Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE @ DIRECT®

ELSEVIER

Computer Vision and Image Understanding 93 (2004) 245–259

Computer Vision and Image Understanding

www.elsevier.com/locate/cviu

Efficient detection under varying illumination conditions and image plane rotations ☆

Margarita Osadchy and Daniel Keren*

Department of Computer Science, University of Haifa, Haifa 31905, Israel

Received 6 September 2001; accepted 7 October 2003

Abstract

This paper focuses on the detection of objects with a Lambertian surface under varying illumination and pose. We offer to apply a novel detection method that proceeds by modeling the different illuminations from a small number of images in a training set; this automatically voids the illumination effects, allowing fast illumination invariant detection, without having to create a large training set. It is demonstrated that the method "fits in" nicely with previous work about modeling the set of object appearances under varying illumination. In the experiments, an object was correctly detected under image plane rotations in a 45° range, and a wide variety of different illuminations, even when significant shadows were present.

Keywords: Detection; Recognition; Varying illumination; Varying pose

1. Introduction

Slight changes in pose and illumination produce large changes in object appearance. Recognition of objects under various classes of geometric transformations or under varying viewpoints was previously studied in [9,16,19,20]. However, these methods offer no solution for the problem of illumination variability, which has a

☆ An earlier version of this work was presented in the ICCV-2001 conference.

* Corresponding author. +972-4-824-9331.

E-mail addresses: rita@nec-labs.com (M. Osadchy), dkeren@cs.haifa.ac.il (D. Keren).

1077-3142/$ - see front matter
doi:10.1016/j.cviu.2003.10.001

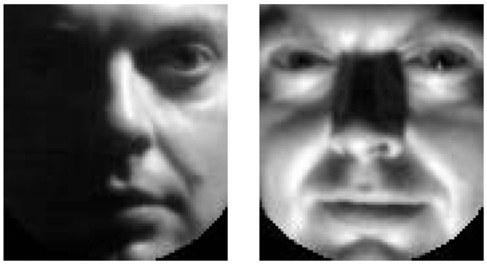

Fig. 1. Variability in appearance due to differences in illumination (the images are from Harvard face database).

very strong effect on the appearance of an object. Fig. 1 shows two images of a person with the same facial expression and photographed from the same viewpoint. There is a significant variability in these two images due to differences in illumination. In fact, it has been observed [12] that in face recognition, the variability due to illumination is often greater than that due to a change in the person's identity. On the other hand, changes in viewpoint also have a dramatic effect on the object appearance. Variation in both illumination and pose results in a complex image set with a very high linear dimension.

In this paper we use the observations from [3] and the *anti-face* method [9] to detect 3D objects under varying illumination and pose. The anti-face method offers an attractive solution, which proceeds by modeling the effects of different illumination conditions in the training set; this automatically voids the illumination effects, allowing fast illumination invariant detection, without having to create a large training set.

The following applications are presented:

1. Detection of an object with a Lambertian surface under varying pose and illumination, without shadows.
2. Detection of an object with a Lambertian surface under varying pose and illumination, with attached shadows.

In the first case, the detection was successful for a rather large class of different poses (360° rotation). In the second case, the existence of shadows results in a considerably more complicated image collection, however the algorithm was still able to correctly detect objects under a 45° rotation range. These results compare favorably to previous work, in which detection over a wide range of pose variation was achieved by individually applying illumination cones to 4×4 degree patches [5].

Theoretically, the algorithm requires that the object be convex, in order to exclude cast shadows. However, good results were obtained for a non-convex object, when significant cast shadows were not present.

1.1. Structure of the paper

Section 1.2 surveys the related work on illumination variability. Section 2 focuses on applying the anti-face method to the *illumination space* and *illumination cone*, and presents the algorithms mentioned above. In Section 3, a Taylor series approximation of rotated images is used to decrease the size of the training set, and Section 4 presents the experimental results.

1.2. Previous work

Appearance-based methods can recognize an object under a particular pose and lighting, if it has been previously seen under similar circumstances: see for example [14]. To extend these methods to handle illumination variability, a large set of images of the object under varying illumination should be used for the learning stage, which is inefficient [13]. Hence it is very popular to represent the set of images that an object produces under varying illumination using a low-dimensional linear subspace. In particular, the image space of a 3D Lambertian surface under changing illumination, without attached shadows, is spanned by a basis of three images [2,7,11,15,18,21,23]. Koenderink and Van Doorn [10] extended these results to allow an ambient component, resulting in a 4D space. The ambient light problem was also considered in [22]. Belhumeur and Kriegman [3] demonstrated that all object appearances produced by illumination changes including attached shadows (without cast shadows) are described by a convex cone which is represented by three images. This representation was used by Georghiades et al. [6] for object recognition and then extended to variation in pose [5]. In this method each "cone" models a 4×4 degree patch of the visibility sphere, hence recognition under a large variability in pose is accomplished by calculating the distance to each cone, which is more computationally expensive than our approach. Another attempt to find a low-dimensional representation of the image space that a Lambertian object can produce due to illumination was proposed by Basri and Jacobs [1] and Ramamoorthi and Hanrahan [17]. They show analytically that illumination variability for a Lambertian object can be very closely approximated by a 9D linear subspace. This result has been used in face recognition. Jacobs et al. [8] proposed a simple local measure of image comparison based on the gradient of image ratio. This measure performed well for face recognition under varying illumination. It is important to mention that this method does not require a training set; it uses only a single image. Chen et al. [4] extended this work by using image gradient distribution for developing illumination insensitive measure of image comparison. This new measure produced better results on the same face database.

2. Illumination invariant detection

In this section we show that unlike other learning techniques, the anti-face method [9] requires only a small number of training images in order to recognize an object under different lighting conditions, and it offers a very fast detection algorithm.

2.1. Anti-faces short overview

Anti-faces [9] is a novel detection method, which works well in case of a rich image collection—for instance, frontal face under a large class of linear transformations, or 3D objects under different viewpoints. Call the collection of images, which should be detected, a *multi-template.* The detection problem is solved by sequentially applying very simple filters (or detectors), which act as inner products with a given image (viewed as a vector) and satisfy the following conditions:

1. The absolute values of their inner product with multi-template images are small.
2. They are smooth, which results in the absolute values of their inner product with "random images" being large; this is the characteristic which enables the detectors to separate the multi-template from random images.
3. They act in an independent manner, which implies that their false alarms are not correlated; hence, the false alarm rate decreases exponentially in the number of detectors.

The detection process is very simple: the image is classified as a member of the multi-template iff the absolute value of its inner product with each detector is smaller than some (detector specific) threshold. Only images which passed the threshold test imposed by the first detector are examined by the second detector, etc. This, in turn, leads to a very fast detection algorithm. Typically, $(1 + \delta)N$ operations are required to classify an N-pixel image, where $\delta < 0.5$.

The anti-face method classifies an image as belonging to the sought image collection (or the multi-template) iff its inner products with certain detectors are small. Hence, if this collection can be described by linear combinations with small coefficients of a small number of basis images, the anti-face method can be trained only on the basis elements, resulting in a very efficient algorithm. This makes it a natural candidate to use for detection under varying illumination.

2.2. The illumination model

The following observations [3,7,15,21] allow to model object appearance under a wide range of illuminations, instead of physically creating them. The following discussion draws from [3].

Consider a convex object with a Lambertian reflectance function, which is illuminated by a single point light source at infinity. Let $B \in R^{n\times 3}$ be a matrix where each row is the product of the albedo with the inward pointing unit normal for a point on the surface corresponding to a particular pixel in the image viewed as a vector of size n. Let $s \in R^3$ denote the product of the light source intensity with the unit vector in the direction of the light source. The resulting image $x \in R^n$ is then given by

$$x = \max(B \cdot s, 0). \tag{1}$$

The pixels set to zero correspond to surface points lying in an *attached shadow.* Convexity of the object is assumed to avoid *cast shadows.* When no part of the object is shadowed, x lies in a 3D subspace L, called the *illumination space*, given by the span of the matrix B:

$$L = \{x | x = B \cdot s \ \forall s \in R^3\}. \tag{2}$$

Hence the illumination subspace can be constructed from just three basis images [7,15,21]. It was shown in [3] that the set C of all possible images of a convex

Lambertian surface, created by varying the direction and strength of an arbitrary number of point light sources at infinity, can be represented as follows:

$$C = \left\{ x | x = \sum_{i=1}^{k} \max(B \cdot s_i, 0) \ \forall s_i \in R^3, \ \forall k \in Z^+ \right\} \tag{3}$$

and C is a convex cone in R^n. Furthermore, it was shown in [3] that any image in the cone C can be represented as a convex combination of *extreme rays* given by

$$x_{ij} = \max(B \cdot s_{ij}, 0), \tag{4}$$

where

$$s_{ij} = b_i \times b_j, \quad i \neq j, \tag{5}$$

where b_i and b_j range over the rows of B. It was proved in [3] that the number of shadowing configurations is at most $m(m-1)+2$, where $m \leqslant n$ is the number of distinct normals, hence there are at most $m(m-1)$ extreme rays. Since there is a finite number of extreme rays, the cone is polyhedral.

The illumination subspace method [6] offers a way to construct the illumination cone. Gather three or more images of the object (with a fixed pose) under varying illumination without shadowing, and use these images to estimate the 3D illumination subspace L by normalizing the images to unit length, and then using singular value decomposition (SVD) to estimate the optimal 3D orthogonal basis B^* in a least square sense. It was proved in [3] that B^* is sufficient for determining the subspace L. Then from B^*. the extreme rays defining the illumination cone C can be computed using Eqs. (4) and (5).

2.3. Application of the anti-face method to illumination invariant detection

To extend anti-faces to handle illumination variability we should find a small number of "basis images" and corresponding *smooth detectors* [9] such that:

(A) After normalization, the different object appearances can all be represented as linear combinations of the basis images, with small combination coefficients.

(B) The detectors have small inner products with the basis images. Because of (A), they will also have small inner products with all the object appearances. This will be formalized in Proposition 2.1.

The following observations [3] support condition (A). Consider a convex object with a Lambertian reflectance function.

- When no part of the object is shadowed, its image lies in the 3D subspace L given by the span of the matrix B; L can be constructed from three basis images.
- The set of images under an arbitrary number of point light sources at infinity is a convex polyhedral cone in R^n, which can be expressed as a convex combination of extreme rays.

In order to satisfy these conditions, let us first analyze the positive set of the anti-face detector (that is, the set of images accepted by the detector).

Assume that the training set consists of orthonormal vectors. This assumption is feasible, because we can replace the original training set with the vectors produced by the following process:
1. Perform singular value decomposition (SVD) on the original training set.
2. Take the eigenvectors that correspond to the eigenvalues with 99% of the energy.

This process also speeds up the computation of the anti-face detectors, as it reduces the size of the training set.

Proposition 2.1. *Let $\{v_i\}_{i=1}^{k}$ be the orthonormal basis produced by SVD from the normalized training set. Let d be an anti-face detector, such that $|(d, v_i)| \leqslant \epsilon_i$ $\forall i = 1, \ldots, k$. Then for each $v = \sum_{i=1}^{k} \alpha_i v_i$, which satisfies $||v|| = 1$, $|(d, v)| \leqslant \sqrt{\sum_{i=1}^{k} \epsilon_i^2}$.*

Proof.

$$1 = |v|^2 = \sum_{i=1}^{k} \sum_{j=1}^{k} \alpha_i \alpha_j v_i^t v_j = \sum_{i=1}^{k} \alpha_i^2$$

and

$$|(d, v)| = \left| \sum_{i=1}^{k} \alpha_i (d, v_i) \right| \leqslant \sum_{i=1}^{k} |\alpha_i (d, v_i)| = \sum_{i=1}^{k} |\alpha_i| |(d, v_i)| \leqslant \sum_{i=1}^{k} |\alpha_i| \epsilon_i.$$

From the Cauchy–Schwarz inequality it follows that

$$\sum_{i=1}^{k} |\alpha_i| \epsilon_i \leqslant \sqrt{\sum_{i=1}^{k} \alpha_i^2} \sqrt{\sum_{i=1}^{k} \epsilon_i^2} = \sqrt{\sum_{i=1}^{k} \epsilon_i^2} \qquad \square$$

From Proposition 2.1 it follows that if the three basis images for the illumination subspace are used as a training set for the detector, it will detect the entire illumination subspace if the threshold is properly chosen.

As was previously mentioned, the illumination cone can be represented by linear combinations of the vectors x_{ij} (Eq. (4)) with non-negative coefficients. In practice the extreme rays of the illumination cone lie near a low-dimensional linear subspace. This observation was theoretically justified by Basri and Jacobs [1] and Ramamoorthi and Hanrahan [17] for convex objects. Thus from the last observation and Proposition 2.1, it follow that if the detector is trained on the basis vectors of the low-dimensional subspace that approximates the illumination cone, it will detect the illumination cone, if the threshold is correctly chosen.

2.4. Detection under varying pose and illumination (without shadows)

We showed in the previous section that if we want to detect an object under fixed pose using anti-faces, we should train the detector on the three basis images for the illumination subspace, and this will allow to detect all images that lie in this subspace. This method can be easily extended to different poses, by training the detector on a linear subspace that contains the basis images for illumination subspaces that correspond to all training poses. The following pseudo code describes the algorithm for detection of a convex object under varying illumination and pose, when no part of the object is shadowed.

1. Find the three basis images for illumination subspace for every sample of object positions:
   (i) Gather three or more images of the object under varying illumination without shadowing.
   (ii) Normalize the images to unit length, apply SVD, and take the three eigenvectors that correspond to the largest eigenvalues.

(Steps 1 will produce 3M images where M is a number of training poses).

2. Replace the training set. containing 3M images produced in the previous step by the eigenvectors $\{v_i\}_{i=1}^{k}$ that correspond to the eigenvalues which capture 99% of the energy. (Obviously k depends on the dimension of this linear subspace.)
3. Find anti-face detectors using the new training set.
4. For each detector d, choose the threshold as $\sqrt{\sum_{i=1}^{k} \epsilon_i^2}$, where $|(d, v_i)| = \epsilon_i$, $i = 1, \ldots, k$.

From Proposition 2.1, it follows that the positive set of such a detector includes the entire illumination space for all object positions on which the detector was trained.

2.5. Detection under varying pose and illumination (allowing shadows)

A similar idea can be used for detection of illumination cones for various poses. Here we find the extreme rays that form the illumination cone at every pose and then train the detector on the linear subspace that contains the union of the illumination cones for all training poses. The following pseudo-code describes an algorithm for detection of a convex object under varying pose and an arbitrary number of point light sources at infinity. Attached shadows are allowed.

1. Find the illumination cone for every sample of object positions:
   (i) Gather three or more images of the object under varying illumination without shadowing.
   (ii) Normalize the images to unit length, and use SVD to estimate the best 3D orthogonal basis B^* in the least square sense.
   (iii) From B^* compute the vectors x_{ij} using Eqs. (4) and (5).
2. Apply SVD to the collection of vectors x_{ij} for all object positions in order to find the eigenvectors $\{v_i\}_{i=1}^{k}$ that correspond to the eigenvalues which capture 99% of the energy. ($\{v_i\}_{i=1}^{k}$ is the basis of the linear subspace that contains the union of illumination cones for all poses).
3. Find anti-face detectors using $\{v_i\}_{i=1}^{k}$ as training set.
4. For each detector d choose the threshold as $\sqrt{\sum_{i=1}^{k} \epsilon_i^2}$, where $|(d, v)| = \epsilon_i$, $i = 1, \ldots, k$.

From Proposition 2.1, it follows that the positive set of the detectors approximates the illumination cones for all object positions. As mentioned in Section

2.1, the number of extreme rays is $m(m-1)$ where $m \leqslant n$ is the number of distinct normals, which is usually large, hence the number of extreme rays needed for construction of the illumination cone can be very large. Therefore, we use the sampling method from [6], that approximates the cone by directly sampling the space of light source directions rather than generating the samples through Eqs. (4) and (5).

3. Incremental pose approximation

The anti-face method, like most other detection and recognition techniques, requires that the multi-template be closely sampled. We showed that for illumination variability this limitation can be overcome by representing the image set using a small number of basis images. For pose variation there is no such representation, however for a small range of image plane rotations (about five degrees) object appearances can be estimated using the Taylor expansion. Let $I(x,y)$ be an image; then the rotated image is a function of x, y, and θ:

$$\tilde{I}(x,y,\theta) = I(x\cos(\theta) - y\sin(\theta), y\cos(\theta) + x\sin(\theta))$$

and for a small θ

$$\tilde{I}(x,y,\theta) \cong \tilde{I}(x,y,0) + \left.\frac{\partial \tilde{I}}{\partial \theta}\right|_{x,y,0} \cdot \theta,$$

where

$$\begin{aligned}\frac{\partial \tilde{I}}{\partial \theta} &= I_x(-x\sin(\theta) - y\cos(\theta))|_{x,y,0} + I_y(-y\sin(\theta) + x\cos(\theta))|_{x,y,0} \\ &= I_x|_{x,y} \cdot x - I_y|_{x,y} \cdot y.\end{aligned}$$

In summary, an image I rotated by a small angle θ can be approximated by

$$I(x,y,\theta) = I(x,y) + (I_y|_{x,y} \cdot x - I_x|_{x,y} \cdot y) \cdot \theta. \qquad (6)$$

Hence the anti-face detectors should be trained on I and $I_T \equiv I_y x - I_x y$. This ensures that the detectors will yield small results on I rotated in the image plane by small angles.

The proposed method can be incorporated into the algorithms described in Sections 2.3 and 2.4 in the case of image plane rotations. Instead of creating the extreme rays for each angle, the five-degree range can be covered by Eq. (6).

4. Experimental results

We have experimented with the algorithms presented above. We have chosen image plane rotations for training and testing the algorithms described in Sections 2.3 and 2.4.

4.1. Experiments

Ten images of a toy tiger were captured under varying illuminations without shadowing (Fig. 2A). The object was illuminated by a single light source, but due to diffusion from the surrounding, ambient light is present in all images. To exclude the ambient component we photographed the object under ambient light only (Fig. 2B) and subtracted this image from the 10 images depicted in Fig. 2A. Using the algorithm from Section 2, we found the three basis images that span the illumination subspace L (Fig. 3).

Fig. 4 presents the results of the detection algorithm, under arbitrary rotations and various illuminations without shadowing (Section 2.3). The detectors were trained on 49 basis images that span the linear subspace of rotations and illuminations without shadowing. Ten detectors were sufficient to recover the toy without false alarms. The anti-face method that was trained on the image of the toy subject to arbitrary rotations and illuminated by ambient light alone failed to detect the object in the scenes depicted in Fig. 4.

The following experiment was designed to test the algorithm for detecting an object under fixed pose, illuminated by an arbitrary number of point light sources at infinity. Attached shadows were allowed. We took the same basis images (Fig. 3) as before, and used the sample method [6] to approximate the cone. It was empirically shown in [3] that the cone is flat (i.e., its elements lie near a low-dimensional linear subspace), and that the subsampled cone provides an approximation that results in good recognition performance. In our experiment we created about 60 images, such that the corresponding light source directions were more or less uniformly distributed on the illumination sphere. Fig. 5 demonstrates the results

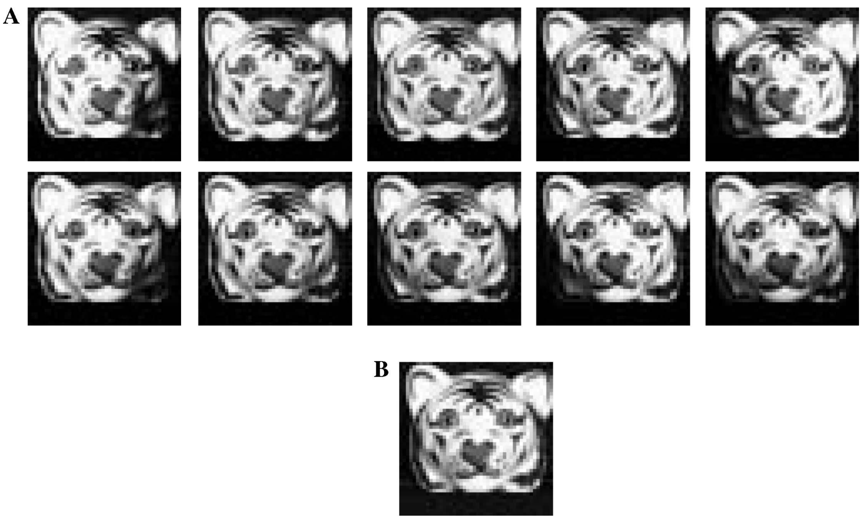

Fig. 2. Initial images for estimation of illumination space. (A) Images illuminated by a single light source and ambient light; (B) image illuminated by ambient light only.

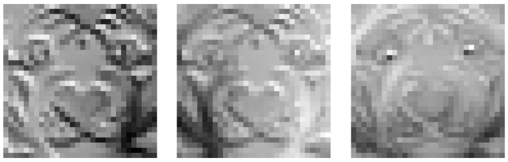

Fig. 3. Basis images (columns of the matrix) that span the illumination subspace L for the toy tiger.

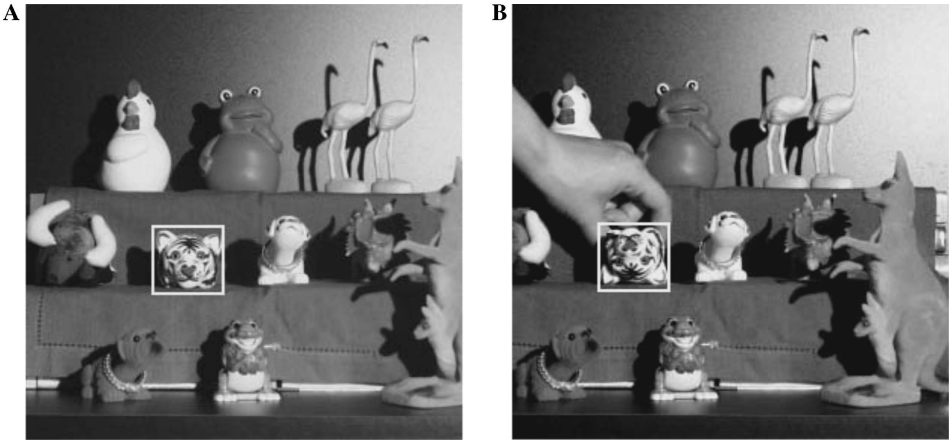

Fig. 4. Results of detection of the toy tiger, subject to image plane rotations and various illuminations without shadowing. The scene was illuminated by point light source and ambient light.

of the detection of the toy tiger in real images under various illuminations. Eight to ten anti-face detectors were used to detect all the instances of the tiger with no false alarms. The detectors were trained on a 16D linear subspace that approximates the cone for the tiger toy.

Since it is very difficult to simulate the light conditions that result in images with significant attached shadows, we tested the algorithm on 200 random samples from the illumination cone of the tiger with one and two light sources. The images were artificially created using the method described in [3]. All 200 samples were recognized as the tiger. Fig. 6 shows some of the images from this test set.

The last experiment was designed to test the algorithm for detecting an object under varying illumination with attached shadows and subject to image plane rotations at a 45° range (Section 2.4). We created the extreme rays that approximate the cone for each rotation angle in the manner described in the previous experiment. Eight sets of anti-face detectors were created, each for a 45° range, thus covering 360°. The images in Fig. 7 depicts the tiger rotated by 180° with different light source directions. Fig. 8 depicts the images rotated by 60° and 100° correspondingly. In these tests, 10 anti-face detectors sufficed to detect the tiger without false alarms. The detectors were trained on 26 basis images that span the linear subspace that approximates the object appearance under both illumination changes (with attached shadows) and plane rotations in a 45° range.

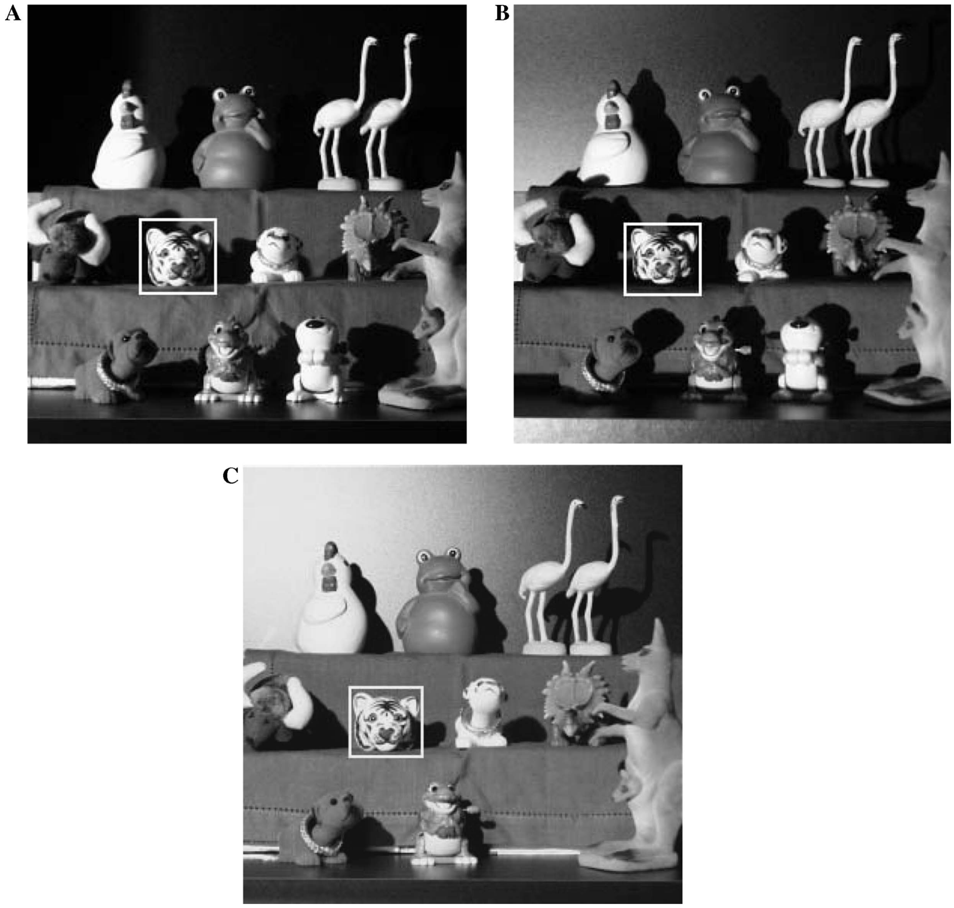

Fig. 5. Detection results in real images; (A) and (B), one light source and ambient light; (C), two light sources and ambient light.

4.2. Detection performance as function of multi-template's structure

During the experiments we observed that the detection performance of the algorithm described in Section 2.4 drops when the range of rotations increases. For instance, using 10 detectors trained on 360° range produced 449 false alarms in a 253×253 pixel image (same image as Fig. 7A). The result can be interpreted as follows. Define the *effective dimension*, as the number of eigenvalues required for 90% of the energy; it is a measure of the image set complexity (by "complexity," we mean the detection complexity, i.e., an empirical measure of how difficult it is to detect the multi-template's images). The effective dimension of the multi-template formed by the illumination cones for various rotations nearly equals the *product* of the effective dimensions of the rotation set and the effective dimension of the illumination cone. This observation is demonstrated by comparing the effective dimensions of all these

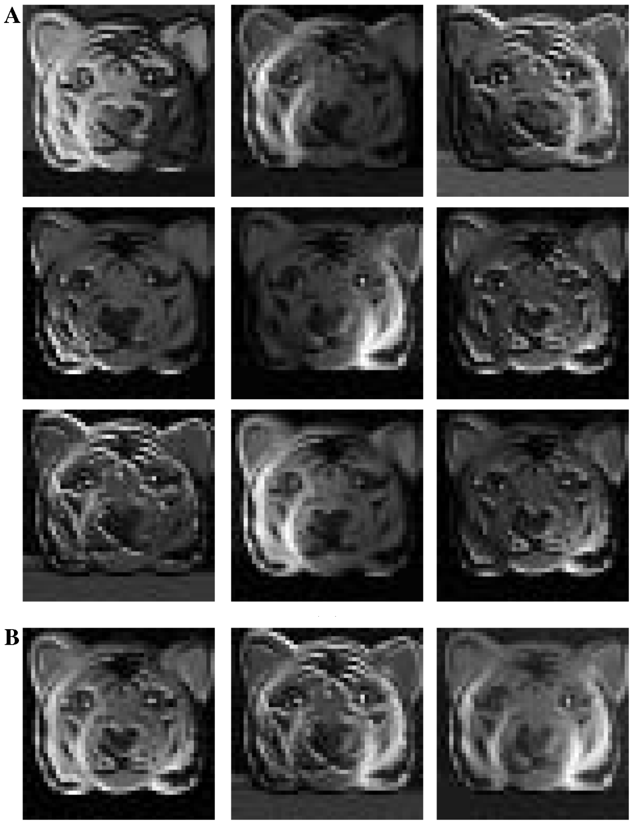

Fig. 6. Random samples from the illumination cone of the tiger: (A) with one and (B) with two light sources. All were correctly detected.

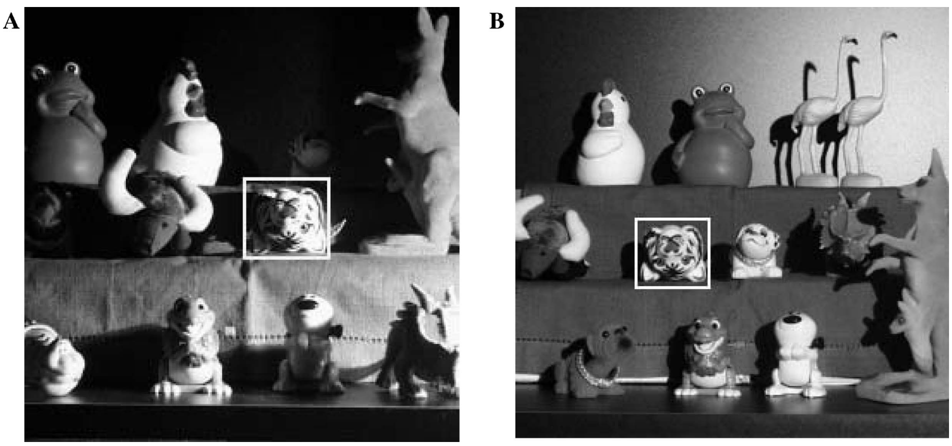

Fig. 7. Detection results for toy tiger rotated by 180° and illuminated by a single point light source and ambient light.

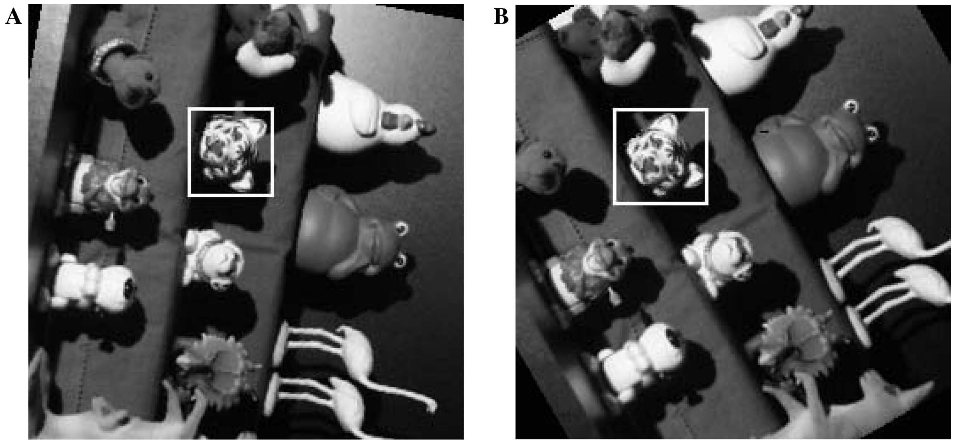

Fig. 8. Detection results for toy tiger under rotation and illumination by a single point light source and ambient light. (A) Image rotated by about 100° degrees; (B) image rotated by about 60°.

Table 1
Approximate linear dimension of different sets with various rotation ranges

Rotation range	Effective dimension for combination of rotation and varying illumination	Effective dimension for rotation only (fixed illumination)
45°	26	3
90°	44	5
180°	65	8

sets. From Table 1 it is clear that for all rotation ranges the ratio between the multi-template for rotations plus illumination and for rotations only is between eight and nine, which is approximately the effective dimension of the illumination cone. Intuitively speaking, "rotations and illuminations don't mix well," and combining them results in a very complex image set—far more complicated than the cases in which only rotations or only illumination changes are allowed. The multiplicative relation between the sets does not change if we vary the effective dimension measure from 90% of energy to 98% with 2% step. Table 2 shows the results for 45° rotation range.

Another interesting observation about the structure of a multi-template is that its complexity depends on the combination of two characteristics: (1) its effective dimension and (2) the smoothness of the principal component (i.e., the vector with the largest eigenvalue in the SVD) of the multi-template in R^n. If the principal direction of

Table 2
Multiplicative relationship of effective dimensions of the sets for different energy thresholds

Effective dimension for combination of rotations and varying illumination	26	36	46	62	93
Varying illumination, fixed pose	8	9	10	12	16
Varying rotations, fixed illumination	3	4	4	5	6

The thresholds start at 90% and increase by 2% between consecutive columns.

Table 3
Estimation of the complexity of various multi-templates vs. number of anti-face detectors required for accurate detection

	Illumination cone for tiger image under fixed pose	Illumination cones for tiger image rotated in 45° range	Illumination cones for tiger image rotated in 360° range	Tiger image rotated in 360° range (fixed illumination)
Roughness of the principal eigenvector	744	581	4	855
Number of detectors	8	10	Impossible to detect	4
Effective dimension	8	26	131	16

the set is smooth (i.e., smooth when viewed as an image), then the anti-face detector that should be orthogonal to the multi-template, is also orthogonal to many other natural smooth images. Thus, both the linear dimension and the principal direction of the multi-template allow to predict the difficulty of the detection problem. This observation can be verified by studying the roughness of the principal eigenvector of the training set for various multi-templates. Table 3 summarizes the results. The roughness is defined as $\sum_{ij}(i^2 + j^2)d_{ij}^2$, where d_{ij} are the coefficients of the DCT of the principal eigenvector, as in [9].

5. Conclusions

In this paper, we have presented a novel algorithm for detection of objects under variable illumination and plane rotations, which accounts for attached shadows. The key element of our approach was to include the effects of different illumination conditions that can be modeled from a small set of images in the training set of the anti-face detectors; this automatically cancels the illumination effects, allowing fast illumination invariant detection. The method was successfully applied to detect an object under variable illumination and rotations in real images with complicated background and simulated images with significant attached shadows.

We have shown empirically that the linear dimension of the multi-template formed by illumination cones for arbitrary rotations, is roughly equal to the product of the linear dimension of the rotation set and the linear dimension of the illumination cone for a fixed pose. We have also shown that the complexity of any multi-template depends on a combination of its linear dimension and the principal direction of the set in R^n. In further research, we plan to test the algorithms on images containing a 3D object under various illuminations and other rotations.

Acknowledgments

The authors are grateful to the three anonymous reviewers for their helpful comments.

References

[1] R. Basri, D. Jacobs, Lambertian reflectance and linear subspaces, Proc. Int'l. Conf. on Comp. Vision vol. II (2001) 283–390.

[2] P.N. Belhumeur, P. Hespanha, D.J. Kriegman, Eigenfaces vs. Fisherfaces: recognition using class specific linear projection, IEEE Trans. Pattern Anal. Machine Intelligence 19 (7) (1997) 711–720.

[3] P.N. Belhumeur, D.J. Kriegman, What is the set of images of an object under all possible lighting conditions?, Internat. J. Comput. Vision 28 (3) (1998) 245–260.

[4] H.F. Chen, P.N. Belhumeur, D.W. Jacobs, In search of illumination invariants, IEEE Conf. Comput. Vision Pattern Recognition (2000) 254–261.

[5] A.S. Georghiades, P.N. Belhumeur, D.J. Kriegman, From few to many: generative models for recognition under variable pose and illumination, IEEE Trans. Pattern Anal. Machine Intelligence 23 (6) (2001) 643–660.

[6] A.S. Georghiades, D.J. Kriegman, P.N. Belhumeur, Illumination cones for recognition under variable lighting: faces, Proc. IEEE Conf. Comput. Vision Pattern Recognition (1998).

[7] H. Hayakawa, Photometric stereo under a light-source with arbitrary motion, JOSA-A 11 (11) (1994) 3079–3089.

[8] D.W. Jacobs, P.N. Belhumeur, R. Basri, Comparing images under variable illumination, Proc. IEEE Conf. Comp. Vision Pattern Recognition (1998) 610–617.

[9] D. Keren, M. Osadchy, C. Gotsman, Anti-faces: a novel, fast method for image detection, IEEE Trans. Pattern Anal. Machine Intelligence 23 (7) (2001) 747–761.

[10] J. Koenderink, A. Van Doorn, The generic billinear calibration-estimation problem, Internat. J. Comput. Vision 23 (3) (1997) 217–234.

[11] Y. Moses, Face recognition: generalization to novel images, Ph.D. Thesis, Weizmann Institute of Science, 1993.

[12] Y. Moses, Y. Adini, S. Ulman, Face recognition: the problem of compensating for changes in illumination direction, European Conf. Comput. Vision (1994) 296–298.

[13] H. Murase, S.K. Nayar, Visual learning and recognition of 3d objects from appearance, Internat. J. Comput. Vision 14 (1) (1995) 5–24.

[14] E. Osuna, R. Freund, F. Girosi, Training support vector machines: an application to face detection, Proc. IEEE Conf. Comput. Vision Pattern Recognition (1997) 130–136.

[15] S. Nayar, H. Murase, Dimensionality of illumination in appearance matching, IEEE Conf. Robotics Automat. (1996).

[16] M. Pontil, A. Verri, Support vector machines for 3d object recognition, IEEE Trans. Pattern Anal. Machine Intelligence 20 (6) (1998) 637–646.

[17] R. Ramamoorthi, P. Hanrahan, On the relationship between radiance and irradiance: determining the illumination from images of a convex Lambertian object, JOSA 18 (10) (2001) 2448–2459.

[18] T. Riklin-Ravin, A. Shashua, The quotient image: class based re-rendering and recognition with varying illuminations, Proc. IEEE Conf. Comput. Vision Pattern Recognition (1999) 566–571.

[19] D. Roobaert, M.M. Van Hulle, Viewbased 3d object recognition with support vector machines, IEEE Internat. Workshop Neural Networks Signal Process. (1999) 77–84.

[20] H.A. Rowley, S. Baluja, T. Kanade, Rotation invariant neural networkbased face detection, IEEE Conf. Comput. Vision Pattern Recognition (1998).

[21] A. Shashua, Geometry and photometry in 3D visual recognition, Ph.D. Thesis, MIT, 1992.

[22] A.L. Yuille, D. Snow, Shape and albedo from multiple images using integrability, IEEE Conf. Comput. Vision Pattern Recognition (1997).

[23] L. Zhao, Y.H. Yang, Theoretical analysis of illumination in PCA-based vision systems, Pattern Recognition 32 (1999) 547–564.