

In Search of Illumination Invariants

Hansen F. Chen
Department of Physics
Yale University
New Haven, CT 06520-8120
hansen.chen@yale.edu

Peter N. Belhumeur*
Departments of EE and CS
Yale University
New Haven, CT 06520-8267
belhumeur@yale.edu

David W. Jacobs
NEC Research Institute
Princeton, NJ 08540
dwj@research.nj.nec.com

Abstract

We consider the problem of determining functions of an image of an object that are insensitive to illumination changes. We first show that for an object with Lambertian reflectance there are no discriminative functions that are invariant to illumination. This result leads us to adopt a probabilistic approach in which we analytically determine a probability distribution for the image gradient as a function of the surface's geometry and reflectance. Our distribution reveals that the direction of the image gradient is insensitive to changes in illumination direction. We verify this empirically by constructing a distribution for the image gradient from more than 20 million samples of gradients in a database of 1,280 images of 20 inanimate objects taken under varying lighting conditions. Using this distribution, we develop an illumination insensitive measure of image comparison and test it on the problem of face recognition.

1 Introduction

Changes in viewpoint and illumination [1] can dramatically alter the appearance of an object. We focus here solely on changes in illumination and ask: Are there discriminative illumination invariants? If not, are there local image measurements that are at least insensitive to illumination changes? (The geometric analogue of this problem has been considered extensively [29, 12, 6, 8].)

We show that even for objects with Lambertian reflectance [23], there are no discriminative functions of images of objects that are invariant to illumination. This differs from earlier findings in that we do not assume a homogeneous BRDF [21, 5], coplanarity [31], or consider invariance based on multiple images [40]. We show that for any two images — whether or not they are of the same object — there is always a family of surfaces, albedo patterns, and light sources that could have produced them. Our proof on the absence of the illumination invariants significantly advances [27] in that we prove this for objects composed of surfaces — not freely floating planar facets in space.

This result suggests that in comparing images for recognition, alignment, or tracking, one must fall back on probabilistic measures of comparison. (Of course, if one has multiple training images of the object under varying illumination, one could extract invariants [40] or construct representations for modeling the illumina-

tion variation [35, 16, 4, 14].) We develop a measure of image comparison by considering illumination to be a random variable which gives rise to the apparent randomness in local image measurements. Specifically, we derive the probability distribution of the image gradient of a point on a surface as determined by the differential geometric and reflectance properties of that point. Our distribution reveals that the direction of the image gradient (which is perpendicular to the flow field [5]) is insensitive to changes in illumination direction. Using the image gradient distribution, we then develop an illumination insensitive measure of image comparison.

Finally, we verify our qualitative theoretical arguments by empirically constructing both our distribution for the image gradient and our illumination insensitive measure of image comparison. We do this using a database of 1,280 images of 20 objects taken under varying illumination direction. We then test our illumination insensitive measure on the problem of face recognition, and compare our performance on 450 images of 10 individuals to other existing methods.

2 Illumination Invariants?

Is it always possible to determine whether two images were created by the same object under two different lighting conditions — or by different objects? While this seems possible, in this section we show the contrary.

To analyze this question, we study the existence of discriminative illumination invariants: functions of an image which are invariant to illumination on the object but vary with the object identity. Formally, let \mathcal{O} be some set of rigid objects, including their optical properties; \mathcal{S} be certain illumination conditions; and \mathcal{I} be the set of all images, or the piecewise smooth nonnegative valued functions of \mathcal{O} under \mathcal{S} . Function $Q : \mathcal{O} \times \mathcal{S} \rightarrow \mathcal{I}$ gives the image $I \in \mathcal{I}$ of object $o \in \mathcal{O}$ under illumination $s \in \mathcal{S}$, i.e., $I = Q(o, s)$. Note, function Q , defined this way, may not be surjective.

We adopt the following definitions:

Definition 1 A function μ on \mathcal{I} is invariant to illumination $\iff \mu(Q(o, s)) = \mu(Q(o, l)), \forall s, l \in \mathcal{S}, o \in \mathcal{O}$.

Definition 2 A illumination invariant μ is nondiscriminative for object set $\mathcal{O} \iff \mu(I) = \mu(J), \forall I \neq J$, where I and J are in the range of Q . A illumination invariant is discriminative iff it is not nondiscriminative.

This definition implies μ does not depend on o for

*P. N. Belhumeur was supported by a Presidential Early Career Award for Scientists and Engineers, an NSF Career Award IRI-9703134, ARO grant DAAH04-95-1-0494, and by the NEI.

nondiscriminative invariants.

Lemma 2.1 *There are no discriminative illumination invariants for \mathcal{O} if for any two images I and J in the range of Q , there is always an object $o \in \mathcal{O}$ which under some two lighting conditions in \mathcal{S} could have produced both image I and J .*

Proof. The proof follows immediately from the above definitions. \square

The existence of discriminative invariants depends on the specific set of illumination and objects. It is obvious that for large enough class of illuminations, there are no such invariants. For example, a movie projector can create arbitrary images by projecting carefully designed patterns on any given object with no zero reflectance. The larger the set of illumination conditions we admit, the smaller the set of invariants will be. Likewise, for purely specular convex surfaces, it is again obvious that there are no discriminative illumination invariants. However these cases are clearly extremes in illumination and reflectance. We next consider in detail the case of Lambertian surfaces under the illumination of point light sources. What is surprising is that there is no discriminative illumination invariants even here. (While it is not clear that the above statement holds when interreflections are considered, our simulations show that their additive effect is often negligible when considered in the context of this problem [7].)

Theorem 2.1 *Discounting interreflection, all illumination invariants for objects with Lambertian reflectance under point light sources at infinity are nondiscriminative.*

Proof. Let μ be an illumination invariant. Given two arbitrary images I, J in the range of Q , by Lemma 2.5 below, there always exists a Lambertian surface of an object o and two light sources at infinity s and l , such that I and J are images of o under s and l , respectively, or $I = Q(o, s)$ and $J = Q(o, l)$. By Lemma 2.1, μ must be a nondiscriminative invariant. \square

Note that for the above theorem and proof we require objects to be composed of surfaces — not freely floating planar facets in space as in [27]. In fact, [27] concedes that “clearly real objects consist of surfaces rather than patches.”

We begin the outline of our proof of Theorem 2.1 with the following notation. Let \vec{r} be a vector, then \hat{r} is the unit vector in the direction of \vec{r} . Set the camera optical axis as the z -axis; the object is viewed from the direction $(0, 0, 1)$. We can therefore represent the object by a graph of a function $z = f(x, y)$. A point light source at infinity is represented by a single 3-vector in the opposite direction of light ray, with the magnitude of the vector equal to the power of the source. Suppose the two images I and J are generated by f under two point

light sources $\vec{s} = (s_x, s_y, s_z)$ and $\vec{l} = (l_x, l_y, l_z)$. Excluding interreflection, the images I and J of the Lambertian surface are respectively

$$I(x, y) = \alpha(x, y) \vec{s} \cdot \hat{n}(x, y), \quad (1)$$

$$J(x, y) = \alpha(x, y) \vec{l} \cdot \hat{n}(x, y) \quad (2)$$

where $\alpha(x, y)$ is the albedo of object o , $\vec{n} = (-\frac{\partial f}{\partial x}, -\frac{\partial f}{\partial y}, 1)$ is a normal vector to the surface f of o , \hat{n} is the unit vector for \vec{n} , and \vec{s} and \vec{l} are chosen to be linearly independent. $I(x, y) = 0$ or $J(x, y) = 0$ if (x, y) is in the cast shadow under either of the light sources.

Our goal is to show that there always exists a continuous and piecewise smooth solution in a compact region to the partial differential equations (PDE's) given in Eq.'s 1 and 2, which in turn leads to Theorem 2.1. We should point out that this section builds on the ideas of [19]. However [19] does not consider the broader question of illumination invariance. Neither does it provide conditions and the proof of the global existence of continuous solutions. The major problems we need to address are the exclusion of the stagnation and bifurcation of the characteristic curves; the treatment of discontinuities and zero-valued regions in the images; the assignment of boundary conditions; and the exclusion of cast shadows so as not to destroy the local differential description of the surface in Eq.'s 1 and 2.

To handle regions in the images that may be in shadow or have albedo equal to zero, we divide the domain of interest into two parts: the region D_1 , where both I and J are equal to zero, and the region D_2 , where they do not vanish simultaneously. As we shall see below, there are many continuous and piecewise smooth solutions in each region. Continuous global solutions over the whole domain, if they exist, are formed by stitching solutions from D_1 and D_2 continuously across their shared boundary, see [7].

Finding a solution to the system of PDE's in region D_1 is simple. If $\alpha \neq 0$, then Eq.'s 1 and 2 in turn become $\vec{s} \cdot \hat{n} = \vec{l} \cdot \hat{n} = 0$ which implies f is any plane parallel to \vec{s} and \vec{l} . If $\alpha = 0$, then f may assume arbitrary values, provided that the surface casts no shadow outside this subregion. Thus, in the region D_1 where the images I and J are both zero, we have a solution to the PDE.

However, finding a solution to the system of PDE's for region D_2 is somewhat more complicated. We multiply the left of Eq. 1 with the right of Eq. 2 and the right of Eq. 1 with the left of Eq. 2. We then divide the resulting equation by α (α cannot be zero by the definition of D_2) to obtain the following first order linear partial differential equation:

$$(\vec{l} - \vec{s}) \cdot \vec{n} = 0. \quad (3)$$

The characteristic curves $\vec{r}(t)$ of Eq. 3 satisfy

$$\frac{d\vec{r}}{dt} = \vec{l} - \vec{s}. \quad (4)$$

We now develop a series of lemmas needed to complete the proof of Theorem 2.1. (The trusting reader may want to skip ahead to Section 3.) We focus our attention on Eq. 4 in region D_2 . We assume that I and J are continuous. (Note that this assumption is not necessary to guarantee continuity of f , but it does slightly simplify the argument.) All the proofs for the following lemmas are relegated to a technical report [7].

We claim that for a vector field uniformly bounded from below, there is a characteristic curve in the XY -plane through each point in D_2 with both ends lying on the boundary ∂D_2 .

Lemma 2.2 *Let Ω be a compact subset of \mathbb{R}^2 . Suppose $I, J \in C(\Omega)$ (continuous on Ω), $I, J \geq 0$, and $|\vec{I} - \vec{J}| > c$ for some $c > 0$. Let \vec{p} denote points in the XY -plane. Then through an arbitrarily given point q in Ω , there exists $\zeta_q \in \mathbb{R}$ and a characteristic curve $\vec{p}(t)$, $t \in [0, \zeta_q]$ in Ω homeomorphic to $[0, \zeta_q]$, such that $\vec{p}(0), \vec{p}(\zeta_q) \in \partial\Omega$.*

The above Lemma 2.2 shows the global existence of the characteristic curves within D_2 .

Lemma 2.3 *Let the hypothesis of Lemma 2.2 be satisfied. In addition, I and J are analytic in Ω . The characteristic curve in the XY -plane passing through an arbitrary point is then unique. Moreover, the XY -plane characteristic curve is analytic with respect to t and its initial point. The characteristic curve is C^k if I and J are C^k .*

The uniqueness of the characteristic curves provided by Lemma 2.3 removes the ambiguity and potential conflict, enabling the later construction of a smooth surface instead of a multivalued function. It also guarantees the needed smoothness for the characteristic curves, although not yet the surface.

The smooth surface can then be fabricated by carefully assigning smooth boundary conditions to appropriate subintervals — so as to cover the whole region Ω and avoid any contradiction — of the boundary.

Lemma 2.4 *Let the hypothesis of Lemma 2.3 be satisfied, moreover I and J are (real) analytic in Ω , and $\partial\Omega$ is a piecewise analytic Jordan curve. There is a C^k or analytic surface f on Ω satisfying Eq. 3.*

Our task would be much simpler if there is a surface that does not cast shadow on itself. Otherwise the local PDE in Eq. 3 does not completely describe the surface. Fortunately Lemma 2.5 dispels that concern. By judiciously choosing the boundary condition, we can generate smooth surfaces with no cast shadows.

Lemma 2.5 *Let the hypothesis of Lemma 2.4 be satisfied. There is, for a family of initial Cauchy data curves g , a family of surfaces f that have no cast shadows.*

3 A Probabilistic Approach

In the previous section, we concluded that for Lambertian and purely specular surfaces there are no illumination invariants. Thus, we must settle for something less: a weaker, probabilistic form of invariance. To this end, we will show that even if the direction of the light source is uniformly random, the direction of the image gradient is not. We first analytically construct the probability distribution for the image gradient as a function of the surface curvature and reflectance. We will eventually verify this empirically by constructing a distribution for the image gradient from a database of real images.

Consider a point on a surface whose tangent plane is parallel to the image plane. Set up a perpendicular coordinate system at that point. Let (u, v) be coordinates on the surface such that $(x(u, \cdot), y(u, \cdot), f(x(u, \cdot), y(u, \cdot)))$ and $(x(\cdot, v), y(\cdot, v), f(x(\cdot, v), y(\cdot, v)))$ are lines of curvature, with u, v being the length of the lines. For subsequent development, let κ_u and κ_v be the principal curvatures of the surface in principal directions \hat{u} and \hat{v} , respectively. We adopt the tangents of the lines of curvature along with the surface normal as the local Cartesian coordinate system and call this the u - v - n coordinate system.

The BRDF α is a function of (u, v) , the direction of the incident light \hat{i} , and the direction of outgoing light \hat{o} in this local coordinate system. The radiance from a point on a piecewise C^2 surface in the camera direction \hat{c} (in u - v - n coordinate system) is then

$$\begin{aligned} L(u, v, \hat{c}) &= \int_{\Omega} \alpha(u, v, \hat{i}(u, v, \hat{s}), \hat{o}(u, v, \hat{c})) \hat{n}(u, v) \cdot \vec{s} \, d\hat{s} \\ &= \hat{n} \cdot \int_{\Omega} \alpha \vec{s} \, d\hat{s} \end{aligned} \quad (5)$$

where Ω denotes the solid angle of light seen at the point of concern and \hat{c} .

We analyze the influence of the differential geometric and reflectance properties by examining the scene radiance under a single light source at infinity. The scene radiance in Eq. 5 becomes

$$L(u, v, \hat{c}) = \alpha(u, v, \hat{i}(u, v, \hat{s}), \hat{o}(u, v, \hat{c})) \vec{s} \cdot \hat{n}(u, v). \quad (6)$$

Suppose we have a patch of a Lambertian surface with constant albedo and principal curvatures k_u and k_v . If $|k_u| \neq |k_v|$ and the direction of light sources are distributed spherically uniformly, the gradient of the scene radiance is most likely in the direction in which the magnitude of the curvature is maximal. As a second example, let us consider a planar patch of surface with nonhomogeneous BRDF. For any BRDF, the direction of the gradient of the scene radiance always points in the direction of the spatial gradient of the BRDF. These observations suggest that if the light source directions are distributed uniformly, the distribution of the gradient of scene radiance will be biased by the underlying geometry and reflectance.

We shall expound on these observations in two steps. First, we derive the relation between the gradient of the

scene radiance $\vec{\nabla}L$ and the surface's local geometry and reflectance. Second, we impose a probability distribution on the light source \vec{s} and determine the resulting distribution on the gradient of the scene radiance $\vec{\nabla}L$. (Note that for simplicity we are doing our analysis in the tangent plane of the surface, whereas the image records the projection of scene radiance from the surface down to the x-y plane. Thus, our analysis ignores the effects of projection.)

In the following derivation, $\vec{\nabla}$ is understood to be the gradient taken in the tangent plane, or u-v plane, at a point on the surface. A short calculation shows that the gradient of the scene radiance

$$\vec{\nabla}L = \alpha(\vec{s} \cdot \vec{\nabla})\hat{n} + (\vec{\nabla}\alpha)\vec{s} \cdot \hat{n}. \quad (7)$$

Equation 7 teases out two factors that determine the gradient of the scene radiance. The first term, called the geometric gradient, is the contribution from geometric changes; the second term, called the reflectance gradient, is the contribution from changes in the BRDF.

Consider the geometric gradient term:

$$(\vec{s} \cdot \vec{\nabla})\hat{n} = \hat{u}\kappa_u s_u + \hat{v}\kappa_v s_v \quad (8)$$

where κ_u and κ_v are the two principal curvatures, and s_u and s_v are the u and v components of the light source in the u-v-n coordinate system.

Consider next the reflectance gradient term:

$$\vec{\nabla}\alpha = \hat{u}\frac{\partial\alpha}{\partial u} + \hat{v}\frac{\partial\alpha}{\partial v} \quad (9)$$

where, after a little calculation,

$$\begin{aligned} \frac{\partial\alpha}{\partial u} &= \vec{\nabla}_{\hat{c}}\alpha \cdot \frac{\partial\hat{c}}{\partial u} + \vec{\nabla}_{\hat{s}}\alpha \cdot \frac{\partial\hat{s}}{\partial u} + \hat{u} \cdot \vec{\nabla}\alpha \\ &= \kappa_u \hat{v} \cdot (\vec{\nabla}_{\hat{c}}\alpha \times \hat{c} + \vec{\nabla}_{\hat{s}}\alpha \times \hat{s}) + \hat{u} \cdot \vec{\nabla}\alpha, \\ \frac{\partial\alpha}{\partial v} &= \kappa_v \hat{u} \cdot (\vec{\nabla}_{\hat{c}}\alpha \times \hat{c} + \vec{\nabla}_{\hat{s}}\alpha \times \hat{s}) + \hat{v} \cdot \vec{\nabla}\alpha \end{aligned}$$

where $\vec{\nabla}_{\hat{c}}$ is the gradient taken with respect to \hat{c} and $\vec{\nabla}_{\hat{s}}$ is gradient taken with respect to \hat{s} .

We combine the geometric and reflectance gradients to get an overall expression for $\vec{\nabla}L$:

$$\begin{aligned} \vec{\nabla}L &= \underbrace{(\hat{u}\kappa_u s_u + \hat{v}\kappa_v s_v)}_{\text{geometric gradient}} + \\ &\underbrace{\left[(\kappa_u \hat{u} \hat{v} + \kappa_v \hat{v} \hat{u}) \cdot (\vec{\nabla}_{\hat{c}}\alpha \times \hat{c} + \vec{\nabla}_{\hat{s}}\alpha \times \hat{s}) + \vec{\nabla}\alpha \right]}_{\text{reflectance gradient}} \vec{s} \cdot \hat{n}. \end{aligned} \quad (10)$$

(Note that $\hat{u} \hat{v}$ and $\hat{v} \hat{u}$ are tensor products of the vectors; the associative rule applies here; and the first term in the square bracket, although classified as part of the




Figure 1: Distribution $\rho_{u,v}$ for constant albedo as described in Eq. 13 in polar coordinates (r, φ) .

reflectance gradient, is induced by the rotation of the u-v-n coordinate system along the lines of curvature.) If the BRDF is Lambertian, the above expression becomes

$$\vec{\nabla}L = \underbrace{(\hat{u}\kappa_u s_u + \hat{v}\kappa_v s_v)}_{\text{geometric}} + \underbrace{(\vec{\nabla}\alpha)\vec{s} \cdot \hat{n}}_{\text{reflectance}} \quad (11)$$

Now that we have expressions for the gradient of the scene radiance in terms of the differential geometric and reflectance properties of the patch of surface, we determine the distribution for scene radiance by imposing a distribution on light sources. We consider only the light sources seen by the surface patch, i.e., $s_n > 0$. A reasonable assumption for the distribution of lighting \vec{s} would be that the direction of the source is symmetric in the upper hemisphere. We felt it is important for this initial discussion to choose a distribution that does not favor any particular direction on the hemisphere, even though it may prove useful at some later point. In addition, let the components of the source s_u , s_v , and s_n be chosen independently. With these assumptions, it can be shown that probability density for \vec{s} is given by

$$\rho_s(\vec{s}) = \frac{1}{(\sqrt{2\pi}\sigma)^3} e^{-\frac{1}{2\sigma^2}(s_u^2 + s_v^2 + s_n^2)}, \quad s_n \in [0, \infty). \quad (12)$$

The expression for the resulting probability density function on $\vec{\nabla}L$ can be found in [7]. Here we consider two special cases which we feel are the two determining factors for the distribution for the image gradient.

Case I: Constant Albedo

If the surface patch has spatially homogeneous reflectance (constant albedo), then $\frac{\partial\alpha}{\partial u} = \frac{\partial\alpha}{\partial v} = 0$. In coordinate system u-v-n, probability density function for $\vec{\nabla}L$ is then

$$\rho_{u,v}(u, v) = \frac{1}{\pi^{\frac{3}{2}}\sigma^2\kappa_u\kappa_v} e^{-\frac{1}{2\sigma^2}\left(\left(\frac{u}{\kappa_u}\right)^2 + \left(\frac{v}{\kappa_v}\right)^2\right)}. \quad (13)$$

Note that the level curves of this function are concentric ellipses and that there is a ridge along the larger




Figure 2: Distribution $\rho_{u,v}$ for planar surface as described in Eq. 14 in polar coordinates (r, φ) .

of the curvature directions. In polar coordinates (r, φ) , where $r = \sqrt{u^2 + v^2}$ and $\tan \varphi = \frac{v}{u}$, there are two equal ridges at $\varphi = 0$ and $\varphi = \pi$ along the direction of x -axis whenever $\frac{\kappa_u}{\kappa_v} \neq 1$ as shown in Fig. 1. The ridges grow sharper as the ratio deviates farther away from 1. The existence of these ridges imply that the gradients are more likely to point in the larger principal curvature direction.

Case II: Local anisotropy

In this case, $\kappa_u = \kappa_v = 0$. The two curvatures are the same, so there is no preordained orientation for the u - v coordinate system. We may simply choose the direction of the albedo gradient as the positive u axis. The probability density function for the gradient is then simply

$$\rho_{u,v}(u, v) = \frac{\sqrt{2}}{\pi\sigma} e^{-\frac{1}{\sigma^2}(\frac{u}{\alpha_u})^2} \delta(v)\Theta(u), \quad (14)$$

where $\alpha_u = \frac{\partial \alpha}{\partial u}$, δ is the Dirac delta function, and Θ is the Heaviside theta function (or step function). The distribution in the u - v coordinate system is one thin slice along the positive u axis. In polar coordinates (r, φ) , there is only one slice at $\varphi = 0$ along r as depicted in Fig. 2. Therefore, the image gradient is always directed along the albedo gradient, regardless of the direction of the light source.

We want to develop an illumination insensitive measure of image comparison encompassing the above two factors. If we are given in advance the directions and magnitudes of principal curvature along with the nature of the BRDF, then we can use this density function to determine, in a probabilistic sense, how faithful an image is to the given values. The problem we are interested in is slightly different: we want to compare two images and determine the likelihood that they have been produced by the same object. For this problem, the magnitudes and directions of surface curvature and the BRDF are unknown. We must then look at the joint density for gradients of scene radiance as given in

two images, (as they are the only observables amongst the previously introduced quantities) and integrate out the unknown nonobservable quantities, i.e., the magnitudes and directions of the principal curvatures and the reflectance parameters.

Furthermore, we do not know the variation in the strength of the source, i.e., the standard deviation σ in Eq. 12. This variation in the strength of the light source is, of course, carried into large variation in the magnitude of the gradient. Because of this, we elect to integrate out the magnitude of the image gradient as well, distilling the discriminative power of the distribution.

The gradients are observed in a fixed x - y coordinate system. Given the angle γ between \hat{x} and \hat{u} , the principal curvatures denoted by κ 's, and the albedo gradient $\vec{\nabla} \alpha$, the probability density of observing a gradient with magnitude r and angle φ from \hat{x} is $\rho_r(r, \varphi | \gamma, \kappa, \vec{\nabla} \alpha) = r \rho_{(u,v)}(r \cos(\varphi - \gamma), r \sin(\varphi - \gamma))$. (The scaling of ρ by r comes from the Jacobian from Cartesian to polar coordinates.) Noticing the angular dependence is only on the difference of the angles $\varphi - \gamma$, we can also write the density as $\rho_r(r, \varphi - \gamma | \kappa, \vec{\nabla} \alpha)$. Now, the joint probability density of observing two scene radiance gradients (r_1, φ_1) and (r_2, φ_2) under two independent and identically distributed light sources is

$$\begin{aligned} & \rho(r_1, \varphi_1, r_2, \varphi_2) \\ &= \int \rho_r(r_1, \varphi_1 - \gamma | \kappa, \vec{\beta}) \rho_r(r_2, \varphi_2 - \gamma | \kappa, \vec{\beta}) \\ & \quad dP(\gamma, \kappa, \vec{\beta}), \end{aligned}$$

where $\vec{\beta} = \vec{\nabla}_{(u,v)} \alpha$, $P(\gamma, \kappa, \vec{\beta})$ is the probability measure on the nonobservable random variables, and the integration is over the whole sample space. Furthermore if azimuthal symmetry holds for P , then the density ρ above can be rewritten as a function of three variables

$$\begin{aligned} & \rho(r_1, \varphi = \varphi_1 - \varphi_2, r_2) \\ &= \int \int_{\gamma=-\pi}^{\pi} \rho_r(r_1, (\varphi_1 - \varphi_2) - \gamma | \kappa, \vec{\nabla}_{(x,y)} \alpha) \\ & \quad \rho_r(r_2, -\gamma | \kappa, \vec{\beta}) d\gamma dP(\kappa, \vec{\beta}). \end{aligned} \quad (15)$$

Azimuthal symmetry is almost intrinsic for any set of reasonably random image samples: the unrestrained relative rotation of the objects and the camera along the optical axis of the lens would almost surely render the azimuthal angle indistinguishable. Equation 15 implies that the joint density depends on the angle between two image gradients or the absolute value of $\varphi = \varphi_2 - \varphi_1$. It is an even function with respect to φ .

Figure 3(a) shows a graph of the joint probability $\rho(r_1, \varphi, r_2)$ where we have pre-chosen the parameters specifying the distributions. For this graph we chose $\sigma = 1.7$, $\kappa_u = 7$, $\kappa_v = 1.7$, $a_u = 17$, $P(\text{Case I}) = 0.2$, and $P(\text{Case II}) = 0.8$. We chose these values simply to illustrate the nature of the shape of the density. Different values will, of course, yield different densities, but their underlying shapes remain qualitatively the same.

Now, integrate ρ over r_1 and r_2 ,

$$\rho_\varphi(\varphi) = \int \int \rho(r_1, \varphi, r_2) dr_1 dr_2. \quad (16)$$

Clearly, ρ_φ inherits from $\rho(r_1, \varphi, r_2)$ the property of having a unique maximum at $\varphi = 0$.

This joint distribution could be used as an illumination insensitive measure. Under this distribution, we would expect high probability assigned to two images of the same object (differing only in illumination, not viewpoint) and low probability assigned to two images of different objects. Yet, we are hamstrung in that we do not know the probability distribution for the nonobservable (curvatures and albedo gradients) and, thus, cannot perform the integration in Eq. 15. Furthermore, all of the above analysis is done in the tangent plane of the surface, and consequently ignores the effects of the projection into the image plane.

To circumvent this difficulty and for pragmatic reasons, in the next section we will empirically construct the distribution in Eq. 15 from real images of objects under varying illumination.

4 Empirical Joint Density

Using a geodesic dome with 64 photographic flashes, we gathered a database of 1,280 images of 20 objects. The 64 flashes are positioned on the dome to cover slightly more than a hemisphere of directions. The objects included folded cloth, a computer keyboard, cups, an umbrella, plants, a styrofoam mannequin, etc.¹ We estimate the density $\rho(r_1, \phi, r_2)$ in Eq. 15 by a histogram of the image gradients.

A slice of the joint probability density as a function of the gradient magnitudes and angle is shown in Fig. 3(b). As expected from Sec. 3, the distribution has all the features depicted in Fig. 3(a). It is symmetric with respect to the plane $\varphi = 0$. There is a prominent ridge at $\varphi = 0$ and it is the unique global maximum on the line of arbitrarily fixed r_1, r_2 . The spike in Fig. 3(b) at $r_2 = r_1, \varphi = 0$ arises from the concentration of lights near the camera's optical axis. With purely uniform sampling of the light source direction the spike disappears. This shows that the statistical regularity of scene radiance gradient does reflect the intrinsic geometric and reflectance properties of surfaces and this regularity can be exploited. In Sec. 5, we will demonstrate this on the problem of face recognition under varying illumination.

5 Application to Face Recognition

Given an object o seen in images I and J under two different lighting conditions, the probability of observing the gradients of I and J is assumed to satisfy

$$P(\angle \vec{\nabla} I, \angle \vec{\nabla} J) = \prod_{i \in M} \rho(\angle \vec{\nabla} I_i, \angle \vec{\nabla} J_i)$$

¹All databases used in this paper are available for download from subdirectories "hrlfaces," "yaleAselected," and "objects" at <ftp://plucky.cs.yale.edu/ftproot>.

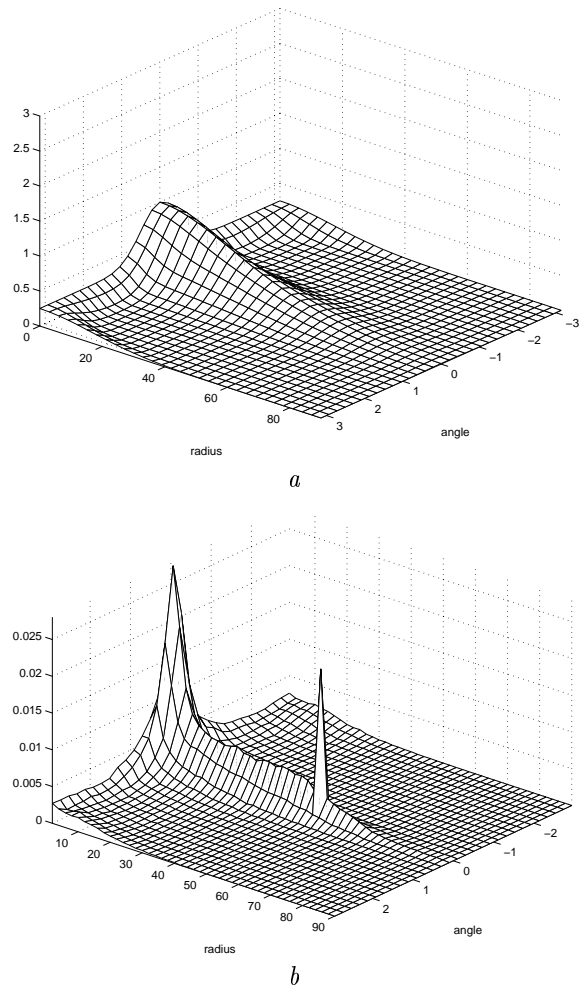



Figure 3: Model (a) and empirical (b) joint probability densities of the two image gradients $\rho(r_1, \varphi, r_2 = 50)$, expressed as a function of the magnitude of one gradient and the angle between the two, with the other's magnitude set to 50.

$$= \prod_{i \in M} \rho_\varphi(\varphi(i)).$$

where, M is the set of pixel indices and φ is the angle between the two gradient vectors. We treat the points on the surface as independent, ignoring correlations arising from spatial proximity.

We apply this scheme to face recognition. Given a testing image I of a face, we compute $P(\angle \vec{\nabla} I, \angle \vec{\nabla} J)$ for every training image using an empirically collected probability database as described in Sec. 4. The one training image having the highest P value is deemed the likeliest to have come from the same face or object as the testing image I . Fig. 5 shows the result of a face recognition test and compares it to other methods. Images of 10 faces each under 45 different lighting conditions are used. One image of each face under frontal illumination is taken as a training image. The recognition test is then performed for the rest 440 images. The results are



COMPARISON OF RECOGNITION METHODS			
Method	Error Rate (%) vs. Illumination		
	Subsets 1 & 2	Subset 3	Subset 4
Correlation	0.0	11.7	65.0
Eigenfaces	0.0	16.7	69.3
Eigenfaces w/o 1st 3	0.0	3.3	57.9
Linear subspace	0.0	1.7	12.9
Cones-attached	0.0	0.8	9.3
Gradient Angle	0.0	0.0	1.4
Cones-cast	0.0	0.0	0.0

Figure 4: Each of the methods is trained on images with near frontal illumination (Subsets 1 and 2). This graph shows the error rates under more extreme light source conditions.




Figure 5: Images of one of the 10 individuals under the 4 subsets of lighting.

grouped into 4 subsets according to the lighting angle with respect to the frontal or camera axis. The first two cover angles $0^\circ - 25^\circ$, third $25^\circ - 50^\circ$, and the fourth $50^\circ - 77^\circ$. We compared our method with those tested and reported in [15]. The image gradient method clearly outperforms all other methods except for that of Cone Cast.

It should be pointed out that all the other methods use all the images of Subset 1 and 2 for training, therefore they have almost by definition zero error rates for Subset 1 and 2. Our method, on the contrary, is at a disadvantage for using only one frontal image for each individual. Yet, the result is still better. Also note that Cone Cast is the Illumination Cone method in [14] in which Subset 1 and 2 images are used to construct person specific illumination representations. In contrast, our method is much simpler using only local image comparisons.

It is worth while to note that the probability distribution used to perform the test is gathered from images of objects rather than human faces. The distributions collected from different categories of objects or faces are remarkably similar. This is expected from the analysis in Sec. 3.

6 Conclusion

In summary, this paper presents two results: first, illumination invariants do not exist for Lambertian surfaces; and second, the angle (or direction) of the image gradient is insensitive to changes in illumination directions. Note that the latter statement is not inconsistent with the conclusion in [1] of the nonexistence of linear filter for image comparison, since the gradient angle is a nonlinear function of the image. However, we hope that the reader does not conclude that image edges are good measures of image comparison under varying illumination — quite the contrary is true. Most edge detection methods are highly sensitive to the magnitude of the image gradient. And, as we can see from the distributions in Fig. 3, the magnitude of the gradient varies quite dramatically with the change in the direction of the light source.

In a sense, this paper looks at two extremes. First, we look for measures of illumination invariance and prove there are none. As a consequence, we back way off and look for illumination insensitivities in local measures, i.e., image gradients. Yet, our assumption about pixel independence implicit in Eq. 17 is, of course, not true — the light sources are fixed for all pixels in a given image, and neighboring surface points tend to have similar geometric and reflectance properties. We hope to remedy the crudeness of our independence assumption by exploring how spatial dependencies [2] in pixel values can be exploited to produce even more discriminating measures of the images.

Acknowledgements

We are grateful to Alan Yuille, David Kriegman, David Forsyth, Steven Zucker, Jonas August, Athinodoros Georgiades, Steven Haker, Patrick Huggins, Fredrik Kahl, and Jie Lu for their help and advice in writing this paper.

References

- [1] Y. Adini, Y. Moses, S. Ullman. "Face Recognition: The Problem of Compensating for Changes in Illumination Direction", *IEEE Trans.PAMI*, Vol.19, No.7: 721-732, July 1997.

- [2] J. August and S. Zucker, "The curve indicator random field: curve organization via edge correlation," in *Perceptual Organization for Artificial Vision Systems*, K. Boyer and S. Sarkar, Eds., pp. 265–288, Kluwer Academic, Boston, 2000.
- [3] S. Baker, S. K. Nayar. "Global Measures of Coherence for Edge Detector Evaluation", *CVPR 99*, Vol.2: 373–379, June 1999.
- [4] P. Belhumeur, D. Kriegman. "What is the Set of Images of an Object Under All Possible Lighting Conditions?", *IEEE CVPR 96*: 270–277, 1996.
- [5] P. Breton, S. Zucker. "Shadows and Shading Flow Fields", *CVPR 96*: 782–789, 1996.
- [6] J. Burns, R. Weiss, E. Riseman. "The Non-Existence of General-Case View-Invariants", *Geometric Invariance in Computer Vision*, edited by J. Mundy A. Zisserman, MIT Press, Cambridge, 1992.
- [7] H. Chen, P. Belhumeur. "Technical Report on Illumination Invariance". Technical Report, Center for Computational Vision, Computer Science Dept., Yale University, Dec., 1999.
- [8] D. Clemens, D. Jacobs. "Space and Time Bounds on Model Indexing", *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol.13, No.10: 1007–1018, 1991.
- [9] E. Coleman, R. Jain. "Obtaining 3-Dimensional Shape of Textured and Specular Surfaces Using Four-Source Photometry", *CGIP 82*, Vol.18, No.4: 309–328, 1982.
- [10] Earl A. Coddington, Norman Levinson. *Theory of Ordinary Differential Equations*, Krieger Publishing Co. Malabar, FL, 1984.
- [11] J. Fan, L. Wolff. "Surface Curvature and Shape Reconstruction from Unknown Multiple Illumination and Integrability", *Computer Vision and Image Understanding*, Vol.65, No.2: 347–359, 1997.
- [12] O. Faugeras, L. Robert. "What Can Two Images Tell Us about a Third One?", *Int. J. of Comp. Vis.*, Vol. 18, No.1: 5–19, 1996.
- [13] W. Freeman, "The Generic Viewpoint Assumption in a Framework for Visual Perception", *Nature* **368**: 542–545, 1994.
- [14] A. Georghiades, D. Kriegman, P. Belhumeur. "Illumination Cones for Recognition Under Variable Lighting: Faces", *CVPR 98*: 52–59, 1998.
- [15] A. Georghiades, P. Belhumeur, D. Kriegman. "From Few to Many: Generative Models for Recognition Under Variable Pose and Illumination", *Int. Conf. on Automatic Face and Gesture Recognition 2000* 2000.
- [16] P. Hallinan. "A Low-Dimensional Representation of Human Faces for Arbitrary Lighting Conditions", *CVPR 94*: 995–999, 1994.
- [17] B. Horn. "Determining Lightness from an Image", *CGIP 74*, Vol.3, No.4: 277–299, 1974.
- [18] D. Jacobs. "Matching 3-D Models to 2-D Images", *Int. J. of Comp. Vis.*, Vol.21, No.1-2: 123–153, 1997.
- [19] D. Jacobs, P. Belhumeur, R. Basri. "Comparing Images Under Variable Illumination", *CVPR 98*: 610–617, 1998.
- [20] J. J. Koenderink. *Solid Shape* The MIT Press Cambridge, MA, 1990.
- [21] J. J. Koenderink, A. J. van Doorn. "Photometric Invariants related to solid Shape", *Optica Acta*, Vol. 27, No.7: 981–996, 1980.
- [22] S. Konishi, A. L. Yuille, J. Coughlan, S. C. Zhu. "Fundamental Bounds on Edge Detection: An Information Theoretic Evaluation of Different Edge Cues", *CVPR 99*, Vol.2: 573–579, 1999.
- [23] J. Lambert. "Photometria Sive de Mensura et Gradibus Luminis, Colorum et Umbrae", Eberhard Klett, 1760.
- [24] D. Lowe. *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, The Netherlands, 1985.
- [25] D. Marr. *Vision*, W.H. Freeman and Company, San Francisco, 1982.
- [26] Y. Moses. *Face recognition: generalization to novel images*, Ph.D. Thesis, Weizmann Institute of Science, 1993.
- [27] Y. Moses, S. Ullman. "Generalization to Novel Views: Universal, Class-based, and Model-based Processing", *Int. J. of Comp. Vis.*, Vol.29, No.3: 233–253, 1998.
- [28] Y. Moses, S. Ullman. "Limitations of Non Model-Based Recognition Schemes", *Sec. Eur. Conf. on Comp. Vis.*: 820–828, 1992.
- [29] J. Mundy, A. Zisserman (eds.). *Geometric Invariance in Computer Vision*, MIT Press, Cambridge, 1992.
- [30] H. Murase, S. Nayar. Visual learning and recognition of 3D objects from appearance. *Int. J. of Comp. Vis.*, Vol.14, No.1: 5–25, 1995.
- [31] S. Nayar, R. Bolle. "Reflectance Based Object Recognition", *Int. J. of Comp. Vis.*, Vol.17, No.3: 219–240, 1996.
- [32] R. Onn, F. Bruckstein. "Integrability Disambiguates Surface Recovery in Two-Image Photometric Stereo", *Int. J. of Comp. Vis.*, Vol.5, No.1: 105–113, 1990.
- [33] D. Ruderman, "The statistics of natural images", *Network: Computation in Neural Systems* **5** (1994), 517–548.
- [34] D. Ruderman and W. Bialek, "Statistics of Natural Images: Scaling in the Woods", *Physical Review Letters* **Vol 73**, No. 6 (1994), 814–817.
- [35] A. Shashua. "On Photometric Issues in 3D Visual Recognition from a Single 2D Image", *Int. J. of Comp. Vis.*, Vol.21, No.1-2: 99–122, 1997.
- [36] M. Turk, A. Pentland. "Eigenfaces for Recognition", *Journal of Cognitive Neuroscience* Vol.3, No.1: 71–96, 1991.
- [37] S. Ullman. "Aligning Pictorial Descriptions: An Approach to Object Recognition", *Cognition*, Vol.32, No.3: 193–254, 1989.
- [38] S. Ullman, R. Basri. "Recognition by Linear Combinations of Models", *IEEE Trans. PAMI*, Vol.13, No.10: 992–1007, 1991.
- [39] L. Wolff, E. Angelopoulou. *Eur. Conf. on Comp. Vis.*: 247–258, 1994.
- [40] L. Wolff, J. Fan. "Segmentation of Surface Curvature with a Photometric Invariant", *J. Opt. Soc. Am. A*, Vol.11, No.11: 3090–3100, Nov. 1994.