 465–484 (1995).
25. C. Fermuller and Y. Aloimonos, "Ordinal representations of visual space," in *Proceedings of the Image Understanding Workshop* (Morgan Kaufmann, San Francisco, Calif., 1996), pp. 897–904.
26. P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille, "The bas-relief ambiguity," *Int. J. Comput. Vis.* **35**, 33–44 (1999).
27. M. P. DoCarmo, *Differential Geometry of Curves and Surfaces* (Prentice-Hall, Englewood Cliffs, N.J., 1976).
28. L. Donati and N. Stolfi, "Singularities of illuminated surfaces," *Int. J. Comput. Vision* **23**, 207–216 (1997).
29. E. Artin, *Geometric Algebra* (Interscience, New York, 1957).
30. P. N. Belhumeur and D. J. Kriegman, "What is the set of images of an object under all possible lighting conditions," in *Proceedings of the IEEE Conference on Computer Vision*

- and *Pattern Recognition* (IEEE Computer Science Press, Los Alamitos, Calif., 1996), pp. 270–277.
31. J. Fan and L. B. Wolff, “Surface curvature and shape reconstruction from unknown multiple illumination and integrability,” *Comput. Vision Image Underst.* **65**, 347–359 (1997).
 32. P. Breton and S. W. Zucker, “Shadows and shading flow fields,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE Computer Society Press, Los Alamitos, Calif., 1996), pp. 782–789.
 33. C. Jiang and M. O. Ward, “Shadow segmentation and classification in a constrained environment,” *CVGIP: Graph. Models Image Process.* **59**, 213–225 (1994).
 34. A. P. Witkin, “Intensity-based edge classification,” in *Proceedings of the American Association for Artificial Intelligence* (AIII Press, Menlo Park, Calif., 1982), pp. 36–41.
 35. H. von Helmholtz, *Treatise on Physiological Optics* (Dover, New York, 1925).
 36. J. J. Koenderink and A. J. van Doorn, “Geometry of binocular vision and a model for stereopsis,” *Biol. Cybern.* **21**, 29–35 (1976).