What Shadows Reveal About Object Structure

David J. Kriegman and Peter N. Belhumeur

Center for Computional Vision and Control, Yale University, New Haven CT 06520
david.kriegman@yale.edu and belhumeur@yale.edu
http://cvc.yale.edu

Abstract. *In a scene observed from a fixed viewpoint, the set of shadow curves 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 an object up to these transformations from the shadows alone.*

1 Introduction

In his fifteenth century *Treatise on Painting* [15], Leonardo da Vinci errs in analysis of shadows while comparing painting and relief sculpture:

> As far as light and shade are concerned low relief fails both as sculpture and as painting, because the shadows correspond to the low nature of the relief, as for example in the shadows of foreshortened objects, which will not exhibit the depth of those in painting or in sculpture in the round.

It is true that – when illuminated by the same light source – a relief surface and a surface "in the round" will cast different shadows. However, Leonardo's statement appears to overlook the fact that 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 later show, applies equally to a greater set of projective transformations.

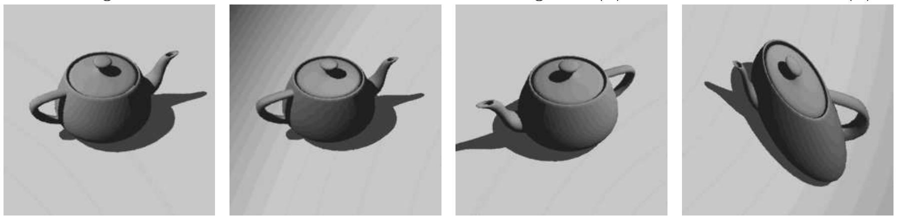

Fig. 1. An illustration of the effect of applying a generalized perspective bas-relief (GPBR) transformation to a scene composed of a teapot resting on a supporting plane. The first image shows the original teapot. The second image shows the teapot after having undergone a GPBR transformation $(a_1, a_2, a_3, a_4) = (.05, .05, .05, 1)$ 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 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 same viewpoint as used for the third image.

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 direction or location of the light source *is 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 [27]. Subsequently, Shafer showed how geometric constraints on surface orientation could be obtained from labeled line drawings using shadow and surface outlines under orthographic projection [23]. The Entry-Exit method was developed to segment and label shadow curves using information about the projection onto the image plane of the light source direction [11]. 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 [12, 17, 29]. Shadows have also been used in the interpretation of aerial images, particularly to locate and reconstruct buildings when the sun direction is known [6, 13, 14, 20].

Here we consider shadows on unknown objects produced by light sources whose directions are also unknown. In the next section 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 Figure 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 [3]. As will be shown in Section 3, the GBR transformation is unique in that any two smooth surfaces which 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 we give it to demonstrate that under ideal conditions information from shadows alone is enough to determine the structure of the object up to a GBR transformation.

2 Shadowing Ambiguity

Let us define two objects as being *shadow equivalent* if there exists two sets of point light sources $\mathcal{S}$ and $\mathcal{S}'$ such that for every light source in $\mathcal{S}$ illuminating one object, there exists a light source in $\mathcal{S}'$ 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 *any* light source illuminating one object, there exists a source illuminating the second object such that shadowing is identical – i.e., $\mathcal{S}$ is the set of all point light sources. 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 halfspace $z > 0$. 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 collection of point light sources, located nearby or at infinity. Note that this is a restriction of the lighting model presented by Langer and Zucker [19] which permits anisotropic light sources whose intensity is a function of direction. In this paper, we will represent surfaces, light sources, and the camera center as lying in either a two or three dimensional real projective space ($\mathbb{RP}^2$ or $\mathbb{RP}^3$). (For a concise treatment of real projective spaces, see [21].) 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 $\mathbf{s} = (s_x, s_y, s_z)$. In projective (homogeneous) coordinates, the light source $\mathsf{s} \in \mathbb{RP}^3$ can be written as $\mathsf{s} = (s_x, s_y, s_z, 1)$. (Note that different fonts are used to distinguish between Euclidean and projective coordinates.) 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 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 $\mathbf{s} = (s_x, s_y, s_z, 0)$. (Note that when the light source at infinity is represented in projective coordinates, the antipodal points from S^2 must be equated.)

For a single point source $\mathbf{s} \in \mathbb{RP}^3$, let us define the set of *light rays* as the lines in $\mathbb{RP}^3$ passing through $\mathbf{s}$. For any $\mathbf{p} \in \mathbb{RP}^3$ with $\mathbf{p} \neq \mathbf{s}$, there is a single light ray passing through $\mathbf{p}$. Naturally it is the intersection of the light rays with the surface f which determine the shadows. We differentiate between two types of shadows: *attached shadows* and *cast shadows* [2, 25]. See Figures 2 and 3. A surface point $\mathbf{p}$ lies on the border of an *attached shadow* for light source $\mathbf{s}$ if and only if it satisfies both a local and global condition:

Local Attached Shadow Condition: The light ray through $\mathbf{p}$ lies in the tangent plane to the surface at $\mathbf{p}$. Algebraically, this condition can be expressed as $\mathbf{n}(\mathbf{p}) \cdot (\mathbf{p} - \mathbf{s}) = 0$ for a nearby light source and as $\mathbf{n}(\mathbf{p}) \cdot \mathbf{s} = 0$ for a distant light source. A point $\mathbf{p}$ which satisfies at least the local condition is called a *local attached shadow boundary point.*

Global Attached Shadow Condition: The light ray does not intersect the surface between $\mathbf{p}$ and $\mathbf{s}$, i.e., the light source is not occluded at $\mathbf{p}$.

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

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 surface is preserved under projective transformations, the line defined by $\mathbf{p}'$ and $\mathbf{s}'$ lies in the tangent plane at $\mathbf{p}'$.

Cast shadows occur at points on the surface that face the light source, but where some other portion of the surface lies between the shadowed points and the light source. A point $\mathbf{p}$ lies on the boundary of a cast shadow for light source

$\mathbf{s}$ if and only if it similarly satisfies both a local and global condition:

Local Cast Shadow Condition: The light ray through $\mathbf{p}$ grazes the surface at some other point $\mathbf{q}$ (i.e., $\mathbf{q}$ lies on an attached shadow). A point $\mathbf{p}$ which satisfies at least the local condition is called a *local cast shadow boundary point.*

Global Attached Shadow Condition: The only intersection of the surface and the light ray between $\mathbf{p}$ and $\mathbf{s}$ is at $\mathbf{q}$.

Lemma 2. *A point $\mathbf{p}$ on a smooth surface is a local cast 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 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. From Lemma 1, $\mathbf{q}'$ is also an attached shadow point.

Taken together, Lemmas 1 and 2 indicate that under a projective transformation of a surface and light source, the set of local shadow curves is a projective transformation of the local shadow curves 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 may not hold.

2.1 Perspective Projection: GPBR

We will further restrict the set of projective transformations. Modeling the camera as a function $\pi : \mathbb{RP}^3 \rightarrow \mathbb{RP}^2$, we require that for any point $\mathbf{p}$ on the surface $\pi(\mathbf{p}) = \pi(a(\mathbf{p}))$ where a is a projective transformation – that is $\mathbf{p}$ and $a(\mathbf{p})$ must project to the same image point. We will consider two specific camera models in turn: perspective projection π_p and orthographic projection π_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. Letting the homogeneous coordinates of an image point be given by $\mathtt{u} \in \mathbb{RP}^2$, then pinhole perspective projection of $\mathtt{p} \in \mathbb{RP}^3$ is given by $\mathtt{u} = \Pi_p \mathtt{p}$ where

$$\Pi_p = \begin{bmatrix} 1\,0\,0\,0 \\ 0\,1\,0\,0 \\ 0\,0\,1\,0 \end{bmatrix} . \tag{1}$$

For $\pi_p(\mathbf{p}) = \pi_p(a(\mathbf{p}))$ 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 $a : \mathtt{p} \mapsto A\mathtt{p}$ where

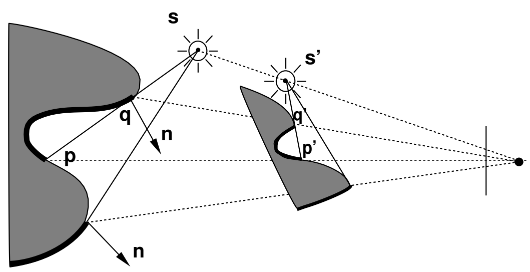

Fig. 2. In this 2-d illustration of the generalized perspective bas-relief transformation (GPBR), the lower shadow is an attached shadow while the upper one is composed of both attached and cast components. A GPBR transformation has been applied to the left surface, yielding the right one. Note that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

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

We call this transformation the Generalized Perspective Bas-Relief (GPBR) transformation. In Euclidean coordinates, the transformed surface and light source are given by

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

where $\mathbf{a} = (a_1, a_2, a_3)^T$. Figure 2 shows a 2-d example of 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 halfspace $z > 0$. 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 which 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 1. *The image of the shadow curves for a surface f and light source $\mathbf{s}$ is identical to the image of the shadow curves for a surface f' and 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 GPBR is a projective transformation, Lemmas 1 and 2 show that the local attached and cast shadow curves on the transformed surface f' from light source $\mathbf{s}'$ are a GPBR of the local shadow curves on f from light source $\mathbf{s}$. For any point $\mathbf{p}$ on the surface and any GPBR transformation A, we have $\Pi_p \mathrm{p} = \Pi_p A \mathrm{p}$, and so the images of the local shadow curves are identical

To show that the global condition for an attached shadow 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 for 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 may not be satisfied. Hence two objects differing by 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 $a_1, \ldots, a_4$ such that $\mathbf{a} \cdot \mathbf{s} + a_4 > 0$ and $\mathbf{a} \cdot \mathbf{p} + a_4 > 0$. Hence, there exist a set of objects which are *shadow equivalent.*

Since the shadow curves of multiple light sources are the union of the shadow curves from the individual light sources, this also holds for multiple light sources. It should also be noted that the occluding contour (silhouette) of f and f' are identical, since the camera center is a fixed point under GPBR and the occluding contour is the same as the attached shadow curve 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 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 on the object's shape is apparent.

This result compliments past work on structure from motion in which the aim of structure recovery is a weaker non-Euclidean representation, such as affine [18, 22, 24, 26], projective [9], or ordinal [10].

2.2 Orthographic Projection: GBR

When a camera is distant and can be modeled as 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 consider the viewing direction to be in the direction of the z-axis

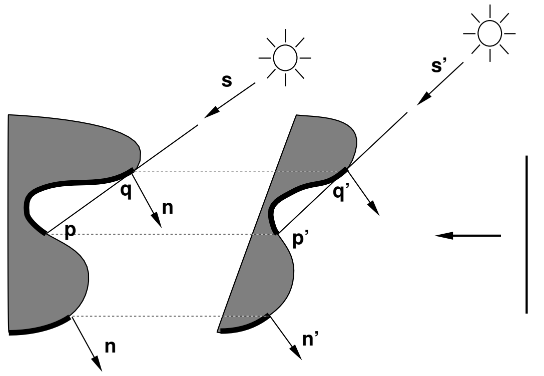

Fig. 3. The image points that lie in shadow for a surface under light source $\mathbf{s}$ are identical to those in shadow for a transformed surface under light source $\mathbf{s}'$. In this 2-d illustration, the lower shadow is an attached shadow while the upper one is composed of both attached and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

and the x and y axes to span the image plane. Again, letting the homogeneous coordinates of an image point be given by $\mathtt{u} \in \mathbb{RP}^2$, orthographic projection of $\mathtt{p} \in \mathbb{RP}^3$ can be expressed as $\mathtt{u} = \Pi_o \mathtt{p}$ where

$$\Pi_o = \begin{bmatrix} 1\,0\,0\,0 \\ 0\,1\,0\,0 \\ 0\,0\,0\,1 \end{bmatrix}. \tag{4}$$

Now, let us consider another set of projective transformations $g : \mathbb{RP}^3 \to \mathbb{RP}^3$. For $\pi_o(\mathbf{p}) = \pi_o(g(\mathbf{p}))$ 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 $g : \mathtt{p} \mapsto G\mathtt{p}$ where

$$G = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ g_1 & g_2 & g_3 & g_4 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{5}$$

with $g_3 > 0$. The mapping g is an affine transformation which was introduced in [3] and was called the generalized bas-relief (GBR) transformation. Consider the effect of applying GBR to a surface parameterized as the graph of a depth function, $(x, y, f(x, y))$. This yields a transformed surface

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \begin{bmatrix} x \\ y \\ g_1 x + g_2 y + g_3 f(x,y) + g_4 \end{bmatrix}.$$

See Figure 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 [3], when $g_1 = g_2 = 0$ and $0 < g_3 < 1$, the resulting transformation is simply a compression of the surface's relief, as in relief sculpture.

Proposition 2. *The image of the shadow curves for a surface f and light source $\mathbf{s}$ are identical to the image of the shadow curves for a surface f' and light source $\mathbf{s}'$ transformed by any GBR.*

Proof. The proof follows that of Proposition 1.

It should be noted that Proposition 2 applies to both nearby light sources and those at infinity. However, in contrast to the GPBR transformation, nearby light source do not move to infinity nor do light sources at infinity become nearby light sources since GBR is an affine transformation which fixes the plane at infinity. Since Proposition 2 holds for *any* light source, all objects differing by a GBR transformation are *strongly shadow equivalent*.

An implication of Propositions 1 and 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 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

Here we prove that under orthographic projection the generalized bas-relief (GBR) transformation is unique in that there is no other transformation of an object's surface which preserves the set of shadows produced by illuminating the object with all possible point sources at infinity. We consider only the simplest case – an object with convex shape casting no shadows on its own surface – and show that the set of attached shadow boundaries are preserved *only* under a GBR transformation of the object's surface.

Recall that an attached shadow boundary is defined as the contour of points $(x, y, f(x, y))$ satisfying $\mathbf{n} \cdot \mathbf{s} = 0$, for some $\mathbf{s}$. For a convex object, the global attached shadow condition holds everywhere. Here the magnitude and the sign of the light source are unimportant as neither effects the location of the attached shadow boundary. Thus, let the vector $\mathbf{s} = (s_x, s_y, s_z)^T$ denote in homogeneous coordinates a point light source at infinity, where all light sources producing the same attached shadow boundary are equated, i.e., $(s_x, s_y, s_z)^T \equiv (ks_x, ks_y, ks_z)^T \; \forall k \in \mathbb{R}, k \neq 0$. With this, the space of light source directions $\mathcal{S}$ 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 the previous section, we represented light sources as points in $\mathbb{RP}^3$; here, we restrict our selves only to distant light sources lying on the plane at infinity of $\mathbb{RP}^3$, (a real projective plane).

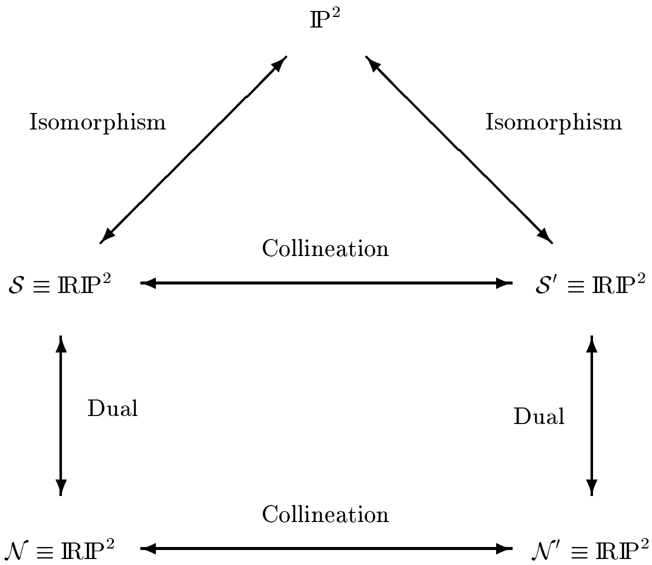

Fig. 4. The relation of different spaces in proof of Proposition 3.

Let $\mathbf{n} = (n_x, n_y, n_z)^T$ denote the direction of a surface normal. Again, the magnitude and 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 of surface normals $\mathcal{N}$ 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 defined by the points (x, y) satisfying $\mathbf{n} \cdot \mathbf{s} = 0$, for some $\mathbf{s}$. These image contours are the 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.

The set of imaged attached shadow boundaries for a convex object forms an abstract projective plane $\mathbb{P}^2$, 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 passing through a common point in the image plane. To see this, note the obvious projective isomorphism between the real projective plane of light source directions $\mathcal{S}$ and the abstract projective plane of imaged attached shadow boundaries $\mathbb{P}^2$. Under this is isomorphism, we have bijections mapping points to points and lines to lines.

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 objects have the same set of imaged attached shadow boundaries as seen in the image plane (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. *If the visible surfaces of two convex objects f and f' are strongly shadow equivalent, then the surfaces are related by a generalized bas-relief transformation.*

Proof. As illustrated in Figure 4, we can construct a projective isomorphism between the set of imaged attached shadow boundaries $\mathbb{P}^2$ and the real projective plane of light source directions $\mathcal{S}$ illuminating surface $f(x, y)$. The isomorphism is chosen to map the collection of imaged attached shadow boundaries passing through a common point (x, y) in the image plane (i.e., a line in $\mathbb{P}^2$) to the surface normal $\mathbf{n}(x, y)$. In the same manner, we can construct a projective isomorphism between $\mathbb{P}^2$ and the real projective plane of light source directions $\mathcal{S}'$ illuminating the surface $f'(x, y)$. The isomorphism is, likewise, chosen to map the same collection of imaged attached shadow boundaries passing through (x, y) in the image plane to the surface normal $\mathbf{n}'(x, y)$. Under these two mappings, we have a projective isomorphism between $\mathcal{S}$ and $\mathcal{S}'$ which in turn is a projective transformation (collineation) [1]. Because $\mathcal{N}$ and $\mathcal{N}'$ are the duals of $\mathcal{S}$ and $\mathcal{S}'$ 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, i.e., the transformation P pointwise fixes the line at infinity of surface normals. Thus, 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 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 generalized bas-relief transformation. Note that the shadows are independent of the translation g_4 along the line of sight under orthographic projection.

4 Reconstruction from Attached Shadows

In the previous section, we showed that under orthographic projection with distant light sources, the only transformation of a surface which preserves the set of imaged shadow contours is the generalized bas-relief transformation. However, Proposition 3 does not provide a prescription for actually reconstructing a surface up to 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 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

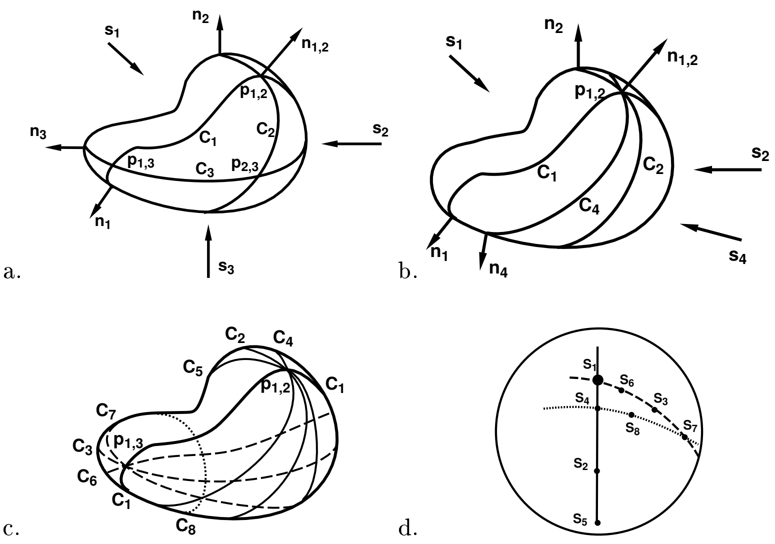

Fig. 5. Reconstruction up to GBR from attached shadows: For a single object in fixed pose, these figures show superimposed attached shadow contours C_i for light source direction $\mathbf{s}_i$. The surface normal where C_i intersects the occluding contour is denoted by $\mathbf{n}_i$. The normal at 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 $\mathbf{s}_1 \ldots \mathbf{s}_8$ and the surface normals at the intersection points can be determined up to GBR. d) The structure of the illumination sphere S^2 for the light source directions generating the attached shadow boundaries in Fig. 5.c.

source directions (for a fixed GBR) and from the measured attached and cast shadow curves [12, 17, 29].

First, consider the occluding contour (silhouette) of a surface which will be denoted 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 Figure 5.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 [8]. 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

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

From the shadows in a single image, it is not possible to further constrain $\mathbf{s}_1$ nor does it seem possible to obtain any further information about points on C_1.

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 only intersect at 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

$$\mathbf{n}_{1,2} = \mathbf{s}_1(\theta_1) \times \mathbf{s}_2(\theta_2). \tag{6}$$

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. 5.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 $\theta_i, i = 1 \ldots n$.

However, the number of unknowns can be reduced when three contours intersect at a common point. Consider Fig. 5.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 $\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_4$ 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 curves (C_1, C_2, C_4 in Fig. 5.b) passing through a common point ($\mathbf{p}_{1,2}$) is generated by light sources ($\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_4$ in Fig. 5.d) located on a great circle of S^2. The light source directions can be determined up to two 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 5.d containing $\mathbf{s}_1, \mathbf{s}_3, \mathbf{s}_6, \mathbf{s}_7$), the corresponding shadow contours (C_1, C_3, C_6, C_7 in Fig 5.c) intersect at another point on the surface ($\mathbf{p}_{1,3}$). Again, we can express the location of light sources ($\mathbf{s}_6, \mathbf{s}_7$) 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 three degrees of freedom θ_1, θ_2, and θ_3. Furthermore, the direction of any other light source ($\mathbf{s}_8$ on Fig. 5.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 using Eq. 6 at all points where the shadow curves intersect. As mentioned earlier, there are $O(n^2)$ such points – observe the number of intersections in Fig. 5.c. It is easy to verify algebraically that the three degrees of freedom θ_1, θ_2 and θ_3 correspond to the degrees of freedom in GBR g_1, g_2 and g_3. The translation g_4 of the surface along the line sight cannot be determined under orthographic projection.

5 Discussion

We have defined notions of shadow equivalence for object, showing that two objects differing by a four parameter family of projective transformations (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 exits a light source illuminating a transformed object such that the shadows are identical. We have proven that GBR is the only transformation having this property. While we have shown that the occluding contour is also preserved under GPBR and 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 [4, 3], we showed that for any set of point light sources, the shading as well as the shadowing on an object with Lambertian reflectance are identical to the shading and shadowing on any generalized bas-relief transformation of the object, i.e., the illumination cones are identical. This is consistent with the effectiveness of well-crafted relief sculptures in conveying a greater sense of the depth than is present. It is clear that shading is not preserved for GPBR or for GBR when the light sources are proximal; the image intensity falls off by the reciprocal of the squared distance between the surface and 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 only vary slightly.

Furthermore, we have shown that it is possible to reconstruct a surface up to GBR from the shadow boundaries in a set of images. To implement a reconstruction algorithm based on the ideas in Section 4 requires detection of cast and attached shadow boundaries. While detection methods have been presented [5, 28], 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 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 which use information about cast shadows to constrain the light source direction may lead to practical implementations.

Leonardo da Vinci's statement that shadows of relief sculpture are "foreshortened" is, strictly speaking, incorrect. However, reliefs are often constructed in a manner such that the cast shadows will differ from those produced by sculpture in the round. Reliefs have been used to depict narratives involving numerous figures located at different depths within the scene. Since the sculpting medium is usually not thick enough for the artist to sculpt the figures to the proper relative depths, sculptors like Donatello and Ghiberti employed rules of perspective to determine the size and location of figures, sculpting each figure to the proper relief [16]. While the shadowing for each figure is self consistent, the shadows cast from one figure onto another are incorrect. Furthermore, the shadows cast onto the background, whose orientation usually does not correspond to that of a wall or floor in the scene, are also inconsistent. Note however, that ancient Greek sculpture was often painted; by painting the background of the Parthenon Frieze a dark blue [7], cast shadows would be less visible and the distortions less apparent. Thus, Leonardo's statement is an accurate characterization of complex reliefs such as Ghiberti's East Doors on the Baptistery in Florence, but does not apply to figures sculpted singly.

Acknowledgments

Many thanks to David Mumford for leading us to the proof of Proposition 3, and to Alan Yuille for many discussions about the GBR.

References

1. E. Artin. *Geometric Algebra.* Interscience Publishers, Inc., New York, 1957.
2. M. Baxandall. *Shadows and Enlightenment.* Yale University Press, New Haven, 1995.
3. P. Belhumeur, D. Kriegman, and A. Yuille. The bas-relief ambiguity. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 1040–1046, 1997.
4. P. N. Belhumeur and D. J. Kriegman. What is the set of images of an object under all possible lighting conditions. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 270–277, 1996.

5. P. Breton and S. Zucker. Shadows and shading flow fields. In *CVPR96*, pages 782–789, 1996.
6. F. Cheng and K. Thiel. Delimiting the building heights in a city from the shadow in a panchromatic spot-image. 1 test of 42 buildings. *JRS*, 16(3):409–415, Feb. 1995.
7. B. F. Cook. *The Elgin Marbles*. Harvard University Press, Cambridge, 1984.
8. L. Donati and N. Stolfi. Singularities of illuminated surfaces. *Int. J. Computer Vision*, 23(3):207–216, 1997.
9. O. Faugeras. Stratification of 3-D vision: Projective, affine, and metric representations. *J. Opt. Soc. Am. A*, 12(7):465–484, 1995.
10. C. Fermuller and Y. Aloimonos. Ordinal representations of visual space. In *Proc. Image Understanding Workshop*, pages 897–904, 1996.
11. L. Hambrick, M. Loew, and R. Carroll, Jr. The entry-exit method of shadow boundary segmentation. *PAMI*, 9(5):597–607, September 1987.
12. M. Hatzitheodorou. The derivation of 3-d surface shape from shadows. In *Proc. Image Understanding Workshop*, pages 1012–1020, 1989.
13. A. Huertas and R. Nevatia. Detection of buildings in aerial images using shape and shadows. In *Proc. Int. Joint Conf. on Art. Intell.*, pages 1099–1103, 1983.
14. R. Irvin and D. McKeown. Methods for exploiting the relationship between buildings and their shadows in aerial imagery. *IEEE Systems, Man, and Cybernetics*, 19(6):1564–1575, 1989.
15. M. Kemp, editor. *Leonardo On Painting*. Yale University Press, New Haven, 1989.
16. M. Kemp. *The Science of Art: Optical Themes in Western Art from Brunelleschi to Seurat*. Yale University Press, New Haven, 1990.
17. J. Kender and E. Smith. Shape from darkness. In *Int. Conf. on Computer Vision*, pages 539–546, 1987.
18. J. Koenderink and A. Van Doorn. Affine structure from motion. *JOSA-A*, 8(2):377–385, 1991.
19. M. Langer and S. Zucker. What is a light source? In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 172–178, 1997.
20. G. Medioni. Obtaining 3-d from shadows in aerial images. In *CVPR83*, pages 73–76, 1983.
21. J. Mundy and A. Zisserman. *Geometric invariance in computer vision*. MIT Press, 1992.
22. R. Rosenholtz and J. Koenderink. Affine structure and photometry. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 790–795, 1996.
23. S. Shafer and T. Kanade. Using shadows in finding surface orientation. *Comp. Vision, Graphics, and Image Proces.*, 22(1):145–176, 1983.
24. L. Shapiro, A. Zisserman, and M. Brady. 3D motion recovery via affine epipolar geometry. *Int. J. Computer Vision*, 16(2):147–182, October 1995.
25. A. Shashua. *Geometry and Photometry in 3D Visual Recognition*. PhD thesis, MIT, 1992.
26. S. Ullman and R. Basri. Recognition by a linear combination of models. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 13:992–1006, 1991.
27. D. L. Waltz. Understanding line drawings of scenes with shadows. In P. Winston, editor, *The Psychology of Computer Vision*, pages 19–91. McGraw-Hill, New York, 1975.
28. A. Witkin. Intensity-based edge classification. In *Proc. Am. Assoc. Art. Intell.*, pages 36–41, 1982.
29. D. Yang and J. Kender. Shape from shadows under error. In *Proc. Image Understanding Workshop*, pages 1083–1090, 1993.