

Symmetric Objects are Hardly Ambiguous*

Gaurav Aggarwal^{††}, Soma Biswas[†] and Rama Chellappa[†]

Center for Automation Research, UMIACS

[†]Dept. of ECE, ^{††}Dept. of Computer Science

University of Maryland, College Park

{gaurav, soma, rama}@cfar.umd.edu

Abstract

Given any two images taken under different illumination conditions, there always exist a physically realizable object which is consistent with both the images even if the lighting in each scene is constrained to be a known point light source at infinity [10]. In this paper, we show that images are much less ambiguous for the class of bilaterally symmetric Lambertian objects. In fact, the set of such objects can be partitioned into equivalence classes such that it is always possible to distinguish between two objects belonging to different equivalence classes using just one image per object. The conditions required for two objects to belong to the same equivalence class are very restrictive, thereby leading to the conclusion that images of symmetric objects are hardly ambiguous. The observation leads to an illumination-invariant matching algorithm to compare images of bilaterally symmetric Lambertian objects. Experiments on real data are performed to show the implications of the theoretical result even when the symmetry and Lambertian assumptions are not strictly satisfied.

1. Introduction

The problem of matching images of an arbitrary scene/object under different illumination conditions has been quite elusive. Lack of information about the geometry and reflectance map makes this problem in its generality, ill-posed. In fact, Jacobs *et al.* [10] show that this problem cannot be solved even under hard constraints of Lambertian reflectance and known single point light sources placed at infinity.

Quite often in vision problems, the intractability of the problem can be reduced significantly by restricting the domain of the problem and using appropriate constraints. In this paper, we analyze the problem of matching symmetric objects across illumination variations. In particular, we

show that unlike general objects, it is almost always possible to distinguish between two bilaterally symmetric objects using just one image per object.

The symmetry assumption eliminates the unknown albedo in the Shape from Shading (SFS) formulation, thereby allowing us to deal with arbitrarily varying albedo maps. Moreover, symmetry leads to a linear constraint on the values of the unknown surface gradients for each point of the object. Though the constraint makes the SFS problem more tractable, it is still not sufficient to recover the surface gradients for general unknown albedo maps.

Unlike the existing work on symmetric SFS, our goal here is illumination-invariant matching rather than shape recovery. We use the linear constraint provided by symmetric SFS to prove the well-posedness of the matching problem for the class of bilaterally symmetric objects. Given two linear constraints from two different images, we solve for the surface gradients. The correctness of the gradients can be checked by substituting them back in the original image irradiance equations for the images and computing albedo from the two separately. We show that the two albedo estimates are identical if the corresponding pixels represent the same physical reality (same shape and albedo). If the points differ physically, the computed albedos almost always differ. We derive the rare condition under which they are same. In fact, the condition partitions the set of symmetric Lambertian objects into equivalence classes such that it is always possible to distinguish between two different objects belonging to different equivalence classes based on just one image per object.

The theoretical analysis leads to an algorithm that can be used to match images of real objects where the symmetry and Lambertian assumptions are not strictly satisfied. Given an image, an illumination-invariant representation is derived that can be used for matching. If the assumptions are strictly satisfied, the algorithm is provably correct (up to the described ambiguity). Experimental results show the usefulness of the approach on real images.

*Partially funded by an NSF-ITR Grant 03-25119

1.1. Organization of the Paper

The rest of the paper is organized as follows. Section 2 discusses the related work. The SFS formulation utilizing the 3D bilateral symmetry is described in Section 3. The theoretical analysis to prove that the images of symmetric objects are hardly ambiguous is outlined in Section 4. In Section 5, we propose an algorithm to perform illumination-invariant matching of such objects. Experiments performed to evaluate the performance of the matching algorithm are described in Section 6. The paper concludes with a summary and discussion in Section 7.

2. Previous Work

There has been a lot of work on the problem of illumination-invariant matching and recognition. Brooks *et al.* [4] discuss the existence and uniqueness of shapes consistent with a given intensity pattern. In [9], a given image is filtered to suppress the lighting effects in order to recover the object reflectance. A method to recover intrinsic properties of an object using multiple images is proposed in [13]. Jacobs *et al.* [10] describe a matching algorithm based on the observation that the ratio of two images of the same object is simpler than that of two different objects. Chen *et al.* [5] utilize the insensitivity of the direction of image gradients to changes in illumination direction in a probabilistic framework to recognize faces across illumination.

Other than these generic methods, a lot of research has been directed towards recognizing faces across illumination variations. Quite often face-specific methods physically model the image formation process which involves illumination sources, albedo and shape. Class specific properties of faces have been utilized to perform reliable reconstruction or recognition in spite of the ill-posed nature of the problem. [3][19][16][8][14][1] are a few remarkable works in this direction.

Yuille *et al.* [16] use singular value decomposition (SVD) to learn generative models of objects from a set of images taken under different unknown illuminations. Shashua and Raviv [14] perform recognition across varying illumination under an ideal-class assumption. All objects belonging to the ideal class are assumed to have the same shape. [8] uses illumination cone models for illumination-invariant face recognition. They require a small number of training images of each face under different illuminations to recover the shape and albedo of the face. Basri and Jacobs [1] propose methods for recovering surface normals in a scene. Result in [2] and [12] forms the basis of their work, which proves that the set of all Lambertian reflectance maps obtained with arbitrary distant illumination sources approximately lie in a 9D linear subspace. In [3], Blanz and Vetter perform face recognition across pose and illumination by fitting a 3D morphable model to the images. Zhou *et*

al. [19] generalize the traditional photometric approach to handle all appearances of all objects in a class. They impose a rank constraint on shape and albedo in a class to separate the two from illumination.

Though SFS approaches for the recovery of shape and albedo have been studied for a long time, it is only recently that attempts have been made to use them for real matching problems. Due to the ill-posed nature of the problem, the SFS research typically makes uniform albedo assumption which often limits the applicability of the approaches. In a recent work [17][18], Zhao and Chellappa present an SFS approach to recover both shape and albedo for a symmetric object from a single image under piecewise constant constraint on albedo. In [17], they use the same approach for generating frontally illuminated prototype images to perform face recognition. They use partial gradient information from a generic 3D model to perform this task. Using the same formulation, Dovgird and Basri [6] make use of class-specific constraints by writing the unknown surface gradients as a linear combination of the surface gradients of a set of known 3D face models to recover the shape.

Though our work is partly motivated by Zhao and Chellappa's work [17][18], we differ in the following aspects

1. We derive precise conditions under which images of two different objects are ambiguous.
2. Our approach for illumination-invariant matching is provably correct for symmetric Lambertian objects.
3. We do not use any class-specific information like generic 3D model as used in [17].

3. Symmetric Shape from Shading

Under the assumptions of orthographic projection and Lambertian reflectance, the perceived intensity of a surface point of an object can be written as

$$I = L\rho \frac{1 - pl - qk}{\sqrt{p^2 + q^2 + 1}\sqrt{l^2 + k^2 + 1}} \quad (1)$$

where ρ is the surface albedo, $\frac{(p,q,1)}{\sqrt{p^2+q^2+1}}$ is the surface normal, L is the intensity of the light source and $\frac{(l,k,1)}{\sqrt{l^2+k^2+1}}$ is the illuminant direction. As done normally in SFS formulations, we assume that the image intensity I is normalized by the known light source intensity to eliminate L from the expression.

The albedo ρ_- and surface normals $\{p_-, q_-\}$ of the bilaterally symmetric point are characterized as follows

$$\rho_- = \rho \quad \{p_-, q_-\} = \{-p, q\} \quad (2)$$

Therefore, its intensity I_- can be written in terms of the albedo and surface normals of its symmetric counterpart as

follows

$$I_- = \rho \frac{1 + pl - qk}{\sqrt{p^2 + q^2 + 1}\sqrt{l^2 + k^2 + 1}} \quad (3)$$

Using (1) and (3), the albedo can be eliminated leading to the following linear constraint on the surface gradients

$$\frac{I_-}{I} = \frac{1 + pl - qk}{1 - pl - qk} \quad (4)$$

$$(I_- - I) - (I_- + I)pl - (I_- - I)qk = 0 \quad (5)$$

$$Slp + Dkq = D \quad (6)$$

where $S = I_- + I$ is the sum of the intensities of the symmetric points and $D = I_- - I$ is the difference of the two. The linear relation implies that the set of possible surface gradients $\{p, q\}$ lie on a straight line in the pq -space, parameterized by the perceived intensity and the lighting condition. Note that the regular reflectance map provides a quadratic constraint on the values surface gradients can take, given the pixel intensity, albedo and illumination conditions. Figure 1 shows the regular quadratic reflectance map and the corresponding linear constraints (6). Even if the albedo is known, there are two possible solutions for the unknown surface gradients. Though enforcing integrability [7] helps in removing the ambiguity completely for constant and piece-wise constant albedo maps, the problem is still ill-posed for the more general case of unknown arbitrary albedo map [18]. Though the shape recovery problem is still ill-posed, the formulation is quite useful for illumination-invariant matching as discussed next.

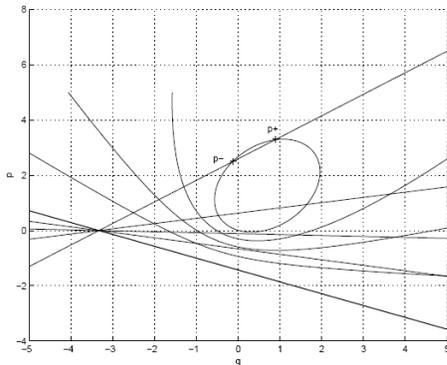


Figure 1. Regular and symmetric reflectance maps [18].

4. Role of Symmetry in Illumination-invariant Matching

In this section, we use the symmetric SFS formulation to analyze the problem of illumination-invariant matching for the class of bilaterally symmetric objects. Given an image

of a bilaterally symmetric object, each pair of symmetric points results in a linear constraint of the form (6). Given a second image of the same surface, we obtain another linear relation for each corresponding point pair which leads to the following Lemma.

Lemma 4.1 *The linear relations for a point with surface gradients $\{p_0, q_0\}$, derived from images taken under different light sources, are concurrent with $\{p_0, q_0\}$ as the point of concurrence.*

Proof Line $Slp + Dkq = D$ in the pq -space has to pass through the point $\{p_0, q_0\}$. This is true for all such lines derived from all possible images of the point under various illumination conditions. As two lines can intersect at only one point, the lines are concurrent with $\{p_0, q_0\}$ as the point of concurrence, which proves the lemma.

Therefore, if two images come from the same object, the corresponding lines intersect at their true surface gradient. Interestingly, even if the two points are not physically same (i.e., they have different surface gradients), the two lines still intersect in the pq -space unless they are parallel. As the points have different surface gradients, the point of intersection can not be the true surface gradient for both of them. These observations help us prove that it is possible to distinguish between two symmetric Lambertian objects using just one image per object as described in the following subsection.

4.1. The Ambiguity in Matching

In a matching scenario, the goal is to determine if the two images come from the same physical object or not. Given two images taken under different illumination conditions, we get an intersection point in the pq -space for each corresponding symmetric point pair, which is a possible solution for the unknown surface gradients.

For each pair of corresponding points from the two images, we get two linear constraints as follows

$$S_1 l_1 p + D_1 k_1 q = D_1 \quad (7)$$

$$S_2 l_2 p + D_2 k_2 q = D_2 \quad (8)$$

where the subscripts 1 and 2 distinguish the quantities corresponding to the two images. Unless they are parallel, the two lines intersect at a point (say $\{\bar{p}, \bar{q}\}$) in the pq -space. Substituting the intersection point back in the image irradiance equations (1) for the two images, following two albedo estimates are obtained

$$\hat{\rho}_1 = \frac{\sqrt{\bar{p}^2 + \bar{q}^2 + 1}\sqrt{l_1^2 + k_1^2 + 1}}{1 - \bar{p}l_1 - \bar{q}k_1} I_1 \quad (9)$$

$$\hat{\rho}_2 = \frac{\sqrt{\bar{p}^2 + \bar{q}^2 + 1}\sqrt{l_2^2 + k_2^2 + 1}}{1 - \bar{p}l_2 - \bar{q}k_2} I_2$$

From Lemma 4.1, if the two points have same surface gradients and albedo, then the two lines intersect at their true surface gradient. Substituting the true surface gradient back in the irradiance equation will always produce the same true albedo. Though not intuitive, it is possible to get $\hat{\rho}_1 = \hat{\rho}_2$ even when the two points are physically different (i.e., they differ either in surface gradients or albedo). The condition on the two points for this to happen is derived in the following theorem.

Theorem 4.2 *The two albedos $\hat{\rho}_1$ and $\hat{\rho}_2$ are same if the following condition is satisfied*

$$\frac{\rho_1}{\rho_2} = \frac{p_2 \sqrt{1 + p_1^2 + q_1^2}}{p_1 \sqrt{1 + p_2^2 + q_2^2}} \quad (10)$$

where ρ_1 and ρ_2 are the true albedos for the two points and $\frac{(p_1, q_1, 1)}{\sqrt{1 + p_1^2 + q_1^2}}$ and $\frac{(p_2, q_2, 1)}{\sqrt{1 + p_2^2 + q_2^2}}$ are the corresponding true surface normals.

Proof Suppose $\frac{(l_1, k_1, 1)}{\sqrt{l_1^2 + k_1^2 + 1}}$ and $\frac{(l_2, k_2, 1)}{\sqrt{l_2^2 + k_2^2 + 1}}$ are the illuminant directions for image 1 and 2 respectively. For image 1, the true surface gradients $\{p_1, q_1\}$ satisfy (7), i.e.,

$$S_1 l_1 p_1 + D_1 k_1 q_1 = D_1 \quad (11)$$

Using (11) and (7), we get

$$q = \frac{1}{k_1} - \frac{1 - k_1 q_1}{k_1 p_1} p \quad (12)$$

Similarly, for image 2, we have

$$q = \frac{1}{k_2} - \frac{1 - k_2 q_2}{k_2 p_2} p \quad (13)$$

These lines intersect at the following point $\{\bar{p}, \bar{q}\}$ in the pq -space

$$\bar{p} = \frac{p_1 p_2 (k_1 - k_2)}{p_1 k_1 (1 - k_2 q_2) - p_2 k_2 (1 - k_1 q_1)} \quad (14)$$

$$\bar{q} = \frac{p_1 (1 - k_2 q_2) - p_2 (1 - k_1 q_1)}{p_1 k_1 (1 - k_2 q_2) - p_2 k_2 (1 - k_1 q_1)} \quad (15)$$

Now the two albedos obtained by substituting $\{\bar{p}, \bar{q}\}$ back in the image irradiance equations for the two points are same if

$$\begin{aligned} & \frac{\sqrt{\bar{p}^2 + \bar{q}^2 + 1} \sqrt{l_1^2 + k_1^2 + 1}}{1 - \bar{p} l_1 - \bar{q} k_1} I_1 \\ &= \frac{\sqrt{\bar{p}^2 + \bar{q}^2 + 1} \sqrt{l_2^2 + k_2^2 + 1}}{1 - \bar{p} l_2 - \bar{q} k_2} I_2 \end{aligned} \quad (16)$$

i.e.,

$$\frac{1 - \bar{p} l_1 - \bar{q} k_1}{1 - \bar{p} l_2 - \bar{q} k_2} \cdot \frac{\sqrt{l_2^2 + k_2^2 + 1}}{\sqrt{l_1^2 + k_1^2 + 1}} = \frac{I_1}{I_2} \quad (17)$$

Substituting \bar{p} and \bar{q} from (14) and (15), the left hand side of (17) simplifies to

$$\frac{p_2}{p_1} \cdot \frac{1 - l_1 p_1 - q_1 k_1}{1 - l_2 p_2 - q_2 k_2} \cdot \frac{\sqrt{l_2^2 + k_2^2 + 1}}{\sqrt{l_1^2 + k_1^2 + 1}} \quad (18)$$

Also, the right hand side of (17) can be written in terms of the true surface gradients and albedos as follows

$$\frac{\rho_1}{\rho_2} \cdot \frac{1 - l_1 p_1 - q_1 k_1}{1 - l_2 p_2 - q_2 k_2} \cdot \frac{\sqrt{l_2^2 + k_2^2 + 1} \sqrt{p_2^2 + q_2^2 + 1}}{\sqrt{l_1^2 + k_1^2 + 1} \sqrt{p_1^2 + q_1^2 + 1}} \quad (19)$$

From (18) and (19), the condition in (17) is true if

$$\frac{p_2}{p_1} = \frac{\rho_1}{\rho_2} \cdot \frac{\sqrt{p_2^2 + q_2^2 + 1}}{\sqrt{p_1^2 + q_1^2 + 1}} \quad (20)$$

which proves the theorem.

Theorem 4.2 leads to a few interesting observations which are described in the following corollaries.

Corollary 4.3 *The condition in Theorem 4.2 is trivially satisfied if the two points have the same surface gradients and albedo.*

Corollary 4.4 *The condition in Theorem 4.2 can be true for points even if they differ either in surface gradients or albedo. This essentially means that the point characterized by surface gradients $\{\bar{p}, \bar{q}\}$ and albedo $\hat{\rho}_1 = \hat{\rho}_2$ can account for both the images, i.e., it is not possible to distinguish between the two points using just one image (of each point) even under hard constraints of bilateral symmetry, Lambertian reflectance and known distant point light sources.*

Corollary 4.4 establishes the ambiguity on a per-point basis. If this is true for all visible points of the two objects, then the two objects are indistinguishable given just one image per object taken under different illumination conditions. As chances of such a condition being satisfied by all the corresponding points of two objects are low, it can be concluded that symmetry helps in disambiguating images across illumination. Note that the condition is on the surface gradients and albedo maps of the objects and not on their particular images.

4.2. Equivalence Classes of Bilaterally Symmetric Objects

We consider the condition in Theorem 4.2 as a relation $R(i, j)$ relating two objects i and j (assuming the condition is satisfied for all corresponding point pairs). Hence, $R(1, 2)$ means that the condition is satisfied for all corresponding points of objects 1 and 2. It is interesting to see that relation R is

1. reflexive, i.e., $R(i, i)$ holds,
2. symmetric, i.e., $R(i, j)$ implies $R(j, i)$, and
3. transitive, i.e., $R(i, j)$ and $R(j, k)$ implies $R(i, k)$.

Therefore, the condition in Theorem 4.2 induces an equivalence relation on the set of all possible bilaterally symmetric objects. In other words, such a set can be partitioned into equivalence classes such that any two objects belonging to the same equivalence class cannot be distinguished using just one image per object. This follows directly from Corollary 4.4. On the other hand, two objects belonging to two different equivalence classes do not satisfy the condition in Theorem 4.2 and thus can always be distinguished using just one image per object.

5. Illumination-invariant Matching

If the assumptions of Lambertian reflectance and bilateral symmetry are reasonably adhered to, the formulation in Section 4.1 can directly be used to reliably match images across illumination. As the chance of getting images of two different objects that belong to the same equivalence class is very low, the algorithm should not make any error in matching.

Unfortunately, in most practical applications, the objects are neither Lambertian nor perfectly symmetric. From Section 4.1, two images are recognized as belonging to the same physical object, if the two estimated albedos $\hat{\rho}_1$ and $\hat{\rho}_2$ are same. $\hat{\rho}_1$ and $\hat{\rho}_2$ depend non-linearly on the estimated surface gradients $\{\bar{p}, \bar{q}\}$. Estimation of surface gradients $\{\bar{p}, \bar{q}\}$ in turn depends on how strictly the assumptions are adhered to. Deviations from the assumptions make the estimation of surface gradients $\{\bar{p}, \bar{q}\}$ and hence $\hat{\rho}_1$ and $\hat{\rho}_2$ quite unstable. The instability in the estimation makes the scheme unsuitable for real data.

Here, we propose a novel algorithm to match images of symmetric objects across illumination which follows naturally from Theorem 4.2. The algorithm does not involve estimation of $\{\bar{p}, \bar{q}\}$ or $\hat{\rho}_1$ and $\hat{\rho}_2$, and thus degrades quite gracefully when the assumptions are not strictly satisfied.

From Theorem 4.2 and Corollaries 4.3 and 4.4, two objects appear similar (given one image per object) iff

$$\frac{\rho_1}{\rho_2} = \frac{p_2 \sqrt{1 + p_1^2 + q_1^2}}{p_1 \sqrt{1 + p_2^2 + q_2^2}} \quad (21)$$

That is, iff

$$p_1 \frac{\rho_1}{\sqrt{1 + p_1^2 + q_1^2}} = p_2 \frac{\rho_2}{\sqrt{1 + p_2^2 + q_2^2}} \quad (22)$$

From the given images, we have the following image irradiance relation for each point on the object

$$I = \rho \frac{1 - pl - qk}{\sqrt{p^2 + q^2 + 1} \sqrt{l^2 + k^2 + 1}} \quad (23)$$

Substituting for ρ_1 and ρ_2 from the image irradiance equations for the two objects in (22)

$$I_1 \frac{\sqrt{1 + l_1^2 + k_1^2}}{1 - p_1 l_1 - q_1 k_1} p_1 = I_2 \frac{\sqrt{1 + l_2^2 + k_2^2}}{1 - p_2 l_2 - q_2 k_2} p_2 \quad (24)$$

For each image, symmetry provides a linear constraint of the form (7) which has to be satisfied by the true surface gradients $\{p_1, q_1\}$, i.e.,

$$S_1 l_1 p_1 + D_1 k_1 q_1 = D_1 \quad (25)$$

For pixels with $D_1 \neq 0$,

$$\frac{S_1}{D_1} l_1 p_1 + k_1 q_1 = 1 \quad (26)$$

From (26) and (24), the condition for the corresponding points of the two objects to appear similar becomes

$$I_1 \frac{\sqrt{1 + l_1^2 + k_1^2}}{l_1 \left(\frac{S_1}{D_1} - 1\right)} = I_2 \frac{\sqrt{1 + l_2^2 + k_2^2}}{l_2 \left(\frac{S_2}{D_2} - 1\right)} \quad (27)$$

Interestingly, the condition in (27) involves only light source directions and image intensities. Thus, given two images, one can use this simple condition for each corresponding pixel to decide whether they come from the same object or not. If the symmetry and Lambertian assumptions are strictly adhered to, the matching decision is provably correct up to the ambiguity in Corollary 4.4. As the condition in (27) does not involve any unstable estimation of surface gradients or albedo, the algorithm degrades gracefully with deviations from the assumptions.

The two sides of the condition in (27) can be treated separately as the illumination-invariant representation of the respective objects as follows

$$I_{1r} = I_1 \frac{\sqrt{1 + l_1^2 + k_1^2}}{l_1 \left(\frac{S_1}{D_1} - 1\right)} \quad (28)$$

$$I_{2r} = I_2 \frac{\sqrt{1 + l_2^2 + k_2^2}}{l_2 \left(\frac{S_2}{D_2} - 1\right)} \quad (29)$$

Two images can be easily compared by generating these *virtually* relighted images.

6. Experiments

The real contribution of this work is the theoretical statement that unlike general objects, it is possible to distinguish between bilaterally symmetric Lambertian objects using just one image. In this section, we describe experiments performed on simulated and real data to evaluate the practical implications of the work.

6.1. Experiments on Simulated Data

First, we use simulated data to verify the correctness of the proposed theoretical result. We use the 3D face models used by Blanz and Vetter in their morphable model [3]. We generate several images of 100 subjects in the database under randomly selected illumination conditions. Here, the faces are made bilaterally symmetric and the images are generated using Lambertian reflectance. As the assumptions made in the theoretical formulation are strictly adhered to, the matching algorithm does not make any error.

6.2. Experiments on Real Data

We also test the performance of the algorithm on PIE dataset [15]. The PIE dataset has 68 subjects with images of each subject in 21 different illumination conditions. The images show deviations from Lambertian and symmetry assumptions. Moreover, the light source direction needs to be estimated which involves some error. Figure 2 shows the *virtually* relighted images obtained from different images of a subject in the dataset. The light source direction in an image is estimated using a simple algorithm recently proposed by Lee and Moghaddam [11]. The relighted images look like flattened frontally illuminated images. As desired, the illumination effects in the original images mostly disappear in the relighted images.

Though the relighted images are not perfect (as the assumptions are not strictly held), they seem promising to be used for matching images across illumination variations. We perform a face recognition experiment using the PIE dataset. A set of commonly used challenging illumination conditions from the PIE dataset are chosen to test our simple relighting based scheme (see Figure 3). In this setting, all images in one illumination scenario are used to form the gallery and another one to form the probe set. Thus, both the gallery and the probe set have one image per subject. The recognition experiment is repeated for all combinations of gallery and probe sets. Similarity between a gallery and a probe image is measured using a simple cross correlation between the corresponding relighted images as follows. Suppose f_g and f_p are two vectorized relighted images, then the similarity of the images is given by

$$S(g, p) = \frac{\langle f_g, f_p \rangle}{|f_g||f_p|} \quad (30)$$

where $\langle f_g, f_p \rangle$ denotes the scalar product of the two vectors. This is a very simple measure and fits well with the goal of stress testing the practical usefulness of the theoretical results. Table 1 shows the recognition results obtained in the experiment. The proposed approach using the relighted images works quite well even with such a simple distance measure. Unlike most face recognition methods, we do not make use of any face-based statistics (like Eigenfaces, 3D

morphable models, etc.). Recognition performance using the intensity images directly is also shown for comparison. Intensity images are normalized before computing the similarity. For most gallery-probe scenarios, relighted images perform better than the normalized intensity images. The improvement is quite significant when the illumination conditions for the gallery and probe scenarios are very different.

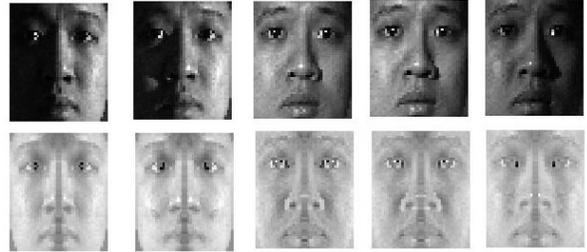


Figure 2. Virtually relighted image examples using images from the PIE dataset.

7. Summary and Discussion

We showed that two bilaterally symmetric objects can almost always be distinguished using just one image per object taken under different illumination conditions. The condition under which they cannot be distinguished, partitions the set of symmetric Lambertian objects into equivalence classes. It is difficult for two objects to satisfy the condition in practice leading to the conclusion that bilaterally symmetric objects are hardly ambiguous.

Based on the theoretical formulation, we proposed a virtual relighting algorithm to recognize real objects that do not strictly satisfy the assumptions made. The algorithm is provably correct for symmetric Lambertian objects up to the ambiguity described in Theorem 4.2. The relighted images obtained on real images seem to be free of any illumination effects. Face recognition experiments using the relighted images showed excellent performance without using any sophisticated classifier or class-based statistics.

There exist a few specific cases where symmetric SFS analysis may not be applied. Shadow pixels do not reveal much information about the surface gradients and have to be excluded from the formulation. Moreover, if $l = 0$ or $p = 0$, two symmetric points have same image intensity, thereby providing no additional information due to symmetry.

References

- [1] R. Basri and D. Jacobs. Photometric stereo with general, unknown lighting. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 374–381, 2001.
- [2] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25:218–233, 2003.

Table 1. Recognition results on the PIE dataset. f_i denotes images taken with i^{th} flash ON as labeled in the PIE dataset. Each $(i, j)^{th}$ entry in the table shows the recognition rate obtained with the images from f_i as gallery and from f_j as probes. The first number is the rank-1 recognition performance using the relighted images while the second number is the performance using the intensity images directly.

Probe	f_{09}	f_{12}	f_{13}	f_{14}	f_{15}	f_{16}	f_{17}	f_{21}	f_{22}
Gallery									
f_{09}	-/-	99/99	97/97	97/94	75/63	60/44	56/34	99/99	85/84
f_{12}	99/99	-/-	99/99	100/99	81/74	62/46	59/31	100/100	85/96
f_{13}	99/94	100/97	-/-	100/100	100/100	94/78	81/54	100/100	100/100
f_{14}	99/91	100/97	100/100	-/-	99/100	94/79	82/59	100/100	100/100
f_{15}	94/35	100/49	100/100	100/100	-/-	99/100	99/96	100/68	100/100
f_{16}	97/38	100/49	100/94	100/96	100/100	-/-	100/100	100/65	100/99
f_{17}	85/37	91/44	97/63	99/71	100/100	100/100	-	94/50	100/90
f_{21}	99/99	100/100	100/100	100/100	87/79	76/51	69/44	-/-	100/97
f_{22}	97/54	100/81	100/100	100/100	100/100	97/96	97/72	100/96	-/-

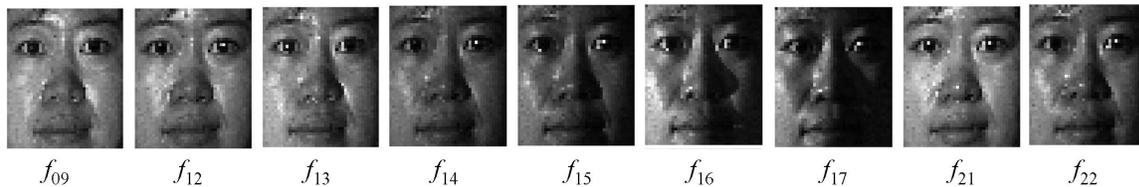


Figure 3. Illumination conditions from the PIE dataset used in the face recognition experiment.

- [3] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(9):1063–1074, Sep 2003.
- [4] M. J. Brooks, W. Chojnacki, and R. Kozera. Impossible and ambiguous shading patterns. *International Journal of Computer Vision*, 7(2):119–126, 1992.
- [5] H. Chen, P. Belhumeur, and D. Jacobs. In search of illumination invariants. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 254–261, 2000.
- [6] R. Dovgand and R. Basri. Statistical symmetric shape from shading for 3d structure recovery of faces. In *European Conf. on Computer Vision*, 2004.
- [7] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading problem. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 10:439–451, July 1988.
- [8] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(6):643–660, June 2001.
- [9] B. K. P. Horn. Determining lightness from an image. *Computer Graphics and Image Processing*, 3(4):277–299, 1974.
- [10] D. Jacobs, P. Belhumeur, and R. Basri. Comparing images under variable illumination. In *IEEE Conf. on Computer Vision and Pattern Recognition*, Santa Barbara, CA, pages 610–617, 1998.
- [11] K. C. Lee and B. Moghaddam. A practical face relighting method for directional lighting normalization. In *International Workshop on Analysis and Modeling of Faces and Gestures*, 2005.
- [12] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: determining the illumination from images of convex Lambertian object. *Jl. of the Optical Society of America A*, pages 2448–2459, Oct 2001.
- [13] A. Shashua. On photometric issues in 3d visual recognition from a single 2d image. *International Journal of Computer Vision*, 21:99–122, 1997.
- [14] A. Shashua and T. R. Raviv. The quotient image: Class-based re-rendering and recognition with varying illuminations. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 23(2):129–139, Feb 2001.
- [15] T. Sim, S. Baker, and M. Bsat. The CMU pose, illumination, and expression database. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(12):1615–1618, Dec 2003.
- [16] A. L. Yuille, D. Snow, R. Epstein, and P. N. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using svd and integrability. *International Journal of Computer Vision*, 35(3):203–222, 1999.
- [17] W. Zhao and R. Chellappa. Illumination-insensitive face recognition using symmetric shape-from-shading. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 1286–1293, 2000.
- [18] W. Zhao and R. Chellappa. Symmetric shape from shading using self-ratio image. *International Journal of Computer Vision*, 45(1):55–75, Oct 2001.
- [19] S. Zhou, R. Chellappa, and D. Jacobs. Characterization of human faces under illumination variations using rank, integrability, and symmetry constraints. In *European Conf. on Computer Vision*, 2004.