

Illumination Face Spaces are Idiosyncratic

Jen-Mei Chang¹, J.R. Beveridge², B.A. Draper², M. Kirby^{1,2}, H. Kley¹, C. Peterson¹ *
Colorado State University, Departments of Computer Science² and Mathematics¹, Fort Collins, CO 80523, USA.

Abstract

Illumination spaces capture how the appearances of human faces vary under changing illumination. This work models illumination spaces as points on a Grassmann manifold and uses distance measures on this manifold to show that every person in the CMU-PIE and Yale data sets has a unique and identifying illumination space. This suggests that variations under changes in illumination can be exploited for their discriminatory information. As an example, when face recognition is cast as matching sets of face images to sets of face images, subjects in the CMU-PIE and Yale databases can be recognized with 100% accuracy.

1. Introduction

Recent work has yielded surprising results about the low dimensionality of linear subspaces associated with face images acquired under varying illumination conditions [7], [3], [2]. Two face data sets, Yale [7] and CMU-PIE [17], are frequently used by researchers to explore how changing illumination alters the appearance of human subjects; they are the largest freely available data sets of face imagery to offer many (21 and 64, respectively) illuminations of each subject. A basic question regarding face images in general, and these data sets in particular, is whether illumination spaces are idiosyncratic. To be more precise, are the illumination spaces associated with individual subjects necessarily distinct, and if so, can they be empirically estimated with enough accuracy to exploit this property?

Answering this question requires two steps. First, we need a theoretical framework for defining what it means for two illumination spaces to be similar. In Section 4 we review Grassmann manifolds, and explain how illumination spaces may be viewed as points on a Grassmann manifold. This connection allows us to consider several standard measures of distance on the Grassmann manifold, and to measure how well they separate illumination

spaces associated with distinct people. The second step is to estimate illumination spaces from real data and validate the theory. Empirical results over the CMU-PIE and Yale data sets are presented in Section 5.

Specifically, we estimate two illumination subspaces for every subject in the Yale and CMU-PIE data sets. The subspaces for each person are estimated from randomly selected and non-intersecting sets of 8 or more images of the subject. For the 67 subjects in the CMU-PIE data set, this creates 67 pairs of matching subspaces and 4,422 pairs of non-matching subspaces. For the 10 subjects in the smaller Yale data set, there are 10 pairs of matching subspaces and 90 pairs of non-matching subspaces. In all of our experiments, pairs of subspaces associated with the same individual are always closer to each other than are pairs of subspaces associated with different individuals. We therefore view these data sets as *Grassmann separable*. Section 6 illustrates the potential of this approach to a harder problem: recognizing people from their illumination spaces when the images were collected under unbalanced sampling conditions. Finally, Section 7 summarizes our results.

2. Background: Illumination Spaces

There are two main schools of thought regarding how to craft algorithms that recognize faces under different illumination conditions. One view promotes the removal of illumination variations, for example, by computing an illumination invariant image. An example of such an approach is [16]. Histogram equalization may be viewed as an *ad hoc* attempt to accomplish the same goal. The other approach is to keep the illumination variations and model the underlying properties of the data set. The most widely used representation is the linear subspace approximation: although the ambient dimension of the image space is very large, the actual illumination data is close to lying on a low-dimension subspace. Some of the works that attempt to model the illumination variations among human face subjects are [7], [2], [12], and [8]. This paper is an extreme example of modeling illumination variations. Instead of trying to eliminate the effects of lighting,

*All authors and this work partially supported by NSF DMS and DCCF, grant MSPA-MCS 0434351. In addition, M. Kirby was partially supported by AFOSR contract FA9550-04-1-0094.

we claim that the way the appearance of a face changes under varying illumination is an identifying property of the person.

Belhumeur and Kriegman have shown that the set of m -pixel monochrome images of a convex, Lambertian object illuminated by an arbitrary number of distinct point light sources forms a convex polyhedral cone in \mathbb{R}^m (*the illumination cone*) [3]. They go on to show that the statement remains true for non-convex objects under relaxed lighting conditions. Unfortunately, for most objects the exact illumination cone is difficult to obtain, especially for non-convex human faces. However, experimental work by Belhumeur et al. [7], and Kriegman et al. [12] shows that this cone lies near a low-dimensional linear subspace in the space of all possible images.

Furthermore, Basri and Jacobs have demonstrated both theoretically and empirically that the set of images of a convex, Lambertian object seen under arbitrary distant light sources can be well approximated by a 9-dimensional linear subspace [2]. In this context, “well approximated” means that over 99% of the energy is captured in the subspace. Basri and Jacobs prove their result by representing reflectance functions as linear combinations of spherical harmonics and further relating the representation of images to the reflectance functions.

In addition, Ramamoorthi transforms the problem of linear approximation with spherical harmonics into linear approximation with principal components which can be shown to be identical to the spherical harmonic basis functions evaluated at the surface normal vectors under the same assumptions of Basri and Jacobs [15]. This prior work is ample evidence for the utility of modeling illumination variation with low-dimensional linear subspaces.

3. Comparing Sets of Images

Traditionally, almost all work on face recognition has focused on comparisons between single images. Recently, however, Experiment 2 of the Face Recognition Grand Challenge considered multi-still to multi-still comparisons [14]. Sets of four images were compared based on four by four (sixteen total) single image comparisons. This was enough to raise the recognition performance of the baseline algorithm from about 66% on single still images to about 88% on the four-way comparison. Works of [6], [18] with images selected from video sequences have promising results supporting the use of multi-still to multi-still image comparisons. As our own work here strongly suggests, future research on face recognition is likely to take much more seriously the question of how best to compare sets of images.

3.1. Estimating Illumination Spaces

The procedure for estimating the illumination subspace for a given person is relatively standard; we review

the procedure here. In all cases, the face imagery is geometrically normalized based upon known eye positions. In addition, the background area outside the face itself has been zeroed. Each image is unrolled into a vector $x_i^{(j)}$, which is the j th image of subject i . A data matrix X_i for subject i is then denoted by $X_i = [x_i^{(1)} | \dots | x_i^{(k)}]$. An orthogonal projection is applied to X_i to serve as a first-step dimensionality reduction.

An illumination subspace representation for the i th subject is then constructed from the k images of its data matrix X_i via *Singular Value Decomposition* (SVD). The q basis vectors for the i th subject’s q -dimensional illumination subspace are the strongest q left singular vectors in the SVD of X_i . In other words, the q -dimensional illumination subspace of X_i is given by the column space $\mathcal{R}(X_i)$ of its first q left singular vectors.

3.2. Reflection Augmentation

Generally, the amount of data available in face recognition data sets is extremely limited. Augmentation of the data set via inclusion of the mirror images effectively increases the available data [11]. Furthermore, this *symmetrization* of the data set imposes even and odd symmetry on the basis functions analogous to sinusoidal expansions. For sets of facial images under varying illumination conditions, reflection augmentation drastically improves the estimated linear representation by both increasing the effective sample set size and introducing novel illumination conditions. As a consequence, the approximation of illumination spaces can be improved without acquiring more data.

4. Subspace Distance Measures

The fundamental thesis of this paper is that the way in which images of an individual responds to varying illumination is characteristic of the individual. Phrased another way, every person has a unique (and therefore idiosyncratic) illumination cone. Formalizing this hypothesis requires that we have a means of measuring the distance between two subspaces. Section 4.1 below reviews the definition of the principal angles between two subspaces which are used in section 4.2 to give expressions for various measures of distance on Grassmann manifolds.

4.1. Principal Angles

A numerically stable algorithm for calculating principal angles between linear subspaces and references is provided in [9]. Briefly, if X and Y are two vector subspaces of \mathbb{R}^m such that $p = \dim(X) \geq \dim(Y) = q \geq 1$, then the *principal angles* $\theta_k \in [0, \frac{\pi}{2}]$, $1 \leq k \leq q$ between X and Y are defined recursively by

$$\cos(\theta_k) = \max_{u \in X} \max_{v \in Y} u^T v = u_k^T v_k$$

subject to $\|u\| = \|v\| = 1$, $u^T u_i = 0$ and $v^T v_i = 0$ for $i = 1 : k - 1$. Clearly, the principal angles satisfy $0 \leq \theta_1 \leq \theta_2 \leq \dots \leq \theta_q \leq \frac{\pi}{2}$. Henceforth, $\theta = (\theta_1, \dots, \theta_q)$ will denote the principal angle vector. Note that a standard way — which we call the *subspace distance*—to define distance between two subspaces using principal angles is (see [9])

$$d_{ss}(X, Y) = \max\{\sin \theta_i\} = \|\sin \theta\|_\infty. \quad (1)$$

Other authors have employed principal angles to compare sets of images to sets of images, see, e.g., [18], [19], but not from the perspective of different geometric embeddings of the Grassmannian.

4.2. Grassmannian Distance Measures

Recall that the (real) *Grassmann manifold* or *Grassmannian* (of q -planes in m -space) is the set $G(q, m)$ of q -dimensional vector subspaces of \mathbb{R}^m (for fixed $q \leq m$). The (differential) topology on $G(q, m)$ can be described in several ways: First, as a quotient (homogeneous space) of the orthogonal group,

$$G(q, m) = O(m)/O(q) \times O(m - q). \quad (2)$$

Next, as a submanifold of projective space,

$$G(q, m) \subset \mathbb{P}(\wedge^q \mathbb{R}^m) = \mathbb{P}^{\binom{m}{q}-1}(\mathbb{R}) \quad (3)$$

via the Plücker embedding. Finally, as a submanifold of Euclidean space,

$$G(q, m) \subset \mathbb{R}^{(m^2+m-2)/2} \quad (4)$$

via a projection embedding described recently in [4].

While the manifold structures on $G(q, m)$ obtained from these three constructions are equivalent (diffeomorphic), they naturally lead to different geometries on the Grassmannian. The standard invariant Riemannian metric on orthogonal matrices $O(m)$ descends via (2) to a Riemannian metric on the homogeneous space $G(q, m)$. The resulting *geodesic* distance function d_g (*arc length*) on the Grassmannian in terms of the principal angles $\theta_1, \dots, \theta_q$ between $X, Y \in G(q, m)$, is (see, e.g., [5])

$$d_g(X, Y) = \left(\sum_{i=1}^q \theta_i^2\right)^{1/2} = \|\theta\|_2.$$

If one prefers the realization (3), then the Grassmannian inherits a Riemannian metric from the Fubini-Study metric on projective space (see, e.g., [10]), and the resulting *Fubini-Study* distance d_{FS} is given in terms of the principal angles by

$$d_{FS}(X, Y) = \cos^{-1} \left(\prod_{i=1}^q \cos \theta_i\right).$$

Finally, one can restrict the usual Euclidean distance function on $\mathbb{R}^{(n^2+n-2)/2}$ to the Grassmannian via (4) to

obtain the *projection* F or *chordal* distance d_c (so called because the image of the Grassmannian under (4) lies in a sphere, so that the restricted distance is simply the distance along a straight-line chord connecting one point of that sphere to another; see [4]) which, in terms of the principal angles, has the expression

$$d_c(X, Y) = \left(\sum_{i=1}^q (\sin \theta_i)^2\right)^{1/2} = \|\sin \theta\|_2.$$

This projection F distance d_c has recently been used in the context of sphere-packing/coding theory in the Grassmannian, where it is significantly more efficient than the “standard” geodesic distance d_g [4], [1]. As a slight variation on the last formula, we may also consider the so-called *chordal Frobenius* distance d_{cF} , given in terms of the principal angles by $d_{cF}(X, Y) = \|2 \sin \frac{1}{2} \theta\|_2$.

It is revealing to consider nested subspaces of $X, Y \in G(q, m)$ by defining the ℓ -truncated principal angle vector $\theta^\ell := (\theta_1, \dots, \theta_\ell)$ where $\theta_1 \leq \dots \leq \theta_q$ are the principal angles between X and Y and $1 \leq \ell \leq q$. We then have ℓ -truncated metrics, e.g., arc length $d_g^\ell(X, Y) := \|\theta^\ell\|_2$. We now proceed to investigate how well these distance measures on $G(q, m)$ separate q -dimensional illumination face spaces of m -pixel images.

5. Balanced Sampling

We begin with the standard assumptions that a person does not change expression or move relative to the camera and hypothesize that given enough noise free samples, any two illumination spaces estimated from images of the same subject should be identical. More importantly, we hypothesize that any two illumination spaces estimated from images of different subjects should be non-intersecting. This is what we mean by the claim that illumination spaces are idiosyncratic.

In practice, however, images are noisy in their high-dimensional ambient space. As a result, two illumination spaces estimated from images of a single subject are not only not identical, they never even intersect in practice. We therefore recast our claim of illumination space idiosyncrasy for the real world: an illumination space estimated from images of one subject should always be “closer” to another illumination space estimated from the same subject than to any illumination space estimated from a different subject. We can test this hypothesis empirically, since the Grassmannian distance measures above give us geometrically motivated methods of measuring distances between subspaces. For all of the empirical analysis which follows, we consider comparisons between pairs of subjects whose image sets X and Y possess a similarity score $S(X, Y)$ defined as $S(X, Y) = 1/D(X, Y)$, where $D(X, Y)$ is one of the Grassmannian distances between $\mathcal{R}(X)$ and $\mathcal{R}(Y)$. The similarities S of different realizations of subspaces for the same person are called match scores while for different people they are called non-match scores.

Given match and non-match scores for a set of people, the *false accept rate (FAR) at a zero false reject rate (FRR)*¹ (defined, e.g., in [13]) indicates how well a similarity score S separates people. Indeed, a zero FAR indicates that image sets X and Y are perfectly separable. In the context of this work, it shows that matching illumination spaces are always closer together than non-matching illumination spaces, because illumination spaces are idiosyncratic. In the experiments which follow, the FAR will be reported and special attention will be drawn to situations where it is zero.

5.1. Baseline Algorithm

Below we present empirical results comparing one set of images to another using Grassmannian distance measures. We are sensitive to the concern that similar outcomes might be observed using direct image to image comparisons. To explore this, we introduce a baseline similarity $S(X, Y)$ for comparing multi-still sets X and Y . Set

$$s(x^{(j)}, Y) = \max_{1 \leq i \leq k} \{\text{Cor}(x^{(j)}, y^{(i)})\},$$

where $x^{(j)} \in X$ and $y^{(i)} \in Y$ and define

$$S(X, Y) = \frac{1}{2} \sum_{j=1}^k (s(x^{(j)}, Y) + s(y^{(j)}, X)).$$

Notice that $S(X, Y)$ is nothing more than a straight across comparison between images, and therefore serves as an excellent baseline in evaluating the performance of our method.

5.2. The separability of the CMU-PIE Database

The PIE database includes imagery of 68 people under different pose, illumination conditions, and expressions. Our work here concerns illumination variations rather than pose, so only frontal (c27) images are used. For these frontal images, there are 21 distinct sources of lights used to illuminate the face. In addition, these 21 sources are sampled both with the background room lights on and the background room lights off.

For each of the two types of imagery, room lights on and room lights off, we have randomly selected two disjoint sets of 10 images of 67 people in the PIE database.² This sampling is “balanced” in so much as it is randomized relative to the specific illumination settings. The use of only 10 images to estimate the illumination space is probably approaching a lower bound on the necessary number of samples. To augment the sample, the mirror

reflection of each image is also included in the image set when estimating illumination subspaces.

Figure 1 summarizes the room lights off and room lights on results, partitioned into the case where all 67 people are considered versus when only 39 people without glasses are considered. Each plot shows the FAR using our baseline algorithm and four ℓ -truncated Grassmannian distance measures, for $1 \leq \ell \leq 15$. The values shown are averages taken over ten trials. In each trial, random pairs of disjoint sets are created for each of the 67 people in the PIE database.

The horizontal axes in Figure 1 indicate the number of principal angles (ℓ) retained when computing the Grassmannian distance. The left plot in Figure 1 shows results for the imagery without subtracting the average image from each image set; the right plot in Figure 1 shows results for the imagery with means subtracted. The results for the mean subtracted imagery leads us to a significant observation: *There is sufficient information on how different people respond to changes in illumination to perfectly separate the 67 people in the PIE database based upon ten sample images.* The evidence for this claim is the fact that all four of the ℓ -truncated Grassmannian distance measures result in a zero FAR over a range of ℓ . While one cannot tell directly from inspection of the right plot in Figure 1, the points for each of the ℓ -truncated Grassmannian distance measures do lie exactly at zero for $2 \leq \ell \leq 10$. In contrast, the baseline algorithm has a FAR as high as 28%. It is also very interesting to observe how sensitive the baseline algorithm is to mean subtraction while the Grassmannian measures are relatively unaffected. Other notable observations from Figure 1 include that the baseline algorithm performs perfectly, with a zero FAR, for the case where room lights are on and mean retained. Where room lights are off, however, the baseline algorithm has about 1% FAR when all 67 subjects are considered. In contrast, all of the Grassmannian measures achieve a zero FAR for many of the settings in all cases summarized in Figure 1.

Given that the Grassmannian distance measures perfectly separate the 67 PIE subjects with lights on and lights off, it is natural next to ask what happens if these two cases are pooled together. Thus, our second experiment is similar to our first, except that 2 random disjoint sets of size 21 are created for each of the 67 people in the PIE database by randomly selecting from among both the room lights on and room lights off imagery. In this case, since there are twice as many sample images from which to estimate illumination subspaces, mirror images are not used. The results of this pooled test are summarized in Figure 2. The left plot in Figure 2 shows results where the average face image for each set is retained, while the right plot in Figure 2 shows the results with the average image subtracted from each set. For the baseline algorithm, this distinction is critical, with the baseline performing perfectly when means are retained, and having a FAR of

¹For the sake of brevity we refer to the false accept rate (FAR) at a zero false reject rate (FRR) simply as FAR.

²Subject 39 was not used because 3 images were missing due to hardware problems during the process of image acquisition.

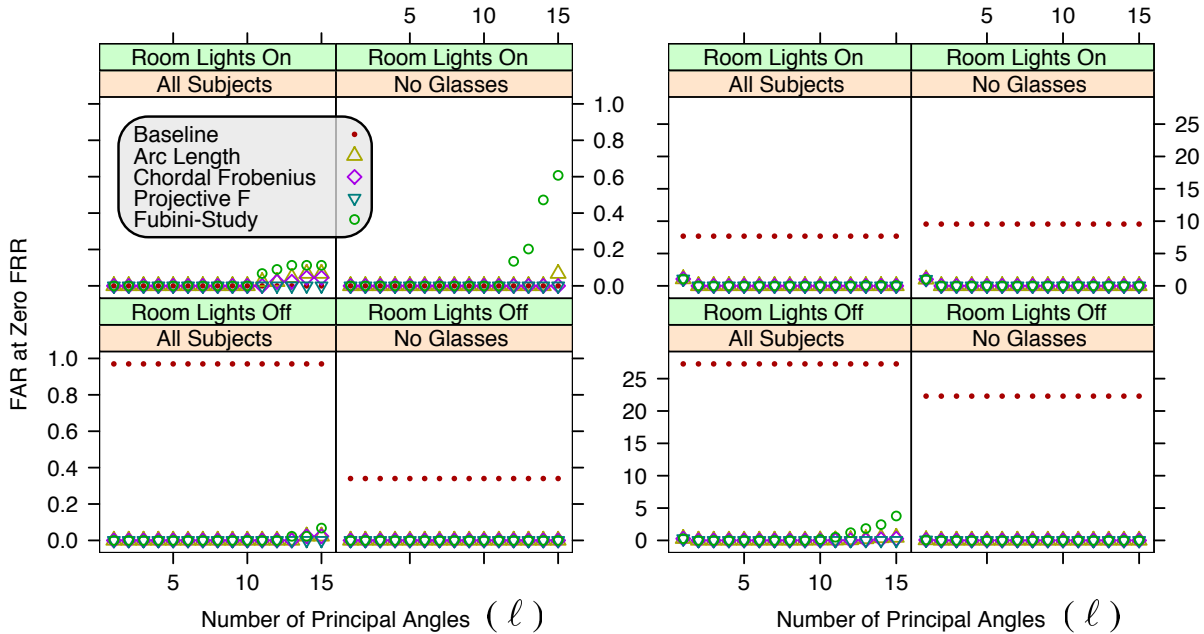


Figure 1. Plots of False Accept Rate (FAR) at a zero False Reject Rate (FRR) in % for the PIE database divided into frontal images with room lights on and room lights off. Between 1 and 15 principal angles are included in the Grassmannian distance computation as shown along the horizontal axis. Left: mean image is not subtracted from sample set. Right: mean image is subtracted from each sample set. In all cases images and their mirror reflections were used to estimate illumination spaces. Note the vertical axis scale changes considerably when means are subtracted.




Figure 2. Plots of FAR at a zero FRR in % for the frontal PIE images with room lights on and off pooled into a single group and then randomly sampled. Left plot is for sample images with the mean retained, versus the right plot where the means have been subtracted. Legend is the same as in Figure 1.

22.8% when means are subtracted. This result reinforces the finding that the Grassmannian distance measures are exploiting how illumination varies, rather than comparing any sort of average or normalized face. This clearly is not true for the baseline algorithm.

5.3. The separability of the Yale Database (YDB)

The YDB [7] has far fewer subjects than PIE. Nonetheless, it is worth studying because it is the oldest and most studied illumination database, and because it has a large

number of images per subject. For each of the ten subjects, two disjoint sets have been created by randomly sampling from the 64 images per person. Then, these sets have been compared using the baseline algorithm and four (ℓ -truncated) Grassmannian distance measures applied to the estimated illumination subspaces.

The results when image sets contain only 8 images are summarized in Figure 3. Results are shown both with and without mirror images added to the set. Also, as before, the FAR rates are averages over ten trials using distinct randomly generated image sets for each subject. There are two significant findings evident in Figure 3. The first concerns the baseline algorithm: *The baseline algorithm performs very poorly on the YDB when only 8 images are used, with FAR rates between 43% and 49%*. This finding is particularly interesting given there are only ten people in the database. It is also interesting that mean subtraction appears to have only a modest influence on how the baseline algorithm performs. The second finding is perhaps more significant: *The YDB is perfectly Grassmann separable using illumination subspaces derived from as few as 8 randomly sampled images*. The strongest evidence for this finding is the mean subtracted case with mirror images. In this case, three of the four ℓ -truncated Grassmannian distance measures have zero FAR when $5 \leq \ell \leq 10$.

The YDB experiment was also carried out for 16, 21 and 32 samples and the results are summarized in Table 1. Observe that the baseline algorithm continues to

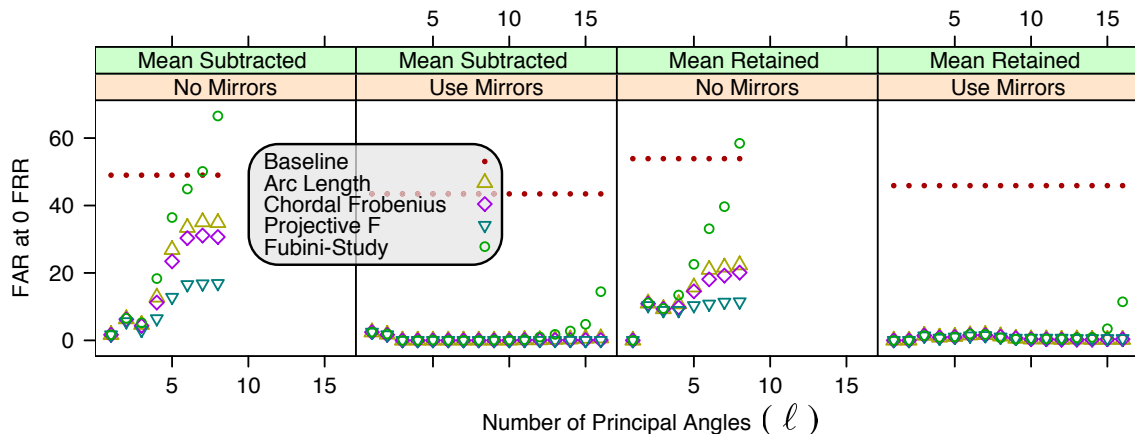


Figure 3. Plots of FAR at a zero FRR in % for the Yale database using 8 sample images.

have difficulty with this data. While it benefits from more images, it never achieves a zero FAR. In contrast, for the 21 and 32 sample case, all four Grassmannian distance measures perfectly separate all ten Yale subjects, as do all four Grassmannian measures when mirrors are included and 16 samples are used.

6. Imbalanced Sampling

In all of the previous tests, the images in the two sets per subject were randomly sampled from the larger set of images without any explicit bias relative to a known factor. This may not always be possible, and it is fair to ask what happens if the image sets are acquired under different conditions. To approach this question, we have

Table 1. FAR at zero FRR for various ℓ -truncated Grassmannian distance measures averaged over $5 \leq \ell \leq 10$, applied to illumination spaces estimated from n disjoint samples from YDB. BL: baseline; AL: arc length; CF: chordal Frobenius; PF: projection F; FS: Fubini-Study.

		With Mirror Images				
n	Mean	BL	AL	CF	PF	FS
8	Retained	45.89	1.07	1.04	1.02	1.06
8	Subtracted	43.45	0.00	0.00	0.00	0.06
16	Retained	29.56	0.00	0.00	0.00	0.00
16	Subtracted	10.22	0.00	0.00	0.00	0.00
21	Retained	16.56	0.00	0.00	0.00	0.00
21	Subtracted	13.44	0.00	0.00	0.00	0.00
32	Retained	6.00	0.00	0.00	0.00	0.00
32	Subtracted	1.11	0.00	0.00	0.00	0.00
		Without Mirror Images				
8	Retained	53.89	20.19	18.00	10.97	38.44
8	Subtracted	49.00	32.58	28.89	15.75	49.50
16	Retained	21.33	0.07	0.06	0.02	0.13
16	Subtracted	22.78	0.00	0.00	0.00	0.04
21	Retained	13.22	0.00	0.00	0.00	0.00
21	Subtracted	9.10	0.00	0.00	0.00	0.00
32	Retained	3.33	0.00	0.00	0.00	0.00
32	Subtracted	1.00	0.00	0.00	0.00	0.00

designed an experiment where one set is made up of the 21 PIE images acquired with the room lights on, and the other image set is made up of the 21 PIE images with the room lights off. Under this scenario, we expect the estimated illumination spaces for a single subject to differ due to the fact that the samples are drawn systematically from very different lighting conditions. How these differences will manifest themselves when comparing illumination subspaces derived from a single person to those derived from different people is not on the surface obvious, and hence the empirical results are all the more intriguing.

Figure 4 summarizes the outcome of this experiment. All four Grassmannian distance measures show a heightened sensitivity to this choice of angles in the experiment, and three of the four appear to become unstable relative to false accept rates as more than 5 principal angles are used – in other words their associated FAR grows and exceeds that of the baseline algorithm. However, there is one extremely striking outcome of this experiment concerning the Fubini-Study (FS) distance measure: *The FS-distance perfectly separates all 67 subjects in PIE in the highly biased case where illumination subspaces estimated with the room lights on are compared to illumination subspaces estimated with the room lights off.* The evidence for this finding is the zero FAR achieved by the FS-distance when 36 or more of the principal angles are used. This is true for all four of the cases plotted in Figure 4. We cannot fully explain this result at this time.

7. Conclusion

This paper addresses how appearance of a face changes with illumination. We introduce the idea of illumination spaces as points on a Grassmannian, and use this formalism to establish distance measures between illumination subspaces. The CMU-PIE and YDB are shown to be Grassmann separable, i.e., the estimated illumination space of a person is always closer to another estimated il-




Figure 4. Plots of FAR at a zero FRR in % for PIE using only room lights on in one sample and room lights off in the other sample.

lumination space of the same person than to the estimated illumination space of any other person in the data set. This work demonstrates that PIE and YDB can be perfectly solved, i.e., are Grassman separable, if the recognition problem is cast as comparing sets of images to sets of images. It also suggests that the appearance of a face under changes in illumination is an identifying characteristic. Rather than trying to remove illumination effects from images, it may be better to acquire multiple images of a person under varying illumination, and base recognition algorithms on the differences among these images.

References

- [1] A. Barg and D. Nogin. Bounds on packings of spheres in the Grassmann manifold. *IEEE Trans. Information Theory*, 48(9):2450–2454, 2002.
- [2] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, 2003.
- [3] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible illumination conditions. *IJCV*, 28(3):245–260, July 1998.
- [4] J. Conway, R. Hardin, and N. Sloane. Packing lines, planes, etc.: Packings in Grassmannian spaces. *Experimental Mathematics*, 5:139–159, 1996.
- [5] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20(2):303–353, 1999.
- [6] A. Fitzgibbon and A. Zisserman. Joint manifold distance: a new approach to appearance based clustering. In *CVPR*, pages 26–36. IEEE Computer Society, 2003.
- [7] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *PAMI*, 23(6):643–660, 2001.
- [8] A. Georghiades, D. Kriegman, and P. Belhumeur. Illumination cones for recognition under variable lighting: Faces. In *CVPR*, pages 52–59. IEEE Computer Society, 1998.
- [9] G. H. Golub and C. F. V. Loan. *Matrix Computations*. Johns Hopkins University Press, third edition, 1996.
- [10] P. Griffiths and J. Harris. *Principles of Algebraic Geometry*. Wiley & Sons, 1978.
- [11] M. Kirby and L. Sirovich. Application of the Karhunen-Loeve procedure for the characterization of human faces. *PAMI*, 12(1):103–108, 1990.
- [12] K.-C. Lee, J. Ho, and D. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *PAMI*, 27(5):684–698, 2005.
- [13] A. Mansfield and J. Wayman. Best practices in testing and reporting of biometric devices: Version 2.01. Technical Report NPL Report CMSC 14/02, Centre for Mathematics and Scientific Computing, National Physical Laboratory, UK, 2002.
- [14] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *CVPR*, pages 947–954. IEEE Computer Society, 2005.
- [15] R. Ramamoorthi. Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object. *PAMI*, 24(10):1322–1333, 2002.
- [16] T. Riklin-Raviv and A. Shashua. The quotient image: Class based re-rendering and recognition with varying illuminations. *PAMI*, 23(2):129–139, 2001.
- [17] T. Sim, S. Baker, and M. Bsat. The CMU pose, illumination, and expression database. *PAMI*, 25(12):1615–1618, 2003.
- [18] L. Wolf and A. Shashua. Kernel principal angles for classification machines with applications to image sequence interpretation. *JMLR*, 4:913–931, 2003.
- [19] O. Yamaguchi, K. Fukui, and K. Maeda. Face recognition using temporal image sequence. In *AFG*, pages 318–323, 1998.