

© 2010 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Pre-print of article that appeared at IECON 2010.

The published article can be accessed from:
http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5675537

Single Image Deblurring for a Real-Time Face Recognition System

Brian Heflin^{*†}, Brian Parks[†], Walter Scheirer^{*†}, and Terrance Boulton^{*†}

^{*}Securics Inc.

[†]University of Colorado at Colorado Springs
Colorado Springs, CO 80918

Abstract—Blur due to motion and atmospheric turbulence is a variable that impacts the accuracy of computer vision-based face recognition techniques. However, in images captured in the wild, such variables can hardly be avoided, requiring methods to account for these degradations in order to achieve accurate results in real time. One such method is to estimate the blur and then use deconvolution to negate or, at the very least, mitigate the effects of blur. In this paper, we describe a method for estimating motion blur and a method for estimating atmospheric blur. Unlike previous blur estimation methods, both methods are fully automated, allowing integration into a real-time facial recognition pipeline. We show experimentally, on datasets processed to include synthetic and real motion and atmospheric blur, that these techniques improve recognition more than prior work. At multiple levels of blur, our results demonstrate significant improvement over related works and our baseline on data derived from both the FERET (fairly constrained data) and Labeled Faces in the Wild (fairly unconstrained data) sets.

I. INTRODUCTION

Facial recognition technology allows for a convenient and non-invasive way to recognize an unknown subject. However, the task of unconstrained face recognition still remains a challenging problem because of its fundamental difficulties concerning various factors in the real world such as pose (in-plane and out-plane rotation), illumination changes, facial expressions, and atmospheric and motion blur. The first three issues are primarily properties of the face in three-dimensions, whereas blur is a variable at the image level. That is to say that until an image is captured, no concrete notion of blur exists between target object and observing body. As soon as the scene is captured, an artifact representing changing environmental conditions over the course of the integration time is apparent. One approach to solving the greater problem of unconstrained environments involves reducing the set of gallery and probe images so that it approximates a constrained environment face recognition problem. This approach lends itself rather well to the case of blurred images, since blur is an artifact at the image level and thus can be reversed (theoretically) simply by manipulating the image. This, then, raises the question of how to manipulate the image such that the recorded blur is reversed or, at the very least and far more likely, mitigated.

Many previous works on the difficulties in facial recognition due to various factors in the real world have focused on pose and illumination variations, but only a few have focused on the issue of motion blur and atmospheric blur. The focus of this paper is automatic motion and atmospheric blur estimation and

deblurring for a real-time facial recognition system. Motion and atmospheric deblurring is a highly ill-posed problem where the observed image $g(x, y)$ is the convolution of the original image $f(x, y)$ and the unknown blur kernel $d(x, y)$ plus additive noise $n(x, y)$.

$$g(x, y) = f(x, y) \otimes d(x, y) + n(x, y) \quad (1)$$

To combat motion blur, many methods using multiple images to perform motion blur estimation have been proposed. They seek to utilize correlation among blurred images, based on the observation that all blur observations come from the same latent image [1]. However, for a real-time facial recognition system this assumption is not satisfied since there can be multiple faces in the image with different levels of blur and there is no guarantee that the same face will be found in subsequent frames. One method using a blur/noisy pair [2] shows good estimation and deblurring results but requires an exposure bracketing function that is not available on most cameras, thereby limiting the application range of this algorithm. Other current methods are focused on space-variant motion blur estimation, nonparametric blur kernel estimation, and blind deconvolution for motion [3], [4], [5]. However, the computational complexity and processing time of these algorithms does not fit into the constraints of a real-time facial recognition system with today's hardware.

Our motion work is based on the previous work of using a single image to estimate the motion or atmospheric blur point spread function (PSF) and then use a deconvolution filter such as a Wiener filter to recover the image. These approaches are based on the fact that some of the blur PSFs have periodic zeros in the frequency domain that can be used to estimate the blur length and direction [6]. Although this approach is somewhat sensitive to noise, it does provide a reasonable tradeoff between algorithm performance, computational complexity, and processing time and is the basis for our current motion blur estimation and compensation algorithm.

On the topic of atmospheric blur, the distortion caused by thermal aberrations in the atmosphere, there has been extensive research on blind deconvolution algorithms over the past 20 years [7], [8]. Blind deconvolution algorithms can be generally categorized into two methods. The first method, which our atmospheric work is based on, separates the atmospheric blur PSF identification as a separate procedure from restoration and is usually a non-iterative process. For example, [9] proposes to

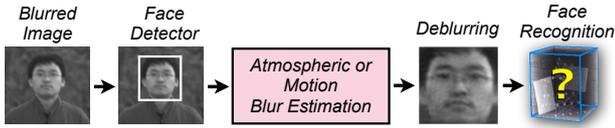


Fig. 1. An overview of the complete face recognition system described in this paper, including atmospheric or motion deblurring.

estimate the PSF directly from the image using automatic best step-edge detection and then use this information to compute the atmospheric modulation transfer function (MTF) to restore the image.

The second method combines the PSF estimation and restoration into one process, and it is usually an iterative process. This method formulates a parametric model for both the atmospheric blur and image at each stage of the process and then uses these models in the subsequent iteration of the algorithm. An isotropic 2D Gaussian function is normally used to model the atmospheric PSF [10], [11], [12]. However, the quality of the restored image and the speed at which the algorithms converge is based on the initial guess of the atmospheric PSF parameters. More recently, a blind deconvolution algorithm that is based on kurtosis minimization has been described in [13]. Using a set choice of blur parameters, the atmospheric blurred image is restored and then a statistical criterion, the minimum kurtosis, is used as a quality metric to select the best restored image.

It is from this perspective of blur reversal or mitigation that the current study treats the effect of motion and atmospheric blur on the face recognition problem. Figure 1 shows the proposed deblurring and facial recognition system framework consisting of 4 blocks: face detection, atmospheric and motion blur PSF estimation, a deblurring block implying deconvolution using the estimated blur PSF and a Wiener filter, and a V1 based recognition core [14].

We organize the rest of the paper as follows: Sections II and III outline the motion and atmospheric blur estimation techniques used in both the preprocessing and recognition phases of testing. Section IV provides further details on the structure and implementation of our recognition pipeline, and presents results from our experiments on both synthetic and real data. We conclude in Section V.

II. IMAGE AND MOTION BLUR MODEL

Motion blur is caused by the relative motion between the camera and the scene during the integration time of the image. The lost spatial frequencies in motion blurred images can be recovered by image deconvolution, provided that the motion is at least locally shift-invariant. The first step to recover an image degraded by motion blur is to determine the original motion function or point spread function (PSF) precisely as possible. More accurately, given a motion blurred image and assuming a linear motion blur, we need to estimate the angle of the motion blur Θ and the blur length L . We consider a linear motion model not because it is the most accurate model of

motion (it isn't), but because it is straightforward to work with algorithmically, and yields excellent results when deblurring for face recognition. According to Equation 1, in order to recover the original image, $f(x, y)$, we need to estimate the motion blur PSF, $d(x, y)$, and convolve the degraded image, $g(x, y)$, with the inverse of the PSF, in addition to estimating the additive noise $n(x, y)$. Since the noise function is usually stochastic, we can just estimate the magnitude or the signal-to-noise ratio (SNR). Assuming a linear motion blur, the PSF can be formulated using the motion blur angle Θ and the motion blur length L using the following equation:

$$d(x, y) = \begin{cases} \frac{1}{L} & \text{if } 0 \leq |x| \leq L \cos \Theta; y = L \sin \Theta \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

The frequency response of $d(x, y)$ is a sinc function and is given by the following equation.

$$D(w_x, w_y) = \text{sinc}(\pi L(w_x \cos \Theta + w_y \sin \Theta)) \quad (3)$$

As shown in Equation 3, the motion blur angle Θ and the motion blur length L are also preserved in the frequency domain representation of the motion blur PSF. The sinc function has periodic zeros according to the blur length L and at an orientation that corresponds to the blur angle Θ . Therefore, we can identify blur parameters by detecting the period and direction of the periodic zeros in the sinc function.

A. Motion Blur Parameter Identification using the Cepstrum of the Image

A method for identifying these periodic zeros in the sinc function, and thus the motion blur parameters, is to use the two-dimensional Cepstrum of the blurred image [6]. The Cepstrum of the blurred image $g(x, y)$ is given by the formula:

$$C(g(x, y)) = F^{-1}(\log |F(g(x, y))|) \quad (4)$$

where F is the Fourier transform.

An interesting aspect of the Cepstrum is that it adheres to the property of additivity under convolution. If we disregard the additive noise, then:

$$C(g(x, y)) = C(f(x, y)) + C(d(x, y)) \quad (5)$$

Biemond et al. show in [15] that $C(d(x, y)) = F^{-1}(\log \|D(x, y)\|)$, where $D(x, y)$ is the Fourier transform of $d(x, y)$, has two large negative spikes at a distance L from the origin which can be used to estimate the motion blur length. However, the motion blur direction must be estimated first.

B. Motion Blur Direction Identification

The identification of blur direction is based on the fact that the Cepstrum of a non-blurred original image is isotropic, and the Cepstrum of a motion blurred image is anisotropic as shown in Figure 2. It can be seen from Figure 2 that the anisotropy in the Cepstrum is perpendicular to the motion blur angle (in this example, 45°). The Hough transform can be used to detect the orientation of the anisotropy in the Cepstrum.

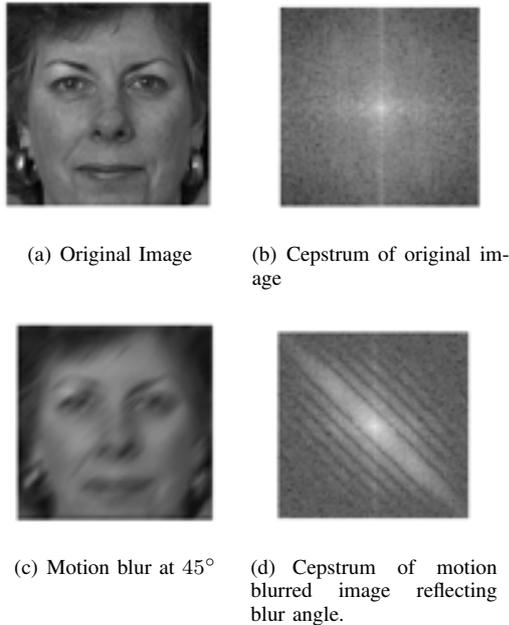


Fig. 2. Images and corresponding Cepstrums.

To reduce the computational time of the Hough transform, the Cepstrum is converted to a binary image by thresholding using the most maximum value of the Cepstrum divided by 128, which can be accomplished using a shift operation.

The Hough transform returns an accumulator array in which the maximum value should correspond to the blur direction. However, our experiments have shown that when the maximum value does not exactly correspond to the blur direction, the 2nd – 5th entry does correspond to the blur direction. Therefore, we return up to 5 angle estimates to increase the accuracy of the system. The inaccurate angle estimates are eliminated during the image restoration phase of the algorithm.

C. Motion Blur Length Identification

Once the candidate motion directions have been determined, the Cepstrum is rotated in the direction opposite to the motion direction. As noted in Section II-A, the Cepstrum of a blurred image will contain two significant negative peaks at a distance L from the origin.

To reduce the noise in the 2D Cepstrum and provide a more accurate estimation we are collapsing the 2D Cepstrum into a 1D signal. The 1D Cepstrum also contains two pronounced negative peaks corresponding to the motion blur length. Most motion blur length estimations that use the Cepstrum for the length estimation either use the location of the largest negative peak, global minimum method, [6], [16] or the position of the first negative value or zero, first crossing method, as the length estimate [17], [18]. However, we have experimentally determined that both estimates can be used to provide more accurate motion blur length estimations. Our research has shown that for blur lengths less than 12 pixels the first negative value or zero for the length estimate is more accurate, while,

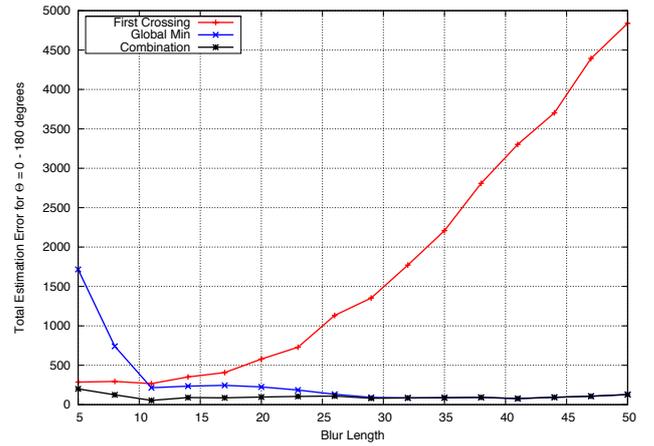


Fig. 3. Results of our Length Estimation Experiments on Motion Blurred Images. Curves express estimation error.

for motion blur lengths greater than 12 pixels the location of the largest negative peak is more accurate as the motion blur length estimation. An extensive review of the motion blur compensation literature did not produce any papers that documented this fact. Figure 3 shows the results of our length estimation experiments on motion blurred images.

For every motion blur length, Θ was iterated from 0° to 180° and 5 length estimates were generated using both the first crossing and global minimum methods for each Θ value. The length estimate that was the closest to the real blur length was used as the estimation error. For example, if the blur length was 5 and the estimated blur length was 4 or 6 then the error is 1. From Figure 3 it can be seen that using both techniques to estimate the motion blur length significantly reduces the error in estimating the motion blur length. The use of both estimates leads to up to 10 total estimates: 5 angle estimates with 2 length estimates per angle estimate. However, all but 2 complete estimates consisting of 1 angle and 1 motion blur length estimate will be eliminated during the image restoration phase of the algorithm.

D. Image Restoration

A Hanning window is applied to the image before restoration to reduce the frequency effects or ringing in the output image due to discontinuities at the edge of the image. A Wiener Filter can be used to perform the image deconvolution once the motion blur angle and length have been estimated. Equation 6 below shows the Wiener Filter formula.

$$R_{Wiener}(u, v) = \frac{H^*(u, v)}{|H(u, v)|^2 + \frac{1}{SNR}} \quad (6)$$

An alternative to the Wiener filter is the Constrained Least Squares (CLS) filter as presented in [19]. The CLS filter replaces the power Cepstrum ratio or $\frac{1}{SNR}$ parameter with a function that varies with frequency. The CLS filter helps eliminate some of the oscillations or “waves” in the output



Fig. 4. The effects of motion blur (top row) and deblurring using the first crossing (middle row) and global min (bottom row) techniques on blurs of 10 pixels (left column), 15 pixels (middle column), and 20 pixels (right column)

image by including a smoothing criterion. The CLS filter formula is shown below in Equation 7:

$$R_{CLS}(u, v) = \frac{H^*(u, v)}{|H(u, v)|^2 + \gamma|P(u, v)|^2} \quad (7)$$

where γ controls how much low-pass filtering occurs and $P(u, v)$ is the Fourier transform of the smoothness criterion function.

Currently, we are using Laplacian cross mask as our smoothness criterion function. The Laplacian cross mask corresponds to a high pass filter; however, since it appears in the denominator of the Wiener filter formula, it acts as a low pass filter. We make use of the SNR-estimation technique presented in [20]. Finally, to eliminate the processing of up to 10 recovered images, we use a computationally effective image metric [21] to select the two “best” images or the images deblurred with the PSF that was closest to the PSF of the blurring function. The two “best” images are then sent to the facial recognition core. The final recognition result comes from the image that produced the highest rank 1 score. Figure 4 shows the effect of motion blur and the results of deblurring using the presented methods.

III. ATMOSPHERIC BLUR MODEL AND RESTORATION

Atmospheric disturbances that can be attributed to the scattering and absorption of particles in the atmosphere and optical turbulence can also have significant impact on facial recognition performance. Our approach for atmospheric blur compensation is a linear systems approach in which a predetermined atmospheric modulation transfer function (AMTF) is deconvolved with the original image. The AMTF that we are using to model the atmospheric blur is based on a PSF that



Fig. 5. The effect of moderately severe atmospheric blur (left) and deblurring using our technique (right)

is within the Levy “stable” density family [22] whose Fourier transform is defined by:

$$F\{h\}(u, v) = e^{-\lambda(u^2+v^2)^\eta}; \quad (8)$$

$$\begin{aligned} (u, v) &\in \mathbb{R}^2, \\ \lambda &> 0, \\ 0 &< \eta \leq 1 \end{aligned}$$

where u and v are the frequency variables, λ controls the severity of the blur, and the value of η is an experimentally determined constant [22]. As the value of λ increases, the blur becomes stronger and when $\lambda = 0$ there is no blur. We are using multiple perturbations of the atmospheric parameters. For each blurred input image we are generating 10 deblurred estimates by perturbing the λ parameter. The images are recovered using the Wiener filter in Equation 6 and the estimated SNR. Finally, to eliminate the processing of all perturbed images, we use the same image metric described in Section II-D to select the “best” image or the image deblurred with the λ that was closest to the λ of the blurring function. A single image statistic, minimum kurtosis, as presented in [13], was also evaluated, however our image metric proved to be a more robust and accurate estimator of the λ parameter. Finally, in [13], the range of atmospheric parameters is set by the user and the range is usually computed by trial and error. The advantage of our algorithm is that the range of λ is recomputed based on the λ of the “best” image, thereby providing an adaptive atmospheric deblurring filter. Figure 5 shows an example of deblurring on an atmospherically blurred image.

IV. RESULTS AND SIGNIFICANCE

A. Facial Recognition

The deblurring methodology described above can be used as a pre-processing module for any facial recognition system. Here we briefly describe the facial recognition approach used to evaluate the impact of blur and our subsequent application of the deblurring algorithms. All images, both probe and gallery, were first geometrically normalized and the faces extracted from the surrounding images using the preprocessing module from the CSU Face Identification Evaluation Toolkit version 5.0 [23]. This resulted in face chips of uniform sizes with uniform orientation to reduce the variation between gallery and probe images due to factors other than those we were testing. For our synthetic data experiments, to blur the

probe images, each face chip was convolved with a filter synthesizing the effects of motion or atmospheric (depending on the test). This final image then became the input to the deblurring phase. The methods proposed in this study were implemented in a high-level prototyping language and run prior to decomposing each image into features to train or test a Support Vector Machine (SVM) classifier. Given a set of blur parameters and an associated model, a deconvolution filter was generated as described in Sections II and III. Each blurred image was then convolved with the filter (a multiplication in frequency space) resulting in a deblurred image. This final image was used as input to the recognition core.

The recognition technique utilized is an augmented form of the technique published by Pinto et al. in [14]. Each gallery image is first filtered by an array of 96 Gabor filters, generating a large array of feature vectors. PCA is used to reduce the dimensionality of these feature vectors prior to using them to train a multiclass SVM. Due to the nature of this method of classification, several gallery images were used for each class so as to increase the accuracy of the SVM’s convergence. In the model of Pinto et al., the probe images are treated the exact same way, with each resulting feature vector classified by the trained SVM. It is in this stage or prior where the authors deblur the probe image before feeding it into the recognition pipeline. This reduces our problem to the same problem as in [14]. This algorithm was chosen for its relative simplicity and excellent baseline performance on popular data sets.

B. Experiments with Synthetic Blur

Using blurred images as probes in a facial recognition situation presents a serious issue that is known to have adverse effects. Several recent studies [24], [25] in facial recognition have treated this issue from various perspectives. Our perspective involves estimating the degree to which the blur has affected the original image and attempting to reverse these effects, as detailed in Sections II and III. These methods are completely automated, requiring no additional human interaction, and demonstrate significant improvement over related works and our baseline, which attempts to match blurred faces against a clean gallery.

To test these deblurring methods, we employed a complete experiment-oriented facial recognition pipeline (shown in Figure 1). The tests were set up such that sets of clean gallery images from public datasets were used to train multiclass SVMs. For each data set, a set of images was chosen to be unique from all gallery images. These images were synthetically blurred as described below. These processed images were then used as probes to test the trained classifier and generate results indicative of the performance of our deblurring methods.

To determine recognition metrics, subsets of two public datasets, FERET [26] and Labeled Faces in the Wild (LFW) [27], were employed. Each was prepared in such a way as to provide the maximum comparability to both the original intent of the dataset as well as with each other. Both datasets were geometrically normalized, but otherwise left unperturbed,

TABLE I
RANK 1 RECOGNITION RESULTS FOR BASELINE AND MOTION DEBLURRED FERET240.

BLUR	None	10px	15px	20px
Baseline blurred	97.50	75.00	39.58	16.67
Deblurred	–	92.89	93.75	86.67

prior to blurring the probes. For motion blur, recognition was tested at blur lengths of 10, 15, and 20 pixels and at integral angles uniformly distributed in the range $0 < \Theta \leq \pi$. For atmospheric blur, recognition was tested for $\lambda = 0.09$.

The FERET subset chosen (dubbed “FERET240”) was determined by choosing the subjects for whom the full set contained four or more images (giving us a sufficient amount of training data for the gallery). Of these, the first three, determined by an alphabetic sort, were utilized as gallery; the fourth in the listing was used as probe. This subset contained 240 subjects and 960 face chips.

In order to use LFW, a protocol identical to that used for FERET was chosen. This varies from the protocol defined in [27] in that ours is tailored to the recognition problem, whereas the original is tailored to the verification problem. Subjects were chosen based on whether or not the original dataset contained four or more images, as with FERET. The first three, given by an alphabetic sort, were chosen as gallery; the fourth was chosen as probe. This subset (dubbed “LFW610”) contained 610 subjects and 2440 face chips.

C. Controlled Motion Blur

For each dataset, a blurred version of each probe was created at lengths of 10, 15, and 20 pixels, representing the range wherein motion blur severely crippled recognition on otherwise unprocessed probes. A suite of experiments utilizing five different sets of parameters was run on each dataset and blur level. For comparison, a baseline test was conducted on each blur length using the method of Pinto et al. [14] without performing any blur correction.

As the FERET data set is a fairly well-behaved data set (consisting of frontal images of unoccluded faces with consistent lighting), results were also fairly well-behaved. The rank 1 results for each experiment are summarized in Table I. As expected, recognition rates dropped as blur increased, due to the increased possibility for error in estimating the blur, with an error in recognition of 4.61%, 3.85%, and 11.12% for blurs of length 10, 15, and 20 pixels, respectively. However, deblurring the probes before attempting recognition demonstrated a marked increase in the percentage of probes recognized, at 23.85% (10 pixels), 136.86% (15 pixels), and 419.92% (20 pixels). The comparative cumulative match curves (representing percentage improvement of the deblurring over the baseline for increasing percentages of the total gallery) are shown in in Figure 6.

Labeled Faces in the Wild (LFW), on the other hand, is an unconstrained dataset by its very nature, so the results of deblurring were considerably lower. In [14], Pinto et al.

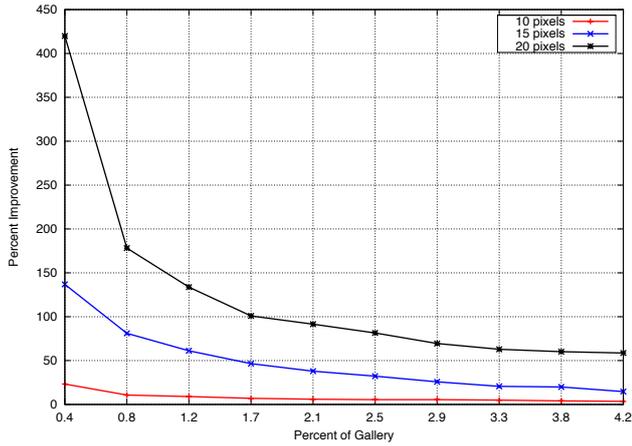


Fig. 6. Comparative cumulative match curves showing percent improvement for our deblurring applied to images from the FERET240 set at three different levels of blur. Note the increasing improvement as the blur levels increase.

TABLE II
RANK 1 RECOGNITION RESULTS FOR BASELINE AND MOTION DEBLURRED LFW610.

BLUR	None	10px	15px	20px
Baseline blurred	41.64	19.34	10.33	5.08
Deblurred	–	33.11	24.26	16.89

treat the set from the perspective of a 1-to-1 face verification problem, using a protocol similar to that laid out in [27]. This protocol is designed to test the problem of whether or not two faces presented are of the same person. Using this technique, they demonstrate a 79.35% rate of success. This result is quite good, considering that chance for the verification protocol is 50%.

However, this paper approaches the dataset from the perspective of a recognition problem, testing the problem of

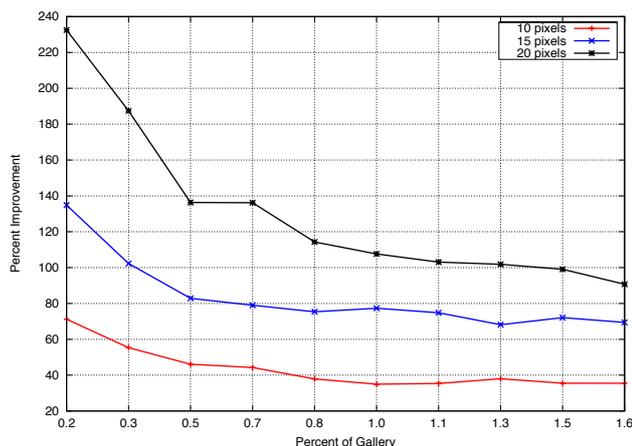


Fig. 7. Comparative cumulative match curves showing percent improvement for our deblurring applied to images from the LFW610 set at three different levels of blur. Note the increasing improvement as the blur levels increase.

TABLE III

RANK 1 RECOGNITION RESULTS ON ATMOSPHERIC DEBLURRED SETS. $\lambda = 0.09$ IN EQUATION 8, REPRESENTING MODERATELY SEVERE BLUR.

	No Blur	Blurred	Deblurred
FERET240	97.50	37.08	90.68
LFW610	41.64	9.34	37.54

classifying a picture as having been taken of a particular subject (out of a gallery of several). Thus, the problem is considerably harder; chance for recognition is 1 in N , where N is the number of subjects in the gallery. As LFW is considered to be one of the hardest set of faces for recognition, it is expected that recognition rates will be considerably lower than those obtained on other, more constrained, datasets. Thus, the numbers presented may appear low compared to the FERET240 results (fulfilling expectations), but are, in fact, demonstrative of a marked performance increase due to deblurring.

Table II summarizes the results obtained from attempting a recognition-style experimental protocol on LFW610. As described above, recognition on a clean copy of the dataset is 41.64%, which is quite a bit lower than the 97.50% obtained on FERET240. However, as the effect of blur further reduces this rank 1 recognition rate, it can easily be seen that deblurring provides a considerable increase in recognition rates at 71.2% for a blur of 10 pixels, 134.85% at 15 pixels, and 232.48% at 20 pixels. Due to the difficulty of the dataset, it is not surprising that the recognition error due to blur and subsequent estimation error is 20.49% (10 pixels), 41.74% (20 pixels), and 59.44% (20 pixels), which is considerably greater than the respective numbers for FERET240. The comparative cumulative match curves are shown in Figure 7.

D. Controlled Atmospheric Blur

For each dataset, the atmospheric blur correction technique was tested with moderately severe synthetic atmospheric blur ($\lambda = 0.09$ in Equation 8). Two sets of tests were performed: one on synthetically blurred probes to establish a baseline and one on deblurred versions of these probes. The rank 1 results are summarized, along with the results of recognition on clean probes, in Table III.

On FERET240, the atmospheric deblurring technique presented demonstrated a relatively minimal recognition error due to blur at a mere 7%, while exhibiting a drastic 144.55% increase in recognition rate over that of the blurred probes. The results for LFW610 were also very strong with a 301.93% increase in recognition rate over the blurred probes, while only suffering a 10.92% error due to the blur. Figure 8 presents comparative cumulative match curves for the atmospheric blur experiments for each of the two datasets.

E. Experiments with Real Blur

In order to further validate our approach, we processed several videos of subjects in motion in an outdoor setting at 100M. In total, three different videos were processed,

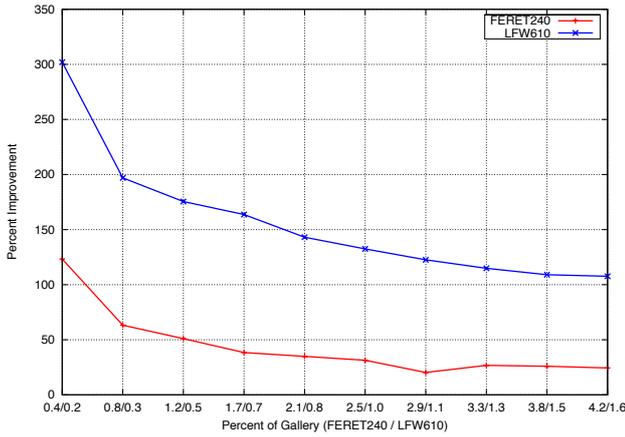


Fig. 8. Comparative cumulative match curves showing percent improvement for atmospheric blur experiments.



Fig. 9. Examples of original imagery from outdoor video (top row) and the resulting deblurred imagery for comparison (bottom row)

containing 85 frames, 125 frames, and 82 frames respectively. The gallery consisted of a small watchlist of 7 subjects (typical of small-scale surveillance applications), including the subjects in each video. Each frame was processed with exactly the same procedure as the above experiments. Only motion deblurring was applied to each frame. Examples of the original source imagery and our deblurred imagery are given in Figure 9. The rank 1 recognition results are given in Table IV. From this table, we see a dramatic performance improvement on real world data in all three cases. Thus, we conclude that our proposed deblurring approach is suitable for live face recognition applications.

TABLE IV
RANK 1 RECOGNITION RESULTS FOR FRAMES FROM VIDEOS FOR THREE DIFFERENT SUBJECTS IN MOTION. THE GALLERY WAS A SIMULATED SURVEILLANCE WATCHLIST OF 7 SUBJECTS.

	Subject 1	Subject 2	Subject 3
Blur	0.0	51.20	4.88
Deblurred	70.59	98.40	70.73

V. CONCLUSION

In this paper, we have presented a set of techniques for dealing with two types of blur evident in the real world insofar as their estimation and correction in the context of a facial recognition pipeline. We have presented a technique for estimating motion blur and a technique for estimating atmospheric blur and tested both on synthetically blurred data generated from two publicly-available datasets, one fairly well-behaved (FERET) and one unconstrained dataset (LFW). We also processed a series of videos containing real motion blur in a live outdoor setting. We have demonstrated a significant increase in recognition rates as a direct result of our deblurring techniques over the baseline recognition on the source images. Future work includes further study into methods of detecting whether or not blur is present, as well as determining which of several deblurring methods is appropriate for a given image. In addition, this study provides opportunity for using similar deblurring techniques with other recognition core implementations.

ACKNOWLEDGMENT

This work was supported by ONR STTR "Biometrics in the Maritime Domain," (Award Number N00014-07-M-0421) and ONR MURI (Award Number N00014-08-1-0638).

REFERENCES

- [1] J. Chen, L. Yuan, C. Tang, and L. Quan, "Robust dual motion deblurring," in *CVPR*, 2008, pp. 1–8.
- [2] L. Yuan, J. Sun, L. Quan, and H.-Y. Shum, "Image deblurring with blurred/noisy image pairs," *ACM Trans. on Graphics*, vol. 26, no. 3, pp. 1–10, 2007.
- [3] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman, "Removing camera shake from a single photograph," *SIGGRAPH*, 2006.
- [4] J. Jia, "Single image motion deblurring using transparency," in *CVPR*, 2007.
- [5] A. Levin, Y. Weiss, F. Durand, and W. Freeman, "Understanding and evaluation blind deconvolution algorithms," in *CVPR*, 2009.
- [6] M. Cannon, "Blind deconvolution of spatially invariant image blurs with phase," *IEEE T. on Acoustics, Speech and Signal Processing*, vol. 24, no. 1, pp. 58–63, 1976.
- [7] D. Kundur and D. Hatzinakos, "Blind image deconvolution," *IEEE Signal Process. Mag.*, vol. 13, pp. 43–64, 1996.
- [8] A. Jalobeanu, J. Zerubia, and L. Blanc-Feraud, "Bayesian estimation of blur and noise in remote sensing imaging," in *Blind Image Deconvolution: Theory and Applications*, P. Campisi and K. Egiazarian, Eds. CRC Press, 2007.
- [9] O. Shacham, O. Haik, and Y. Yitzhaky, "Blind restoration of atmospherically degraded images by automatic best step-edge detection," *PRL*, vol. 28, no. 15, pp. 2094–2103, 2007.
- [10] A. E. Savakis and H. J. Trussell, "Blur identification by residual spectral matching," *IEEE Trans. Image Process.*, pp. 2141–2151, 1993.
- [11] G. Pavlovic and A. M. Tekalp, "Maximum likelihood parametric blur identification based on a continuous spatial domain model," *IEEE Trans. Image Process.*, pp. 496–504, 1992.
- [12] D. G. Sheppard, H. Bobby, and M. Michael, "Iterative multi-frame super-resolution algorithms for atmospheric turbulence-degraded imagery," *J. Opt. Soc. Am.*, pp. 978–992, 1998.
- [13] L. Dalong and S. Simske, "Atmospheric turbulence degraded-image restoration by kurtosis minimization," *IEEE Geoscience and Remote Sensing Letters*, vol. 6, no. 2, pp. 244–247, 2009.
- [14] N. Pinto, J. J. DiCarlo, and D. D. Cox, "How far can you get with a modern face recognition test set using only simple features?" in *IEEE CVPR*, 2009.
- [15] J. Biemond, R. Lagendijk, and R. Mersereau, "Iterative methods for image deblurring," *Proceedings of the IEEE*, vol. 78, no. 5, pp. 856–883, 1990.

- [16] W. Fawwaz A M, T. Shimahashi, M. Matsubara, and S. Sugimoto, "PSF estimation and image restoration for noiseless motion blurred images," in *SSIP*, 2008, pp. 1–7.
- [17] R. Lokhande, K. Arya, and P. Gupta, "Identification of parameters and restoration of motion blurred images," in *ACM. Symp. on Applied Computing*, 2006, pp. 301–315.
- [18] S. Wu, Z. Lu, E. Ong, and W. Lin, "Blind image blur identification in cepstrum domain," in *ICCCN*, 2007, pp. 1166–1171.
- [19] K. Rajeswari, K. Krishna, V. Naveen, and A. Vamsidhar, "Performance comparison of wiener filter and cls filter on 2d signals," in *ITNG*, 2009, pp. 1244–1249.
- [20] J. Thong, K. Sim, and J. Phang, "Single-image signal to noise ratio estimation," *Scanning*, vol. 23, pp. 328–336, 2001.
- [21] T. Boulton and W. Scheirer, "Long range facial image acquisition and quality," in *Handbook of Remote Biometrics*, M. Tistarelli, S. Li, and R. Chellappa, Eds. Springer, 2009.
- [22] P. Qiu, "A nonparametric procedure for blind image deblurring," *Computational Statistics and Data Analysis*, vol. 52, pp. 4828–4841, 2008.
- [23] D. Bolme, R. Beveridge, M. Teixeira, and B. Draper, "The CSU Face Identification Evaluation System: Its purpose, features and structure," in *ICVS*, 2003, pp. 304–311.
- [24] Q. Han, M. Zhang, C. Song, Z. Wang, and X. Niu, "A face image database for evaluating out-of-focus blur," in *ISDA*, vol. 2, 2008, pp. 277–282.
- [25] M. Nishiyama, H. Takeshima, J. Shotton, T. Kozakaya, and O. Yamaguchi, "Facial deblur inference to improve recognition of blurred faces," in *CVPR*, 2009, pp. 1115–1122.
- [26] P. J. Phillips, H. Moon, P. J. Rauss, and S. Rizvi, "The FERET evaluation methodology for face recognition algorithms," *IEEE TPAMI*, vol. 22, no. 10, 2000.
- [27] G. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," University of Massachusetts, Amherst, Tech. Rep. 07-49, 2007.