

Submitted to the International Journal of Computer Vision.
Earlier version appeared in CVPR, 1997, pp. 1060--1066.

The Bas-Relief Ambiguity*

Peter N. Belhumeur[†]

David J. Kriegman[‡]

Alan L. Yuille[§]

Center for Computational Vision and Control
Yale University
New Haven, CT 06520-8267

Smith-Kettlewell Eye Research Institute
San Francisco, CA 94115

Abstract

When an unknown object with Lambertian reflectance is viewed orthographically, there is an implicit ambiguity in determining its 3-d structure: we show that the object's visible surface $f(x, y)$ is indistinguishable from a "generalized bas-relief" transformation of the object, $\bar{f}(x, y) = \lambda f(x, y) + \mu x + \nu y$. For each image of the object illuminated by an arbitrary number of distant light sources, there exists an identical image of the transformed object illuminated by similarly transformed light sources. This result holds both for the illuminated regions of the object as well as those in cast and attached shadows. Furthermore, neither small motion of the object, nor of the viewer will resolve the ambiguity in determining the flattening (or scaling) λ of the object's surface. Implications of this ambiguity on structure recovery and shape representation are discussed.

*Please direct all correspondence about this submission to Peter N. Belhumeur, Dept. of Electrical Engineering, Yale University, New Haven, CT 06520-8267, E-mail belhumeur@yale.edu, Phone (203) 432-4249, Fax (203) 432-7481.

[†]P. N. Belhumeur was supported by an NSF Career Award IRI-9703134 and ARO grant DAAH04-95-1-0494.

[‡]D. J. Kriegman was supported by NSF under an NYI, IRI-9257990.

[§]A. L. Yuille was supported by ARO grant DAAH04-95-1-0494 and NSF grant IRI-93-17670.



Figure 1: Frontal and side views of a pair of marble bas-relief sculptures: Notice how the frontal views appear to have full 3-dimensional depth, while the side views reveal the flattening – the sculptures rise only 5 centimeters from the background plane. While subtle shading is apparent on the faces, the shadows on the women’s pleats are the dominant perceptual cue in the body.

1 Introduction

Since antiquity, artisans have created flattened forms, i.e., so-called “bas-reliefs,” which when viewed from a particular vantage point are difficult, if not impossible, to distinguish from full reliefs. See Figure 1. As the sun moves through the sky, the shading and shadows change, yet the degree of flattening cannot be discerned on well sculpted bas-reliefs. Even if an observer’s head moves by a small amount, this ambiguity cannot be resolved. This paper does not simply present an explanation for the effectiveness of relief sculpture, but demonstrates that the ambiguity is implicit in recovering the structure of any object.

Consider the set of images produced by viewing an object from a fixed viewpoint, but under all possible combinations of distant light sources. An ambiguity in determining the object’s structure arises if there exist other objects that differ in shape yet produce the same set of images. We show that there exists a whole family of transformations, termed “generalized bas-relief transformations,” for which this is true.

A generalized bas-relief (GBR) transformation changes both the surface shape and albedo pattern. In particular, if (x, y) denotes the coordinates of points in an image plane, and $z = f(x, y)$ denotes the distance from an object’s surface to the image plane, a generalized bas-relief transformation of the surface shape is given by $\bar{f}(x, y) = \lambda f(x, y) + \mu x + \nu y$ with $\lambda > 0$, and the corresponding generalized bas-relief transformation of the surface albedo is given by Eq. 3. Relief sculptures are constructed using a subset of the transformation on shape with $0 < \lambda < 1$ and $\mu = \nu = 0$ but without – to the best of our knowledge – the corresponding transformation on albedo. Bas-reliefs (low reliefs) are usually defined as having $\lambda < 0.5$.

Yet the subtleties of the bas-relief ambiguity may have eluded Renaissance artists. Leonardo da Vinci, while comparing painting and sculpture, criticized the realism afforded by reliefs [11]:

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 ($\lambda < 1$) and a surface “in the round” ($\lambda = 1$) will cast different shadows. However, Leonardo’s comment appears to overlook the fact that for any classical bas-relief transformation of the surface, there is a corresponding transformation of the light source direction such that the shadows are the same. This is not restricted to classical reliefs but, as we will show, applies equally to the greater set of generalized bas-relief transformations.

The fact that an object and every GBR transformation of the object produce the same shadow regions arises from an implicit duality. For each image of a Lambertian [15, 8] surface f viewed under orthographic projection (parallel lines of sight) and illuminated by a distant light source \mathbf{s} (e.g., the sun), there exists an identical image of a GBR surface \bar{f} with transformed albedo produced by a transformed light source $\bar{\mathbf{s}}$. This equality holds not only for the shadowed regions of the surfaces, but for the shading in the illuminated regions as well. Furthermore, due to superposition, the equality holds for an arbitrary – possibly infinite – number of light sources. It will be shown in the Appendix that for objects with convex shapes the generalized bas-relief transformation is the *only* transformation with this property: no other such ambiguity exists.

Thus, from a single viewpoint, there is an ambiguity in determining the 3-d Euclidean geometry of a surface: one can – at best – determine the relief of the surface up to a three parameter family of linear transformations. No information in either the shadowing or shading of the surface can resolve this. Yet, if the viewer moves relative to the surface, or the surface moves relative to the viewer, can this ambiguity be resolved?

As discussed by Helmholtz [27], image changes produced by an observer’s motion reveal both the depth and shape of viewed objects. For an object undergoing rigid motion and viewed under perspective projection, the object’s Euclidean structure can be determined from as few as two images [17, 20, 28]. If the object is viewed orthographically in two images, its structure can only be recovered up to a one parameter family of affine distortions. For infinitesimal motion under orthographic projection, there is a genuine bas-relief ambiguity: the shape of the surface can only be recovered up to a scale factor in the direction of the camera’s optical axis, i.e., a classical bas-relief transformation for which λ is unknown [13].

Beyond explaining the effectiveness of relief sculptures, the generalized bas-relief ambiguity also has implications for our understanding of human surface perception and for the development of computational vision algorithms. Our results support the recent psychophysical findings of [14] that for a variety of surfaces this ambiguity exists and is often unresolved in the human visual system. Likewise, the results suggest that the aim of structure recovery might be a weaker non-Euclidean representation, such as an affine representation [13, 21, 22, 26], a projective representation [5], or an ordinal representation [6]; for many applications, machine

vision systems need not resolve this ambiguity.

A summary of these and other results follow:

- The set of cast and attached shadows produced by a surface and a GBR transformed surface are identical, irrespective of the material type.
- If the material can modeled as having Lambertian reflectance, then the set of possible images including shadowing under any lighting condition (illumination cone [2]) for a surface and its transformed surface are identical. Therefore, these objects cannot be distinguished by any recognition algorithm.
- The generalized bas-relief transformation is the *only* transformation which has these first two properties.
- Under orthographic projection, the set of motion fields produced by a surface and its classical bas-relief are identical. Therefore, an object and its relief cannot be distinguished from small unknown camera motion.
- For photometric stereo where the light source directions are unknown, the structure can only be determined up to a generalized bas-relief transformation, and shadows do not provide further information. Using prior information about the albedo and light source strength, the structure can be determined up to a reflection in depth. Cast shadows can be used to distinguish these two cases.

Illustrating the GBR ambiguity, Fig. 2 shows four graphically rendered human heads: a “normal” head and three distorted heads obtained through a GBR transformation of the original. When the heads are observed frontally and under appropriately positioned light sources, the resulting images are so similar that their differences provide few cues to the true structure. Even when the head is rotated by a small amount, the ambiguity cannot be resolved. Only through a large motion (e.g., the side views) is the bas-relief transformation revealed.

2 Bas-Relief Ambiguity: Illumination

In this section we present details explaining the complicity of factors that give rise to the generalized bas-relief ambiguity. In particular, we show that there is a duality between a particular set of transformations of an object’s shape and the light sources which illuminate it. Here we consider distant illumination (parallel illuminating rays) of objects viewed under orthographic projection (parallel lines of sight).

Consider a surface observed under orthographic projection and define a coordinate system attached to the image plane such that the \mathbf{x} and \mathbf{y} axes span the image plane. In this coordinate system, the depth of every visible point in the scene can be expressed as

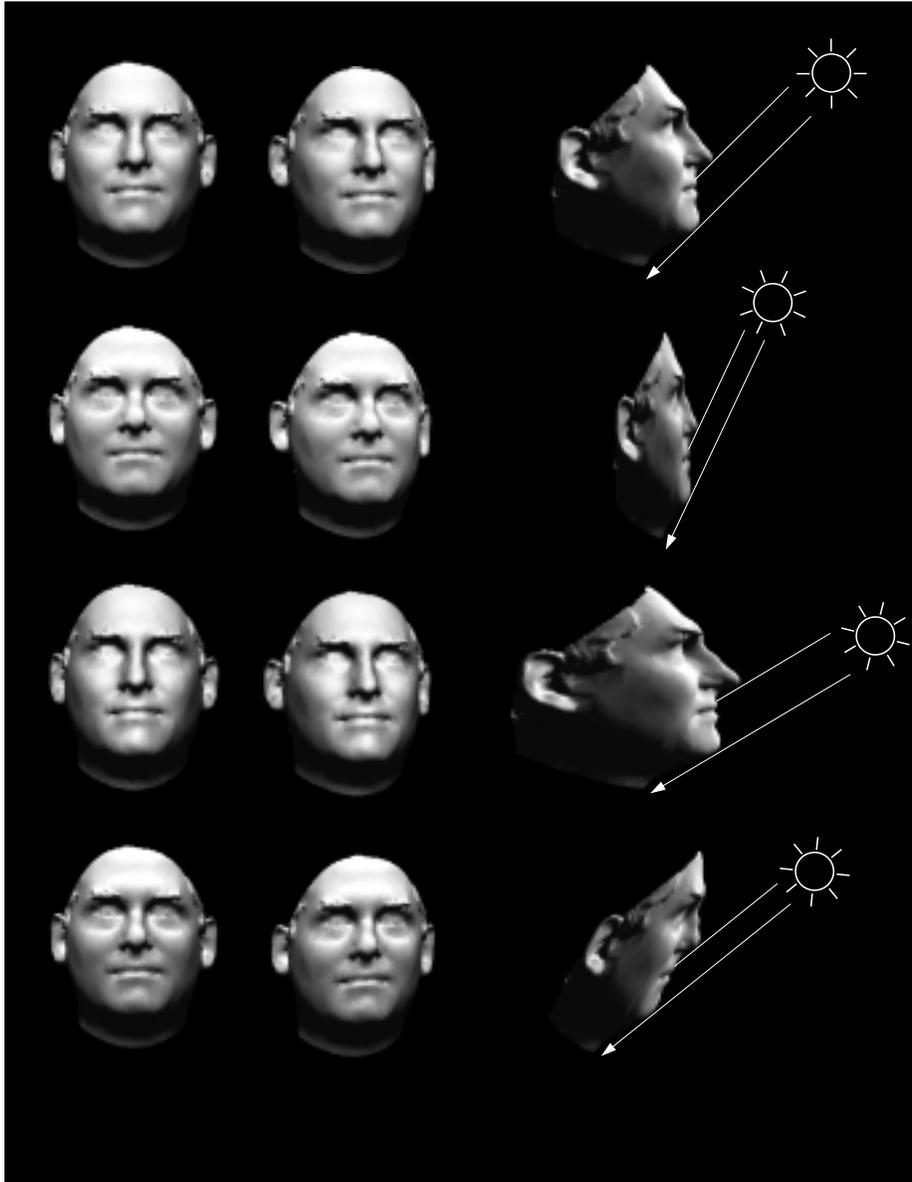


Figure 2: Three-dimensional data for the human head (top row) was obtained using a laser scan (Cyberwave) and rendered as a Lambertian surface with constant albedo (equal grey values for all surface points). The subsequent three rows show images of heads whose shapes have been transformed by different generalized bas-relief transformations, but whose albedos have not been transformed. The profile views of the face in the third column reveal the nature the individual transformations and the direction of the light source. The top row image is the true shape; the second from top is a flattened shape ($\lambda = 0.5$) (as are classical bas-reliefs); the third is an elongated shape ($\lambda = 1.5$); and the bottom is a flattened shape plus an additive plane ($\lambda = 0.7$, $\nu = 0.5$, and $\mu = 0.0$). The first column shows frontal views of the faces in the third column. From this view the true 3-d structure of the objects cannot be determined; in each image the shadowing patterns are identical, and even though the albedo has not been transformed according to Eq. 3, the shading patterns are so close as to provide few cues as to the true structure. The second column shows near frontal views of the faces from the same row, after having been separately rotated to compensate for the degree of the flattening or elongation. The rotation about the vertical axis is 7° for the first row of the second column; 14° for the second row; 4.6° for the third; and 14° for the fourth row. To mask the shearing produced by the additive plane, the fourth row has also been rotated by 5° about the line of sight.

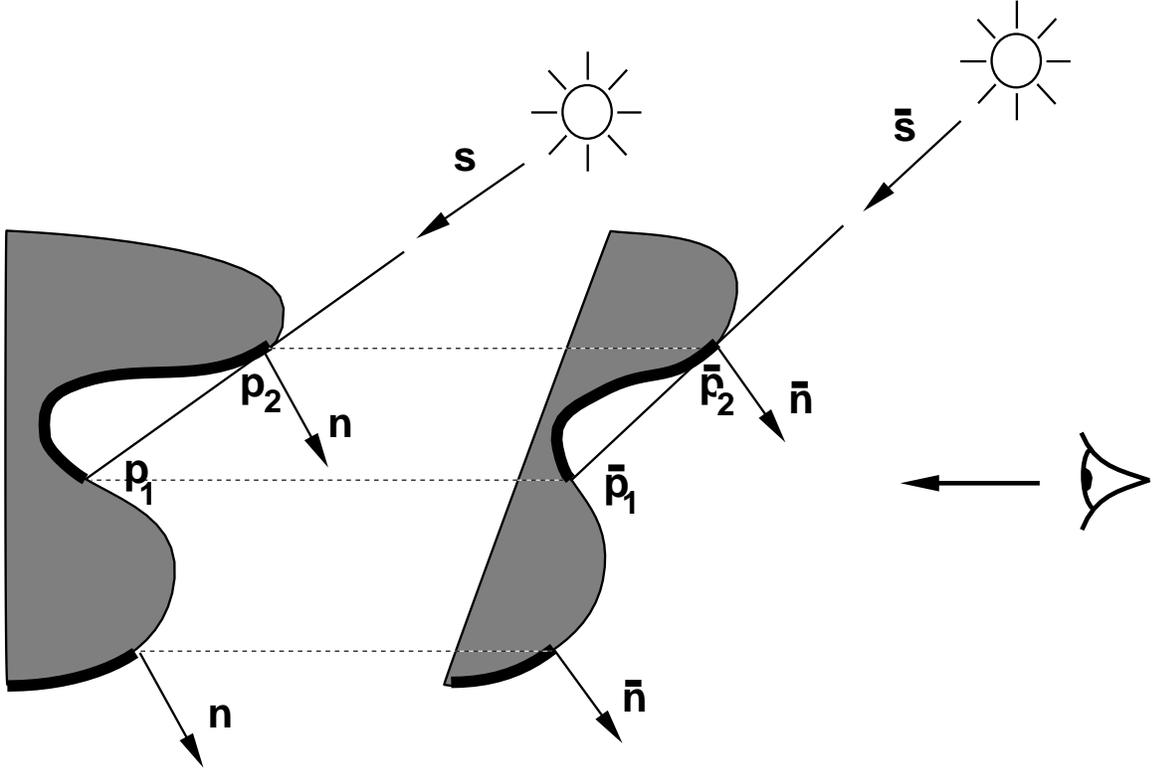


Figure 3: The image points that lie in shadow for a surface under light source \mathbf{s} are identical to those in shadow for a transformed surface under light source $\bar{\mathbf{s}} = G\mathbf{s}$. In this 2-d illustration, the lower shadow is an attached shadow while the upper one is composed of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

$$z = f(x, y)$$

where f is a piecewise differentiable function. The graph of $f(x, y)$, i.e., $(x, y, f(x, y))$, defines a surface which will also be denoted by f . The direction of the inward pointing surface normal $\mathbf{n}(x, y)$ can be expressed as

$$\mathbf{n}(x, y) = \begin{bmatrix} -f_x \\ -f_y \\ 1 \end{bmatrix} \quad (1)$$

where f_x and f_y denote the partial derivatives of f with respect to x and y respectively.

Consider transforming the surface f to a new surface \bar{f} in the following manner. We first flatten (or scale) it along the \mathbf{z} axis and then add a plane, i.e.,

$$\bar{f}(x, y) = \lambda f(x, y) + \mu x + \nu y$$

where $\lambda \neq 0$ [3]. We call this transformation the generalized bas-relief (GBR) transformation. See Figures 2 and 3. When $\mu = 0$ and $\nu = 0$, we call this transformation the classical bas-relief transformation, since for $\lambda < 1$ the surface is flattened like classical bas-relief sculptures.

Note that if $\mathbf{p} = (x, y, f(x, y))$ and $\bar{\mathbf{p}} = (x, y, \bar{f}(x, y))$, then $\bar{\mathbf{p}} = G\mathbf{p}$ where

$$G = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{bmatrix}. \quad (2)$$

Under the matrix product operation, the set $GBR = \{G\}$ forms a subgroup of $GL(3)$ with

$$G^{-1} = \frac{1}{\lambda} \begin{bmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ -\mu & -\nu & 1 \end{bmatrix}.$$

Also, note that for image point (x, y) , the relation between the direction of the surface normal of \bar{f} and f is given by $\bar{\mathbf{n}} = G^{-T}\mathbf{n}$ where $G^{-T} \equiv (G^T)^{-1} = (G^{-1})^T$. As will be seen in Section 4, this is the only linear transformation of the surface's normal field which preserves integrability.

Let the vector \mathbf{s} denote a point light source at infinity, with the magnitude of \mathbf{s} proportional to the intensity of the light source. (For a more general model of illumination, e.g., one that does not restrict light sources to be at infinity, see [16].) We first show that shadowing on a surface f for some light source \mathbf{s} is identical to that on a GBR transformed surface \bar{f} with an appropriate light source $\bar{\mathbf{s}}$; we then show that if the surfaces are Lambertian, the set of all possible images of both surfaces are identical.

We can identify two types of shadows: *attached shadows* and *cast shadows* [1, 23]. See Figure 3. A surface point $\mathbf{p} = (x, y, f(x, y))$ lies in an *attached shadow* for light source direction \mathbf{s} iff $\mathbf{n}(x, y)^T\mathbf{s} < 0$. This definition leads to the following lemma.

Lemma 2.1 *A point $\mathbf{p} = (x, y, f(x, y))$ lies in an attached shadow for light source direction \mathbf{s} iff $\bar{\mathbf{p}} = (x, y, \bar{f}(x, y))$ lies in an attached shadow for light source direction $\bar{\mathbf{s}} = G\mathbf{s}$.*

Proof. If a point \mathbf{p} on f lies in an attached shadow, then $\mathbf{n}^T\mathbf{s} < 0$. On the transformed surface, the point $\bar{\mathbf{p}} = (x, y, \bar{f}(x, y))$ also projects to (x, y) , and for this point $\bar{\mathbf{n}}^T\bar{\mathbf{s}} = (G^{-T}\mathbf{n})^TG\mathbf{s} = \mathbf{n}^T\mathbf{s}$. Therefore, $\bar{\mathbf{p}}$ is also in an attached shadow. The converse clearly holds as well. ■

A necessary condition for a point $\mathbf{p}_1 = (x_1, y_1, f(x_1, y_1))$ on the surface to fall on the *cast shadow boundary* for light source direction \mathbf{s} is that there exists another point $\mathbf{p}_2 = (x_2, y_2, f(x_2, y_2))$ on the surface such that the light ray in the direction \mathbf{s} passing through \mathbf{p}_2 grazes the surface at \mathbf{p}_2 and intersects the surface at \mathbf{p}_1 . The point \mathbf{p}_2 is the boundary of an attached shadow.

Lemma 2.2 *A point $\mathbf{p} = (x, y, f(x, y))$ satisfies the necessary condition for lying on a cast shadow boundary for light source direction \mathbf{s} iff $\bar{\mathbf{p}} = (x, y, \bar{f}(x, y))$ satisfies the condition for light source direction $\bar{\mathbf{s}} = G\mathbf{s}$.*

Proof. The condition for a point \mathbf{p}_1 to be on a shadow boundary cast by \mathbf{p}_2 is that

$$\begin{cases} \mathbf{n}_2^T \mathbf{s} = 0 \\ \mathbf{p}_2 - \mathbf{p}_1 = \gamma \mathbf{s} \end{cases}$$

for some $\gamma < 0$. For the transformed surface, the first condition for a point to be on the shadow boundary is

$$\bar{\mathbf{n}}_2^T \bar{\mathbf{s}} = (G^{-T} \mathbf{n}_2)^T G \mathbf{s} = \mathbf{n}_2^T \mathbf{s} = 0.$$

Under the GBR transformation $\bar{\mathbf{p}} = \lambda G \mathbf{p}$, and the second condition can be expressed for the relief surface as

$$\begin{aligned} \bar{\mathbf{p}}_2 - \bar{\mathbf{p}}_1 - \bar{\gamma} \bar{\mathbf{s}} &= G(\mathbf{p}_2 - \mathbf{p}_1) - \bar{\gamma} G \mathbf{s} \\ &= (\mathbf{p}_2 - \mathbf{p}_1) - \bar{\gamma} \mathbf{s} = 0. \end{aligned}$$

This condition clearly holds when $\bar{\gamma} = \gamma$. The converse of this lemma can be similarly proven. ■

This lemma becomes both necessary and sufficient for a point to lie on a shadow boundary when the ray from \mathbf{p}_1 passing through \mathbf{p}_2 does not intersect any other portion of the surface for both f and \bar{f} . In general, this is true when $\lambda > 0$.

Taking these two lemmas together, it follows that if some portion of the surface f is in a cast or attached shadow for a light source direction \mathbf{s} , then if the surface is subject to a GBR transformation G , there exists a lighting direction $\bar{\mathbf{s}} = G \mathbf{s}$ such that the same portion of the transformed surface is also shadowed. Let us specify these shadowed regions – both attached and cast – through a binary function $\Psi_{f,\mathbf{s}}(x, y)$ such that

$$\Psi_{f,\mathbf{s}}(x, y) = \begin{cases} 0 & \text{if } (x, y) \text{ is shadowed} \\ 1 & \text{otherwise.} \end{cases}$$

Using this notation and the above two lemmas, we can then write $\Psi_{f,\mathbf{s}}(x, y) = \Psi_{\bar{f},\bar{\mathbf{s}}}(x, y)$.

We should stress that the shadowing $\Psi_{f,\mathbf{s}}(x, y)$ from direct illumination by a light source is a function of the object's geometry – it is unaffected by the reflectance properties of the surface.¹ For any surface, any light source direction, and any GBR transformation of that surface, there exists a light source direction such that the shadowing will be identical. Furthermore, the GBR transformation is the *only* transformation for which this is true. See the Appendix for a proof of this for objects with convex shape.

We now show that if the surface reflectance is Lambertian [8, 15], then the sets of images produced by a surface (i.e., the surface's illumination cone [2]) and a transformed surface under

¹It should be noted that the image location of a specularity is not preserved under GBR, and so shadows arising due to indirect illumination of the surface from the virtual light source of a specularity will not be preserved under GBR.

all possible lighting conditions are identical. Letting the albedo of a Lambertian surface f be denoted by $a(x, y)$, the intensity image produced by a light source \mathbf{s} can be expressed as

$$\mathbf{I}_{f,a,\mathbf{s}}(x, y) = \Psi_{f,\mathbf{s}}(x, y)\mathbf{b}(x, y)^T\mathbf{s}$$

where $\mathbf{b}(x, y)$ is the product of the albedo $a(x, y)$ of the surface and the inward pointing unit surface normal $\hat{\mathbf{n}}(x, y)$.

The set of images produced by f and \bar{f} are identical when \bar{f} has albedo $\bar{a}(x, y)$ given by

$$\bar{a} = \frac{a}{\lambda} \sqrt{(\lambda\hat{n}_1 - \mu\hat{n}_3)^2 + (\lambda\hat{n}_2 - \nu\hat{n}_3)^2 + \hat{n}_3^2} \quad (3)$$

where $\hat{\mathbf{n}} = (\hat{n}_1, \hat{n}_2, \hat{n}_3)^T$. The effect of applying Eq. 3 to a classical bas-relief transformation $0 < \lambda < 1$ is to darken points on the surface where \mathbf{n} points away from the optical axis.

This transformation on the albedo is subtle and warrants discussion. For λ close to unity, the transformation on albedo is nearly impossible to detect. That is, if you transform the shape of a surface by a GBR transformation, but leave the albedo unchanged, then the differences in the images produced under varying illumination are too small to reveal the structure. In Fig. 2, we left the albedo unchanged, $\bar{a}(x, y) = a(x, y)$, and even though λ ranges from 0.5 to 1.5, the differences in shape cannot not be discerned from the frontal images. However, when the albedo is unchanged and the flattening is more severe, e.g., tenfold ($\lambda = 0.1$), the shading patterns can reveal the flatness of the surface. This effect is often seen on very low relief sculptures (e.g., Donatello's *rilievo schiacciato*) which reproduce shadowing accurately, but shading poorly.

For the set of images to be identical it is necessary that the albedo be transformed along with the surface and $\lambda > 0$. When $\lambda < 0$, the surface \bar{f} is inverted (as in *intaglio*); for a corresponding transformation of the light source $\bar{\mathbf{s}}$, the illuminated regions of the original surface f and the transformed surface \bar{f} will be the same. This is the well known ‘‘up/down’’ (convex/concave) ambiguity. However, the shadows cast by \bar{f} and f may differ quite dramatically.

Lemma 2.3 *For each light source \mathbf{s} illuminating a Lambertian surface $f(x, y)$ with albedo $a(x, y)$, there exists a light source $\bar{\mathbf{s}}$ illuminating a surface $\bar{f}(x, y)$ (a GBR transformation of f) with albedo $\bar{a}(x, y)$ (as given in Eq. 3), such that $I_{f,a,\mathbf{s}}(x, y) = I_{\bar{f},\bar{a},\bar{\mathbf{s}}}(x, y)$.*

Proof. The image of f is given by

$$I_{f,a,\mathbf{s}}(x, y) = \Psi_{f,\mathbf{s}}(x, y)\mathbf{b}^T(x, y)\mathbf{s}$$

For any 3×3 invertible matrix A , we have that

$$I_{f,a,\mathbf{s}}(x, y) = \Psi_{f,\mathbf{s}}(x, y)\mathbf{b}^T(x, y)A^{-1}A\mathbf{s}.$$

Since GBR is a subgroup of $GL(3)$ and $\Psi_{f,\mathbf{s}}(x, y) = \Psi_{\bar{f},\bar{\mathbf{s}}}(x, y)$,

$$\begin{aligned} I_{f,a,\mathbf{s}}(x, y) &= \Psi_{f,\mathbf{s}}(x, y)\mathbf{b}^T(x, y)G^{-1}G\mathbf{s} \\ &= \Psi_{\bar{f},\bar{\mathbf{s}}}(x, y)\bar{\mathbf{b}}^T(x, y)\bar{\mathbf{s}} \\ &= I_{\bar{f},\bar{a},\bar{\mathbf{s}}}(x, y) \end{aligned}$$

where $\bar{\mathbf{b}}(x, y) = G^{-T}\mathbf{b}(x, y)$ and $\bar{\mathbf{s}} = G\mathbf{s}$. ■

With the above three lemmas in hand, we can now state and prove the central proposition of this section:

Proposition 2.1 *The set of images under all possible lighting conditions produced by a Lambertian surface f with albedo $a(x, y)$ and those surfaces \bar{f} differing by any GBR transformation with albedo $\bar{a}(x, y)$ given by Eq. 3 are identical.*

Proof. From Lemmas 2.1, 2.2, and 2.3, we have that the image of a surface f produced by a single light source \mathbf{s} is the same as the image of a GBR transformed surface \bar{f} produced by the transformed light source $\bar{\mathbf{s}} = G\mathbf{s}$, i.e., $I_{f,a,\mathbf{s}}(x, y) = I_{\bar{f},\bar{a},\bar{\mathbf{s}}}(x, y)$. When the object is illuminated by a set of light sources $\{\mathbf{s}_i\}$, then the image is determined by the superposition of those images that would be formed under the individual light sources. Similarly, the same image can be produced from the transformed surface if it is illuminated by the set of light sources given by $\{\bar{\mathbf{s}}_i\}$, where $\bar{\mathbf{s}}_i = G\mathbf{s}_i$. ■

Taken together, the above results demonstrate that when both the surface and light source direction are transformed by G , both the shadowing and shading are identical in the images of the original and transformed surface. An implication of this result is that given any number of images taken from a fixed viewpoint, neither a computer vision algorithm nor biological process can distinguish two objects that differ by a GBR transformation. Knowledge (or assumptions) about surface shape, surface albedo, light source direction, or light source intensity must be employed to resolve this ambiguity. See again Fig. 2.

3 Bas-Relief Ambiguity: Motion

While neither the shading nor shadowing of an object, seen from a single viewpoint, reveals the exact 3-d structure, motion does provide additional cues [27]. If the surface undergoes a rigid motion and is viewed under perspective projection, the object's Euclidean structure can be determined from as few as two images [17, 20, 28]. If the object is viewed orthographically, the object's structure can only be determined up to a one parameter family of affine distortions from two images [13]. To determine the Euclidean structure under orthographic projection, at least three images, taken from separate viewpoints, are needed.

Yet, complications arise when the object's motion is small. For infinitesimal motion under perspective projection, the structure estimates are sensitive to noise, producing an implicit

error in the estimate of the relief of the surface [19, 25]. For small (infinitesimal) unknown motion under orthographic projection, there is a genuine bas-relief ambiguity: the shape of the surface can only be recovered up to a scale factor in the direction of the camera’s optical axis, i.e., a classical bas-relief transformation ($\lambda > 0, \mu = \nu = 0$).

To see this, let us assume that the surface does, in fact, undergo an arbitrary infinitesimal motion. The velocity $\dot{\mathbf{p}} = (\dot{x}, \dot{y}, \dot{z})$ of a point $\mathbf{p}(x, y, z)$ on the surface f induces a velocity (\dot{x}, \dot{y}) in the image plane. The collection of velocities for all points in the image plane is often called the motion field [8]. In the following proposition, we show that the set of motion fields induced by all 3-d infinitesimal motions of a surface f is the same, under orthographic projection, as the set of all motion fields of a surface differing by a classical bas-relief transformation (not a generalized bas-relief transformation).

Proposition 3.1 *The set of motion fields induced by all 3-D infinitesimal motions of a surface f is the same, under orthographic projection, as the set of all motion fields of a surface differing by a bas-relief transformation $\bar{f}(x, y) = \lambda f(x, y)$ where $\lambda \neq 0$.*

Proof. Let $\Omega = (\Omega_x, \Omega_y, \Omega_z)$ and $\mathbf{v} = (\mathbf{v}_x, \mathbf{v}_y, \mathbf{v}_z)$ respectively denote the angular and linear velocity of the surface f with respect to the observer. The 3-d velocity of a point \mathbf{p} on f is $\dot{\mathbf{p}} = \Omega \times \mathbf{p} + \mathbf{v}$. For the transformed surface, $\bar{\mathbf{p}} = G\mathbf{p}$ with $\mu = \nu = 0$, and the velocity of the point is given by $\dot{\bar{\mathbf{p}}} = \bar{\Omega} \times (G\mathbf{p}) + \bar{\mathbf{v}}$. It is easy to show that the motion fields for both surfaces will be identical (i.e., $\dot{x} = \dot{\bar{x}}$ and $\dot{y} = \dot{\bar{y}}$) when $\bar{\mathbf{v}}_x = \mathbf{v}_x$, $\bar{\mathbf{v}}_y = \mathbf{v}_y$, $\bar{\Omega}_z = \Omega_z$, $\bar{\Omega}_x = \Omega_x/\lambda$ and $\bar{\Omega}_y = \Omega_y/\lambda$. That is, the component of the angular velocity parallel to the image plane is scaled inversely with respect to the relief. Thus, for every motion f there is a motion of \bar{f} that will yield the same motion field and so the set of motion fields is identical. ■

This proof follows the results in [13, 21]. An implication of Proposition 3.1 is that under orthographic projection, a small motion of either the object or the observer cannot resolve the bas-relief ambiguity. Furthermore, since the motion field is linear in $f(x, y)$, the classical bas-relief transformation is the *only* transformation of f that will preserve the set of motion fields.

Revisit the second column in Figure 2. The image produced by the “normal” relief after a rotation of 7° from frontal is nearly identical to the images produced by a rotation of 14° for the flattened head and 4.6° for the elongated head. No rotation, however, will completely disguise the distortions produced by the bottom image in the second row where $\nu \neq 0$; here we rotated the head by 10° around the vertical axis and then by 5° about the line of sight.

4 Integrability, Reconstruction, and the Bas-Relief Ambiguity

In this section, we investigate the role of the generalized bas-relief ambiguity on surface reconstruction using photometric stereo. Let us assume that a Lambertian surface is illuminated by

a point light source at infinity. When there is no shadowing (i.e., $\Psi_{f,\mathbf{s}}(x, y) = 1$), the intensity image produced by a light source \mathbf{s} can be expressed as

$$\mathbf{I}_{f,a,\mathbf{s}}(x, y) = \mathbf{b}(x, y)^T \mathbf{s} \quad (4)$$

where $\mathbf{b}(x, y)$ is the product of the albedo $a(x, y)$ of the surface and the inward pointing unit surface normal $\hat{\mathbf{n}}(x, y)$. From multiple images of the object seen from a fixed viewpoint but with different light source direction, we can solve Eq. 4 for \mathbf{b} when the light source strengths and directions are known. This, of course, is the standard photometric stereo technique, see [8, 24, 29].

However, if the light source strengths and directions are *not known*, then we can only determine the vector field $\mathbf{b}(x, y)$ of surface normals and albedos up to a 3×3 linear transformation. For any invertible 3×3 linear transformation $A \in GL(3)$ [7, 3, 21]

$$\mathbf{b}^T \mathbf{s} = (A\mathbf{b})^T A^{-T} \mathbf{s}. \quad (5)$$

If $\mathbf{b}(x, y)$ is the true vector field of surface normals then the recovered vector field $\mathbf{b}^*(x, y)$ is any vector field in the orbit of $\mathbf{b}(x, y)$ under the group $GL(3)$. For a pixelated image with no surface point in shadow, \mathbf{b}^* can be estimated from a collection of images using singular value decomposition; when some of the surface points are shadowed, Jacobs' method can be used to estimate \mathbf{b}^* [10]. Note, however, that not all vector fields $\mathbf{b}^*(x, y)$ correspond to continuous (or even piecewise continuous) surfaces. We will use this observation to restrict the group of allowable transformations on $\mathbf{b}(x, y)$ [3].

If \mathbf{b} is transformed by an arbitrary $A \in GL(3)$ (i.e., any vector field $\mathbf{b}^*(x, y)$ in the orbit of \mathbf{b} under $GL(3)$), then in general, there will be no surface $f^*(x, y)$ with unit normal field $\hat{\mathbf{n}}^*(x, y)$ and albedo $a^*(x, y)$ that could have produced the vector field $\mathbf{b}^*(x, y)$. For $f^*(x, y)$ to be a surface, it must satisfy the following integrability constraint [9]:

$$f_{xy}^* = f_{yx}^*$$

which, in turn, means $\mathbf{b}^*(x, y)$ must satisfy

$$\begin{pmatrix} b_1^* \\ b_3^* \end{pmatrix}_y = \begin{pmatrix} b_2^* \\ b_3^* \end{pmatrix}_x \quad (6)$$

where $\mathbf{b}^* = (b_1^*, b_2^*, b_3^*)^T$ and the subscripts x and y denote partial derivatives.

Proposition 4.1 *If $\mathbf{b}(x, y)$ corresponds to a surface $f(x, y)$ with albedo $a(x, y)$, then the set of linear transformation $\mathbf{b}^*(x, y) = A\mathbf{b}(x, y)$ which satisfy the integrability constraint in Eq. 6 are $A = G^{-T}$ where the generalized bas-relief transformations G is given in Eq. 2.*

Proof. The integrability constraint given in Eq. 6 can be written as $(b_{1_y}^* - b_{2_x}^*)b_3^* + b_{3_x}^*b_2^* - b_{3_y}^*b_1^* = 0$. Letting A_{ij} be the i, j -th element of A , and recalling that $\mathbf{b}^* = A\mathbf{b}$, the left hand side is a function of $b_i(x, y)$, $b_{i_x}(x, y)$ and $b_{i_y}(x, y)$ for $i = 1, 2, 3$. Since these functions are generally independent, the coefficients of these function must all vanish for the integrability constraint to hold for all (x, y) . This leads to the following algebraic constraints on the elements of A .

$$\begin{cases} A_{22}A_{31} - A_{21}A_{32} = 0 \\ A_{21}A_{33} - A_{23}A_{31} = 0 \\ A_{12}A_{33} - A_{13}A_{32} = 0 \\ A_{12}A_{31} - A_{11}A_{32} = 0 \\ A_{22}A_{33} - A_{11}A_{33} + A_{13}A_{31} - A_{32}A_{23} = 0 \end{cases}$$

Since this system is homogeneous, for any A satisfying this system, ρA also satisfies the system; varying ρ corresponds to changing the light source intensity while making a corresponding global scaling of the albedo function. It can be shown that if $A_{33} = 0$, the matrix A satisfying the constraints is singular. So we can let $A_{33} = 1$, and solve for the remaining coefficients. The only nonsingular solution is $A_{11} = A_{22}$ and $A_{12} = A_{21} = A_{31} = A_{32} = 0$. That is, A must be a generalized bas-relief transformation. ■

The choice of $\mathbf{b}^*(x, y)$ is, of course, not unique since $\mathbf{b}^*(x, y) = G\mathbf{b}$ satisfies the integrability constraint for any $G \in GBR$. Yet, every \mathbf{b}^* has a corresponding surface f^* with a corresponding albedo $a(x, y)$, and these surfaces differ by a GBR ambiguity. Thus, if we have at least three images – each acquired under different light source directions – of a Lambertian surface $f(x, y)$, then by imposing the integrability constraint in Eq. 6, we can recover the surface $f(x, y)$ up to a GBR transformation $\tilde{f}(x, y) = \lambda f(x, y) + \mu x + \nu y$. See [30] for a method to estimate an integrable b^* from image data. Note that no information given in the image shadows can resolve this ambiguity, as Section 2 showed that the set of all possible images of a surface $f(x, y)$ is invariant under the GBR transformation. It should be noted that Fan and Wolff showed that the Hessian of $f(x, y)$ can be determined from three images up to scale factor using the integrability constraint to enforce equality of the off-diagonal elements, $f_{xy} = f_{yx}$ [4]. The unknown scale of the Hessian corresponds to the parameter λ and the unknown initial conditions of Fan and Wolff’s differential equation correspond to the parameters μ and ν .

If, however, we have additional information about the albedo or the strength of the light sources we can further restrict the ambiguity.

Corollary 4.1 *If the albedo $a(x, y)$ is constant (or known), or the light sources \mathbf{s}_i all have the same intensity, then the GBR ambiguity G is restricted to the binary subgroup given by $\lambda = \pm 1, \mu = 0$, and $\nu = 0$.*

Proof. If $a(x, y) = |\mathbf{b}(x, y)|$ is constant (or known), then for $|\mathbf{b}(x, y)| = |\mathbf{b}^*(x, y)| = |A\mathbf{b}(x, y)|$, A must preserve length for any \mathbf{b} . The only matrices that preserve length are the orthonormal matrices. The only orthonormal matrices that are also GBR transformations

correspond to $\lambda = \pm 1, \mu = 0$, and $\nu = 0$. A similar argument holds when the light source intensities are known. ■

Thus, we can determine the true surface up to a sign, i.e., $\bar{f}(x, y) = \pm f(x, y)$. This is the classical in-out ambiguity that occurs in shape from shading [8, 18]. Note however, that the shadowing configurations change when λ changes sign, and if shadowing is present, this ambiguity can be resolved.

5 Conclusion

We have shown that under any lighting condition, the shading and shadowing on an object with Lambertian reflectance are identical to the shading and shadowing on any generalized bas-relief transformation of the object. The GBR transformation is unique in that it is the only transformation of the surface having this property. Thus, from a single viewpoint, there is an ambiguity in the recovery of the surface: we can only determine the relief of the surface up to a three parameter family of linear transformations. No information in either the shadowing or shading can resolve this. Furthermore, not even the motion fields produced by small motions of the viewer (or object) can resolve the surface relief.

Leonardo da Vinci’s statement in the introduction 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 slab 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 [12]. Barring the effects of constant albedo, the shading and shadowing for each figure is self consistent; however, the shadows cast from one figure onto another are incorrect. Furthermore, the shadows cast onto the background slab, whose orientation usually does not correspond to that of a wall or floor in the scene, are also inconsistent. 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 such as the ones shown in Figure 1.

Putting the subtleties of relief sculpture aside, we should point out that while shadowing is preserved exactly under GBR transformations of an object, there are certain shading effects which are not. Specularities arising from non-Lambertian (glossy) surfaces and the effect of inter-reflection of light from one part of a surface onto another depend on the surface relief. Nevertheless, these effects may be secondary in that they may not allow a human observer to resolve the GBR ambiguity – even when viewing a known object. Recently, Koenderink, Van Doorn and Christon performed a series of psychophysical experiments in

which subjects observe images of sculptures (one abstract and one torso) and provide estimates of the orientation of the surface normals at about 300 points [14]. Since the estimated vector field of normals satisfies an integrability constraint for all subjects, this experiment provides evidence that humans maintain a surface representation of an observed scene. Furthermore, while the subjects correctly estimate the overall shape of the surface, they consistently misestimate its relief and slant.

The nature of GBR ambiguity and its apparent presence in the human visual system suggest that machine vision systems will also be similarly impaired. One objective of computer vision is the recovery of models of surfaces from multiple images. In photometric stereo, the Euclidean structure is estimated from multiple images of a scene taken from a fixed viewpoint, but under different lighting conditions for which the illuminant directions are known [8]. While it had been thought that photometric stereo with unknown light source directions could be solved by first estimating the light source directions and then estimating the surface structure, this paper has shown that these estimates are coupled through a GBR transformation. The only way to resolve these ambiguities is to use additional information beside that contained in the image data [7].

Recently, it has been proposed that structure recovery methods should be stratified according to the available information about the image formation process [5, 13, 21]. So for example, the Euclidean structure of a scene observed under perspective can be recovered from two images when the camera’s intrinsic parameters (e.g., focal length, principal point) are known. When they are unknown, the structure can only be recovered up to a projective transformation. Here, we introduce a new layer of the stratification between affine and Euclidean structure. These non-Euclidean representations (affine [13, 21, 22, 26], projective [5] or ordinal [6]), can still be used to solve numerous vision-based tasks such as object recognition, vehicle navigation, robotic manipulation and synthetic image generation without resolution of these ambiguities.

Appendix: Uniqueness of the Generalized Bas-Relief Transformation

Here we prove that the generalized bas-relief transformation is unique in that there is no other transformation of the object’s surface which preserves the set of images 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 and, thus, the set of images 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 $\mathbf{n}^T \mathbf{s} = 0$, for some \mathbf{s} . 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 $\mathbf{s} = (s_1, s_2, s_3)^T$ denote a point light source at infinity, but let us equate all light sources producing the same attached shadow boundary, i.e., $(s_1, s_2, s_3)^T = (ks_1, ks_2, ks_3)^T \forall k \in \mathbb{R}, k \neq 0$. 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 $(s_1, s_2, 0)$.

Let the triple $\mathbf{n} = (n_1, n_2, n_3)^T$ denote a surface normal. Again, the magnitude and sign of the surface normal are unimportant, so we equate $(n_1, n_2, n_3)^T = (kn_1, kn_2, kn_3)^T \forall k \in \mathbb{R}, k \neq 0$. Thus, the space of surface normals \mathcal{N} is, likewise, equivalent to \mathbb{RP}^2 . Note that under the equation $\mathbf{n}^T \mathbf{s} = 0$, 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 $\mathbf{n}^T \mathbf{s} = 0$, 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{IP}^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{IP}^2 . Under this isomorphism, we have a bijection 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 $\tilde{f}(x, y)$. If the objects have the same set of imaged attached shadow boundaries as seen in the image plane (i.e., if the set of image contours produced by orthographically projecting the attached shadow boundaries is the same for both objects), then the question arises: How are the two surfaces $f(x, y)$ and $\tilde{f}(x, y)$ related?

Proposition 5.1 *If two surfaces $f(x, y)$ and $\tilde{f}(x, y)$ produce the same set of attached shadow boundaries as seen in the image plane, 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{IP}^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{IP}^2) to the surface normal $\mathbf{n}(x, y)$. In the same manner, we can construct a projective isomorphism between \mathbb{IP}^2 and the real projective plane of light source directions $\tilde{\mathcal{S}}$ illuminating the surface $\tilde{f}(x, y)$. The isomorphism is, likewise, chosen to map the same

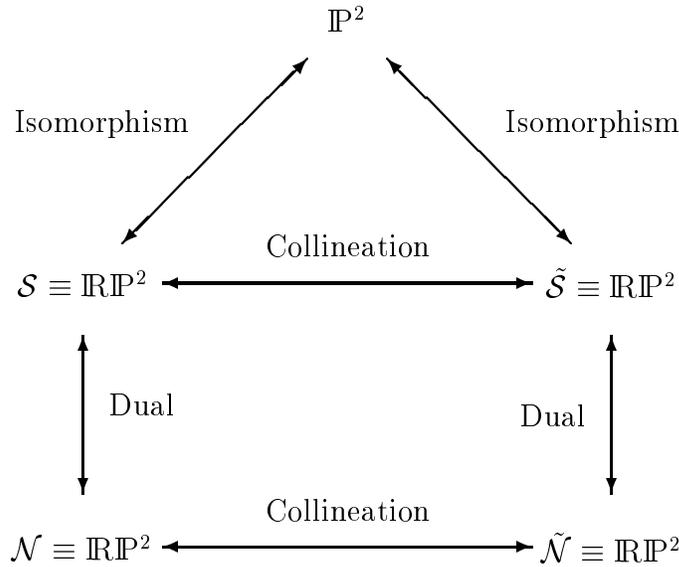


Figure 4: The relationship of different spaces in proof of Proposition 5.1.

collection of imaged attached shadow boundaries passing through (x, y) in the image plane to the surface normal $\tilde{\mathbf{n}}(x, y)$. Under these two mappings, we have a projective isomorphism between \mathcal{S} and $\tilde{\mathcal{S}}$ which in turn is a projective transformation (collineation). Because \mathcal{N} and $\tilde{\mathcal{N}}$ are the duals of \mathcal{S} and $\tilde{\mathcal{S}}$ respectively, the surface normals of $f(x, y)$ are also related to the surface normals of $\tilde{f}(x, y)$ by a projective transformation, i.e., $\tilde{\mathbf{n}}(x, y) = P\mathbf{n}(x, y)$ where P is a 3×3 matrix in the general projective group $\text{GP}(3)$.

The transformation P is further restricted in that the surface normals along the occluding contour of f and \tilde{f} are equivalent, i.e., the transformation P pointwise fixes the line at infinity of surface normals. Thus, P is of the form given by the generalized bas-relief transformation in Eq. 2, and the surfaces, in turn, must be related by a generalized bas-relief transformation. ■

Acknowledgments

Many thanks to David Mumford for leading us to the proof of Proposition 5.1 and to Niko Troje from Max-Planck-Institute, Tuebingen for the data used in Fig. 2.

References

- [1] M. Baxandall. *Shadows and Enlightenment*. Yale University Press, New Haven, 1995.
- [2] 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.

- [3] R. Epstein, A. Yuille, and P. N. Belhumeur. Learning and recognizing objects using illumination subspaces. In *Proc. of the Int. Workshop on Object Representation for Computer Vision*, 1996.
- [4] J. Fan and L. Wolff. Surface curvature and shape reconstruction from unknown multiple illumination and integrability. *Computer Vision and Image Understanding*, 65(2):347–359, February 1997.
- [5] O. Faugeras. Stratification of 3-D vision: Projective, affine, and metric representations. *J. Opt. Soc. Am. A*, 12(7):465–484, 1995.
- [6] C. Fermuller and Y. Aloimonos. Ordinal representations of visual space. In *Proc. Image Understanding Workshop*, pages 897–904, 1996.
- [7] H. Hayakawa. Photometric stereo under a light-source with arbitrary motion. *JOSA-A*, 11(11):3079–3089, Nov. 1994.
- [8] B. Horn. *Computer Vision*. MIT Press, Cambridge, Mass., 1986.
- [9] B. Horn and M. Brooks. The variational approach to shape from shading. In *Computer Vision, Graphics and Image Processing*, volume 33, pages 174–208, 1986.
- [10] D. Jacobs. Linear fitting with missing data: Applications to structure from motion and characterizing intensity images. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, 1997.
- [11] M. Kemp, editor. *Leonardo On Painting*. Yale University Press, New Haven, 1989.
- [12] M. Kemp. *The Science of Art: Optical Themes in Western Art from Brunelleschi to Seurat*. Yale University Press, New Haven, 1990.
- [13] J. Koenderink and A. Van Doorn. Affine structure from motion. *JOSA-A*, 8(2):377–385, 1991.
- [14] J. J. Koenderink, A. J. Van Doorn, and C. Christon. Shape constancy in pictorial relief. In J. Ponce, A. Zisserman, and M. Hebert, editors, *Object Representation in Computer Vision II*, pages 151–164. Springer, 1996.
- [15] J. Lambert. *Photometria Sive de Mensura et Gradibus Luminus, Colorum et Umbrae*. Eberhard Klett, 1760.
- [16] M. Langer and S. Zucker. What is a light source? In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 172–178, 1997.

- [17] H. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, 1981.
- [18] J. Oliensis. Uniqueness in shape from shading. *Int. J. Computer Vision*, 6(2):75–104, June 1991.
- [19] J. Oliensis. Structure from linear or planar motions. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 335–342, June 1996.
- [20] J. Roach and K. Aggarwal, J. Computer tracking of objects from a sequence of images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 1(2):127–135, 1979.
- [21] R. Rosenholtz and J. Koenderink. Affine structure and photometry. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 790–795, 1996.
- [22] L. Shapiro, A. Zisserman, and M. Brady. 3D motion recovery via affine epipolar geometry. *Int. J. Computer Vision*, 16(2):147–182, October 1995.
- [23] A. Shashua. *Geometry and Photometry in 3D Visual Recognition*. PhD thesis, MIT, 1992.
- [24] W. Silver. *Determining Shape and Reflectance Using Multiple Images*. PhD thesis, MIT, Cambridge, MA, 1980.
- [25] R. Szeliski and S. Kang. Shape ambiguities in structure from motion. In *European Conf. on Computer Vision*, pages I:709–721, 1996.
- [26] S. Ullman and R. Basri. Recognition by a linear combination of models. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 13:992–1006, 1991.
- [27] H. von Helmholtz. *Handbuch der Physiologischen Optik*. Verlag von Leopold Voss, Hamburg, Germany, 1910.
- [28] A. Waxman and S. Ullman. Surface structure and three-dimensional motion from image flow kinematics. *Int. J. Robotics Research*, 4:72–92, 1985.
- [29] R. Woodham. Analysing images of curved surfaces. *Artificial Intelligence*, 17:117–140, 1981.
- [30] A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pages 158–164, 1997.