Whitening for Photometric Comparison of Smooth Surfaces under Varying Illumination

Margarita Osadchy
NEC Labs
4 Independence Way,
Princeton, NJ 08540

Michael Lindenbaum
Dept. of Computer Science
The Technion
Haifa, Israel

David Jacobs
Dept. of Computer Science
The University of Maryland
College Park, Maryland

Whitening for Photometric Comparison of Smooth Surfaces under Varying Illumination

Abstract

We consider the problem of image comparison in order to match smooth surfaces under varying illumination. In a smooth surface nearby surface normals are highly correlated. We model such surfaces as Gaussian processes and derive the resulting statistical characterization of the corresponding images. Supported by this model, we treat the difference between two images, associated with the same surface and different lighting, as colored Gaussian noise, and use the whitening tool from signal detection theory to construct a measure of difference between such images. While our Gaussian assumption is a simplification of reality, we find that this measure functions well in practice for both synthetic and real smooth objects. Our model does not apply for objects associated with rougher geometry and sharp albedo changes. Existing methods handle these situations well, suggesting that a combined method would function best for general objects.

1 Introduction

Image comparison is a fundamental component in many computer vision tasks such as recognition, alignment and tracking. Various measures for comparing images under varying illumination have been proposed[6, 13, 15, 1, 3]. These methods can be shown to work well on objects containing discontinuities or places of rapid change in albedo or shape. However, comparing images of smooth surfaces with no edges or texture under varying illumination remains a challenging problem. This problem is important since most real surfaces are complex and contain smooth, untextured areas along with rough and textured ones. The more accurately we can handle all types of shape, the better we can hope to achieve accurate recognition or dense registration or tracking. In this paper we propose a new measure for image comparison especially designed for smooth surfaces. We demonstrate this comparison method on the problem of recognition under varying illumination and fixed pose.

Previous approaches have focused on comparing representations of images that are invariant, or quasi-invariant to lighting. Edges are a classic example. [3] discuss the quasi-invariance of derivative operators to lighting changes. Gabor jets are also widely used for image comparison, in part because they are also considered to be insensitive to lighting changes (eg., [15]). [6] point out that the direction of the gradient is relatively insensitive to lighting changes. However, it is well-known that quasi-invariance to lighting changes is difficult to achieve for smooth objects. This is made explicit in the analysis of [6], which shows that gradient direction is truly invariant to lighting direction for surfaces with discontinuities, and varies more rapidly the smoother an object is.

We approach this problem by constructing a statistical model of smooth shapes and then using this to model the effect that lighting changes have on the appearance of a smooth object. Specifically, let I_1 and I_2 be two images of the same surface taken under different lighting conditions, and define the *difference image* to be $I_d = I_1 - I_2$. With a statistical characterization of I_d in hand, we can then design an appropriate detection strategy for image comparison.

For example, the simplest image comparison method is to just take $\|I_d\|$, which is equivalent to taking the sum-of-squared-differences (SSD) between I_1 and I_2. This is optimal if I_d is independent, identically distributed Gaussian noise. However, nearby points on smooth surfaces have highly correlated normals that produce highly correlated gray levels. So nearby values of I_d are not independent. We tackle this problem by finding operators to decorrelate the pixels of I_d prior to computing its magnitude. Signal detection theory tells us that this is the optimal approach when I_d consists of colored (non-independent) Gaussian noise. We show that for a simple model of smooth surfaces, this is a good characterization of I_d.

We use whitening to lessen dependencies in I_d. Whitening is a common technique in statistics and signal processing. It relies on a known second order statistical description of I_d. Whitening has also been used for decorrelation of images in different image processing tasks such as watermarking [8, 7], image restoration [19, 4, 5], and texture feature extraction [9, 16]. Many methods have used some differential operators or the Laplacian [18] to approximate the whitening filter, though Lin and Attikiouzel [16] used a 2D causal linear prediction model to derive whitening filters in a feature extraction framework.

We start by specifying a simple model of smooth, Lambertian objects and derive the covariance structure of the images they produce. This allows us to theoretically justify the use of whitening. Then we propose a way to implement the whitening process and to use it for image comparison. The method is tested in comparison to other methods on synthetic and real data and it is shown to be very effective for smooth objects. In a complete system our approach can be combined with other methods that perform well on rough surfaces.

2 The Whitening Approach

Our goal is to define a distance between images that will be low between two images, I_1, I_2 that were created by the same object under two different lighting conditions, and large otherwise. A perfect solution is impossible, because two different objects can look the same under different lighting conditions [2] and, in fact, for any two images there always exists a single object that could produce both, under two different lighting conditions, [6].

We will take a statistical approach, regarding the difference image, $I_d = I_1 - I_2$ as a random variable. We analyze this considering a Lambertian surface illuminated by distant point sources. Neglecting shadows, we can model the images as: $I_1 = \rho N s_1$ and $I_2 = \rho N s_2$, where N are surface normals, ρ is albedo, and s_1, s_2 are light sources in two images. The difference image then is an image itself, associated with the same object geometry and a difference lighting

$$I_d = \rho N s_1 - \rho N s_2 = \rho N(s_1 - s_2) \tag{1}$$

Dependencies that exist between the nearby surface normals of an object lead to dependencies in I_d. To handle these we model I_d as *colored* Gaussian noise. Colored Gaussian noise captures noise with dependencies, whereas white noise is independent. Note that the difference between two images associated with different objects would be transformed to white noise as well, since each image individually would be turned into white noise. However, the difference image produced by two different objects would be of much larger magnitude.

While this model is not strictly true, we will show that it is a valuable approximation that opens the way to using linear filtering to reduce dependency in the difference image. We will next review whitening, from signal processing, which removes correlations in Gaussian noise. Then we will present a simple model of smooth surfaces that allows us to analyze the dependencies in I_d. We then show how to whiten difference images. Finally, we describe how to use the results for recognition.

2.1 Whitening in Signal Processing

Let $\mathbf{n}$ represent a set of pixels in the difference image, as a vector. Assume that $\mathbf{n}$ is Gaussian colored noise. This implies that it is fully characterized by its first and second order statistics. In particular, the whitening filter may be designed using the covariance matrix. Let $C = E[\mathbf{n}\mathbf{n}^T]$ be the covariance matrix characterizing the distribution of $\mathbf{n}$ (E denotes expected value). Let W be a matrix composed of the scaled eigenvectors of C, $\frac{1}{\sqrt{\lambda_i}}\mathbf{e}_i$ as rows. Then, the components of $\mathbf{y} = W\mathbf{n}$ are independent, as apparent from their Gaussianity and their covariance:

$$E[\mathbf{y}\mathbf{y}^T] = diag(\lambda_1, \lambda_2, \ldots \lambda_m)$$

That is, the multiplication by the matrix W "whitens" the vector $\mathbf{n}$.

2.2 A Covariance Model for Natural Images - Rough Plane Covariance

Consider a surface, characterized by its normal vectors, which make a small random perturbation about a common direction (without loss of generality the z axis). We refer to such a surface as *roughly planar* and assume that locally a smooth surface behaves like a roughly planar surface. This is a generalization of the common facet model [12]. The surface is given by a function $z = f(x)$. The normals at every point x are random (but not independent!) and every one of them is specified by a single parameter θ, which is its angle relative to the z axis (Figure 1). Quantitatively we characterize the function $\theta(x)$ as a wide sense (w.s.) stationary Gaussian random process [17]. That is, we assume that the expected value at every point is constant $\mu_\theta = 0$, that the variance $C_\theta(x, x) = \sigma_\theta^2$ is constant as well,

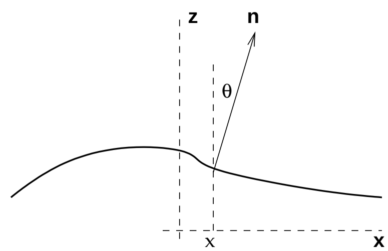

Figure 1: The roughly planar (random) surface is specified (in 2D approximation) by the angle $\theta(x)$ that the normal makes with the z direction.

and that the auto-correlation

$$C_\theta(x_1, x_2) = r(x_1, x_2)\sigma_\theta^2 = r(|x_1 - x_2|)\sigma_\theta^2$$

depends only on the distance between two points. $r(|x_1 - x_2|)$ is a correlation coefficient. We also assume that the surface is Lambertian, and that its albedo ρ, is constant, at least locally. For a distant light source, illuminating the surface at angle ϕ (relative to the z axis), the reflected light at the location x is

$$I(x) = \rho(sin\theta(x), cos\theta(x))(sin\phi, cos\phi)^T$$

Proposition 1: Under the above assumptions the reflected light function $I(x)$ is a random w.s. stationary process. Its expected value, variance and auto-correlation are:

$$E[I(x)] = \rho cos\phi e^{-\sigma_\theta^2/2}$$

$$\sigma_I^2 = \frac{1}{2}\rho^2(sin^2\phi(1 - e^{-2\sigma_\theta^2}) + cos^2\phi(1 - e^{-\sigma_\theta^2})^2)$$

$$\begin{aligned} C_I(x_1, x_2) &= \frac{1}{2}\rho^2(sin^2\phi e^{-\sigma_\theta^2}(e^{r\sigma_\theta^2} - e^{-r\sigma_\theta^2}) \\ &+ cos^2\phi(e^{-\sigma_\theta^2}(e^{r\sigma_\theta^2} + e^{-r\sigma_\theta^2}) - 2e^{-\sigma_\theta^2})) \end{aligned} \quad (2)$$

where x_1, x_2 are the two points for which the correlation coefficient of the tangent direction is $r = r(|x_1 - x_2|)$.

Proof: The reflected light function $I(x)$ is a random process. Its expected value is:

$$E[I(x)] = \rho(sin\phi\, E[sin\theta(x)] + cos\phi\, E[cos\theta(x)]) = \rho cos\phi E[cos\theta(x)] = \rho cos\phi e^{-\sigma_\theta^2/2}. \tag{3}$$

In computing the variance, we make use of the Gaussian integral $\int_\infty cosx\, e^{-x^2/2a^2} dx = \sqrt{2\pi}|a|e^{-a^2/2}$ [11].

The variance of $I(x)$ is:

$$\begin{aligned} \sigma_I^2 &= E[(I(x) - E[I(x)])^2] \\ &= \rho^2 E[(sin\phi sin\theta + cos\phi cos\theta - cos\phi E[cos\theta])^2] \\ &= \rho^2 (sin^2\phi E[sin^2\theta] + cos^2\phi E[(cos\theta - E[cos\theta])^2] \\ &\qquad + sin\phi cos\phi E[sin\theta(cos\theta - E[cos\theta])]) \\ &= \frac{1}{2}\rho^2(sin^2\phi(1 - e^{-2\sigma_\theta^2}) + cos^2\phi(1 - e^{-\sigma_\theta^2})^2) \end{aligned} \tag{4}$$

(The last term in the third line integrates to zero. Simple trigonometric expressions like $sin^2\theta = \frac{1}{2}(1 - cos2\theta)$ and the Gaussian integral above suffice to derive this expression.).

Consider now the autocorrelation. Let x_1, x_2 be two points for which the correlation coefficient of the tangent direction is $r = r(\|x_1, x_2\|)$. Then,

$$\begin{aligned} C_I(x_1, x_2) &= E[(I(x_1) - E[I(x)])(I(x_2) - E[I(x)])] \\ &= \rho^2 E[(sin\phi sin\theta_1 + cos\phi cos\theta_1 - cos\phi E[cos\theta_1]) \cdot \\ &\qquad (sin\phi sin\theta_2 + cos\phi cos\theta_2 - cos\phi E[cos\theta_2])] \\ &= \rho^2(sin^2\phi E[sin\theta_1 sin\theta_2] + cos^2\phi E[cos\theta_1 cos\theta_2] - cos^2\phi E[cos\theta]^2 \\ &= \frac{1}{2}\rho^2(sin^2\phi e^{-\sigma_\theta^2}(e^{r\sigma_\theta^2} - e^{-r\sigma_\theta^2}) + cos^2\phi(e^{-\sigma_\theta^2}(e^{r\sigma_\theta^2} + e^{-r\sigma_\theta^2}) - 2e^{-\sigma_\theta^2})) \end{aligned}$$

Note that all $sin\theta_i cos\theta_j$ terms vanish due to symmetry. The rest of the derivation requires us to change variables, to the sum and difference of θ_1 and θ_2, which are independent. We write their explicit (2D Gaussian) distribution, and use the Gaussian integral.

□

As expected, for $r = 1$ (same pixel), the correlation reduces to the variance expression, and for $r = 0$ (distant pixels), the correlation becomes 0. For rougher surfaces (larger σ_θ^2) the correlation is lower for the same distance $|x_1 - x_2|$, and for the (impossible) white surface (independent normals, $r = 0$), the image is white as well.

Note that the covariance in eq. 2 may be written as

$$\begin{aligned} C_I(x_1, x_2) &= sin^2\phi f_1(r(|x_1 - x_2|)) \\ &+ cos^2\phi f_2(r(|x_1 - x_2|)) \end{aligned} \tag{5}$$

where

$$f_1(r(|x_1 - x_2|)) = \frac{1}{2}\rho^2 e^{-\sigma_\theta^2}(e^{r\sigma_\theta^2} - e^{-r\sigma_\theta^2})$$

and

$$f_2(r(|x_1 - x_2|)) = \frac{1}{2}\rho^2 (e^{-\sigma_\theta^2}(e^{r\sigma_\theta^2} + e^{-r\sigma_\theta^2}) - 2e^{-\sigma_\theta^2})$$

The covariance in eq. 5 is not a constant and it varies with ϕ. Note that ϕ is the angle between the light and the predominant surface normal in a planar patch of the object, and so varies across the surface of an object. Therefore, it seems that it is more difficult to learn the covariance and less justified to use a single filter (as described below, in section 2.3) to implement whitening. Note, however, that $f_1()/f_2() \approx (2 + \sigma_\theta^2)/\sigma_\theta^2$, implying that for say, $\sigma_\theta = 0.1$, and for all $\phi > 12$ degrees, the first term in eq. 5 is ten times larger than the second term and thus dominant. Therefore, we can conclude:

Covariance characterization for rough Lambertian plane: the second order statistical behavior of a rough Lambertian plane surface, illuminated by a single source, is characterized by an autocorrelation function which, for nearly every illumination, is approximately invariant of the illumination direction up to a multiplicative factor.

For empirical verification of this result see Section 3.4. For a real object, the normals may change substantially about the average direction, but still the approximation of the autocorrelation by one function (and a slowly varying multiplicative factor) is good for many objects and most locations on them. This fails only where the normal roughly points at the illumination source, or where there are significant effects of perspective foreshortening. Note that for rougher objects, associated with high normal direction variance, the image correlation is reduced and the image gets closer to being white itself.

For image comparison, we are mostly concerned with the difference images, which for Lambertian objects, behave as a single image does.

2.3 Whitening using AR models

The model described above gives the autocorrelation of reflected light images, but usually its parameters (σ_θ^2, r) are unknown. Estimating the autocorrelation from data is considered to be a noisy unstable process. Fitting the parametric Autoregressive (AR) model, is a way to get the whitening filter directly without explicitly calculating the covariance [14].

A sequence $x(n)$ is called an AR process of order p if it can be generated as the output of the recursive causal linear system

$$x(n) = \sum_{k=1}^{p} a(k)x(n-k) + \varepsilon(n), \forall n \tag{6}$$

where $\varepsilon(n)$ is white noise. The estimate

$$\bar{x}(n) = \sum_{k=1}^{p} a(k)x(n-k) \tag{7}$$

is the best linear MS predictor of $x(n)$ based on the previous p samples. For Gaussian signals the prediction error sequence: $\varepsilon(n) = x(n) - \bar{x}(n)$ is white. Therefore the prediction error filter is a whitening filter for $x(n)$,

$$W = (1, -a_1, \ldots, -a_p). \tag{8}$$

The AR model can be applied to the whitening of the difference image (eq. 1). We adopted a 2D "causal" model described in [14], where a gray level $x(n)$ is predicted from the previous gray levels in a $p \times p$ neighborhood in column by column scan (Figure 2). We denote this set of pixels by N^n. As the image is not truly a Gaussian AR process, we determine the size of the whitening filter based on the image correlation. See section 3. Note that using a non-causal neighborhood leads to a lower SSD, but the prediction error sequence is not white [14].

Since $\bar{x}(n)$ is the best linear MS predictor of $x(n)$, the coefficients a_k in eq.7 can be determined by solving

$$\arg\min_{\{a_k\}} \sum_{n | N^n \in image} (x(n) - \bar{x}(n))^2$$

This is an over-determined least squares problem that can be solved using SVD. We propose to estimate the whitening filter in this way, using image differences from the domain in which we will operate.

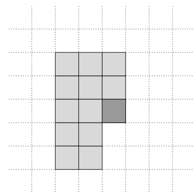

Figure 2: A typical 2D “Causal” neighborhood used in our experiments. The darkest pixel is the one where the gray level is estimated from the grey levels in the other dark pixels.

Note that we could estimate the covariance matrix from subimage samples and then use the Yule-Walker method [17] to get the AR coefficients. There are some subtle difficulties with this approach: First, as shown above, the covariance is not stationary, and therefore the covariance matrix is not specified. Moreover, the average is not constant as well, and an estimate of the local average is required for the Yule-Walker method.

Note that scaling all the grey levels by the same factor, would give a correlation function which is the same up to a multiplicative constant. This is essentially what happens when the angle between the average normal and the illumination direction is changed, in the model described in section 2.2. Note however, that this does not change the AR coefficients at all. Therefore, the whitening filter is invariant to a multiplicative change in the correlation, implying that a space invariant filter may be used for whitening the image.

The whitening filter depends on the image statistics. Intuitively, for smoother images the correlation is larger and decorrelating it requires a wider filter. For images which are not so smooth the decorrelation is done over a small range, and the filter looks very much like the Laplacian, which is also known to have some whitening effect. Therefore, for rougher images, we do not expect to perform better than an alternative procedure using the Laplacian. As we shall see later, for smooth objects the performance difference is significant.

2.4 Recognition

We use the following recognition scheme as an application of our approach to image comparison: A set of objects is represented in a library containing one image for every object. Let $I_{M_1}, I_{M_2}, \ldots, I_{M_j}, \ldots$ be reference images in the library. Let I_Q be the query image of one of the objects from this set, taken with the same pose, but different illumination. The task is to decide which of the objects is the one in the query image.

The proposed basic algorithm is very simple:

1. For every reference image, I_{M_j}, whiten the difference image to get $W(I_Q - I_{M_j})$. Calculate the L_2 norm $E_j = \|W(I_Q - I_{M_j})\|$.

2. Choose the model associated with the smallest whitened error norm, E_j.

This algorithm is used in the context of communication, where the given signal (image) is a sum of the (known) clean signal (model) and uniform noise. The situation here is a bit different: First, the reference image I_{M_j} was taken with a different illumination intensity than the test image. Therefore, every scaled version of it is a valid model as well. Minimizing the SSD over all scaled versions is equivalent to taking the SSD between the normalized whitened images, $\|\frac{W(I_{M_j})}{\|W(I_{M_j})\|} - \frac{W(I_Q)}{\|W(I_Q)\|}\|$. This normalization also compensates for the fact that some objects are rougher than others, which makes the difference between two differently illuminated images of them larger.

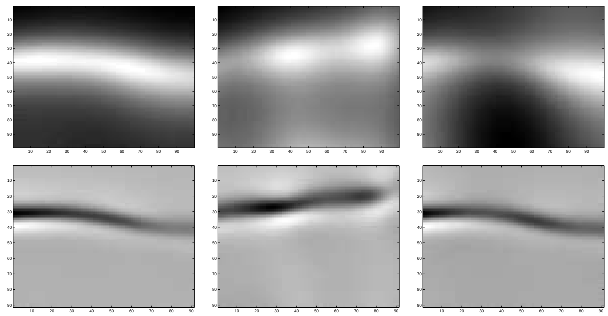

Figure 3: Whitening often reveals hidden differences: The left and middle image in the top row correspond to different surfaces and one illumination. The third one is created from the same scene used for the left image, but with a different illumination. The bottom row shows the same images after whitening.

3 Experiments

We tested the whitening approach against other methods in several experiments. The whitening approach was developed to handle smooth textureless surfaces on which other methods perform poorly. Most of the experiments were designed with this goal in mind, but we also experimented with the Yale database of human faces [10], as an example of surfaces that contain albedo discontinuities and regions of high curvature.

3.1 Synthetic images

The first set of experiments was done using synthetic images. Every scene was created as a sum of random harmonic functions, with (roughly) fixed amplitudes but random directions and phases. This way, we get an ensemble of different images with similar statistical properties. These were rendered as Lambertian surfaces with point sources.

For every set of such images we learned a whitening filter as the 2D causal

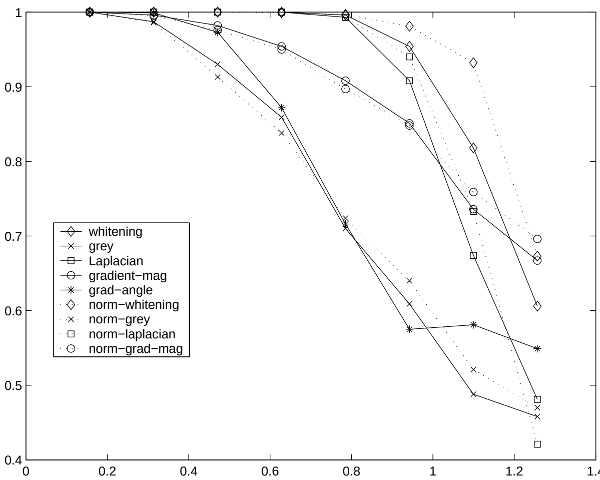

Figure 4: Recognition success rate for simulations on synthetic smooth surfaces using whitening, grey levels, Laplacian, Gradient Magnitude and gradient angles. The success rate is plotted against the angle (in radians) between the illumination source and the average surface normal. The solid curves correspond to the unnormalized quantities and the dotted curves are for the normalized quantities.

filter that minimizes the MS prediction error. A typical filter had 265 coefficients inside a 23×23 window. For learning we used a fixed illumination, deviating $3\pi/8$ degrees from the z direction and 1000-5000 random images. The learning set was independent of the test set.

A test was done as follows: two random scenes were illuminated by the same nearly vertical illumination to create two references images I_r, I'_r. The test image I_t was synthesized from the first scene, with a different illumination, making an angle ϕ with the z axis. A typical triple of such images is shown in Figure 3.

For comparison we also tested other algorithms using the SSD of the gray level image, of the result of a Laplacian, of the magnitude of the gradient and of its angle. Both non-normalized and normalized versions were tested for all representations (except the gradient angle, for which normalization seems unnecessary

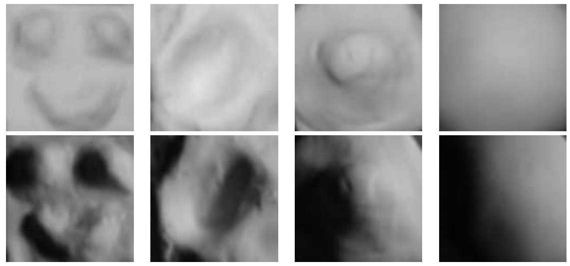

Figure 5: Samples from the smooth real objects data set; top – frontal illumination, bottom – side illumination.

and no normalization seems reasonable). See Figure 4 for the results.

We came to several conclusions:

1. For the relatively smooth images that we used, whitening was the best method.

2. The advantage of the whitening method was higher when the filter was larger but even for filters which were as small as 7×7 the advantage was substantial and consistent over all viewing angles except the very extreme ones.

3. In particular whitening was always better than the Laplacian, even when a 3×3 filter was used, implying that both large distance correlations and causality are important.

4. In most cases, the normalized version worked better than the un-normalized version.

3.2 Real Smooth Objects

The experiments above show that for smooth objects, Lambertian reflection and no shadows, the whitening approach performs significantly better than other mea-

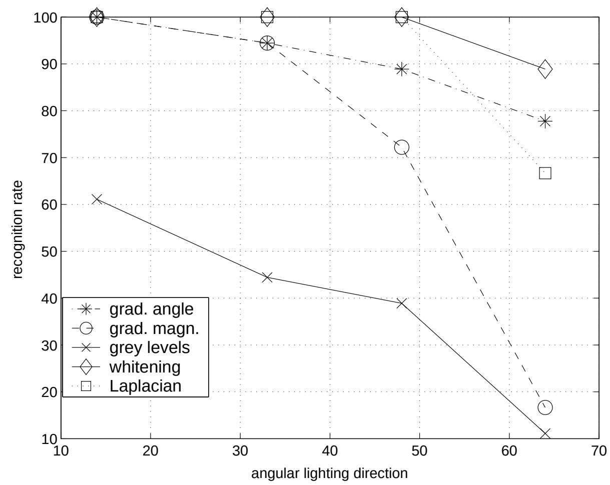

Figure 6: Recognition success rate for experiments done with real smooth objects, with several methods. The success rate is plotted against the average angle (in degrees) between the illumination source and the average surface normal.

sures. We now relax these assumptions and repeat the experiment on real, smooth objects that produce images with substantial shadows (Figure 5). We created eighteen objects from clay and illuminated them by a single light source moving along a half circle, so that its distance from the object is roughly fixed. We used a camera placed vertically above the object, and took 14 images of every object with different lighting directions at angles in the range $[-70, 70]$ degrees to the vertical axis. The images associated with a nearly vertical illumination (one for every object) were chosen as the reference images.

The whitening filter was trained on the difference images between reference images and corresponding images associated with the same object and six other illuminations. Only twelve images associated with 2 objects (out of 18) were used. We learned the whitening filter, as a 2D causal filter with 25 coefficients inside 7×7 windows.

The task was to identify an object in a given image with unknown lighting direction. All images of the 18 objects except the reference images were used as query images (234 images). We divided the query images into four groups according to their angular lighting direction:$10^o - 25^o$, $26^o - 40^o$, $41^o - 55^o$, and $56^o - 70^o$. The plot in Figure 6 shows the performance of the whitening approach in comparison to other tested measures (normalized versions).

Whitening again performed better than the other methods. Taking a closer look at the data, we observed that for a few of the clay objects that were roughest, the Laplacian, whitening and gradient angle performed equally well. For smoother objects, however, whitening worked considerably better. The Laplacian couldn't whiten the smooth surfaces, because its size was insufficient to handle the high correlations between the grey levels of the smooth surfaces.

3.3 Images of Face

To illustrate the limitations of the whitening method we tested it also on the Yale face database [10], corresponding to rougher objects and associated with abrupt albedo changes. We replicated the results in [6], (Figure 7) showing that the gradient angle was indeed a winner of all the methods considered here. Whitening does not perform well on this data, because it fails on edges, which can not be accounted for by the model (section 2.2). The Laplacian works better then whitening, but worse than gradient angle, because it performs as an edge detector, which is still sensitive to large illumination changes.

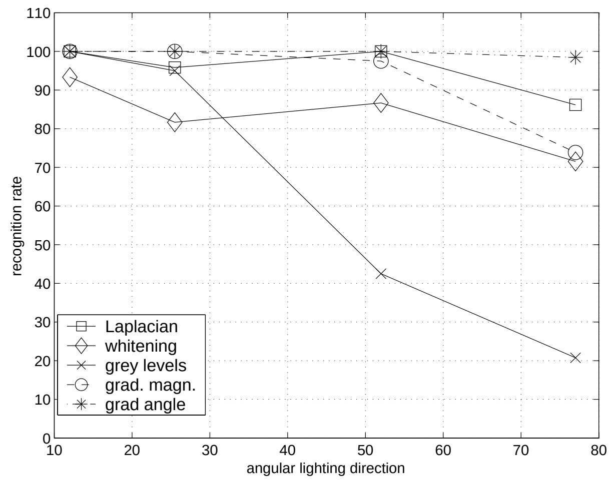

Figure 7: Recognition success rate for experiments done with real faces (Yale database), with several methods. The success rate is plotted against the average angle (in degrees) between the illumination source and the average face normal.

3.4 Verifying Covariance Model

To confirm the multiplicative behavior of the covariance model described in Section 2.2, we took a high resolution image of a real, approximately Lambertian sphere, illuminated by a point source. We divided the image into 50×50 pixel patches, and calculated covariance in every patch using 5×5 neighborhoods. Figure 8 shows the estimated covariance as a function of the distance. Each curve represents a different patch. The plots confirm that covariance in different patches differs by a multiplicative factor as claimed in Section 2.2.

4 Conclusions

In this work we have proposed a measure for image comparison of smooth surfaces under varying illumination. The measure was motivated by a simple statistical model of smooth surfaces. This model showed that the error between two

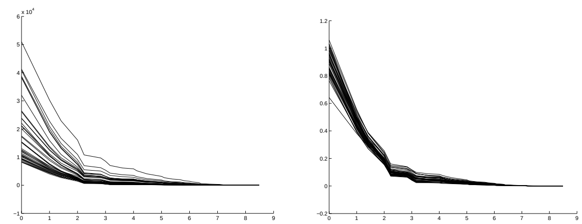

Figure 8: Covariance estimates for different patches of a real objects: left – non-normalized covariances differ by multiplicative factors; right – covariances which are normalized by the variance are almost the same for all angles.

images associated with the same object under different lighting may be modelled as colored noise. We adapted well-known techniques of whitening to perform matching of images corrupted by such noise.

We found that whitening was more effective than other representations for comparing images of smooth surfaces taken under varying illumination conditions. Previous methods have commonly used the Laplacian or the magnitude of gradient, as whitening approximations. This seems to be adequate for rough images but leads to inferior results for smoother ones.

We believe that recognition (or image comparison in general) should use all the image information. Many current methods neglect photometric information and thus cannot handle smooth objects. It seems that combined methods, using both the (more geometric) information in edges and the (more photometric) information in smooth patches, would yield superior results, especially in hard tasks.

References

[1] P.N. Belhumeur, J.P. Hespanha, and D.J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class-specific linear projection. *PAMI*, 19(7):711–720, July 1997.

[2] P.N. Belhumeur, D.J. Kriegman, and A.L. Yuille. The bas-relief ambiguity. *IJCV*, 35(1):33–44, November 1999.

[3] R. Brunelli and T. Poggio. Face recognition: Features versus templates. *PAMI*, 15(10):1042–1062, 1993.

[4] B. Bundschuh. A linear predictor as a regularization function in adaptive image restoration and reconstruction. In *5th Int, Conf. on Computer Analysis of Images and Patterns*, 1993.

[5] H. Bundschuh, B. Schulz and D. Schneider. Adaptive least squares image restoration using whitening filters of short length. In *Second HST Image Restoration Workshop*, 1993.

[6] H.F. Chen, P.N. Belhumeur, and D.W. Jacobs. In search of illumination invariants. In *CVPR00*, pages I: 254–261, 2000.

[7] M. L. Cox, I. J. Miller and Bloom J. A. *Digital Watermarking*. Morgan Kaufmann, 2002.

[8] T. Depovere, G. Kalker and J.P. Linnartz. Improved watermark detection using filtering before correlation. In *IEEE Int. Conf. on Image Processing*, pages I: 430–434, 1998.

[9] O.D. Faugeras and W.K. Pratt. Decorrelation methods of texture feature extraction. *PAMI*, 2(4):323–332, July 1980.

[10] P.N. Georghiades, A.S. Belhumeur and D.J. Kriegman. From few to many: Generative models for recognition under variable pose and illumination. *PAMI*, 23(6):643–660, 2001.

[11] I.S. Gradshteyn and I.M. Ryzhik. *Table of Integrals, Series and Products*. New York, Academic, 1980.

[12] R.M. Haralick and L.G. Shapiro. Computer and robot vision. In *Addison-Wesley*, 1992.

[13] D.W. Jacobs, P.N. Belhumeur, and R. Basri. Comparing images under variable illumination. In *CVPR98*, pages 610–617, 1998.

[14] A.K. Jain. Fundamentals of digital image processing. In *Prentice Hall*, 1989.

[15] M. Lades, J.C. Vorbruggen, J. Buhmann, J. Lange, C. von der Malsburg, R.P. Wurtz, and W. Konen. Distortion invariant object recognition in the dynamic link architecture. *TC*, 42(3):300–311, March 1993.

[16] Z. Lin and Y. Attikiouzel. Two-dimensional linear prediction model-based decorrelation method. *PAMI*, 11(6):661–665, June 1989.

[17] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw Hill, 3rd edition, 1991.

[18] W.K. Pratt. *Digital Image Processing (First Edition)*. Wiley, 1978.

[19] L.P. Yaroslavsky. *Digital Picture Processing. An Introduction.* Springer Verlag,Berlin, Heidelberg, 1985.