Appeared in the 2002 International Conference on Automatic Face and Gesture Recognition

Eigen Light-Fields and Face Recognition Across Pose

Ralph Gross, Iain Matthews, and Simon Baker
The Robotics Institute, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213
Email: {rgross,iainm,simonb}@cs.cmu.edu

Abstract

In many face recognition tasks the pose of the probe and gallery images are different. In other cases multiple gallery or probe images may be available, each captured from a different pose. We propose a face recognition algorithm which can use any number of gallery images per subject captured at arbitrary poses, and any number of probe images, again captured at arbitrary poses. The algorithm operates by estimating the *eigen light-field* of the subject's head from the input gallery or probe images. Matching between the probe and gallery is then performed using the eigen light-fields. We present results on the CMU PIE and the FERET face databases.

Keywords: Face recognition across pose, eigen light-fields, appearance-based object recognition, image re-rendering.

1. Introduction

In many face recognition scenarios the pose of the probe and gallery images are different. For example, the gallery image might be a frontal "mug-shot" and the probe image might be a 3/4 view captured from a camera in the corner of the room. The number of gallery and probe images can also vary. For example, the gallery may consist of a pair of images for each subject, a frontal mug-shot and full profile view (like the images typically captured by police departments). The probe may be a similar pair of images, a single 3/4 view, or even a collection of views from random poses.

Face recognition across pose (i.e. face recognition where the gallery and probe images do not have the same poses) has received very little attention. Algorithms have been proposed which can recognize faces [12] (or more general objects [10]) at a variety of poses. Most of these algorithms require gallery images at every pose, however. Algorithms have been proposed which do generalize across pose, for example [6], but this algorithm computes 3D head models using a gallery containing a large number of images per subject captured with substantial controlled illumination variation. It cannot be used with arbitrary gallery and probe sets.

We propose an algorithm for face recognition across pose. Our algorithm can use any number of gallery images captured at arbitrary poses, and any number of probe images also captured with arbitrary poses. A minimum of 1 gallery and 1 probe image are needed, but if more images are available the performance of our algorithm generally gets better.

Our algorithm operates by estimating (a representation of) the light-field [9] of the subject's head. First, generic training data is used to compute an eigen-space of head light-fields, similar to the construction of eigen-faces [16]. Light-fields are simply used rather than images. Given a collection of gallery or probe images, the projection into the eigen-space is performed by setting up a least-squares problem and solving for the projection coefficients similarly to approaches used to deal with occlusions in the eigenspace approach [4, 8]. This simple linear algorithm can be applied to any number of images, captured from any poses. Finally, matching is performed by comparing the probe and gallery eigen light-fields.

We evaluate our algorithm on the pose subset of the CMU PIE database [15] and a subset of the FERET database [13]. We demonstrate that our algorithm is able to reliably recognize people across pose and that our algorithm performs better if more gallery or probe images are used. We also investigate the variation in performance of our algorithm with the poses of the gallery and probe images. Finally, we investigate whether it is better to have a large gallery and a singe probe image or equal sized gallery and probe sets.

The remainder of this paper is organized as follows. We proceed in Section 2 to introduce light-fields, eigen light-fields, our algorithm for estimating an eigen light-field from any number of images, and derive some of the properties of the algorithm. In Section 3 we describe how the theoretical eigen light-field estimation algorithm can be used for face recognition across pose. After presenting our experimental results in Section 4 we conclude in Section 5.

2. Theory

2.1. Object Light-Fields

The *plenoptic function* [1] or *light-field* [9] is a function which specifies the radiance of light in free space. It is a

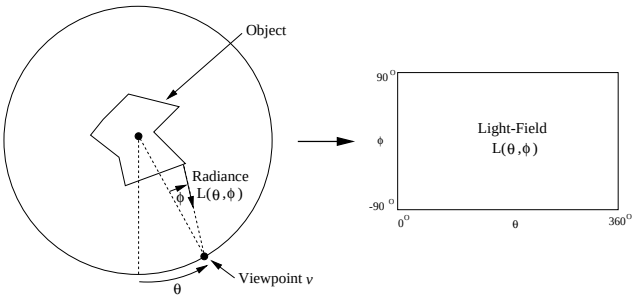

Figure 1. An illustration of the 2D light-field of a 2D object [9]. The object is conceptually placed within a circle. The angle to the viewpoint v around the circle is measured by the angle θ, and the direction that the viewing ray makes with the radius of the circle is denoted ϕ. For each pair of angles θ and ϕ, the radiance of light reaching the viewpoint from the object is then denoted by $L(\theta, \phi)$, the *light-field*. Although the light-field of a 3D object is actually 4D, we will continue to use the 2D notation of this figure in this paper for ease of explanation.

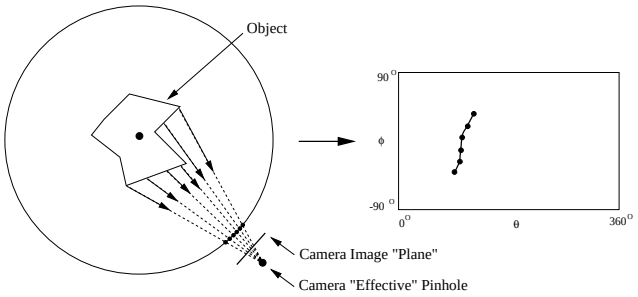

Figure 2. The 1D image of a 2D object corresponds to a curve (surface for a 2D image of a 3D object) in the light-field. Each pixel in the image corresponds to a ray in space through the camera pinhole and the location of the pixel on the image plane. In general this ray intersects the light-field circle at a different point for each pixel. As the pixel considered "moves" in the image plane, the point on the light-field circle therefore traces out a curve in θ-ϕ space. This curve is a straight vertical line iff the "effective pinhole" of the camera lies on the circle used to define the light-field.

5D function of position (3D) and orientation (2D). In addition, it is also sometimes modeled as a function of time, wavelength, and polarization, depending on the application in mind. Assuming that there is no absorption or scattering of light through the air [11], the light-field is actually only a 4D function, a 2D function of position defined over a 2D surface, and a 2D function of direction [7, 9]. In 2D, the light-field of a 2D object is actually 2D rather, than the 3D that might be expected. See Figure 1 for an illustration.

2.2. Eigen Light-Fields

Suppose we are given a collection of light-fields $L_i(\theta, \phi)$ where $i = 1, \ldots, N$. See Figure 1 for the definition of this notation. If we perform an eigen-decomposition of these vectors using Principal Components Analysis (PCA), we obtain $d \leq N$ eigen light-fields $E_i(\theta, \phi)$ where $i = 1, \ldots, d$. Then, assuming that the eigen-space of light-fields is a good representation of the set of light-fields under consideration, we can approximate any light-field $L(\theta, \phi)$ as:

$$L(\theta, \phi) \approx \sum_{i=1}^{d} \lambda_i E_i(\theta, \phi) \tag{1}$$

where $\lambda_i = \langle L(\theta, \phi), E_i(\theta, \phi) \rangle$ is the inner (or dot) product between $L(\theta, \phi)$ and $E_i(\theta, \phi)$. This decomposition is analogous to that used in face and object recognition [10, 16]; it is just performed on the entire light-field rather than on images.

2.3. Estimating Light-Fields from Images

Capturing the complete light-field of an object is a difficult task, primarily because it requires a huge number of images [7, 9]. In most object recognition scenarios it is unreasonable to expect more than a few images of the object; often just one. As shown in Figure 2, however, any image of the object corresponds to a curve (for 3D objects, a surface) in the light-field. One way to look at this curve is as a highly occluded light-field; only a very small part of the light-field is visible.

Can the eigen coefficients λ_i be estimated from this highly occluded view? Although this may seem hopeless, consider that light-fields are highly redundant, especially for objects with simple reflectance properties such as Lambertian. An algorithm is presented in [8] to solve for the unknown λ_i for eigen-*images*. A similar algorithm was proposed in [4]. Rather than using the inner product $\lambda_i = \langle L(\theta, \phi), E_i(\theta, \phi) \rangle$, Leonardis and Bischof [8] solve for λ_i as the least squares solution of:

$$L(\theta, \phi) - \sum_{i=1}^{d} \lambda_i E_i(\theta, \phi) = 0 \tag{2}$$

where there is one such equation for each pair of θ and ϕ that are un-occluded in $L(\theta, \phi)$. Assuming that $L(\theta, \phi)$ lies *completely within the eigen-space* and that enough pixels are un-occluded, then the solution of Equation (2) will be exactly the same as that obtained using the inner product:

Theorem 1 *Assuming that $L(\theta, \phi)$ is in the linear span of $\{E_i(\theta, \phi) \mid i = 1, \ldots d\}$, then $\lambda_i = \langle L(\theta, \phi), E_i(\theta, \phi) \rangle$ is always an exact solution of Equation (2).*

Since there are d unknowns $(\lambda_1 \ldots \lambda_d)$ in Equation (2), at least d un-occluded light-field pixels are needed to over-constrain the problem, but more may be required due to linear dependencies between the equations. In practice, $2 - 3$

times as many equations as unknowns are typically required to get a reasonable solution [8]. Given an image $I(m,n)$, the following is then an algorithm for estimating the eigen light-field coefficients λ_i:

Algorithm 1: Eigen Light-Field Estimation

1. For each pixel (m,n) in $I(m,n)$ compute the corresponding light-field angles $\theta_{m,n}$ and $\phi_{m,n}$. (This step assumes that the camera intrinsics are known, as well as the relative orientation of the camera to the object.)
2. Find the least-squares solution (for $\lambda_1 \ldots \lambda_d$) to the set of equations:

$$I(m,n) - \sum_{i=1}^{d} \lambda_i E_i(\theta_{m,n}, \phi_{m,n}) \;=\; 0 \qquad (3)$$

   where m and n range over their allowed values. (In general, the eigen light-fields E_i need to be interpolated to estimate $E_i(\theta_{m,n}, \phi_{m,n})$. Also, all of the equations for which the pixel $I(m,n)$ does not image the object should be excluded from the computation.)

Although we have described this algorithm for a single image $I(m,n)$, any number of images can obviously be used (so long as the camera intrinsics and relative orientation to the object are known for each image.) The extra pixels from the other images are simply added in as additional constraints on the unknown coefficients λ_i in Equation (3).

2.4. Properties of the Algorithm

Algorithm 1 can be used to estimate a light-field from a collection of images. Once the light-field has been estimated, it can then be used to render new images of the same object under different poses. (See [17] for a related algorithm.) In this section we show that, if the objects used to create the eigen-space of light-fields all have the same shape as the object imaged to create the input to the algorithm, then this re-rendering process is in some sense "correct," assuming that all the object are Lambertian. As a first step, we show that the eigen light-fields $E_i(\theta,\phi)$ capture the shape of the objects in the following sense:

Lemma 1 *If $\{L_i(\theta,\phi) \,|\, i = 1,\ldots,N\}$ is a collection of light-fields of Lambertian objects with the same shape, then all of the eigen light-fields $E_i(\theta,\phi)$ have the property that if (θ_1,ϕ_1) and (θ_2,ϕ_2) define two rays which image the same point on the surface of any of the objects then:*

$$E_i(\theta_1,\phi_1) \;=\; E_i(\theta_2,\phi_2) \qquad \forall i = 1 \ldots d. \qquad (4)$$

Proof: The property in Equation (4) holds for all of the light-fields $\{L_i(\theta,\phi) | i = 1,\ldots,N\}$ used in the PCA because they are Lambertian. Hence, it also holds for any linear combination of the L_i. Therefore it holds for the eigenvectors because they are linear combinations of the L_i. □

The property in Equation (4) clearly also holds for all linear combinations of the eigen light-fields. It therefore also holds for the light-field recovered in Equation (3) in Algorithm 1, assuming that the light-field from which the input image is derived lies completely in the eigen-space and so Theorem 1 applies. This fact means that Algorithm 1 estimates the light-field in a way that is consistent with the object being Lambertian and of the appropriate shape:

Theorem 2 *Suppose $\{E_i(\theta,\phi) \,|\, i = 1,\ldots,d\}$ are the eigen light-fields of a set of Lambertian objects with the same shape and $I(m,n)$ is an image of another Lambertian object with the same shape. If the light-field from which $I(m,n)$ is derived lies in the light-field eigen-space, then the light-field recovered by Algorithm 1 has the property that if $\theta_{m,n}, \phi_{m,n}$ is* any *pair of angles which image the same point in the scene as the pixel (m,n) then:*

$$I(m,n) \;=\; E(\theta_{m,n}, \phi_{m,n}). \qquad (5)$$

where $E(\theta_{m,n}, \phi_{m,n})$ is the light-field estimated by Algorithm 1; i.e. Algorithm 1 correctly re-renders the object under the Lambertian reflectance model.

Theorem 2 implies that Algorithm 1 is acting reasonably in estimating the light-field, a task which is in general impossible without a prior model on the shape of the object. (The shape model here is contained in the eigen-space.) Theorem 2 assumes that faces are approximately the same shape, but that is a common assumption [14]. Theorem 2 also assumes that faces are Lambertian and that the light-field eigenspace accurately approximates any face light-field. The extent to which these assumptions are valid is demonstrated by the empirical results obtained by our algorithm. See Section 4.

3. Face Recognition Across Pose

We propose to use Algorithm 1 to perform face recognition across pose. Although the derivation in Section 2 is in terms of the entire light-field, the results also clearly hold for any subset of rays (θ,ϕ) in the light-field. We will evaluate our algorithm on a subset of the CMU PIE database [15] and a subset of the FERET database [13]. In the CMU PIE database 68 subjects are imaged under 13 different poses totalling 884 images (figure 3). We furthermore report results on 75 subjects from the FERET database which are recorded in 9 different poses (figure 5). The remainder of this section describes how we map the abstract algorithms in Section 2 onto the data in Figure 3.

3.1. Gallery, Probe, & Generic Training Data

There are 68 subjects in the PIE database. We randomly select $N = 34$ of these subjects and use the images of them as generic training data to construct the eigen light-fields.

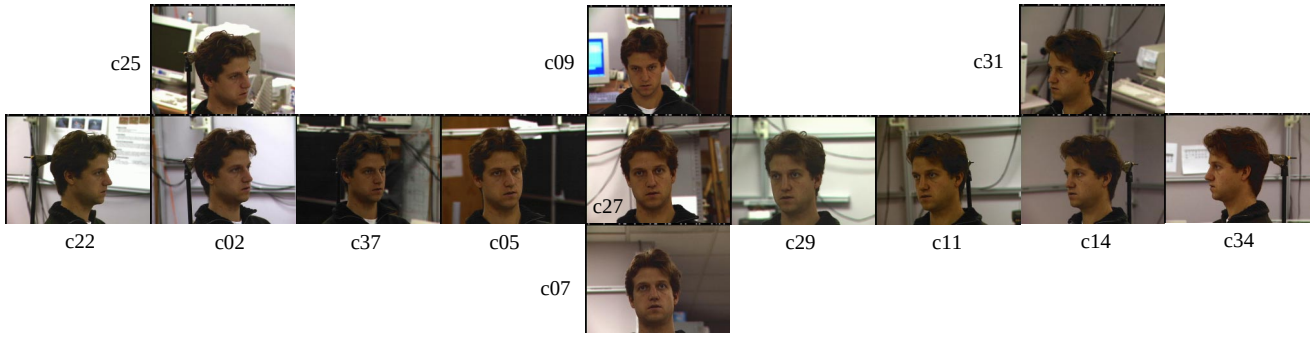

Figure 3. The pose variation in the PIE database [15]. The pose varies from full left profile (c34) to full frontal (c27) and on to full right profile (c22). The 9 cameras in the horizontal sweep are each separated by about 22.5°. The 4 other cameras include 1 above (c09) and 1 below (c07) the central camera, and 2 in the corners of the room (c25 and c31), typical locations for surveillance cameras.

Figure 4. Face normalization. The original images from camera views c27 and c37 are shown together with the normalized and crop face region.

Of the remaining 34 subjects, the images are divided into *completely disjoint* gallery and probe sets based on the camera they are captured from. For example, if we consider the scenario described in the introduction where the gallery consists of a pair of images for each subject, a frontal mugshot and full profile view, we might use the images from cameras c22 and c27 for the gallery (see figure 3). If the probe images are 3/4 views, we would use the images from camera c37 (or c11.)

3.2. Extracting the Face Region

We hand-labeled the x-y positions of both eyes (pupils) and the tip of the nose in all 884 images of the PIE database and the 675 images of the FERET database. Within each pose separately the face images are normalized for rotation, translation, and scale. The face region is then tightly cropped using the normalized feature point distances. Figure 4 shows the result of face region extraction for two cameras (c27 and c37) of the PIE database.

3.3. Constructing the Light-Field Eigenspace

Face region extraction is performed on every image from every camera; generic training data, gallery images, and probe images. Suppose that the results are the images:

$$\mathrm{cropped\_face\_id\_c} \tag{6}$$

where $\mathrm{id} \in \{1, 2, \ldots, 68\}$ is the identity of the subject and $\mathrm{c} \in \{02, 05, 07, 09, 11, 14, 22, 25, 27, 29, 31, 34, 37\}$ is the camera number. Suppose $\mathrm{Training} \subset \{1, 2, \ldots, 68\}$ is the set of generic training subjects, $\mathrm{Gallery}$ is the set of gallery cameras, and Probe is the set of probe cameras. We then form the light-field eigenspace as follows. For each $\mathrm{id} \in \mathrm{Training}$, the images:

$$\{\mathrm{cropped\_face\_id\_c} \mid \mathrm{c} \in \mathrm{Gallery} \cup \mathrm{Probe}\} \tag{7}$$

are raster-scanned and concatenated. PCA is performed on these $N = 34$ vectors to form the eigen-vectors E_i.

3.4. Processing the Gallery Images

For each E_i ($i = 1, \cdots, d$, the dimension of the light-field eigenspace) we extract the elements corresponding to the gallery images and form a shorter vector E^{G}. For each non-training subject $\mathrm{id} \notin \mathrm{Training}$ we raster-scan the images:

$$\{\mathrm{cropped\_face\_i\_c} \mid \mathrm{c} \in \mathrm{Gallery}\} \tag{8}$$

to form a vector of the same length. We solve Equation (3) for these shortened vectors. Suppose the result is λ_i^{id}.

3.5. Processing the Probe Images

We process the probe images similarly. For each eigenvector E_i we extract the elements corresponding to the probe images and form a shorter vector E^{P}. For each non-training subject $\mathrm{id} \notin \mathrm{Training}$ we raster-scan the images:

$$\{\mathrm{cropped\_face\_i\_c} \mid \mathrm{c} \in \mathrm{Probe}\} \tag{9}$$

Figure 5. The pose variation in the FERET database [13]. The pose varies from +60 (bb) to full frontal (ba) and on to -60 (bi).

to form a vector of the same length. We solve Equation (3) for these shortened vectors. Suppose the result is μ_i^{id}; i.e. we use μ as the equivalent of λ for the probe images.

3.6. The Classification Algorithm

We use a nearest neighbor classification algorithm for simplicity. For each probe subject id we find the closest matching gallery eigen light-field as:

$$\arg\min_{\mathrm{id}*} \sum_{i=1}^{d} \left(\mu_i^{\mathrm{id}} - \lambda_i^{\mathrm{id}*}\right)^2. \tag{10}$$

If the nearest neighbor eigen light-field $\mathrm{id}* = \mathrm{id}$ the algorithm has correctly recognized the subject. If $\mathrm{id}* \neq \mathrm{id}$ the algorithm has recognized the subject incorrectly.

4. Experimental Results

4.1. Comparison with Other Algorithms

We first compared our algorithm with eigenfaces [16], as implemented by Beveridge *et. al* [3], and FaceIt, the commercial face recognition system from Visionics. FaceIt finished as the top performer in the Face Recognition Vendor Test 2000 [5]. In Figure 6 we compare the recognition rate of the three algorithms for the gallery $E^{\mathrm{G}} = \{05, 27, 29\}$. We plot the recognition rate for each of the 10 cameras not in the gallery and the overall average. As would be expected for a simple appearance based algorithm, eigenfaces performs poorly in all cases (except for c07, one of the cameras closest to the gallery). FaceIt performs comparably to our algorithm on the 4 cameras closest to the gallery images c07, c09, c11, and c37. For the profile views c02, c22, c25, c31, and c34, our algorithm outperforms both of the other algorithms by a huge margin. On average our algorithm achieves a recognition accuracy of 69% vs. 42% for FaceIt (version 2) and 18% for eigenfaces. In terms of relative performance the eigen light-fields improve the FaceIt result by 60.7% (the error rate improves from 58% to 31%, an improvement of $(58-31)/((58+31)/2) = 60.7$).

We observe similar results on the FERET database. When trained on frontal and near frontal images our algorithm outperforms FaceIt by far on the poses closest to the profile views (92% vs. 83%). See Figure 7. Here our algorithm improves the FaceIt result by 72% in the same sense as above. The similarity in the improvements across two databases suggests that our results are general in nature and not specific to one database.

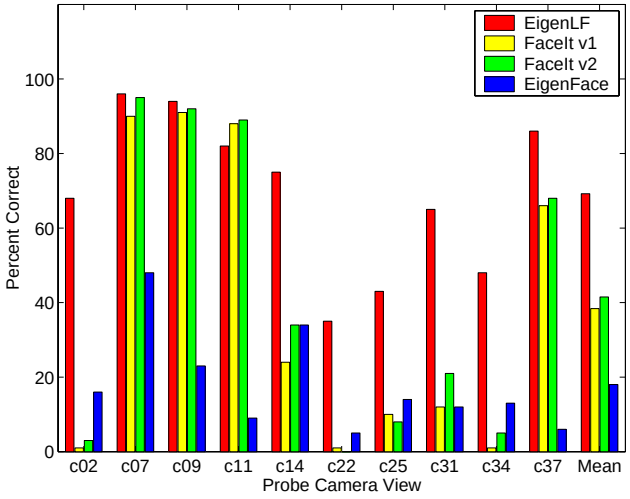

Figure 6. A comparison of our algorithm with eigenfaces [3, 16] and two FaceIt versions on the gallery $E^{\mathrm{G}} = \{05, 27, 29\}$. FaceIt version 1 was available before Visionics had access to the PIE database. FaceIt version 2 came out after Visionics had a copy of PIE for approximately one year. The recognition rate of our algorithm and FaceIt is similar for the 4 cameras $\{07, 09, 11, 37\}$ closest to the gallery. For the profile views $\{02, 22, 25, 31, 34\}$ our algorithm outperforms FaceIt by far. In all cases, eigenfaces performs poorly.

4.2. Improvement with the Number of Images

In Figure 8(a) we plot an example of the improvement of the performance of our algorithm with the number of gallery images. We plot the recognition rate, computed on average over all of the other images in the database, for 5 different galleries $\{25\}$, $\{25, 27\}$, $\{25, 27, 34\}$, $\{11, 25, 27, 34\}$, and $\{05, 11, 25, 27, 34\}$. For each of these cases, only one probe is used at a time. On the same graph we also plot the results using a single gallery image, but varying the num-

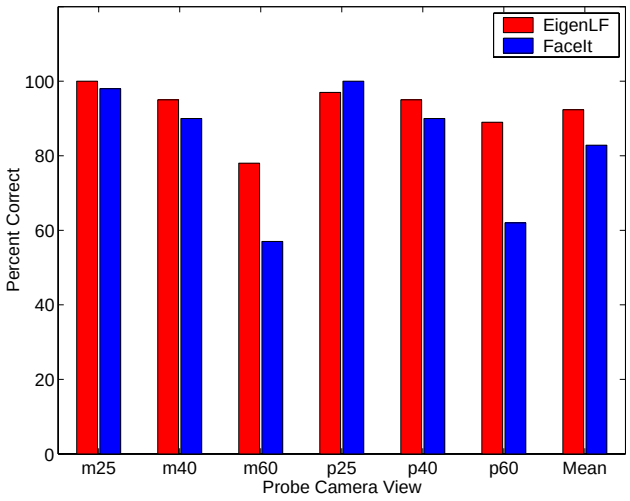

Figure 7. A comparison of our algorithm with FaceIt on the gallery $E^{\mathrm{G}} = \{ba, bf, be\}$ of the FERET database. Our algorithm outperforms FaceIt by far on the poses closest to the profile views $\{bb, bi\}$.

ber of probe images. We find that the performance of our algorithm increases rapidly with the number of gallery images and also that the role of the gallery and probe sets are approximately interchangeable.

4.3. Variation with the Pose of the Gallery

In Figure 8(b) we compare using a gallery containing only frontal images in $\{05, 07, 09, 11, 27, 29, 37\}$ and one containing only profile views in $\{02, 11, 14, 22, 25, 31, 34, 37\}$. We compute the recognition rate for different gallery sizes. For each gallery size, we randomly generate a large number of frontal galleries of that size and profile galleries of that size. We then compute the average recognition rate over all single probe images not in the gallery. We then repeat this process for 99 other randomly chose gallery pairs and average the results. The graphs in Figure 8(b) show that the performance of the frontal galleries is far superior to that of the profile galleries.

4.4. Division between the Gallery and Probe

A natural question which occurs at this point is whether it is better to have a large gallery and a single probe image, or a gallery and a probe half the size of the large gallery. We randomly generated pairs of galleries and probes. In the first case we choose a large gallery and a single probe. In the second case we choose an equal size gallery and probe. We then compute the average recognition rate. In Figure 8(c) we plot two curves against the total size of the gallery and probe sets combined. The results clearly show that it is better to divide the images equally between the gallery and probe sets than to have a large gallery and a single probe.

5. Discussion

5.1. Summary

We have proposed an algorithm for face recognition across pose based on an algorithm to estimate an eigen light-field from a collection of images. The algorithm can use any number of gallery images captured form arbitrary poses and any number of probe images also captured from arbitrary poses. The gallery and probe poses do not need to overlap, and any number of gallery and probe images can be used. We have shown that our algorithm can reliably recognize faces across pose. We have also shown that our algorithm can take advantage of the additional information contained in widely separated views to improve recognition performance if more than 1 gallery or probe image is available.

5.2. Current Limitations of the Algorithm

In this first paper describing our face recognition algorithm we have concentrated on showing that our algorithm can: (1) recognize people across pose and (2) take advantage of widely spaced views to yield improved face recognition performance. To get preliminary results, we have simplified the task in several ways: (1) the poses of the cameras are known and fixed, (2) the locations of the eyes and the nose used to extract the face region are marked by hand, and (3) the generic training data is captured with the same cameras that are used to capture the gallery and probe images. All of these factors make face recognition easier and are limitations on the current algorithm. We are continuing to develop our algorithm to remove these limitations, while retaining the desirable properties of the algorithm.

We recently conducted preliminary experiements using PIE images as generic training data and FERET images as gallery and probe images. Our algorithm achieves a recognition accuracy of 81.3%, which compares very well to the performance of FaceIt over the same dataset (84.4%).

5.3. Future Work: Illumination

The question "what is the set of images of an object under all possible illumination conditions?" was recently posed and answered in [2]. Is there an analogous result for light-fields? Since images consist of subsets of rays from the light-field, it is not surprising that the analogous result does hold for light-fields:

Theorem 3 *The set of n-pixel light-fields of any object, under all possible lighting conditions, is a convex cone in $\mathbf{R}^n$.*

An illumination invariant face recognition algorithm was proposed in [2] based on the equivalent of this "convex

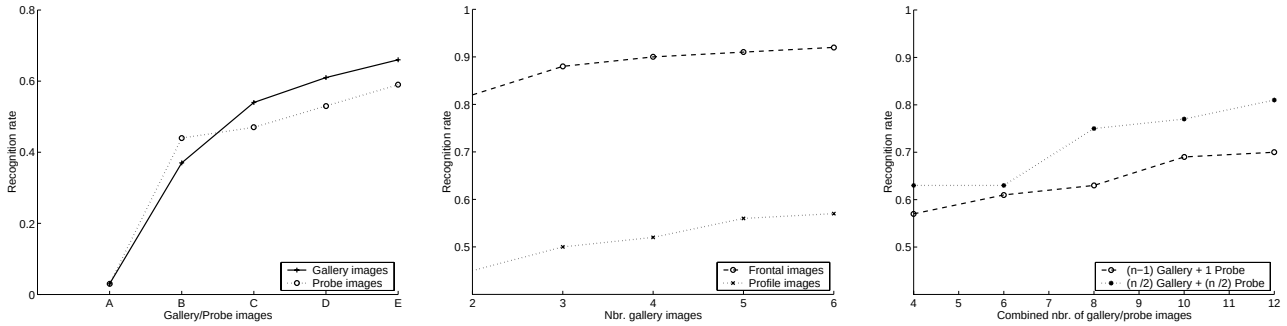

(a) Variation with the Number of Images (b) Variation with the Pose of the Gallery (c) Division between the Gallery and Probe

Figure 8. (a) The variation in the performance of our algorithm with the number of gallery and probe images. The camera sets A $= \{27\}$, B $= \{25, 27\}$, C $= \{25, 27, 34\}$, D $= \{11, 25, 27, 34\}$, and E $= \{05, 11, 25, 27, 34\}$. Our algorithm takes advantage of more gallery and probe images. (b) The variation in performance between frontal and profile galleries. Our algorithm operates better when the gallery images are frontal rather than profile. (c) A comparison of 2 ways of dividing the images between the gallery and probe set. Our algorithm operates better if the images are divided equally between gallery and probe sets rather than having a large gallery and a single probe.

cone" property for images. We are currently working on extending our algorithm to be able to recognize faces across both pose and illumination.

Acknowledgements

We wish to thank Terence Sim and Takeo Kanade for preliminary discussions on the light-field estimation algorithm. The research described in this paper was supported by U.S. Office of Naval Research contract N00014-00-1-0915. Portions of the research in this paper use the FERET database of facial images collected under the FERET program.

References

[1] E. Adelson and J. Bergen. The plenoptic function and elements of early vision. In Landy and Movshon, editors, *Computational Models of Visual Processing*. MIT Press, 1991.

[2] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible lighting conditions? *IJCV*, 28(3):1–16, 1998.

[3] R. Beveridge Evaluation of face recognition algorithms. www.cs.colostate.edu/evalfacerec, 2001.

[4] M. Black and A. Jepson. Eigen-tracking: Robust matching and tracking of articulated objects using a view-based representation. *International Journal of Computer Vision*, 36(2):101–130, 1998.

[5] D. Blackburn, M. Bone, and P. Phillips. Facial recognition vendor test 2000: Evaluation report, 2000.

[6] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Generative models for recognition under variable pose and illumination. In *Proceedings of the 4th ICAFGR*, 2000.

[7] S. Gortler, R. Grzeszczuk, R. Szeliski, and M. Cohen The lumigraph In *SIGGRAPH*, 1996.

[8] A. Leonardis and H. Bischof. Dealing with occlusions in the eigenspace approach. In *Proceedings of CVPR*, 1996.

[9] M. Levoy and M. Hanrahan. Light field rendering. In *Proc. of SIGGRAPH*, 1996.

[10] H. Murase and S. Nayar. Visual learning and recognition of 3-D objects from appearance. *Int. J. of Computer Vision*, 14:5–24, 1995.

[11] S. Nayar and S. Narasimhan. Vision in bad weather. In *Proc. 7th ICCV*, 1999.

[12] A. Pentland, B. Moghaddam, and T. Starner. View-based and modular eigenspaces for face recognition. In *Proceedings of CVPR*, 1994.

[13] P. J. Phillips, H. Wechsler, J. Huang, and P. Rauss. The FERET database and evaluation procedure for face recognition algorithms. *Image and Vision Computing J*, 16(5):295–306, 1998.

[14] T. Riklin-Raviv and A. Shashua. The Quotient image: Class based recognition and synthesis under varying illumination. In *Proc. of the IEEE Conf. on Comp. Vis. and Patt. Rec.*, 1999.

[15] T. Sim, S. Baker, and M. Bsat. The CMU pose, illumination, and expression (PIE) database. In *Proc. of the 5th IEEE International Conference on Automatic Face and Gesture Recognition*, 2002.

[16] M. Turk and A. Pentland. Face recognition using eigenfaces. In *Proc. of CVPR*, 1991.

[17] T. Vetter and T. Poggio. Linear object classes and image synthesis from a single example image. *IEEE Trans. on PAMI*, 19(7):733–741, 1997.