

Facial Deblur Inference to Improve Recognition of Blurred Faces

Masashi Nishiyama¹, Hidenori Takeshima¹, Jamie Shotton², Tatsuo Kozakaya¹, Osamu Yamaguchi¹

¹Corporate Research & Development, Toshiba Corporation, Japan

²Microsoft Research Cambridge, UK

masashi.nishiyama@toshiba.co.jp

Abstract

This paper proposes a novel method for deblurring facial images to recognize faces degraded by blur. The main problem is how to infer a Point Spread Function (PSF) representing the process of blur. Inferring a PSF from a single facial image is an ill-posed problem. To make this problem more tractable, our method uses learned prior information derived from a training set of blurred facial images of several individuals. We construct a feature space such that blurred faces degraded by the same PSF are similar to one another and form a cluster. During training, we compute a statistical model of each PSF cluster in this feature space. For PSF inference we compare a query image of unknown blur with each model and select the closest one. Using the PSF corresponding to that model, the query image is deblurred, ready for recognition. Experiments on a standard face database artificially degraded by focus or motion blur show that our method substantially improves the recognition performance compared with state-of-the-art methods. We also demonstrate improved performance on real blurred images.

1. Introduction

Identification using face recognition technology allows a convenient, and non-intrusive biometric authentication system [15]. Identification performance is significantly influenced both by the variation in appearance due to facial pose and illumination, and by image degradation due to sampling, blur, and noise. Many existing methods handle facial pose and illumination, but only recently has research started to focus on handling image degradation [1, 10, 11]. It is important to solve this problem for real-life face recognition applications such as watch-list monitoring and video surveillance. However, these papers mainly address sampling of face images. The focus of this paper is thus blur, and in particular, automatic deblurring for face recognition.

Blur changes the appearance of faces in images, causing two problems for face recognition: (i) the facial appearance

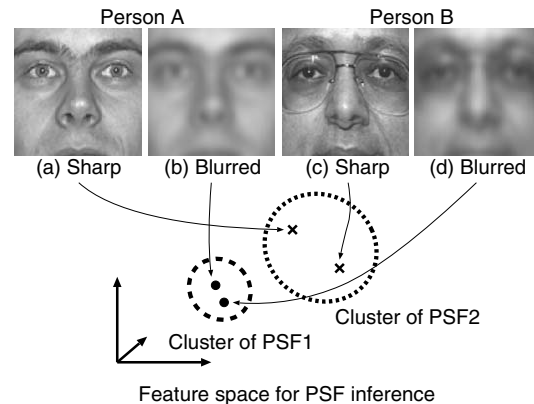


Figure 1. Variation in facial appearance caused by blur. We construct a new frequency magnitude-based feature space that clusters similar levels of blur together, regardless of identity. These clusters are used to find an appropriate PSF to deblur query images for accurate face recognition.

of an individual changes drastically, e.g. Figure 1 (a) and (b); and (ii) different individuals tend to appear more similar when blurred, e.g. Figure 1 (b) and (d). The first problem can be alleviated by matching a query image to artificially blurred copies of the original sharp target images registered for identification [23]. However, the second problem remains since the blurred copies become similar to each other. Moreover, in a real application the target images may themselves already be blurred. Instead, we choose an approach that *removes* blur from facial appearances using blind image deconvolution [14]. The deblurred images can then be used to perform more robust identification. The approach solves both problems (i) and (ii) simultaneously. The deconvolution requires a *Point Spread Function* (PSF) that represents the blurring process.

The majority of previous methods (see Section 2) use the smoothness of intensity changes around edges to infer the PSF from a single image, without prior knowledge of the image contents. These methods thus often infer a poor quality PSF as the problem is ill-posed: it is difficult to distinguish between blurred edges and smooth object sur-

faces. Our method instead exploits prior knowledge containing how facial appearances are changed by blur.

To this end, we propose a new method called Facial Deblur Inference for inferring PSFs using learned statistical models of the variation in facial appearance caused by blur. At training time, we choose a representative set of PSFs for a particular application. For each PSF in this set, we artificially blur the sharp training images and use to build a statistical model of facial appearance under that PSF. At test time, we infer a PSF by comparing the query image (with unknown identity and amount of blur) to each model.

Our method maps blurred images to a feature space for learning statistical models. Simply vectorizing the images does not generalize well for the feature space since the vectorized images are not clearly separated, and indeed our experiments show poor performance (‘Baseline A’ in Section 5). We instead design a new frequency magnitude-based feature space, illustrated in Figure 1, in which the variation of facial appearance caused by blur is larger than the variation of facial appearance between individuals. In this new feature space, faces blurred by the same PSF are similar to one another and form a cluster. Our method learns the statistical models by approximating each cluster as a low-dimensional linear subspace using principal component analysis (PCA). At test time, we compare query images to each subspace. The most similar subspace gives an accurate inferred PSF which is used to deblur the query image for input to a standard face recognition algorithm.

The salient contributions of this paper are the following:

- Learning statistical models of the variation in facial appearance caused by blur to accurately infer PSFs from blurred images.
- A new frequency magnitude-based feature space in which these statistical models are learned. This feature space is designed to emphasize the different appearances of different levels of blur but be invariant to facial identity.
- A demonstration that these statistical models generalize well, maintaining PSF inference accuracy: the models are trained on faces that are completely different to target images for identification.

Our experiments on real and artificially blurred face images show high PSF inference and face recognition accuracy that improves on prior work.

The rest of this paper is organized as follows: Section 2 describes related work, Section 3 describes our method for PSF inference, and Section 4 demonstrates the effectiveness of our method through experiments. We conclude our discussion in Section 5.

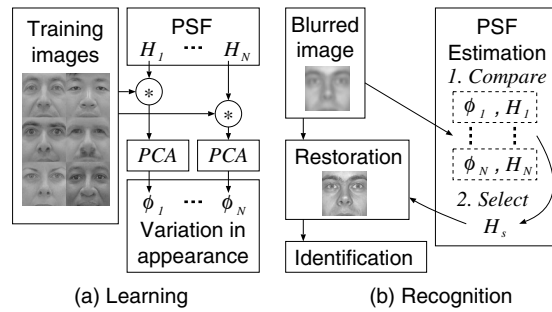


Figure 2. Our method is composed of two steps: (a) learning statistical models of facial blur appearance, and (b) using these models to recognize individuals in query images of unknown blur.

2. Related work

In the field of blind image deconvolution, many methods have been proposed for deblurring from a single image. Chan & Wong [3] simultaneously infer a PSF and deblur an image using Total Variation regularization. Other methods attempt to model the smoothness of intensity changes around edges. A PSF is inferred using information derived from this smoothness using the variation of Gaussian scale [12, 5], wavelet coefficient [20, 24], the summation of image derivatives [27], or alpha values representing the object boundary transparency [13]. These methods have to solve an ill-posed problem because they cannot use prior knowledge of the image contents, and so the quality of a deblurred image using an inferred PSF is often poor. As we will see in our experiments, the deblurred images using these methods are insufficient for accurate face recognition. Yuan et al. [28] infer a PSF using multiple images captured from the same scene with long and short exposure times. This setting limits face recognition applications. Fergus et al. [7] infer a PSF using heavy-tailed natural image priors, but this prior is very generic and so fairly weak, and the method requires much computation time for PSF inference.

3. Facial Deblur Inference

3.1. Background

The degradation process caused by blur is defined as

$$\mathbf{g} = H\mathbf{f} + \mathbf{n}, \quad (1)$$

where vector \mathbf{g} represents the blurred image $g(u, v)$, matrix H the PSF, vector \mathbf{f} the original sharp image, and vector \mathbf{n} the noise. Equation (1) represents an explicit appearance model for blur. Note that \mathbf{g} and \mathbf{f} consist of only the facial region given by a face detector.

If the PSF H and noise \mathbf{n} are known, a sharp image \mathbf{f} can be exactly recovered from the blurred image \mathbf{g} . This paper focuses on accurately inferring a PSF H from a blurred facial image \mathbf{g} , though we also compare in our experiments




Figure 3. Comparison of feature vectors. The blue curves are feature vectors generated from the top-left sharp image and the red curves are ones from the bottom-left blurred image. The feature vectors in our feature space (d) are much more clearly separated than in (a) using Baseline A. (a) One scan-line through a simple vectorization of the image (Baseline A). (b) One row of the the 2D Fourier transform of the images. (c) The same row after taking logarithms. (d) The same row with normalization. These transformations create our feature space. The horizontal axis of (a) represents position in space domain. The horizontal axes of (b),(c),(d) represent frequency in the Fourier domain, with low frequency in the center and higher frequencies further away from the center.

how, given our inferred PSF, two conventional deblurring methods cope with the noise.

3.2. Overview

We start with an overview of our method, illustrated in Figure 2. This paper assumes that a representative fixed set $\Omega = \{H_i\}_{i=1}^N$ of N PSFs are constituted as described in [21]. These PSFs are chosen for a particular application. We represent statistical models for PSF inference as subspaces in the frequency magnitude-based feature space (see Section 3.3). We denote a set of statistical models as

$$\Phi = \{ (H_i, \phi_i) \}_{i=1}^N, \quad (2)$$

where ϕ_i represents the subspace modeling the variation in facial appearance induced by PSF H_i . We learn subspaces ϕ_i from a set of M training images $\Psi = \{\mathbf{f}'_k\}_{k=1}^M$. A training image \mathbf{f}'_k is a sharp image that is artificially blurred by each H_i to learn the subspace ϕ_i (see Section 3.4). Note that the individuals in Ψ can be completely different to target images for identification.

A query image \mathbf{g} of unknown blur and identity is compared to each subspace using the subspace method [25] which can be easily implemented and gives an accurate and stable similarity measure (see Section 3.5). The closest subspace ϕ_s is selected and the corresponding PSF H_s is the result of our PSF inference. The query image \mathbf{g} is deblurred using H_s (see Section 3.6), and finally identification of the deblurred image is performed.

3.3. Frequency Magnitude-based Feature Space

We construct a feature space that is sensitive to the appearance variations of different blurs but insensitive to the difference between individuals. We base the feature space on frequency domain amplitude, because it has a phenomenon that high frequency amplitudes for blurred images become smaller than those for sharp images [9]. Since we

want a feature space invariant to identity, we deliberately ignore phase information.¹

Inferring a PSF using frequency domain amplitude is well-known technique [21, 8, 2, 4]. The method [21] selects a PSF from a representative fixed set Ω using amplitude of a sharp image and noise. But, it is difficult to correctly estimate these amplitudes from the blurred image because the sharp image is unknown. Other methods [8, 2, 4] use the relationship between PSF parameters and the positions of the zero crossings of amplitude. However, detecting these positions is difficult for real blurred images because the zero crossings are affected by the noise amplitude. Our method instead uses the whole frequency amplitude domain for constructing the feature space. The variation caused by noise is included in the subspaces ϕ_i learned in our feature space by adding noise to training images (see Section 3.4).

We emphasize discriminative high frequency amplitude to improve the PSF inference accuracy. Our method first transforms a blurred image $g(u, v)$ in space domain to a feature image $x(\xi', \eta')$ as

$$x(\xi', \eta') = [C(|g(\xi, \eta)|)] \downarrow, \quad (3)$$

where $g(\xi, \eta)$ is the 2D Fourier transform of $g(u, v)$, $|\cdot|$ takes the amplitude, C takes dynamic range compression, and $[\cdot] \downarrow$ represents downsampling. Dynamic range compression is performed to emphasize high frequency values because non-discriminative low frequency values are much larger than discriminative high frequency values. For the function C , we simply take logarithms that shifts from a large value to a small value while maintaining order relation. Downsampling helps reduce noise and speed-up subspace learning. Bilinear interpolation is used for downsampling. We choose the procedures C and $[\cdot] \downarrow$ from the viewpoints of computation time and inference accuracy. The feature vector \mathbf{x} is generated by raster-scanning $x(\xi', \eta')$.

¹Without blur, of course, phase information is a good feature for face recognition [22].




Figure 4. Visualization of clusters in two feature spaces. Artificially blurred facial images (left) are vectorized either directly (center, Baseline A) or using the frequency domain amplitude (right, our method). Colors represent the particular PSFs used to blur the faces. Our evaluation highlights the importance of the much cleaner separation of the PSF clusters using our method.

Finally, as pre-processing for the subspace method, vector \mathbf{x} is normalized so that $\|\mathbf{x}\|_2 = 1$.

Figure 3 compares feature vectors generated from a sharp image and its blurred version. The vectors in Figure 3 (a) to (d) are the outputs of each step in Equation (3) and the normalization. Note that the waveforms of the vectors in Figure 3 (a) are almost the same, but the difference between sharp and blurred images appears most cleanly in our feature space in Figure 3 (d).

Further, Figure 4 uses PCA to visualize two feature spaces containing images blurred by three different PSFs. The Baseline A feature space (center) directly vectorizes the blurred facial images. Observe that our frequency magnitude-based feature space (right) gives clusters that are much better separated. Our experiments below compare these two feature spaces and confirm that our feature space gives more accurate PSF inference and face recognition.

3.4. Learning Subspaces

We generate subspaces for the statistical models to represent clusters in the feature space. Training images are blurred by applying each PSF H_i to sharp training images \mathbf{f}'_k in a training set Ψ as $\mathbf{g}'_{ik} = H_i \mathbf{f}'_k + \mathbf{n}'$, where \mathbf{n}' is the noise which we assume to be white Gaussian.

Next, we apply PCA to blurred training images $\{\mathbf{g}'_{ik}\}_{k=1}^M$ in our feature space, as follows. Having transformed image \mathbf{g}'_{ik} to feature vector \mathbf{x}'_{ik} , the correlation matrix [17] is defined as $A_i = \frac{1}{M} \sum_{k=1}^M \mathbf{x}'_{ik} (\mathbf{x}'_{ik})^T$. Note that the mean vector of $\{\mathbf{x}'_{ik}\}_{k=1}^M$ is not subtracted for the correlation matrix because our implementation of subspace method uses an angle for similarity measure (see Equation (4)). Eigenvectors and eigenvalues of the correlation matrix A_i are calculated. The first D eigenvectors, sorted by decreasing eigenvalue, form the basis vectors of the subspace $\phi_i = \{\mathbf{b}_{ij}\}_{j=1}^D$. The low-dimensional subspace spanned by the basis vectors ϕ_i is defined by maximizing the variance of training feature vectors $\{\mathbf{x}'_{ik}\}_{k=1}^M$ in this subspace.




Figure 5. Examples of synthesized images of camera focus blur from the subset ‘fb’ of FERET. Blurred images (b), (c), (d), and (e) are synthesized from the sharp original image (a) using Gaussian PSFs with given standard deviations σ .

3.5. Determining the PSF

Given a blurred query image \mathbf{g} of unknown blur, we find the closest subspace ϕ_s in a set of statistical models Φ as

$$\begin{aligned} s &= \arg \max_i \cos^2 \theta_i \\ &= \arg \max_i \sum_{j=1}^D (\mathbf{b}_{ij}^T \mathbf{x})^2, \end{aligned} \quad (4)$$

where \mathbf{x} is the feature vector of \mathbf{g} , θ_i is the angle between vector \mathbf{x} and subspace ϕ_i , and $i \in \{1, \dots, N\}$. A small angle θ_i indicates that the blurred query image \mathbf{g} is similar to a cluster of the training images blurred by PSF H_i that generated the subspace. The maximization of $\cos^2 \theta_i$ over i thus implicitly infers the most appropriate PSF H_s .

3.6. Restoration

Given PSF H_s , we can then remove blur from the degraded query image \mathbf{g} . Simple matrix inverse based deconvolutions give a very poor quality result because of existing noise \mathbf{n} in Equation (1). We thus use Wiener filters [26] or the BTV regularization from [6]. These differ in their treatment of noise and computational expense. The detail of BTV is described in Appendix A.

4. Experiments

To demonstrate the effectiveness of our method, we evaluated on the FERET [19] and FRGC 1.0 [18] face databases. We investigate camera focus blur and camera

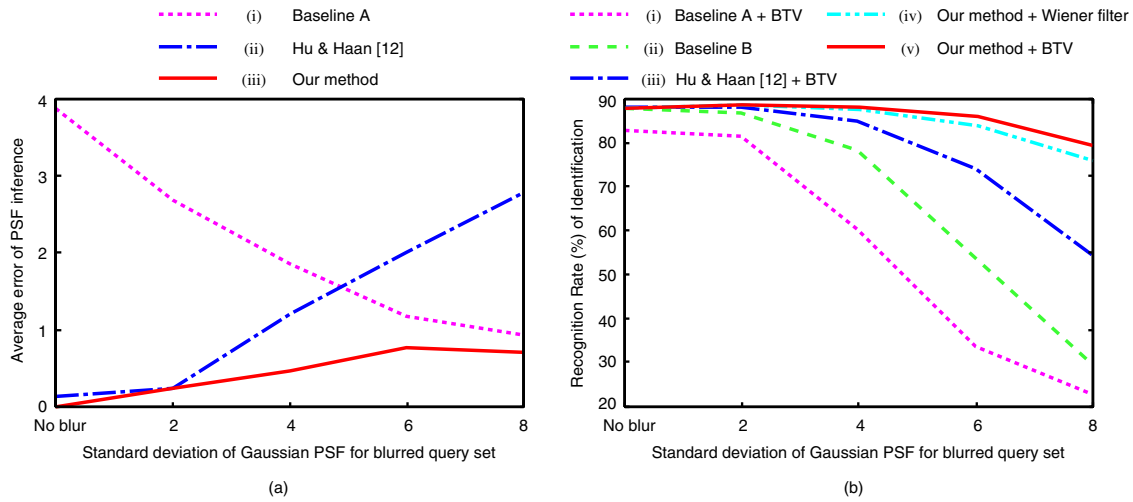


Figure 6. Performance on FERET of artificially camera focus blur. Query set are blurred with Gaussian PSF. The curves in (a) show PSF inference accuracy and the curves in (b) show identification performance .

motion blur. Camera focus blur arises when the camera focal length is not correctly adjusted to the subject. Camera motion blur arises when the camera is moved while capturing an image.

4.1. Evaluation on an artificially blurred database

4.1.1 Camera focus blur

The first experiment tests on synthesized images by blurring sharp query faces from FERET. Sharp faces are blurred by shift-invariant Gaussian PSFs as $H(u, v) = \frac{1}{Z} \exp\left(\frac{-(u^2+v^2)}{2\sigma^2}\right)$, where σ is the standard deviation, and Z is a normalization term as $\int \int H(u, v) dudv = 1$. This PSF is defined in the space domain. White Gaussian noise of 30dB is added to the synthesized images. Examples are shown in Figure 5. The facial images are registered using facial feature points (eyes, nose, etc.). Image $g(u, v)$ is of size 128×128 , and downsampled $x(\xi', \eta')$ is of size 64×64 . For the representative fixed set Ω of PSFs, we use $N = 18$ PSFs, including 17 Gaussians and 1 ‘no blur’ delta function. The parameters of the Gaussian PSFs are set from $\sigma = 1$ to 9 in increments of 0.5. The database includes three subsets: ‘bk’, ‘fa’, and ‘fb’. Each subset contains a single image per person. Subset bk is used as training set Ψ to learn the subspaces in a set of statistical models Φ . The dimension of each subspace is $D = 20$. The number of training images is $M = 200$.

We evaluate performance on the 1001 images that remain in subsets fa and fb after removing the individuals present in subset bk. The identification target set is fa and the query set is fb. The deblurred face images are identified using Euclidean nearest-neighbor matching of the raster-scan vectorized images. This matching can assess standard identifica-

tion performance under the situation where only one query image is used per person. We evaluate performance using another matching method in Section 4.2.

Before reporting identification accuracy, we show PSF inference accuracy on the query images in Figure 6 (a). Average error is defined as $\|\sigma_c - \sigma_s\|_1$ where σ_c is the standard deviation of the PSF used for synthesizing the blurred image (ground-truth), and σ_s is that of the inferred PSF. We compare our method (iii) to two alternatives: (i) Baseline A uses a directly vectorized feature space (cf. Figure 4); and (ii) Hu & Haan [12] estimate the standard deviation of Gaussian PSF from the smoothness of intensity changes around edges. This method can handle only camera focus blur. As we see, our result (iii) is superior to the baseline and [12]. An interesting anomaly is that the PSF inference accuracy of Baseline A increases with blur. This is probably because increasing the blur reduces variation in the Baseline A feature space which can thus be more adequately represented by a low-dimensional subspace.

In Figure 6 (b), we report identification accuracy as recognition rate: the probability that a query image is matched to the correct target image of the same individual. We compare our method (iv) and (v) to three alternatives: (i) Baseline A uses a directly vectorized feature space; (ii) Baseline B does no PSF inference or deblurring, and instead directly matches blurred query images against sharp target images; and (iii) the PSF inference of Hu & Haan [12] is used. Identification accuracy is evaluated by combining restoration (cf. Section 3.6) with PSF inference. In (i), (iii), and (v), BTV regularization is used for restoration. In (iv), the Wiener filter is used for restoration. The results in (iv) and (v) show that using BTV regularization is slightly more accurate than using the Wiener filter. This comparison is only indicative however, as it only used additive white




Figure 7. Examples of images artificially blurred with random standard variation of Gaussian PSF in $(0, 8]$ on FERET. (a) Target images from the subset ‘fa’. (b) Query images for the subset ‘fb’. Note that both target and query images are blurred to simulate real-life face capture conditions.

Table 1. Identification performance on FERET of artificially camera focus blur with random standard variation of Gaussian PSF.

Method	Recognition rate
Baseline A + BTV	60.24%
Baseline B	72.43%
Hu & Haan [12] + BTV	76.02%
Our method + BTV	82.82%

Gaussian noise to artificial blurred images, and does not compare impulse noise tolerance. From Figure 6 (b), our results (iv), (v) are significantly better than both baselines (i), (ii) and (iii), especially in the presence of large blur. In our unoptimized implementation on a single core 3.4GHz processor, learning takes about 5 minutes, and PSF inference takes about 0.5 seconds per image. The BTV regularization [6] takes up to 1.5 seconds per image, and the Wiener filter [26] takes up to 50 milliseconds per image. Fast Fourier transform is used for the Wiener filter.

We also show recognition performance on FERET when the query set images were blurred by an unknown random Gaussian sigma in the range $(0, 8]$. Figure 7 shows examples used in this experiment. Some faces are blurred only in target images, and other faces are blurred only in query images. Note that facial deblur inference is performed on *both* target and query images. From Table 1, we again see an improvement on the other methods, and demonstrate consistent performance for unknown blurs that do not precisely coincide with the training PSFs.

4.1.2 Camera motion blur

The second experiment tests on synthesized images by blurring sharp query faces from FERET using shift-invariant linear motion blur PSFs as $H(u, v) = 1/Z$ if $\|(u, v)\|_2 < b$ and $v = u \tan \theta$, otherwise $H(u, v) = 0$, where b is the length of camera motion, θ is the angle, and Z is a normalization term. White Gaussian noise of 30dB is added to the synthesized images. Identification query images are



Figure 8. Examples of synthesized images of camera motion blur from the subset ‘fb’ of FERET. Blurred images are synthesized from Figure 5 (a) using linear motion blur PSFs with given b, θ .

Table 2. Identification performance on FERET of artificially camera motion blur.

Method	Recognition rate
Baseline A + BTV	26.4%
Baseline B	59.6%
Yitzhaky & Kopeika [27] + BTV	60.0%
Our method + BTV	82.0%

blurred by the PSF in the range $b = 5, 10, 15, 20$ and $\theta = 0, 0.25\pi, 0.5\pi, 0.75\pi$. Examples are shown in Figure 8. For the representative fixed set Ω of PSFs, we use $N = 41$ PSFs, including $10 \times 4 = 40$ motion blur function and 1 ‘no blur’ delta function. The parameters are set from $b = 2.5$ to 25 in increments of 2.5 and from $\theta = 0$ to 0.75π in increments of 0.25π . The other parameters for learning statistical models are the same as in Section 4.1.1.

Table 2 reports average recognition rate. We compare our method to three alternatives. Baseline A and Baseline B are the same as in Section 4.1.1. The PSF inference of Hu & Haan [12] cannot be applied because it is only applicable to focus blur. We instead use the method of Yitzhaky & Kopeika [27]. As we see, our result is significantly better than both baselines and [27].

4.2. Evaluation on a real blurred database

The second experiment uses real blurred images from FRGC 1.0. We evaluated identification performance in terms of verification rate under two setups, ‘Exp1’ and ‘Exp4’. Exp1 is evaluated on controlled still query images while Exp4 is evaluated on uncontrolled still query images




Figure 10. (a) Example images from FRGC 1.0 with real blur. (b), (c) Faces deblurred by combining our method to infer appropriate PSFs with the Wiener filter and the BTV regularization, respectively.




Figure 9. Example images from FRGC 1.0. (a) Target image. (b) Query image in Exp4. Real camera focus blur arises in (b).

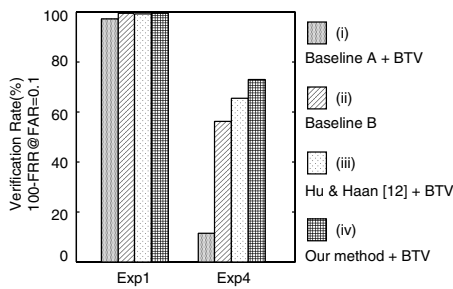


Figure 11. Identification performance on FRGC 1.0 (real blurred images).

including blurred faces. Each query set consists of 608 images of 152 individuals. In Exp 4, we count that 366 query images are degraded by camera focus blur. Target images of Exp1 and Exp4 are collected under a controlled still condition. A single image is captured per person for target. The number of target images is 152. Figure 9 shows example images from FRGC 1.0. Camera focus blur arises in Figure 9 (b) as the camera focal length is adjusted to the background. The amount of blur in both the target and query images of Exp4 is not constant and unknown. Note that facial deblur inference is performed on *both* target and query images, since real face target databases already contain considerable focus blur. For the representative fixed set Ω of PSFs, we use $N = 8$ PSFs, including 7 Gaussians and 1 ‘no blur’ delta function. The parameters of the Gaussian

PSFs are set from $\sigma = 1$ to 4 in increments of 0.5. Three subsets ‘bk’, ‘fa’, and ‘fb’ in FERET are used as training set Ψ to learn the subspaces in a set of statistical models Φ . The dimension of each subspace is $D = 20$. The number of training images is $M = 2591$.

We show some deblurred images from Exp4 in Figure 10. The deblurred images look almost the same between BTV regularization and the Wiener filter, probably because the noise included in FRGC 1.0 can be approximated well by white Gaussian noise instead of impulse noise.

In this experiment, the method from [16] is used to identify the deblurred face. This copes better with illumination and pose changes than nearest-neighbor, but requires multiple images for each individual. We therefore generate several examples for a single training image per person by perturbing the facial feature points.

We evaluated recognition performance in terms of verification rate calculated from False Acceptance Rate (FAR) and False Rejection Rate (FRR). As shown in Figure 11, our method again obtains higher recognition performance than previous methods and the same baselines. The recognition performance using the BTV regularization and the Wiener filter were almost the same in this experiment. We conclude that our method is effective at deblurring real images to allow improved face recognition.

5. Conclusion

This paper proposed a novel method to deblur facial images for face recognition. Our algorithm inferred Point Spread Functions using learned models of facial appearance variation under different amounts of blur. The inferred PSFs were used to sharpen both query and target images, and thereby improve face recognition accuracy. Our experiments on both real and artificially blurred face images demonstrated substantially more accurate PSF inference and face recognition than two baseline methods and previous works.

We believe that our deblurring method has relevance not only for face recognition, but also for other restricted classes of images, such as character, hand and body. As fu-

ture work we intend to expand our evaluation to include tolerance to noise and severe motion blur e.g. camera shake.

References

- [1] O. Arandjelović and R. Cipolla. A manifold approach to face recognition from low quality video across illumination and pose using implicit super-resolution. *ICCV*, 2007.
- [2] M. Cannon. Blind deconvolution of spatially invariant image blurs with phase. *IEEE Trans. Acoust. Speech Signal Process*, ASSP-24:58 – 63, 1976.
- [3] T. F. Chan and C.-K. Wong. Total variation blind deconvolution. *IEEE Trans. Image Processing*, 7(3):370 – 375, 1998.
- [4] M. M. Chang, A. M. Tekalp, and A. T. Erdem. Blur identification using the bispectrum. *IEEE Trans. Signal Processing*, 39(10):2323 – 2325, 1991.
- [5] J. H. Elder. Local scale control for edge detection and blur estimation. *PAMI*, 20(7):699 – 716, 1998.
- [6] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar. Fast and robust multiframe super resolution. *IEEE Trans. Image Processing*, 13(10):1327 – 1344, 2004.
- [7] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graphics*, 2006.
- [8] D. B. Gennery. Determination of optical transfer function by inspection of the frequency-domain plot. *J. Opt. Soc. Amer.*, 63:1571 – 1577, 1973.
- [9] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Prentice Hall, 2007.
- [10] B. K. Gunturk, Y. Batur, A. U. Altunbasak, M. H. Hayes, and R. M. Mersereau. Eigenface-domain super-resolution for face recognition. *IEEE Trans. Image Processing*, 12(5):597 – 606, 2003.
- [11] P. H. Hennings-Yeomans, S. Baker, and B. V. K. V. Kumar. Simultaneous super-resolution and feature extraction for recognition of low-resolution faces. *CVPR*, 2008.
- [12] H. Hu and G. de Haan. Low cost robust blur estimator. *ICIP*, 2006.
- [13] J. Jia. Single image motion deblurring using transparency. *CVPR*, 2007.
- [14] D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Magazine*, pages 43 – 64, May 1996.
- [15] S. Z. Li and A. K. Jain. *Handbook of Face Recognition*. Springer, 2005.
- [16] M. Nishiyama, M. Yuasa, T. Shibata, T. Wakasugi, T. Kawahara, and O. Yamaguchi. Recognizing faces of moving people by hierarchical image-set matching. *Proc. IEEE Workshop on Biometrics*, 2007.
- [17] E. Oja. Subspace methods of pattern recognition. *Research Studies Press*, 1983.
- [18] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. *CVPR*, 1:947 – 954, 2005.
- [19] P. J. Phillips, H. Moon, P. J. Rauss, and S. Rizvi. The feret evaluation methodology for face recognition algorithms. *PAMI*, 22(10):1090 – 1104, 2000.
- [20] F. Rooms, A. Pizurica, and W. Philips. Estimating image blur in the wavelet domain. *ACCV*, pages 210 – 215, 2002.
- [21] A. E. Savakis and H. J. Trussell. Blur identification by residual spectral matching. *IEEE Trans. Image Processing*, 2(2):141 – 151, 1993.
- [22] M. Savvides, B. V. K. V. Kumar, and P. K. Khosla. Eigen-phases vs. eigenfaces. *ICPR*, 3:810 – 813, 2004.
- [23] I. Stainvas and N. Intrator. Blurred face recognition via a hybrid network architecture. *ICPR*, 2:805 – 808, 2000.
- [24] H. Tong, M. Li, H. Zhang, and C. Zhang. Blur detection for digital images using wavelet transform. *IEEE Proc. Int. Conf. Multimedia and Expo*, 1:17 – 20, 2004.
- [25] S. Watanabe and N. Pakvasa. Subspace method of pattern recognition. *Proc. Int. Joint Conf. Pattern Recognition*, pages 25 – 32, 1973.
- [26] N. Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. The MIT Press, 1964.
- [27] Y. Yitzhaky and N. S. Kopeika. Identification of blur parameters from motion blurred images. *Graphical Models and Image Processing*, 59(5):310 – 320, 1997.
- [28] L. Yuan, J. Sun, L. Quan, and H. Y. Shum. Image deblurring with blurred/noisy image pairs. *ACM Trans. Graphics*, 2007.

A. Bilateral Total Variation regularization

The BTV regularization is robust to white noise and impulse noise including outlier e.g. salt-and-pepper noise. In BTV, deblurred image $\hat{\mathbf{f}}$ is defined as

$$\hat{\mathbf{f}} = \arg \max_{\mathbf{f}} [\|H_s \tilde{\mathbf{f}} - \mathbf{g}\|_2^2 + \lambda_b \sum_{l=-P}^P \sum_{m=-P}^P \alpha_b^{|m|+|l|} \|\tilde{\mathbf{f}} - S_x^l S_y^m \tilde{\mathbf{f}}\|_1], \quad (5)$$

where λ_b, α_b, P are constant, matrix S_x^l shifts an image horizontally by l pixels, and matrix S_y^m shifts an image vertically by m . Our experiments use L2 norm instead of L1 norm for $\|H_s \tilde{\mathbf{f}} - \mathbf{g}\|$ to reduce iterations. The deblurred image $\hat{\mathbf{f}}$ is calculated using a gradient method as

$$\hat{\mathbf{f}}_{l+1} = \hat{\mathbf{f}}_l - \beta_b \{ 2H_s^T (H_s \hat{\mathbf{f}}_l - \mathbf{g}) + \lambda_b \sum_{l=-P}^P \sum_{m=-P}^P \alpha_b^{|m|+|l|} (I - S_y^{-m} S_x^{-l}) \text{sign}(\hat{\mathbf{f}}_l - S_x^l S_y^m \hat{\mathbf{f}}_l) \}, \quad (6)$$

where β_b is constant, matrix S_x^{-l}, S_y^{-m} are the transposes of S_x^l, S_y^m respectively, and function $\text{sign}(\cdot)$ transforms positive values to 1 and negative values to -1 for each element of the vector.