

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 but differ for other objects. This result leads us to adopt a probabilistic approach in which we determine a probability distribution for image gradients as a function of the differential geometric (i.e., curvature) and reflectance properties (i.e., bidirectional reflectance distribution function) of the object's surface. We then empirically constructed the distribution from more than 20 million samples of image gradients in a database of 1,280 images of 20 inanimate objects taken under varying lighting condition. This image gradient distribution then serves as the basis for a probabilistic illumination insensitive measure of image comparison. Finally, we apply the developed illumination insensitive measure to the problem of face recognition and compare the results of our method to others on a database of 450 images of human faces.

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 [25, 10, 5, 6].)

We show that even for objects with Lambertian reflectance [20], there are no discriminative functions of images of objects that are invariant to illumination but distinguish different objects. This differs from earlier findings in that we do not assume a homogeneous BRDF [19, 4], coplanarity [27] or consider invariants based on multiple images [34]. 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. Inter-reflection destroy some of the compatibility of two arbitrary images. However, our computation shows its symmetry breaking effect is negligible.

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 [34] or construct representations for modeling the illumination variation [29, 14, 3, 12].) We take a step in this direction

by considering illumination to be a random function or variable which gives rise to the apparent randomness in local image measurements. Specifically, we consider a patch of surface and show how the probability distribution of the image gradient of the patch is determined by the differential geometric and reflectance properties of the surface. Assuming spherically symmetric distributions on light sources, the “sharpness,” or discriminating power of the image gradient distribution is governed by curvature of the patch and the rate of change of the bidirectional reflectance distribution function. Using the image gradient distribution, we then develop a probabilistic illumination insensitive measure of image comparison.

Finally, we implement the theoretically set up probability scheme for image gradients by empirically constructing a distribution from a database of 1,280 images of 20 objects taken under varying illumination direction. Using our illumination insensitive measure, we conduct a face recognition experiment on 450 images of 10 individuals and compare the performance to existing methods.

2 Illumination Invariants?

Is it possible to determine whether two images were created by the same or different objects? While this seems obviously possible, we show in this section the contrary.

We analyze this question by studying the existence of discriminative illumination invariants: functions of an image which are invariant to illumination but vary with the object identity. Formally, let \mathcal{O} be some certain set of rigid objects (including their surface reflectance properties), \mathcal{S} some certain set of lighting conditions¹, and \mathcal{I} the set of all images, defined to be the piecewise smooth nonnegative functions, from $\text{cal}\mathcal{O}$ under \mathcal{I} . 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., $I = Q(o, s)$.

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 An illumination invariant μ is nondiscriminative for object set $\mathcal{O} \iff \mu(I) = \mu(J), \forall I \neq J, I, J \in \mathcal{I}$.

This definition implies μ does not depend on o for nondiscriminative invariants.

¹Clearly we would like to restrict possible lighting conditions, lest our claim be trivial: movie projectors can give rise to all images!

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

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

For the case of purely specular surfaces, it is obvious that there are no discriminative illumination invariants. Yet, what is surprising is that this result holds for Lambertian surfaces as well.

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 $I, J \in \mathcal{I}$, by the 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

For the above proof we need to show that given any two images I and J , there always exists a Lambertian surface of an object o and a pair of lighting conditions s, l that could have produced I and J : $I = Q(o, s)$ and $J = Q(o, l)$. To do this, we define \mathcal{S} and \mathcal{Q} as specified by the Lambertian model of reflectance. We restrict the lighting in each scene to be a point source at infinity. The contribution from any light source is then represented by a single 3-vector in the opposite direction of the ray, with magnitude of the vector equal to the power of the source. We have $\vec{s} = (s^x, s^y, s^z)$ and $\vec{l} = (l^x, l^y, l^z)$.

We assume that the object o is viewed from the direction $(0, 0, 1)$, and that the surface of o can be written as $(x, y, z = f(x, y))$. Since the surface of the object has Lambertian reflectance, we write the equations for image formation Q as

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

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

where $\alpha(x, y)$ is the albedo of object o , $\vec{n} = (-f_x, -f_y, 1)$ is a normal vector to the surface f of o , \hat{n} is the unit vector of \vec{n} , and \vec{s} and \vec{l} are chosen to be linearly independent.

Our goal is to show that there always exists a solution to the partial differential equation (PDE) given in Eq.'s 1 and 2. We should point out that the proof of this fills in all of the details left out in the argument presented in [17], yet, when completed we have a proof of Theorem 2.1.

To handle regions in the images that may be in shadow or have albedo equal to 0, we divide the domain of interest into two parts: the region D_1 , where both I and J are equal to 0, and the region D_2 , where they do not vanish simultaneously.

Finding a solution to the PDE in region D_1 is simple. There are two cases. For Case 1, $\alpha \neq 0$ and 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} . For Case 2, $\alpha = 0$ and 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 0, we have a solution to the PDE.

Finding a solution to the PDE for region D_2 is somewhat more complicated. We recast the the PDE. First, 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. Then, we divide by α (α cannot be 0 by the definition of D_2) to obtain the following first order linear partial differential equation:

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

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

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

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.) We first show in Lemma 2.2 the global existence of the characteristic curves within D_2 . We then show in Lemma 2.3 and Lemma 2.4 that the surface f is as smooth as I and J , given smooth initial curves. Finally, we show in Lemma 2.5 that the surface f does not cast a shadow on itself for either light source s or l . All the proofs are relegated to the Appendix.

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

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

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 an C^k or analytic surface f on Ω satisfying Eq. 3.

We can select from amongst the smooth surfaces the ones having no attached or cast shadows by choosing the appropriate initial Cauchy data curve.

Lemma 2.5 *Let the hypothesis of Lemma 2.4 be satisfied. Adopting the notations in the proof of Lemma 2.4, 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

We have concluded previously that, at least for Lambertian or purely specular surfaces, there are no discriminative illumination invariants. In this section, we take a modest step toward understanding how local image measurements vary with changing illumination. Our goal here is to understand how local variance in geometry and reflectance give rise to that in photometry.

A surface can be made up of materials that reflect light differently at different points on the surface (non-homogeneous BRDF) and for different incident and outgoing angles (non-Lambertian). We intend to do our analysis in the tangent plane of the surface, characterizing the surface by its differential geometric and reflectance properties.

We set up a perpendicular coordinate system on the surface. 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 later 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 a local Cartesian coordinate system; we call this the u-v-n coordinate system.

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

$$L(u, v, \hat{c}) = \int_{\Omega} \alpha(u, v, \hat{s}, \hat{c}) \hat{n} \cdot \vec{s} d\hat{s} = \hat{n} \cdot \int_{\Omega} \alpha \vec{s} d\hat{s}. \quad (5)$$

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

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

The discussion of the previous section indicates the futility of focusing on individual pixels. However, if we look at statistics of the interplay between neighboring pixels, we start to see some statistical regularity. First, let us consider a patch of a Lambertian surface with constant albedo characterized it by its principal curvatures

k_u and k_v . If $k_u = k_v = 0$, the gradient of the scene radiance is 0 regardless of the direction of the illuminant. If $|k_u| > 0$ and $k_v = 0$, the gradient of the scene radiance lies solely in the direction of the nonzero curvature axis, although its magnitude changes with changes in light source strength and direction. If $|k_u| \neq |k_v|$ and the direction of light sources are distributed uniformly, the gradient of the scene radiance is most likely in the direction in which the magnitude of the curvature is maximal. Second, let us consider a planar patch of surface with nonhomogeneous BRDF. For any BRDF, the direction of the gradient of the scene radiance is always in the direction of the spatial gradient of the BRDF. These observations suggest that even if the light source directions are distributed uniformly, the distribution of the gradient of scene radiance is not.

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 local geometry and reflectance of the surface. 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, while 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 taken in the tangent plane, or u-v plane, at a point on the surface. The gradient of the scene radiance is given by

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

(Note that the algebraic steps have been skipped to get to the above expression.) Equation 7 teases out two factors that determine the gradient of the scene radiance. The first term (which we call the geometric gradient) is the contribution from geometric changes; the second term (which we call the reflectance gradient) is the contribution from changes in the BRDF.

Consider the geometric gradient term: Let κ_u and κ_v be the two principal curvatures, then

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

where 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:

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

and

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

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 $\vec{\nabla}L$:

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

(Note that $\hat{u}\hat{v}$ is a tensor product of the vectors; the associative rule applies here; and the first term in the square bracket, although classified as part of the reflectance gradient, is induced by the rotation of the u-v-n coordinate system along the lines of curvature.) We can simplify the above expression by assuming the underlying BRDF is Lambertian:

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

This simplification allows us to push through the calculations for the probability distributions on $\vec{\nabla}L$. It is, of course, possible to compute this for the most general case, but space limitations prevent us.

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 light sources seen by surface patch, i.e., $s_n > 0$. A reasonable assumption for distribution of lighting \vec{s} would be 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. Finally, 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 $\vec{\nabla}L$ is complicated. We consider first two special cases before moving to the general case.

Case I: Homogeneous Reflectance

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

$$\begin{aligned} \rho_{u,v}(u, v) &= \frac{1}{\kappa_u \kappa_v} \int_0^\infty \rho_s \left(\left(\frac{u}{\kappa_u} \right)^2 + \left(\frac{v}{\kappa_v} \right)^2 + s_n^2 \right) ds_n \\ &= \frac{1}{\pi^{\frac{3}{2}} \sigma^2 \kappa_u \kappa_v} e^{-\frac{1}{2\sigma^2} \left(\left(\frac{u}{\kappa_u} \right)^2 + \left(\frac{v}{\kappa_v} \right)^2 \right)}. \end{aligned} \quad (13)$$

Note the level curves of this function are concentric ellipses. In polar coordinate systems, there is a ridge along the direction of $r = \sqrt{u^2 + v^2}$ whenever $\frac{|\kappa_u|}{|\kappa_v|} \neq 1$. The ridge grow sharper the more the ratio deviates from 1.

Case II: $\kappa_u = \kappa_v = 0$

For $\kappa_u = \kappa_v = 0$, the probability density function for the gradient is simply

$$\begin{aligned} \rho_{u,v}(u, v) &= 2\pi\delta\left(\frac{u}{a} - \frac{v}{b}\right)\Theta\left(\frac{u}{a}\right) \int_0^\infty \rho_s \left(s^2 + \left(\frac{u}{a} \right)^2 \right) s ds \\ &= \frac{\sqrt{2}}{\pi\sigma} \delta\left(\frac{u}{a} - \frac{v}{b}\right)\Theta\left(\frac{u}{a}\right) e^{-\frac{1}{2\sigma^2} \left(\frac{u}{a} \right)^2}, \end{aligned} \quad (14)$$

where $s = \sqrt{s_u^2 + s_v^2}$, $a = \frac{\partial \alpha}{\partial u}$, $b = \frac{\partial \alpha}{\partial v}$, and δ is the Dirac delta function and Θ is the Heaviside theta function.

General Case:

The geometry and reflectance factors are mixed. The elliptical level curves in Case I are rotated and distorted by the reflectance factor in Case II if the gradient of albedo does not align with either axis. The distribution in u-v coordinate is

$$\begin{aligned} \rho_{u,v}(u, v) &= \frac{1}{\kappa_u \kappa_v} \int_0^\infty \rho_s \left(\left(\frac{u - a s_n}{\kappa_u} \right)^2 + \left(\frac{v - b s_n}{\kappa_v} \right)^2 + s_n^2 \right) ds_n \\ &= \frac{1}{\pi^{\frac{3}{2}} \sigma^2 \kappa_u \kappa_v c} e^{-\frac{1}{2\sigma^2} \left(\left(1 - \frac{a^2}{c^2}\right) \xi^2 - 2 \frac{ab}{c^2} \xi \zeta + \left(1 - \frac{b^2}{c^2}\right) \zeta^2 \right)} \\ &\quad \int_{-\frac{a\xi + b\zeta}{\sqrt{2}\sigma c}}^\infty e^{-\tau^2} d\tau. \end{aligned} \quad (15)$$

where $a = \frac{\partial \alpha}{\partial u}$, $b = \frac{\partial \alpha}{\partial v}$, $c = \sqrt{\left(\frac{a}{\kappa_u}\right)^2 + \left(\frac{b}{\kappa_v}\right)^2 + 1}$, $\xi = \frac{u}{\kappa_u}$, and $\zeta = \frac{v}{\kappa_v}$.

From this density function we want to develop an illumination insensitive measure of image comparison. If we are given in advance the directions and magnitudes of principal curvature along with the nature of the BRDF, then we could use this density function to determine, in a probabilistic sense, how faithful any image was 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 direction 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., magnitudes and direction of principal curvature and reflectance properties.

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 , and the albedo gradient $\vec{\nabla} \alpha$, the probability density of observing a gradient with magnitude r and angle φ from \hat{x} is $\rho_r(r, \varphi | \hat{u}, \kappa, \vec{\nabla} \alpha) = r \rho_{(u,v)}(r \cos(\varphi - \gamma), r \sin(\varphi - \gamma))$. (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

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

and if azimuthal symmetry holds for P , ρ

$$\begin{aligned} &= \int \int_{\gamma=-\pi}^{\pi} \rho_r(r_1, (\varphi_1 - \varphi_2) - \gamma | \kappa, \vec{\nabla}_{(x,y)} \alpha) \\ & \quad \rho_r(r_2, -\gamma | \kappa, \vec{\beta}) d\gamma dP(\kappa, \vec{\beta}). \end{aligned} \quad (16)$$

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. 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.

Furthermore, Eq.16 implies the angular dependence of $\rho(r_1, \varphi_1, r_2, \varphi_2)$ is only upon the absolute value of their difference $\varphi = \varphi_2 - \varphi_1$. We may therefore proceed to rewrite ρ as $\rho(r_1, \varphi, r_2)$. It is an even function with respect to φ .

Because $\rho_r(\cdot, \varphi | \gamma, \cdot, \cdot)$ depends only on the difference $\varphi - \gamma$, using Cauchy-Schwartz inequality and the fact that ρ often has no period less than 2π , $\rho(r_1, \varphi, r_2)$ has a maximum at $\varphi = 0$ for all fixed r_1 and r_2 and the maximum is unique. The sharpness of the ridge at $\varphi = 0$ is determined by the narrowness of the peak in φ of single distribution $\rho_r(\gamma, \kappa, \vec{\beta})$, the polar coordinate form of $\rho_{u,v}$ in Eq. 15. More often than not, transition from one material (reflectance) to another is accompanied by sharp changes even edges in geometry. Geometry and reflectance factors then reinforce each other and sharpen the angular peak in the single distribution, and in turn narrow the ridge at $\varphi = 0$ in the joint distribution.

Bear in mind that specifications on ρ_r or what is the same $\rho_{u,v}$ could be relaxed without endangering the basic features such as the ridge, and its uniqueness. Nevertheless, azimuthal symmetry of the probability measure P is crucial to the existence of the ridge.

With this distribution in hand, we could construct an illumination insensitive measure of image comparison. To do this, we could write down the joint probability of the scene radiance gradients at each pixel in both images. 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 image of different objects. Yet, we are hamstrung in that we do not know




Figure 1: Empirical joint probability density of the two image gradients $\rho(r_1, \varphi, r_2 = 50)$ under two random lighting conditions, expressed as a function of the magnitude of one gradient and the angle between the two, with the other's magnitude set to 50.

the probability distribution for the nonobservables and, thus, cannot perform the integration in Eq. 16.

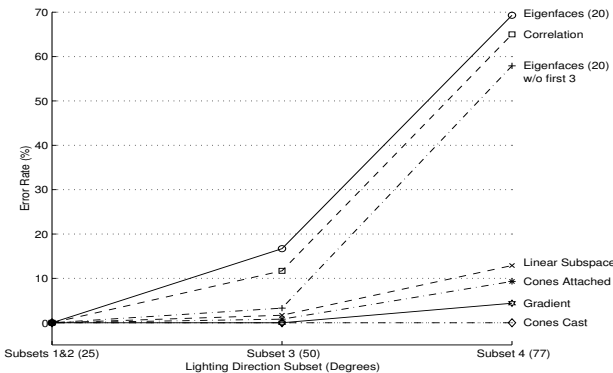
To get around this, in the next section we will construct the distribution in Eq. 16 from real images of object under varying illumination.

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. (The database is available for download.)

We are to find the joint probability density of two image gradients for the same point on the surface. We use empirically gathered image gradients to approximate the density ρ given in Eq. 16. We cannot perform the summation (in lieu of integral, since the samples are discrete) directly as in Eq. 16, since we do not know the distribution of curvatures and reflectance gradient a priori. However the sample frequency by which the image gradient at a prescribed value appears is an approximation of the true distribution. Therefore, we need only let the summation run over the whole set of pixel positions.

A slice of the joint probability density is shown in Fig. 4. As expected from Sec. 3, there is a prominent ridge at $\varphi = 0$ and it is the only global maximum on the line r_1, r_2 constant. For a given r_1 , the distribution peaks sharply at $r_2 = r_1, \varphi = 0$. We see that the experimentally gathered database confirms the angular symmetry (even-ness) of the distribution function ρ and the existence of a unique prominent ridge along $\varphi = 0$. 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 section 5, we will demonstrate this on



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	0.0	0.0	4.3
Cones-cast	0.0	0.0	0.0

Figure 2: 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.

the problem of face recognition under varying illumination.

5 Application to Face and Object Recognition

Armed with the locally illumination insensitive joint distribution density function for two image gradients, we venture to see what can be said globally.

Given a single arbitrary object o , the probability of observing the gradients of two images I and J is assumed to satisfy

$$\begin{aligned}
 P(\vec{\nabla}I, \vec{\nabla}J) &= \prod_{i \in M} \rho(\vec{\nabla}I_i, \vec{\nabla}J_i) \\
 &= \prod_{i \in M} \rho(r_1(i), \varphi(i), r_2(i)).
 \end{aligned}$$

where, M is the set of pixel indices, $r_1(i) = |\vec{\nabla}I_i|$, $r_2(i) = |\vec{\nabla}J_i|$ and φ is the angle between the two gradient vectors. Under this product form of probabilities of individual gradients, we treat the points on the surface independently, ignoring constraints arise from the smooth geometric and reflectance transition of closely spaced points.

We now apply this scheme to face or object recognition. Given a testing image I of a face or an object, we compute $P(\vec{\nabla}I, \vec{\nabla}J)$ for every training image using an empirically collected probability database as described

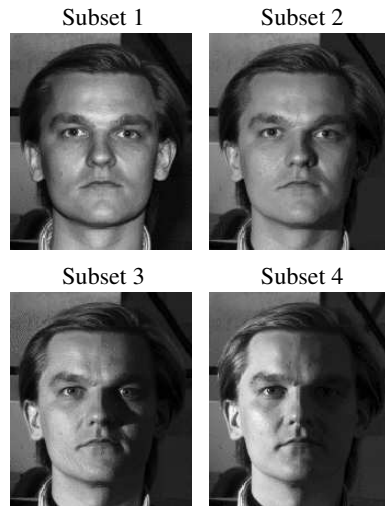


Figure 3: Images of one of the 10 individuals under the 4 subsets of lighting.

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. 2 shows the result of one of the face recognition tests and compares it to other methods. Images of 10 faces are taken using the geodesic dome each under 45 different lighting conditions. The frontal images of each face is taken as a training image. The recognition test is then performed for the 440 images. The results are grouped into 4 subsets according to the lighting angle with respect to the frontal or camera axis. The first two covers 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 [13]. The image gradient method clearly out performs 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 [12] in which Subset 1 and 2 images are used to construct a person specific illumination representations. In contrast, our method is much simpler using only local image comparisons.

It is worthwhile noting that the probability database used to perform the test is gathered from images of objects rather than human faces. The databases collected from different categories of objects or faces are remarkably similar. It is expected from the analysis in Sec. 3.

6 Appendix

Proof of Lemma 2.2.

The following observation on the characteristic curve Eq. 4 is crucial to the proofs of the lemmas.

Observation 6.1 *Let q be an arbitrary point in R^3 , S_q the plane containing q parallel to \vec{l} and \vec{s} , and \vec{r}_q the*

characteristic curve passing through q . It is clear that \vec{r}_q — if it exists — remains in the plane S_q . Moreover, the characteristic vector at every point resides in the planar cone $C = \{a\vec{l} - b\vec{s}, a, b \geq 0\}$ generated by \vec{l} and $-\vec{s}$ and the negative of C . Consequently the curve \vec{r}_q is also confined in the two-sided cone $q + C \cup q - C$. Especially, the projection $\vec{\rho}_p$ of \vec{r}_q onto XY -plane is contained in a cone C_p , the projection of C on the XY -plane.

Without loss of generality, we assume the given point is at the origin. Since I and J are continuous on the compact set Ω , they are bounded. Recall that the characteristic vectors in the XY -plane are confined in the cone C_p , as noted in Observation 6.1. We denote by \vec{h} the vector bisecting the angle subtended by \vec{l}_p and $-\vec{s}_p$. Since the vector field is bounded from below by c , the projection of the vector field on \vec{h} is bounded from below by $c \cos \theta$, where θ is half of the angle between \vec{l}_p and \vec{s}_p .

If $\theta < \frac{\pi}{2}$, any curve starting from the origin with these velocity (or tangential) vectors which leave Ω in time $T = \frac{R}{c \cos \theta}$. If $\theta = \frac{\pi}{2}$, $I\vec{l}_p - J\vec{s}_p$ is on a straight line. The leaving time is then simply bounded by $T = \frac{R}{c}$.

We recursively construct a sequence of functions $\{\vec{\rho}_n\}_{n=0}^{\infty} : [0, T] \rightarrow \Omega$ using Picard iteration:

$$\begin{aligned} \vec{\rho}_0 &= 0, \\ \vec{\rho}_{n+1}(t) &= \int_0^t [I(\vec{\rho}_n(\tau))\vec{l} - J(\vec{\rho}_n(\tau))\vec{s}] d\tau. \end{aligned} \quad (17)$$

In the above process, once $\vec{\rho}$ reaches the boundary of Ω , it is set to stay at the same point until T .

Clearly $\vec{\rho}_n(0) = 0$, $\vec{\rho}_n(\zeta) \in \partial\Omega$, $\forall n \in \mathbf{N}$. We also observe $\vec{\rho}_n(t) \in \Omega$, or $\{\vec{\rho}\}_{n=0}^{\infty}$ is uniformly bounded. Since $|I\vec{l} - J\vec{s}| < M$ for every point in Ω for some $M > 0$, $\{\vec{\rho}\}_{n=0}^{\infty}$ is equicontinuous. There then exists a subsequence uniformly converging to a $C^1[0, \zeta]$ function $\vec{\rho}(t)$. $\vec{\rho}$ can be shown to satisfy Eq. 3 projecting on XY -plane, and $\vec{\rho}(0) = 0$, $\vec{\rho}(\zeta) \in \partial\Omega$.

As for the other part of the curve, replace t with $-t$, every step of the above argument carries through. Re-parameterize the curve, so that the curve has the required parameterization as in the Lemma. Note that this proof of global existence is not the same of the classical one for local existence. \square

Proof of Lemma 2.3. Obviously, I and J satisfy a Lipschitz condition. The proof for these global properties is similar to that for the corresponding local ones classic in many texts (e.g., [8]). \square

Proof of Lemma 2.4. Denote Γ the subset of the boundary having only field vectors that point strictly inward of Ω , not including those that are parallel to tangents of the boundary. Γ is composed of disjoint components each of which is connected. Consider first the characteristic curves in the XY plane. By Lemma 2.3, there is a unique analytic characteristic curve emanating

from each boundary point in Γ . Because of the analyticity of both the curves and (piecewise) of the boundary, the number of the isolated intersections of the curves and the boundary is finite. Thus the number of curves having besides their initial and final points more than two isolated intersection points with the boundary is also finite. Denote the set of those curves by S . Those intersection points then have no accumulation points.

Assign the initial point of each aforementioned curve in S an arbitrary value and integrate the 3-d characteristic curve. Then define $g \in C^k(\Gamma)$ (analytic on Γ), such that g connects "C^k-ly" (analytically) with the 3-d characteristic curves on S .

Finally, integrate the rest of the 3-d characteristic curves for the initial Cauchy data g , we obtain the desired C^k (analytic) surface f . \square

Proof of Lemma 2.5. Construct C^k or analytic g such that there are no two points on its graph and no tangent of its graph lie on the same lighting plane S as defined in Observation 6.1.

Point a casts a shadow on point b if and only if a is between one of the light sources and b . By the way we construct g , it is then obvious from Observation 6.1 that there are no two 3-d characteristic curves which lie on the same lighting plane S . a and b are either on two different light planes as described in Observation 6.1 or on the same 3-d characteristic curve. For the former situation, it is obvious the two points can cast no shadow on each other. As for the latter, a and b stay in the cone C described in Observation 6.1 and can cast no shadows on each other — except when the curve between a and b is a straight line parallel to one of the lighting direction, however it is still not a cast shadow.

As for the attached shadow, parameterize the surface by (s, t) , where $z(s, t) = f(x(s, t), y(s, t))$ and $\{(x(s, 0), y(s, 0)) : s\} = \Gamma$. We need to prove that the dot product of the normal vector of the constructed parameterized surface and any of the light source vector do not change sign as the point moves all over the surface. By the hypothesis, no tangent of the graph of the initial curve is parallel to the lighting plane S , i.e., $\frac{\partial \vec{r}}{\partial s}$ thus does not lie in S , and certainly it will not be proportional to $\frac{\partial \vec{r}}{\partial t}$. A normal vector of the surface at a point \vec{r} is thus $\frac{\partial \vec{r}}{\partial s} \times \frac{\partial \vec{r}}{\partial t}$. Since the 3-d characteristic curves stay in their respective lighting planes S 's which are parallel to each other, vector $\frac{\partial \vec{r}}{\partial s}$ will stay on the same side of S , spanned by \vec{s} and \vec{l} , as the curve evolves from the initial point. Without loss of generality, let us assume

$$-\vec{s} \times \vec{l} \cdot \frac{\partial \vec{r}(s, 0)}{\partial s} > 0. \quad (18)$$

Then the inequality also holds for positive t or points in the interior of Ω , or

$$-\vec{s} \times \vec{l} \cdot \frac{\partial \vec{r}(s, t)}{\partial s} > 0. \quad (19)$$

Thus

$$\frac{\partial \vec{r}(s, t)}{\partial s} \times \frac{\partial \vec{r}(s, t)}{\partial t} \cdot \vec{s} = \frac{\partial \vec{r}(s, t)}{\partial s} \cdot (\vec{I}\vec{l} - J\vec{s}) \times \vec{s} \geq 0, \quad (20)$$

since $I(x, y)$ and $J(x, y)$ are nonnegative. The expression in the parenthesis is Eq. 4. The equality holds if and only if $I = 0$ for some (x, y) . The same holds for the inner product of the normal vector and \vec{l} . \square

References

- [1] Y. Adini, Y. Moses, S. and Ullman, 1997. "Face Recognition: The Problem of Compensating for Changes in Illumination Direction," *IEEE Trans. PAMI* **19**(7):721–732.
- [2] S. Baker and S. K. Nayar, 1999. "Global Measures of Coherence for Edge Detector Evaluation", *IEEE CVPR 99*: 1373-379.
- [3] P. Belhumeur and D. Kriegman, 1996. "What is the Set of Images of an Object Under All Possible Lighting Conditions?," *IEEE Conf. on Comp. Vis. and Pat. Rec.*:270–277.
- [4] P. Breton and S. Zucker, 1996. "Shadows and Shading Flow Fields" *ECCV96*
- [5] J. Burns, R. Weiss, and E. Riseman, 1992, "The Non-Existence of General-Case View-Invariants," *Geometric Invariance in Computer Vision*, edited by J. Mundy, and A. Zisserman, MIT Press, Cambridge.
- [6] D. Clemens and D. Jacobs, 1991, "Space and Time Bounds on Model Indexing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **13**(10):1007-1018.
- [7] E. Coleman and R. Jain, 1982, "Obtaining 3-Dimensional Shape of Textured and Specular Surfaces Using Four-Source Photometry," *CGIP* **18**(4):309–328.
- [8] Earl A. Coddington, Norman Levinson, *Theory of Ordinary Differential Equations*. Krieger Publishing Co. Malabar, Fl.
- [9] J. Fan and L. Wolff, 1997, "Surface Curvature and Shape Reconstruction from Unknown Multiple Illumination and Integrability," *Computer Vision and Image Understanding* **65**(2):347–359.
- [10] O. Faugeras and L. Robert, 1996. What Can Two Images Tell Us about a Third One?. *Int. Journal of Computer Vision*, **18**(1):5–19.
- [11] W. Freeman, 1994, "The Generic Viewpoint Assumption in a Framework for Visual Perception," *Nature* **368**:542–545.
- [12] A. Georghiades, D. Kriegman, and P. Belhumeur, 1998. "Illumination Cones for Recognition Under Variable Lighting: Faces", *IEEE Conf. CVPR*.
- [13] A. Georghiades, P. Belhumeur and D. Kriegman, 2000. "From Few to Many: Generative Models for Recognition Under Variable Pose and Illumination", *FG2000*
- [14] P. Hallinan, 1994. "A Low-Dimensional Representation of Human Faces for Arbitrary Lighting Conditions," *IEEE Conf. CVPR*:995–999.
- [15] B. Horn, 1974, "Determining Lightness from an Image", *CGIP*, **3**(4):277–299.
- [16] Jacobs, D., 1997, "Matching 3-D Models to 2-D Images", *International Journal of Computer Vision* **21**(1/2):123–153.
- [17] D. Jacobs, P. Belhumeur and R. Basri, 1998, "Comparing Images Under Variable Illumination", *IEEE Conf. CVPR*
- [18] J. J. Koenderink, *Solid Shape*. The MIT Press Cambridge, MA.
- [19] J. J. Koenderink and A. J. van Doorn, 1980, "Photometric Invariants related to solid Shape", *Optica Acta*, **27**(7):981-996.
- [20] J.H. Lambert, 1760, "Photometria Sive de Mensura et Gradibus Luminis Colorum et Umbrae", Eberhard Klett.
- [21] D. Lowe, 1985, *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, The Netherlands.
- [22] D. Marr, 1982, *Vision*, W.H. Freeman and Company, San Francisco.
- [23] Y. Moses, 1993. *Face recognition: generalization to novel images*, Ph.D. Thesis, Weizmann Institute of Science.
- [24] Y. Moses and S. Ullman, 1992, "Limitations of Non Model-Based Recognition Schemes," *Second European Conference on Computer Vision*:820-828.
- [25] J. Mundy and A. Zisserman (eds.), 1992, *Geometric Invariance in Computer Vision*, MIT Press, Cambridge.
- [26] H. Murase and S. Nayar, 1995. Visual learning and recognition of 3D objects from appearance. *Int. Journal of Computer Vision*, **14**(1):5–25.
- [27] S. Nayar and R. Bolle, forthcoming, "Reflectance Based Object Recognition," *Int. J. of Comp. Vis.*
- [28] R. Onn, and F. Bruckstein, 1990, "Integrability Disambiguates Surface Recovery in Two-Image Photometric Stereo," *Int. J. of Comp. Vis.* **5**(1):105–113.
- [29] A. Shashua, 1997. On Photometric Issues in 3D Visual Recognition from a Single 2D Image. *IJCV*, **21**(1/2):99–122.
- [30] M. Turk, and A. Pentland, 1991. "Eigenfaces for Recognition", *Journal of Cognitive Neuroscience*, **3**, 71-86.
- [31] S. Ullman, 1989, "Aligning Pictorial Descriptions: An Approach to Object Recognition," *Cognition* **32**(3):193-254.
- [32] S. Ullman and R. Basri, 1991, "Recognition by Linear Combinations of Models," *IEEE Trans. PAMI*, **13**(10):992-1007.
- [33] L. Wolff and E. Angelopoulou, 1994, *Eur. Conf. on Comp. Vis.*:247–258.
- [34] L. Wolff and J. Fan, Nov. 1994, "Segmentation of Surface Curvature with a Photometric Invariant", *J. Opt. Soc. Am. A*, **11**(11):3090-3100
- [35] S. Konishi, A. L. Yuille, J. Coughlan and S. C. Zhu, 1999. "Fundamental Bounds on Edge Detection: An Information Theoretic Evaluation of Different Edge Cues", *IEEE CVPR 99*:II573-579.