Blur and Illumination Robust Face Recognition via Set-Theoretic Characterization

Priyanka Vageeswaran, *Student Member, IEEE,* Kaushik Mitra, *Member, IEEE,* and Rama Chellappa, *Fellow, IEEE*

***Abstract*—We address the problem of unconstrained face recognition from remotely acquired images. The main factors that make this problem challenging are image degradation due to blur, and appearance variations due to illumination and pose. In this paper we address the problems of blur and illumination. We show that the set of all images obtained by blurring a given image forms a convex set. Based on this set-theoretic characterization, we propose a blur-robust algorithm whose main step involves solving simple convex optimization problems. We do not assume any parametric form for the blur kernels, however, if this information is available it can be easily incorporated into our algorithm. Further, using the low-dimensional model for illumination variations, we show that the set of all images obtained from a face image by blurring it and by changing the illumination conditions forms a bi-convex set. Based on this characterization we propose a blur and illumination-robust algorithm. Our experiments on a challenging real dataset obtained in uncontrolled settings illustrate the importance of jointly modeling blur and illumination.**

I. Introduction

FACE recognition has been an intensely researched field of computer vision for the past couple of decades [1]. Though significant strides have been made in tackling the problem in controlled domains (as in recognition of passport photographs) [1], significant challenges remain in solving it in the unconstrained domain. One such scenario occurs while recognizing faces acquired from distant cameras. The main factors that make this a challenging problem are image degradations due to blur and noise, and variations in appearance due to illumination and pose [2] (see Figure 1). In this paper, we specifically address the problem of recognizing faces across blur and illumination.

An obvious approach to recognizing blurred faces would be to deblur the image first and then recognize it using traditional face recognition techniques [3]. However, this approach involves solving the challenging problem of blind image deconvolution [4], [5]. We avoid this unnecessary step and propose a direct approach for face recognition. We show that the set of all images obtained by blurring a given image forms a convex set, and more specifically, we show that this set is the convex hull of shifted versions of the original image. Thus with each gallery image we can associate a corresponding convex set. Based on this set-theoretic characterization, we propose a blur-robust face recognition algorithm. In the basic version of our algorithm, we compute the distance of a given probe image (which we want to recognize) from each of the convex sets, and assign it the identity of the closest gallery image. The distance-computation steps are formulated as convex optimization problems over the space of blur kernels. We do not assume any parametric or symmetric form for the blur kernels; however, if this information is available, it can be easily incorporated into our algorithm, resulting in improved recognition performance. Further, we make our algorithm robust to outliers and small pixel mis-alignments by replacing the Euclidean distance by weighted L_1-norm distance and comparing the images in the LBP (local binary pattern) [6] space.

Fig. 1: *Face images captured by a distant camera in unconstrained settings.* The main challenges in recognizing such faces are variations due to blur, pose and illumination. In this paper we specifically address the problems of blur and illumination.

It has been shown in [7] and [8] that all the images of a Lambertian convex object, under all possible illumination conditions, lie on a low-dimensional (approximately nine-dimensional) linear subspace. Though faces are not exactly convex or Lambertian, they can be closely approximated by one. Thus each face can be characterized by a low-dimensional subspace, and this characterization has been used for designing illumination robust face recognition algorithms [7], [9]. Based on this illumination model, we show that the set of all images

Manuscript received April 30, 2012; revised August 29, 2012.

Priyanka Vageeswaran and Rama Chellappa are with the Department of Electrical and Computer Engineering, and the Center for Automation Research, UMIACS, University of Maryland, College Park.E-mail: {svpriyanka,rama}@umiacs.umd.edu

Kaushik Mitra is with the Department of Electrical and Computer Engineering, Rice University. E-mail: Kaushik.Mitra@rice.edu.

This work was partially supported by a MURI from the Office of Naval Research under the Grant N00014-08-1-0638.

of a face under all blur and illumination variations is a bi-convex one. That is- if we fix the blur kernel then the set of images obtained by varying the illumination conditions forms a convex set; and if we fix the illumination condition then the set of all blurred images is also convex. Based on this set-theoretic characterization, we propose a blur and illumination robust face recognition algorithm. The basic version of our algorithm computes the distance of a given probe image from each of the bi-convex sets, and assigns it the identity of the closest gallery image. The distance computations steps can be formulated as 'quadratically constrained quadratic programs' (QCQPs), which we solve by alternately optimizing over the blur kernels and the illumination coefficients. Similar to the blur-only case, we make our algorithm robust to outliers and small pixel mis-alignments by replacing the Euclidean norm by the weighted L_1-norm distance and comparing the images in the LBP space.

To summarize, the main technical contributions of this paper are:

- We show that the set of all images obtained by blurring a given image forms a convex set. More specifically, we show that this set is the convex hull of shifted versions of the original image.
- Based on this set-theoretic characterization, we propose a blur-robust face recognition algorithm, which avoids solving the challenging and unnecessary problem of blind image deconvolution.
- If we have additional information on the type of blur affecting the probe image, we can easily incorporate this knowledge into our algorithm, resulting in improved recognition performance and speed.
- We show that the set of all images of a face under all blur and illumination variations forms a bi-convex set. Based on this characterization, we propose a blur and illumination robust face recognition algorithm.

A. Related Work

Face recognition from blurred images can be classified into four major approaches. In the first approach, the blurred image is first deblurred and then used for recognition. This is the approach taken in [10] and [3]. The drawback of this approach is that we first need to solve the challenging problem of blind image deconvolution. Though there have been many attempts at solving the blind deconvolution problem [11], [4], [12], [13], [5], it is an avoidable step for the face recognition problem. Also, in [3] statistical models are learned for each blur kernel type and amount; this step might become infeasible when we try to capture the complete space of blur kernels.

In the second approach, blur invariant features are extracted from the blurred image and then used for recognition; [14] and [15] follow this approach. In [14], the local phase quantization (LPQ) [16] method is used to extract blur invariant features. Though this approach works very well for small blurs, it is not very effective for large blurs [3]. In [15], a (blur) subspace is associated with each image and face recognition is performed in this feature space. It has been shown that the (blur) subspace of an image contains all the blurred version of the image. However, this analysis does not take into account the convexity constraint that the blur kernels satisfy, and hence the (blur) subspace will include many other images apart from the blurred images. The third approach is the direct recognition approach. This is the approach taken in [17] and by us. In [17], artificially blurred versions of the gallery images are created and the blurred probe image is matched to them. Again, it is not possible to capture the whole space of blur kernels using this method. We avoid this problem by optimizing over the space of blur kernels. Finally, the fourth approach is to jointly deblur and recognition the face image [18]. However, this involves solving for the original sharp image, blur kernel and identity of the face image, and hence it is a computationally intensive approach.

Set theoretic approaches for signal and image restoration have been considered in [19], [20], [21]. In these approaches the desired signal space is defined as an intersection of closed convex sets in a Hilbert space, with each set representing a signal constraint. Image de-blurring has also been considered in this context [20], where the non-negativity constraint of the images has been used to restrict the solution space. We differ from these approaches as our primary interest lies in recognizing blurred and poorly illuminated faces rather than restoring them.

There are mainly two approaches for recognizing faces across illumination variation. One approach is based on the low-dimensional linear subspace model [7], [8]. In this approach, each face is characterized by its corresponding low-dimensional subspace. Given a probe image, its distance is computed from each of the subspaces, and it is then assigned to the face image with the smallest distance [7], [9]. The other approach is based on extracting illumination insensitive features from the face image and using them for matching. Many features have been proposed for this purpose such as self-quotient images [22], correleration filters [23], Eigenphases method [24], image preprocessing algorithms [25], gradient direction [26], [27] and albedo estimates [28].

The organization of the rest of the paper is as follows: In section II we provide a set-theoretic characterization of the space of blurred images and subsequently propose our approach for recognizing blurred faces, in section III we incorporate the illumination model in our approach and in section IV we perform experiments to evaluate the efficacy of our approach on many synthetic and real datasets.

II. Direct Recognition of Blurred Faces (DRBF)

We first review the convolution model for blur. Next, we show that the set of all images obtained by blurring a given image is convex and finally we present our algorithm for recognizing blurred faces.

A. Convolution Model for Blur

A pixel in a blurred image is a weighted average of the pixel's neighborhood in the original sharp image. Thus, blur is modeled as a convolution operation between the original image and a blur filter kernel which represents the weights [29]. Let I be the original image and H be the blur kernel of

size $(2k+1) \times (2k+1)$, then the blurred image I_b is given by

$$I_b(r,c) = I * H(r,c) = \sum_{i=-k}^{k} \sum_{j=-k}^{k} H(i,j) I(r-i, c-j) \quad (1)$$

where $*$ represents the convolution operator and r, c are the row and column indices of the image. Blur kernels also satisfy the following properties- their coefficients are non-negative, $H \geq 0$, and sum up to 1 (i.e. $\sum_{i=-k}^{k} \sum_{j=-k}^{k} H(i,j) = 1$). The blur kernel may possess additional structure depending on the type of blur (such as circular-symmetry for out-of-focus blurs), and these structures could be exploited during recognition.

B. The Set of All Blurred Images

We want to characterize the set of all images obtained by blurring a given image I. To do this we re-write (1) in a matrix-vector form. Let $\mathbf{h} \in \mathbb{R}^{(2k+1)^2}$ be the vector obtained by concatenating the columns of H, i.e., $\mathbf{h} = H(:)$ in MATLAB notation, and similarly $\mathbf{i_b} = I_b(:) \in \mathbb{R}^N$ be the representation of I_b in the vector form, where N is the number of pixels in the blurred image . Then we can write (1), along with the blur kernel constraints, as

$$\mathbf{i_b} = \mathbf{A}\mathbf{h} \text{ such that } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1 \quad (2)$$

where $\mathbf{A}$ is a $N \times (2k+1)^2$ matrix, obtained from I, with each row of $\mathbf{A}$ representing the neighborhood pixel intensities about the pixel indexed by the row. From 2, it is clear that the set of all blurred images obtained from I is given by

$$\mathcal{B} \triangleq \{\mathbf{A}\mathbf{h} | \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1\} \quad (3)$$

We have the following result about the set $\mathcal{B}$.

Proposition II.1. *The set of all images $\mathcal{B}$ obtained by blurring an image I is a convex set. Moreover, this convex set is given by the convex hull of the columns of matrix $\mathbf{A}$, where the columns of $\mathbf{A}$ are various shifted versions of I as determined by the blur kernel.*

Proof: Let $\mathbf{i_1}$ and $\mathbf{i_2}$ be elements from the set $\mathcal{B}$. Then there exists $\mathbf{h_1}$ and $\mathbf{h_2}$, with both satisfying the conditions $\mathbf{h} \geq 0$ and $\|\mathbf{h}\|_1 = 1$, such that $\mathbf{i_1} = \mathbf{A}\mathbf{h_1}$ and $\mathbf{i_2} = \mathbf{A}\mathbf{h_2}$. To show that the set $\mathcal{B}$ is convex we need to show that for any λ satisfying $0 \leq \lambda \leq 1$, $\mathbf{i_3} = \lambda\mathbf{i_1} + (1-\lambda)\mathbf{i_2}$ is an element of $\mathcal{B}$. Now

$$\begin{aligned} \mathbf{i_3} &= \lambda\mathbf{i_1} + (1-\lambda)\mathbf{i_2} \\ &= \mathbf{A}(\lambda\mathbf{h_1} + (1-\lambda)\mathbf{h_2}) \\ &= \mathbf{A}\mathbf{h_3}. \end{aligned} \quad (4)$$

Note that $\mathbf{h_3}$ satisfies both the non-negativity and sum conditions and hence $\mathbf{i_3}$ is an element of $\mathcal{B}$. Thus, $\mathcal{B}$ is a convex set. $\mathcal{B}$ is defined as

$$\{\mathbf{A}\mathbf{h} | \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1\}, \quad (5)$$

which, by definition, is the convex hull of the columns of $\mathbf{A}$. ■

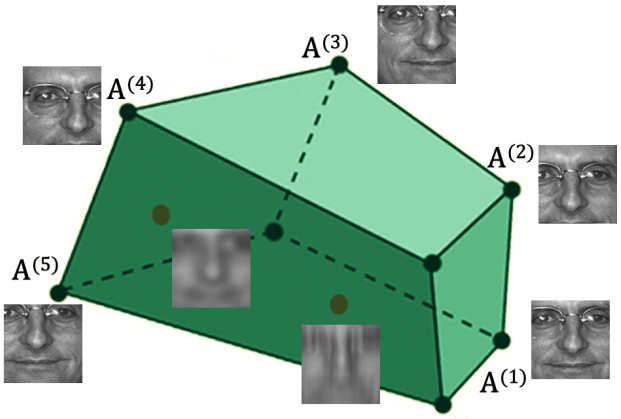

Fig. 2: *The set of all images obtained by blurring an image I is a convex set.* Moreover, this convex set is given by the convex hull of the columns of the matrix $\mathbf{A}$, which represents the various shifted versions of I as determined by the blur kernel.

C. A Geometric Face Recognition algorithm

We first present the basic version of our blur-robust face recognition algorithm. Let $I_j, j = 1, 2, \ldots, M$ be the set of M sharp gallery images. From the analysis above, every gallery image I_j has an associated convex set of blurred images $\mathcal{B}_j$. Given the probe image I_b, we find its distance from the set $\mathcal{B}_j$, which is the minimum distance between I_b and the points in the set $\mathcal{B}_j$. This distance r_j can be obtained by solving:

$$r_j = \min_{\mathbf{h}} ||\mathbf{i_b} - \mathbf{A_j}\mathbf{h}||_2^2 \text{ subject to } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1 \quad (6)$$

This is a convex quadratic program which can be solved efficiently. For $\mathbf{i_b} \in \mathbb{R}^N$ and $\mathbf{h} \in \mathbb{R}^K$, the computational complexity is $\mathcal{O}(NK^2)$. We compute r_j for each $j = 1, 2, \ldots, M$ and assign I_b the identity of the gallery image with the minimum r_j. If there are multiple gallery images per class (person), we can use the k-nearest neighbor rule, i.e. we arrange the $r_j's$ in ascending order and find the class which appears the most in the first k instances. In this algorithm we can also incorporate additional information about the type of blur. The most commonly occurring blur types are the out-of-focus, motion and the atmospheric blurs [29]. The out-of-focus and the atmospheric blurs are circularly-symmetric, i.e. the coefficients of H at the same radius are equal; whereas the motion blur is symmetric about the origin, i.e. $H(i,j) = H(-i,-j)$ [29]. Thus, having knowledge of the blur type, we solve (6) with an additional constraint on the blur kernel:

$$\begin{aligned} r_j &= \min_{\mathbf{h}} ||\mathbf{i_b} - \mathbf{A_j}\mathbf{h}||_2^2 \\ &\quad \text{subject to } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1, \mathcal{C}(\mathbf{h}) = 0, \end{aligned} \quad (7)$$

where $\mathcal{C}(\mathbf{h}) = 0$ represents equality constraints on $\mathbf{h}$. Imposing these constraints reduces the number of parameters in the optimization problem giving better recognition accuracy and faster solutions.

Algorithm *Direct Recognition of Blurred Faces*
Input: (Blurred) probe image I_b and a set of gallery images I_j
Output: Identity of the probe image
1. For each gallery image I_j, find the optimal blur kernel $\mathbf{h_j}$ by solving either (9) or its robust version (10).
2. Blur each gallery image I_j with its corresponding $\mathbf{h_j}$ and extract LBP features.
3. Compare the LBP features of the probe image I_b with those of the gallery images and find the closest match.

Fig. 3: *Direct Recognition of Blurred Faces (DRBF/rDRBF) Algorithm:* Our proposed algorithm for recognizing blurred faces.

D. Making the Algorithm Robust to Outliers and Misalignment

By making some minor modifications to the basic algorithm, we can make it robust to outliers and small pixel misalignments between the gallery and probe images. It is well known in face recognition literature [30] that different regions in the face have different amounts of information. To incorporate this fact we divide the face image into different regions and weigh them differently when computing the distance between the probe image I_b and gallery sets $\mathcal{B}_j$. That is, we modify the distance functions r_j as

$$\begin{aligned} r_j &= \min_{\mathbf{h}} ||W(\mathbf{i_b} - \mathbf{A_j h})||_2^2 \\ &\quad \text{subject to } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1, \mathcal{C}(\mathbf{h}) = 0. \end{aligned} \quad (8)$$

We learn the weight W, a diagonal matrix, using a training dataset. The training procedure is described in the APPENDIX.

Face recognition is also sensitive to small pixel misalignments and, hence, the general consensus in face recognition literature is to extract alignment insensitive features, such as Local Binary Patterns (LBP) [6], [14], and then perform recognition based on these features. Following this convention, instead of doing recognition directly from r_j, we first compute the optimal blur kernel $\mathbf{h_j}$ for each gallery image by solving (8), i.e.

$$\begin{aligned} \mathbf{h_j} &= \arg\min_{\mathbf{h}} ||W(\mathbf{i_b} - \mathbf{A_j h})||_2^2 \\ &\quad \text{subject to } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1, \mathcal{C}(\mathbf{h}) = 0. \end{aligned} \quad (9)$$

We then blur each of the gallery images with the corresponding optimal blur kernels $\mathbf{h_j}$ and extract LBP features from the blurred gallery images. And finally, we compare the LBP features of the probe image with those of the gallery images to find the closest match.

To make our algorithm robust to outliers, which could arise due to variations in expression, we propose to replace the L_2 norm in (9) by the L_1 norm, i.e. we solve the problem:

$$\begin{aligned} \mathbf{h_j} &= \arg\min_{\mathbf{h}} ||W(\mathbf{i_b} - \mathbf{A_j h})||_1 \\ &\quad \text{subject to } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1, \mathcal{C}(\mathbf{h}) = 0. \end{aligned} \quad (10)$$

Note that the above optimization problem is a convex L_1-norm problem, which we formulate and solve as a Linear Programing (LP) problem. The computational complexity of this problem is $\mathcal{O}((K+N)^3)$. The overall algorithm is summarized in Figure 3.

III. Incorporating the Illumination Model

The facial images of a person under different illumination conditions can look very different, and hence for any recognition algorithm to work in practice, it must account for these variations. First, we discuss the low-dimensional subspace model for handling appearance variations due to illumination. Next, we use this model along with the convolution model to define the set of images of a face under all possible lighting conditions and blur. We then propose a recognition algorithm based on minimizing the distance of the probe image from such sets.

A. The Low-Dimensional Linear Model for Illumination Variations

It has been shown in [7], [8] that when an object is convex and Lambertian, the set of all images of the object under different illumination conditions can be approximately represented using a nine-dimensional subspace. Though the human face is not exactly convex or Lambertian, it is often approximated as one; and hence the nine-dimensional subspace model captures its variations due to illumination quite well [31]. The nine-dimensional linear subspace corresponding to a face image I can be characterized by 9 basis images. In terms of these nine basis images $I_m, m = 1, 2, \ldots, 9$, an image I of a person under any illumination condition can be written as

$$I = \sum_{m=1}^{9} \alpha_m I_m \quad (11)$$

where $\alpha_m, m = 1, 2, \ldots, 9$ are the corresponding linear coefficients. To obtain these basis images, we use the "universal configuration" of lighting positions proposed in [9]. These are a set of 9 lighting positions $s_m, m = 1, 2, \ldots, 9$ such that images taken under these lighting positions can serve as basis images for the subspace. These basis images are generated using the Lambertian reflectance model:

$$I_m(r, c) = \rho(r, c) \max(\langle s_m, n(r, c) \rangle, 0) \quad (12)$$

where $\rho(r, c)$ and $n(r, c)$ are the albedo and surface-normal at pixel location (r, c). We use the average 3-D face normals from [32] for n and we approximate the albedo ρ with a well-illuminated gallery image under diffuse lighting. In the absence of a well-illuminated gallery image, we could proceed by estimating the albedo from a poorly lit image using approaches presented in [28],[33] and [34].

B. The set of all images under varying lighting and blur

For a given face characterized by the nine basis images $I_m, m = 1, 2, \ldots, 9$, the set of images under all possible lighting conditions and blur is given by

$$\mathcal{BI} \triangleq \{ \sum_{m=1}^{9} \alpha_m \mathbf{A_m h} | \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1, \mathcal{C}(\mathbf{h}) = 0 \}, \quad (13)$$

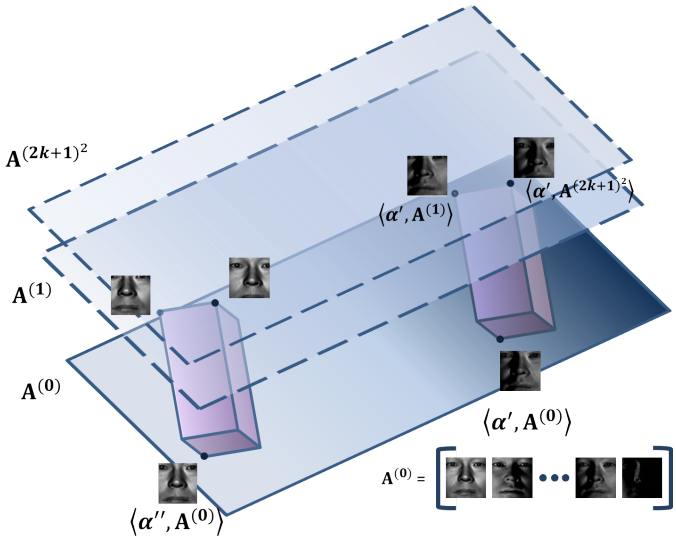

Fig. 4: *The set of all images under varying lighting and blur for a single face image.* This set is a bi-convex set, i.e. if we fix either the filter kernel $\mathbf{h}$ or the illumination condition α, the resulting subset is convex. Each hyperplane in the figure represents the illumination subspace at different blur. For example, all points on the bottom-most plane are obtained by fixing the blur kernel at $\mathbf{h}^{(0)}$ (the impulse function centered at 0, i.e. the no-blur case) and varying the illumination conditions α. On this plane two data-points (faces), corresponding to illumination conditions α' and α'', are explicitly marked. Both these data points are associated with their corresponding blur convex hulls, see figure 2.

Algorithm *Illumination-robust Recognition of Blurred Faces*
Input: (Blurred and poorly illuminated) probe image I_b and a set of gallery images I_j
Output: Identity of the probe image

1. For each gallery image I_j, obtain the nine basis images $I_{j,m}, m = 1, 2, \ldots, 9$.
2. For each gallery image I_j, find the optimal blur kernel $\mathbf{h_j}$ and illumination coefficients $\alpha_{j,m}$ by solving either (14) or its robust version (15).
3. Transform (blur and re-illuminate) the gallery images I_j using the computed $\mathbf{h_j}$ and $\alpha_{j,m}$ and extract LBP features.
4. Compare the LBP features of the probe image I_b with those of the transformed gallery images and find the closest match.

Fig. 5: *Illumination-Robust Recognition of Blurred Faces (IRBF/rIRBF) Algorithm:* Our proposed algorithm for jointly handling variations due to illumination and blur.

where the matrix $\mathbf{A_m}$ is constructed from I_m and represents the pixel neighborhood structure. This set is not a convex set though if we fix either the filter kernel $\mathbf{h}$ or the illumination condition α_m the set becomes convex, see figure 4.

C. Illumination-robust Recognition of Blurred Faces (IRBF)

Corresponding to each sharp well-lit gallery image $I_j, j = 1, 2, \ldots, M$, we obtain the nine basis images $I_{j,m}, m = 1, 2, \ldots, 9$. Given the vectorized probe image $\mathbf{i_b}$, for each gallery image I_j we find the optimal blur kernel $\mathbf{h_j}$ and illumination coefficients $\alpha_{j,m}$ by solving:

$$[\mathbf{h_j}, \alpha_{j,m}] = \arg\min_{\mathbf{h},\alpha_m} \|W(\mathbf{i_b} - \sum_{m=1}^{9} \alpha_m \mathbf{A_{j,m}} \mathbf{h})\|_2^2$$
$$\text{subject to } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1, \mathcal{C}(\mathbf{h}) = 0. \quad (14)$$

We then transform (blur and re-illuminate) each of the gallery images I_j using the computed blur kernel $\mathbf{h_j}$ and the illumination coefficients $\alpha_{j,m}$. Next, we compute the LBP features from these transformed gallery images and compare it with those from the probe image I_b to find the closest match, see Figure 5. The major computational step of the algorithm is the optimization problem of (14), which is a non-convex problem. To solve this problem we use an alternation algorithm in which we alternately minimize over $\mathbf{h}$ and α_m, i.e. in one step we minimize over $\mathbf{h}$ keeping α_m fixed and in the other step we minimize over α_m keeping $\mathbf{h}$ fixed and we iterate till convergence. Each step is now a convex problem: the optimization over $\mathbf{h}$ for fixed α_m reduces to the same problem as (9) and the optimization of α given $\mathbf{h}$ is just a linear least squares problem. The complexity of the overall alternation algorithm is $\mathcal{O}(T(N + K^3))$ where T is the number of iterations in the alternation step, and $\mathcal{O}(N)$ is the complexity in the estimation of the illumination coefficients. We also propose a robust version of the algorithm by replacing the L_2-norm in (14) with the L_1-norm:

$$[\mathbf{h_j}, \alpha_{j,m}] = \arg\min_{\mathbf{h},\alpha_m} \|W(\mathbf{i_b} - \sum_{m=1}^{9} \alpha_m \mathbf{A_{j,m}} \mathbf{h})\|_1$$
$$\text{subject to } \mathbf{h} \geq 0, \|\mathbf{h}\|_1 = 1, \mathcal{C}(\mathbf{h}) = 0. \quad (15)$$

Again, this is a non-convex problem and we use the alternation procedure which reduces each step of the algorithm to a convex L_1-norm problem. We formulate these L_1-norm problems as Linear Programing (LP) problems. The complexity of the overall alternation algorithm is $\mathcal{O}(T(N^3 + (K + N)^3))$. The algorithm is summarized in Figure 5.

IV. Experimental Evaluations

We evaluate the proposed algorithms: the 'blur-only' formulation DRBF of section II and the 'blur and illumination' formulation IRBF of section III on synthetically blurred datasets-FERET [35] and PIE [36], and a real dataset of remotely acquired faces with significant blur and illumination variations [2]- which we will refer to as the REMOTE dataset, see Figure 1. In section IV-A, we evaluate the performance of the DRBF algorithm in recognizing faces blurred by different types and amounts of blur. In section IV-B, we evaluate the effectiveness of the IRBF algorithm in recognizing blurred and poorly illuminated faces. Finally, in section IV-A, we evaluate our algorithms, DRBF and IRBF, on the real and challenging dataset of REMOTE.

A. Face Recognition across Blur

To evaluate our algorithm DRBF on different types and amounts of blur, we synthetically blur face images from the

FERET dataset with four different types of blurs: out-of-focus, atmospheric, motion and general non-parametric blur. We use Gaussian kernels of varying standard deviations to approximate the out-of focus and atmospheric blurs [29], and rectangular kernels with varying lengths and angles for the motion blur. For the general blur we use the blur kernels used in [5]. Figure 6 shows some of the blur kernels and the corresponding blurred images. We compare our algorithm with the FADEIN approach [3] and the LPQ approach [16]. As discussed in section I-A, the FADEIN approach first infers the deblurred image from the blurred probe image and then uses it for face recognition. On the other hand in the LPQ (local phase quantization) approach a blur insensitive image descriptor is extracted from the blurred image and recognition is done on this feature space. We also compare our algorithms with 'FADEIN+LPQ' [3], where LPQ features extracted from the deblurred image produced by FADEIN is used for recognition.

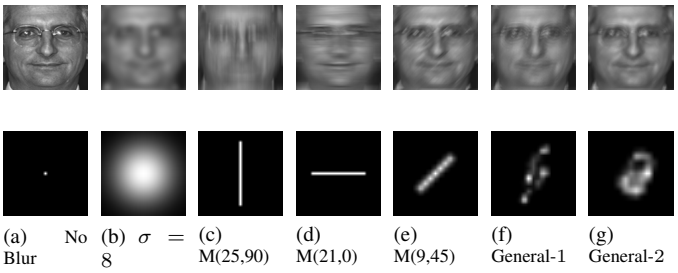

Fig. 6: Examples of blur kernels and images used to evaluate our algorithms. The *General* blurs shown above have been borrowed from [5]

1) Out-of-Focus and Atmospheric Blurs: We synthetically generate face images from the FERET dataset using Gaussian kernels of various standard deviations for evaluation. We use the same experimental set-up as used in FADEIN [3], i.e. we chose our gallery set as 1001 individuals from the *fa* folder of the FERET dataset. The gallery set so constructed has one face image per person and the images are frontal and well-illuminated. We construct the probe set by blurring images of the same set of 1001 individuals from the *fb* folder of FERET (the images in this folder has slightly different expressions from the fa folder). We blur each individual image by Gaussian kernels of σ values $0, 2, 4$, and 8 and kernel size $4\sigma + 1$.

To handle small variations in illumination, we histogram-equalize all the images in the gallery and probe datasets. We then perform recognition using the DRBF algorithm and its robust (L_1) version rDRBF, with the additional constraint of circular symmetry imposed on the blur kernel. Figure 7 shows the recognition results obtained using the above approach along-side the recognition results from the FADEIN, LPQ and FADEIN+LPQ algorithms. Our algorithms-DRBF and its robust version rDRBF, show significant improvement over the other algorithms, especially for large blurs. rDRBF performs even better than DRBF owing to the more robust modeling of expressions and misalignment, as shown in Figure 8

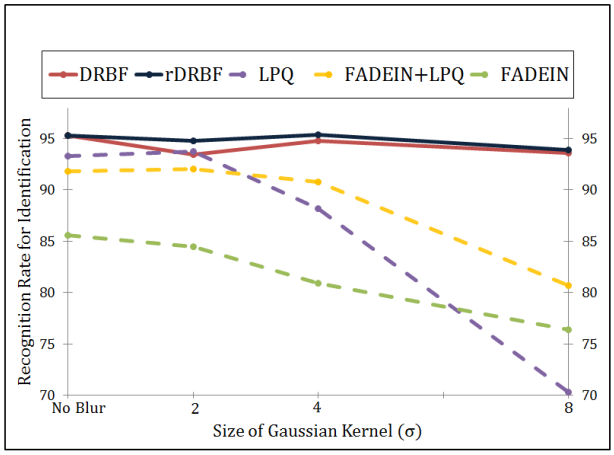

Fig. 7: *Face recognition across Gaussian blur.* Recognition results by different algorithms as the amount of Gaussian blur is varied. Our algorithms, DRBF and its robust (L_1-norm) version rDRBF, shows significant improvement over the algorithms FADEIN, LPQ and FADEIN+LPQ, especially, for large blurs.

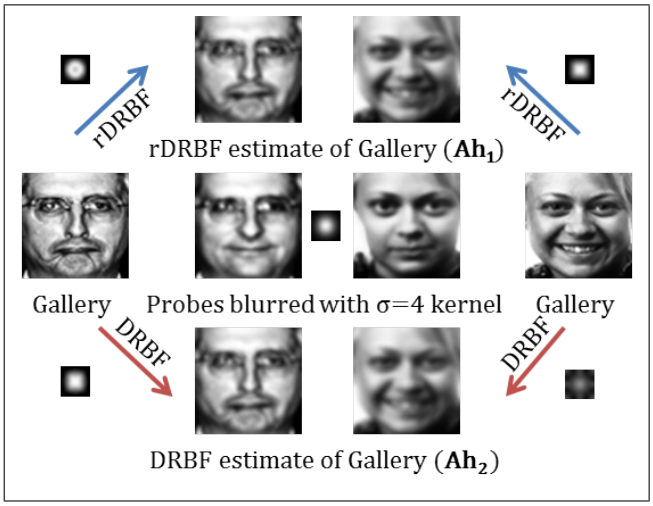

Fig. 8: *Comparison of DRBF with its robust version rDBRF*: The robust version rDBRF can handle outliers, such as those due to expression variations, more effectively. Two gallery images along with their corresponding probes are shown in the center row. The probes have been blurred by a Gaussian blur of $\sigma = 4$. Note that the probe images have a different expression than the gallery images. The blur kernels estimated by the two algorithms rDRBF and DRBF are shown on the top and bottom rows respectively. As can be seen from the figure, the kernels estimated by rDBRF is closer to the actual kernel (at the center). The gallery images blurred by the estimated kernels further illustrate this fact, as the blurred gallery on the bottom row (corresponding to DBRF) looks significantly more blurred than the blurred gallery images in the top row (corresponding to rDBRF).

2) Motion and General Blurs: We now demonstrate the effectiveness of our algorithms, DRBF and rDRBF, on datasets degraded by motion and general non-parametric blurs. For this experiment we use the *ba* and *bj* folders in FERET, both of which contains 200 subjects with one image per subject. We use the *ba* folder as the gallery set. The probe set is formed by blurring the images in the *bj* folder by different motion and general blur kernels, some of them are shown in Figure 6. When we perform recognition using DRBF and rDRBF, we impose appropriate symmetry constraints for the blur types. That is, when we solve for the motion blur case, we impose the 'symmetry about the origin' constraint on the blur kernel, whereas, when we solve for the general or non-parametric blur case we do not impose any constraint. Figure 9 shows that DRBF and rDRBF perform consistently better than LPQ and FADEIN+LPQ. Hence, we can say that our method generalizes well to all forms of blur.

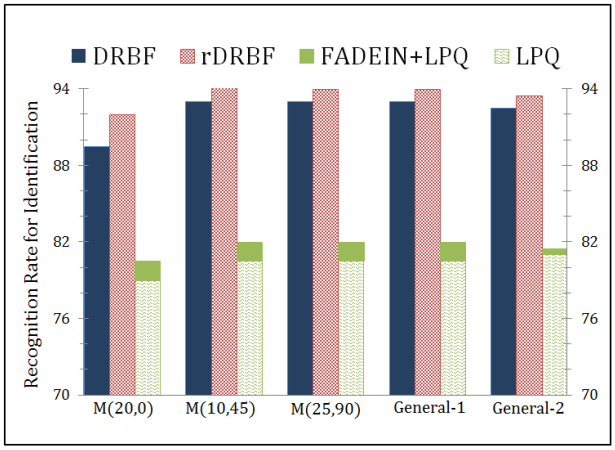

(a)

Fig. 9: *Recognition result for motion and general blurs:* Performance of different algorithms on some selected motion and non-parametric blurs, see Figure 6. Our algorithms, DRBF and rDRBF, perform much better than LPQ and FADEIN+LPQ, with the robust version rDRBF always better than DRBF.

3) Effect of Blur Kernel-Size and Symmetry Constraints On DRBF: In all the experiments described above we have assumed that we know the type and size of the blur kernel, and have used this information while estimating the blur kernel in (9) or (10). For example, for images blurred by a Gaussian blur of standard deviation σ, we impose a kernel-size of $4\sigma+1$ and circular symmetry. Though in some applications we may know the type of blur or the amount of blur, it may not be known for all applications. Hence, to test the sensitivity of our algorithm to blur kernel-size and blur type (symmetry constraint), we perform a few experiments.

We use the *ba* folder of FERET as the gallery set and we create the probe set by blurring the images in the *bj* folder by a Gaussian kernel of $\sigma = 4$ and size $4\sigma + 1 = 17$. We then perform recognition via DRBF with choices of kernel size ranging from 1 to $32\sigma+1$. We consider both the cases of imposing the circular symmetry constraint and not imposing any constraint. The experimental results are shown in Figure

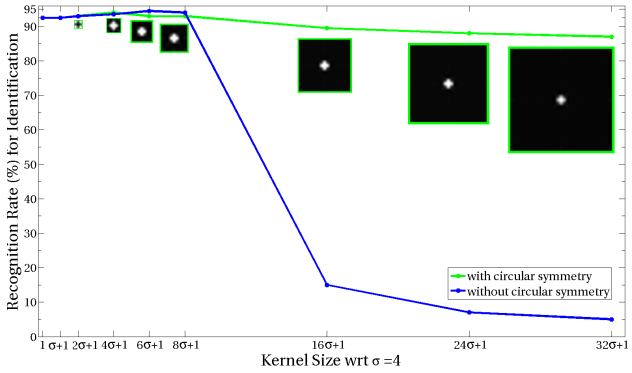

Fig. 10: *Effect of blur kernel size and symmetry constraints on DRBF.* For this experiment, we use probe images blurred by a Gaussian kernel of $\sigma = 4$ and size $4\sigma + 1 = 17$ and perform recognition using DRBF with choices of kernel size ranging from 1 to $32\sigma + 1$. When we impose appropriate symmetry constraints, the recognition rate remains high even when we over-estimate the kernel size by a large margin (we have also shown the estimated blur kernels). This is because imposing symmetry constraints reduces the solution space, and makes it a more regularized problem. On the other hand when no symmetry constraints are imposed, the recognition rate falls drastically after a certain kernel size. However, as long as we do not over-estimate the kernel size by a large margin, we can expect a good performance from the algorithm.

10. We can see that the recognition rates are fairly stable for the case when we impose appropriate symmetry constraints. This is because imposing symmetry constraints reduces the solution space and makes it a more regularized problem. We also show the mean estimated kernels, from which it is clear that the algorithm works well even for large kernel sizes. On the other hand when no symmetry constraints are imposed, the recognition rate falls drastically after a certain kernel size. However, as long as we do not over-estimate the kernel size by a large margin, we can expect good results from the algorithm. We conclude from these experiments that: 1) our algorithm exhibits a stable performance for a wide range of kernel-sizes as long as we do not over-estimate them by a large margin, 2) it is better to under-estimate the kernel size than over-estimate it and 3)if we know the blur type then we should impose the corresponding symmetry constraints because imposing them further relaxes the need for an accurate estimate of the kernel-size.

B. Recognition across Blur and Illumination

We study the effectiveness of our algorithms in recognizing blurred and poorly illuminated faces. We use the PIE dataset, which consists of images of 68 individuals under different illumination conditions. To study the effect of blur and illumination together, we synthetically blur the images with Gaussian kernels of varying σ's. We use face images with a frontal pose (c_{27}) and good illumination (f_{21}) as our gallery and the rest of the images in c_{27} as probe. We further divide the probe dataset into two categories- 1) Good Illumination (*GI*) consisting of subsets f_{09}, f_{11}, f_{12} and f_{20} and 2) Bad Illumination (*BI*) consisting of $f_{13}, f_{14}, f_{15}, f_{16}, f_{17}$ and f_{22}, see

Kernel Size(σ)	0		0.5		1.0		1.5		2.0		2.5		3.0	
Illumination	GI	BI	GI	BI	GI	BI	GI	BI	GI	BI	GI	BI	GI	BI
DRBF	99.63	95.10	99.63	84.55	**99.63**	78.67	**99.63**	77.95	**99.63**	77.45	97.79	58.58	95.58	42.40
IRBF	99.63	93.56	99.63	91.42	**99.63**	90.44	**99.63**	90.68	**99.63**	85.78	98.9	81.13	96.32	77.69
rIRBF	**99.7**	95.1	**99.7**	92.7	**99.63**	92.7	**99.63**	**91.6**	**99.63**	**88.2**	**99.63**	**84.78**	**97.45**	**81.36**
LPQ	99.63	**99.1**	99.63	**97.79**	**99.63**	**96.08**	**99.63**	88.97	97.05	73.04	79.42	58.08	46.32	27.7
FADEIN+LPQ	98.53	91.5	95.6	87.7	93.6	81.8	91.2	69.11	89.8	62.74	88.60	56.37	87.13	44.61

TABLE I: *Recognition across Blur and Illumination on the PIE dataset*. GI and BI represent the 'good illumination' and 'bad illumination' subsets of the probe-set. IRBF and its robust version rIRBF out-perform the other algorithms for blurs of sizes greater than $\sigma = 1$. LPQ performs quite well for small blurs, but for large blurs its performance degrades significantly. This experiment clearly validates the need for modeling illumination and blur in a principled manner.

(a) *Good Illumination*(GI) (b) *Bad Illumination* (BI) (c) *Illumination Basis Images*

Fig. 11: To study the effect of blur and illumination we use the PIE dataset which shows significant variation due to illumination. For each of the 68 subjects in PIE, we choose a well illuminated and frontal image of the person as the gallery set. The probe set, which is obtained from all the other frontal images, is divided into two categories: 1) Good Illumination (*GI*) set and 2) Bad Illumination (*BI*) set. Figures 11(a) and (b) shows some images from the *GI* and *BI* sets respectively. Figure 11(c) shows the 9 illumination basis images generated from a gallery image.

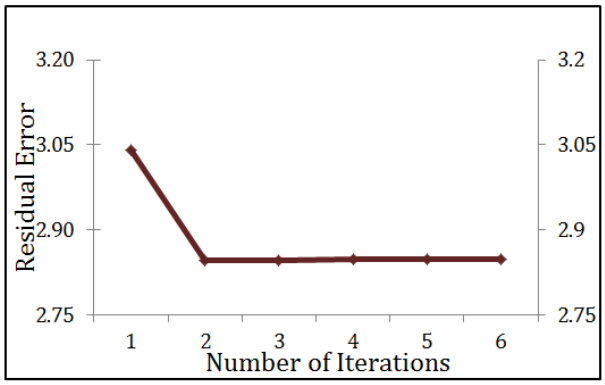

Fig. 12: *Convergence of the IRBF algorithm*- Note that the IRBF algorithm minimizes a bi-convex function which, in general, is a non-convex problem. However, since we alternately optimize over the blur kernel and illumination coefficients, we are guaranteed to converge to a local minimum. The plot shows the average convergence behavior of the algorithm. Based on this plot, we terminate the algorithm after six iterations.

Figure 11. We then blur all the probe images with Gaussian blurs of σ $0, 0.5, 1.0, 1.5, 2.0, 2.5$ and 3.

To perform recognition using our 'blur and illumination' algorithm IRBF, we first obtain the nine illumination basis images for each gallery image as described in section III-A. We impose the circular symmetry constraints while solving the recognition problem using DRBF and IRBF. For comparison, we use LPQ and a modified version of FADEIN+LPQ. Since FADEIN does not model variations due to illumination, we preprocess the intensity images with the self-quotient method [22] and then run the algorithm. Table I shows the recognition results for the algorithms. We see that our algorithms IRBF and rIRBF out-perform the comparison algorithms for blurs of sizes $\sigma = 1.5$ and greater. Moreover, with a 8-core 2.9GHz procesor running MATLAB, it takes us $2.28s$, $9.53s$ and $17.74s$ per query image with a blur of $\sigma = 2$, for a gallery of 68 images for the DRBF, IRBF and rIRBF algorithms respectively. Thus we can conclude that our algorithms are able to maintain a consistent performance across increasing blur in a reasonable amount of time.

As discussed is section III-C, the main optimization step in the IRBF (14) and rIRBF (15) is a bi-convex problem, i.e. it is convex w.r.t. to blur and illumination variables individually, but it is not jointly convex. Thus, the global optimality of the solution is not guaranteed. However, since we alternately optimize over the blur kernel and illumination coefficients, we are guaranteed to converge to a local minimum. Figure 12 plots the average residual error of the cost function in (14) with increasing number of iterations. Note that the algorithm converges in a few iterations. Based on this plot, in our experiments, we terminate the algorithm after six iterations.

C. Recognition in Unconstrained Settings

Finally, we report recognition experiments on the REMOTE dataset where the images have been captured in an unconstrained manner [2], [37]. The images were captured in two settings: from ship-to-shore and from shore-to-ship. The distance between the camera and the subjects ranges from 5meters to 250 meters. Hence, the images suffer from varying amounts of blur, variations in illumination and pose, and even some occlusion. This dataset has 17 subjects. We design our gallery set to have one frontal, sharp and well-illuminated image. The probe is manually partitioned into three categories: 1) the *illum* folder containing 561 images with illumination variations as the main problem, 2) the *blur* folder containing 75 images with blur as the main problem and 3) the *illum-blur* folder containing 128 images with both

(a) *Sample Images* (b) *Basis Images*

Fig. 13: Examples from the REMOTE dataset showing three probe partitions in subplot (a). The top row shows images from the *Illum* folder where variations in illumination are the only problem. The middle and bottom rows have images which exhibit variations in blur alone (from the *Blur* folder), and variations in blur and illumination (from the *Illum-Blur* folder) respectively. Subplot (b) shows the basis images generated from a gallery.

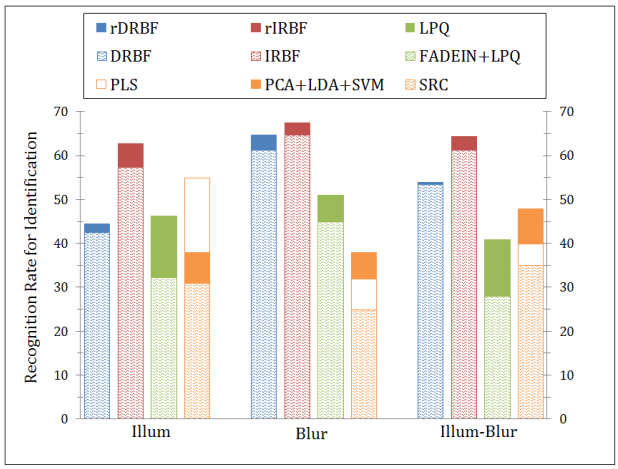

Fig. 14: *Recognition results on the unconstrained dataset REMOTE.* We compare our algorithms, DRBF and IRBF, with LPQ, modified FADEIN+LPQ, a sparse representation-based face recognition algorithm [38] (SRC), PCA+LDA+SVM [2] and a PLS-based (Partial least squares) face recognition algorithm [37]. The good performance by rIRBF and IRBF further confirms the importance of jointly modeling blur and illumination variations.

problems, see Figure 13(a). All three subsets contain near-frontal images of the 17 subjects, as set in the protocol in [37]. We register the images as a pre-processing step and normalize the size of the images to 120×120 pixels. We then run our algorithms-DRBF and IRBF, on the dataset. We assume symmetry about the origin as most of the blur arises due to out-of-focus, atmospheric and motion blur, all of which satisfy this symmetry constraint. For the *illum* folder we assume a blur kernel size of 5, and for the other two folders we assume kernel size of 7. For the IRBF and rIRBF algorithms, we generate the 9 illumination basis images for each image in the gallery, see Figure 13(b). We compare our algorithm with LPQ, modified FADEIN+LPQ (as described in the previous section IV-B). Apart from these algorithms, we also compare our algorithms with the algorithms presented in [37]. These algorithms are- sparse representation based face recognition algorithm [38] (SRC), PCA+LDA+SVM [2] and a PLS-based (Partial least squares) face recognition algorithm [37]. The results are shown in Figure 14. The good performances by rIRBF and IRBF further confirms the importance of jointly modeling blur and illumination variations.

V. Conclusion and Discussion

Motivated by the problem of remote face recognition, we have addressed the problem of recognizing blurred and poorly-illuminated faces. We have shown that the set of all images obtained by blurring a given image is a convex set given by the convex hull of shifted versions of the image. Based on this set-theoretic characterization, we proposed a blur-robust face recognition algorithm DRBF. In this algorithm we can easily incorporate prior knowledge on the type of blur as constraints. Using the low-dimensional linear subspace model for illumination, we then showed that the set of all images obtained from a given image by blurring and changing its illumination conditions is a bi-convex set. Again, based on this set-theoretic characterization, we proposed a blur and illumination robust algorithm IRBF. We also demonstrated the efficacy of our algorithms in tackling the challenging problem of face recognition in uncontrolled settings.

Our algorithm is based on a generative model followed by nearest-neighbor classification between the query image and the gallery space, which makes it difficult to scale it to real-life datasets with millions of images. This is a common issue with most algorithms based on generative models. Broadly speaking, only classifier-based methods have been shown to scale well to very large datasets; this is because the size of the gallery largely affects the training stage, the testing stage remains relatively fast. Hence we believe that incorporating a discriminative-learning based approach like SVM into this formulation would be a very promising direction for future work. We would also like to model pose-variation under the same framework.

Appendix
The Use of Weights in DRBF and IRBF

In the algorithms proposed above, we use the weight-matrix $\mathbf{W}$ to make them robust to outliers due to non-rigid variability (like facial expressions) and misalignment. They also help reduce the importance of pixels in the low-frequency regions of the face in the kernel-estimation step. This is desirable as the effects of blur are not really perceivable in these regions.

To train the weights, we used a method similar to the one used in [39]. We used the 'ba' and 'bj' folders of FERET as gallery and probe, respectively. We blurred the query with a Gaussian kernel of $\sigma = 4$ and partitioned the face images into patches as shown in 15(a). We then used DRBF/rDRBF to get the recognition rate for each patch independently. Finally the weights were assigned to each patch proportional to the recognition rate observed. 15(b) shows the weights obtained by this method for DRBF. These weights were then used for all 3 datasets, namely FERET('fa' and 'fb'), PIE and REMOTE.

(a) Image Patches (b) $\mathbb{W}$

Fig. 15: *The use of weights in DRBF and IRBF:* The weights have been trained on the 'ba' and 'bj' folders of FERET to allow for different regions of the face to contribute differently to the overall cost function. This enables us to give low weights to regions of the face that show high variability like the ears and the mouth. The trained weights are shown in Figure15(b), with white representing the most weight.

This can be verified from Figure 15(b), where the weights obtained for the outer regions of the face(hair, ears, neck etc) and the mouth are very small; as these regions are more prone to show non-rigid variability. The weights for the regions with the cheeks are also relatively small, validating our hypothesis that less textured (low-frequency) regions of the face should contribute less towards the estimation problem. Lastly, regions around the eyes are weighed the most which re-affirms the common understanding that they are the more distinguishable features of the human face.

REFERENCES

[1] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face recognition: A literature survey," *ACM Comput. Surv.*, 2003. 1

[2] J. Ni and R. Chellappa, "Evaluation of state-of-the-art algorithms for remote face recognition," in *International Conference on Image Processing*, 2010. 1, 5, 8, 9

[3] M. Nishiyama, A. Hadid, H. Takeshima, J. Shotton, T. Kozakaya, and O. Yamaguchi, "Facial deblur inference using subspace analysis for recognition of blurred faces," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2010. 1, 2, 6

[4] D. Kundur and D. Hatzinakos, "Blind image deconvolution," *Signal Processing Magazine, IEEE*, vol. 13, May 1996. 1, 2

[5] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding and evaluating blind deconvolution algorithms," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 1964–1971. 1, 2, 6

[6] T. Ojala, M. Pietikainen, and T. Maenpaa, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, July 2002. 1, 4

[7] R. Basri and D. W. Jacobs, "Lambertian reflectance and linear subspaces," *IEEE Transactions Pattern Anal. Mach. Intell.*, 2003. 1, 2, 4

[8] R. Ramamoorthi and P. Hanrahan, "A signal-processing framework for inverse rendering," in *SIGGRAPH*, 2001. 1, 2, 4

[9] K. C. Lee, J. Ho, and D. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2005. 1, 2, 4

[10] H. Hu and G. D. Haan, "Low cost robust blur estimator," in *International Conference on Image Processing*, 2006. 2

[11] W. H. Richardson, "Bayesian-based iterative method of image restoration," 1972. 2

[12] A. Levin, "Blind motion deblurring using image statistics," in *In Advances in Neural Information Processing Systems (NIPS*, 2006. 2

[13] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman, "Removing camera shake from a single photograph," *ACM Trans. Graph.*, vol. 25, July 2006. 2

[14] T. Ahonen, E. Rahtu, V. Ojansivu, and J. Heikkila, "Recognition of blurred faces using local phase quantization," in *International Conference on Pattern Recognition*, 2008. 2, 4

[15] R. Gopalan, S. Taheri, P. Turaga, and R. Chellappa, "A blur-robust descriptor with applications to face recognition," *IEEE Trans Pattern Anal. Mach. Intell.*, 2012. 2

[16] V. Ojansivu and J. Heikkil, "Blur insensitive texture classification using local phase quantization," in *Image and Signal Processing*. Springer Berlin / Heidelberg, 2008. 2, 6

[17] I. Stainvas and N. Intrator, "Blurred face recognition via a hybrid network architecture," in *International Conference on Pattern Recognition*, 2000. 2

[18] H. Zhang, J. Yang, Y. Zhang, N. M. Nasrabadi, and T. S. Huang, "Close the loop: Joint blind image restoration and recognition with sparse representation prior," in *IEEE International Conference on Computer Vision*, 2011. 2

[19] P. L. Combettes, "The convex feasibility problem in image recovery," *Advances in Imaging and Electron Physics*, vol. 95, 1996. 2

[20] H. Trussell and M. Civanlar, "The feasible solution in signal restoration," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 32, no. 2, pp. 201 – 212, apr 1984. 2

[21] P. L. Combettes and J. C. Pesquet, "Image restoration subject to a total variation constraint," *IEEE Trans. Image Process.*, vol. 13, no. 9, pp. 1213 –1222, sept. 2004. 2

[22] H. Wang, S. Z. Li, and Y. Wang, "Face recognition under varying lighting conditions using self quotient image," in *IEEE International Conference on Automatic Face and Gesture Recognition*, may 2004, pp. 819 – 824. 2, 8

[23] B. V. K. V. Kumar, M. Savvides, and C. Xie, "Correlation pattern recognition for face recognition," *Proceedings of the IEEE*, vol. 94, nov. 2006. 2

[24] M. Savvides, B. V. K. V. Kumar, and P. K. Khosla, "Eigenphases vs. eigenfaces," in *International Conference on Pattern Recognition*, 2004. 2

[25] R. Gross and V. Brajovic, "An image preprocessing algorithm for illumination invariant face recognition," in *Proceedings of the 4th international conference on Audio- and video-based biometric person authentication*, 2003. 2

[26] H. F. Chen, P. N. Belhumeur, and D. W. Jacobs, "In search of illumination invariants," in *Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on*, 2000. 2

[27] M. Osadchy, D. W. Jacobs, and M. Lindenbaum, "Surface dependent representations for illumination insensitive image comparison," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, 2007. 2

[28] S. Biswas, G. Aggarwal, and R. Chellappa, "Robust estimation of albedo for illumination-invariant matching and shape recovery," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2009. 2, 4

[29] R. L. Lagendijk and J. Biemond, "Basic methods for image restoration and identication," *The Essential Guide to Image Processing. Ed. A. C. Bovic*, 2009. 2, 3, 6

[30] P. Sinha, B. Balas, Y. Ostrovsky, and R. Russell, "Face recognition by humans: Nineteen results all computer vision researchers should know about," November 2006. 4

[31] R. Epstein, P. Hallinan, and A. Yuille, "5 plus or minus 2 eigenimages suffice: an empirical investigation of low-dimensional lighting models," in *Physics-Based Modeling in Computer Vision, Proceedings of the Workshop on*, 1995. 4

[32] V. Blanz and T. Vetter, "Face recognition based on fitting a 3d morphable model," *IEEE Trans Pattern Anal. Mach. Intell.*, vol. 25, no. 9, 2003. 4

[33] W. A. P. Smith and E. Hancock, "Estimating the albedo map of a face from a single image," in *IEEE International Conference on Image Processing*, vol. 3, 2005. 4

[34] X. Zou, J. Kittler, M. Hamouz, and J. R. Tena, "Robust albedo estimation from face image under unknown illumination," in *Proceedings of SPIE*, vol. 6944, 2008. 4

[35] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss, "The feret evaluation methodology for face-recognition algorithms," *IEEE Trans. Pattern Anal. and Mach. Intell.*, 2000. 5

[36] T. Sim, S. Baker, and M. Bsat, "The cmu pose, illumination, and expression (pie) database," 2002. 5

[37] R. Chellappa, J. Ni, and V. M. Patel, "Remote identification of faces: Problems, prospects, and progress," *Pattern Recognition Letters*, 2011. 8, 9

[38] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. M., "Robust face recognition via sparse representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, feb. 2009. 9

[39] T. Ahonen, A. Hadid, and M. Pietikainen, "Face description with local binary patterns: Application to face recognition," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 28, no. 12, pp. 2037–2041, 2006. 9

PLACE
PHOTO
HERE

Priyanka Vageeswaran received her B.S. degree in Electrical Engineering from the University of Maryland, College Park in 2010. She is currently a graduate student at the Computer Vision Lab in UMCP. Her research interests include unconstrained recognition for biometrics and geolocalization.

Kaushik Mitra is currently a postdoctoral researcher in the Electrical and Computer Engineering department of Rice University. Before this he received his Ph.D. in the Electrical and Computer Engineering department at University of Maryland, College Park in 2011 under the supervision of Prof. Rama Chellappa. His areas of research interests are Computational Imaging, Computer Vision and Statistical Learning Theory.

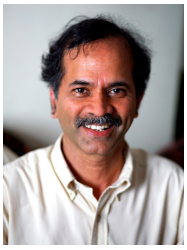

Rama Chellappa received the B.E. (Hons.) degree from the University of Madras, India, in 1975 and the M.E. (Distinction) degree from the Indian Institute of Science, Bangalore, in 1977. He received the M.S.E.E. and Ph.D. Degrees in Electrical Engineering from Purdue University, West Lafayette, IN, in 1978 and 1981 respectively. Since 1991, he has been a Professor of Electrical Engineering and an affiliate Professor of Computer Science at University of Maryland, College Park. He is also affiliated with the Center for Automation Research and the Institute for Advanced Computer Studies (Permanent Member). In 2005, he was named a Minta Martin Professor of Engineering. Prior to joining the University of Maryland, he was an Assistant (1981-1986) and Associate Professor (1986-1991) and Director of the Signal and Image Processing Institute (1988-1990) at the University of Southern California (USC), Los Angeles. Over the last 31 years, he has published numerous book chapters, peer-reviewed journal and conference papers in image processing, computer vision and pattern recognition. He has co-authored and edited books on MRFs, face and gait recognition and collected works on image processing and analysis. His current research interests are face and gait analysis, markerless motion capture, 3D modeling from video, image and video-based recognition and exploitation, compressive sensing, sparse representations and domain adaptation methods. Prof. Chellappa served as the associate editor of four IEEE Transactions, as a co-Guest Editor of several special issues, as a Co-Editor-in-Chief of Graphical Models and Image Processing and as the Editor-in-Chief of IEEE Transactions on Pattern Analysis and Machine Intelligence. He served as a member of the IEEE Signal Processing Society Board of Governors and as its Vice President of Awards and Membership. Recently, he completed a two-year term as the President of IEEE Biometrics Council. He has received several awards, including an NSF Presidential Young Investigator Award, four IBM Faculty Development Awards, an Excellence in Teaching Award from the School of Engineering at USC, two paper awards from the International Association of Pattern Recognition (IAPR), the Society, Technical Achievement and Meritorious Service Awards from the IEEE Signal Processing Society and the Technical achievement and Meritorious Service Awards from the IEEE Computer Society. He has been selected to receive the K.S. Fu Prize from IAPR. At University of Maryland, he was elected as a Distinguished Faculty Research Fellow, as a Distinguished Scholar-Teacher, received the Outstanding Faculty Research Award from the College of Engineering, an Outstanding Innovator Award from the Office of Technology Commercialization, the Outstanding GEMSTONE Mentor Award and the Poole and Kent Teaching Award for Senior Faculty. He is a Fellow of IEEE, IAPR, OSA and AAAS. In 2010, he received the Outstanding ECE Award from Purdue University. He has served as a General and Technical Program Chair for several IEEE international and national conferences and workshops. He is a Golden Core Member of the IEEE Computer Society and served a two-year term as a Distinguished Lecturer of the IEEE Signal Processing Society.