Removing Atmospheric Turbulence via Space-Invariant Deconvolution

Xiang Zhu, *Student Member, IEEE*, and Peyman Milanfar, *Fellow, IEEE*

Abstract—To correct geometric distortion and reduce space and time-varying blur, a new approach is proposed in this paper capable of restoring a single high-quality image from a given image sequence distorted by atmospheric turbulence. This approach reduces the space and time-varying deblurring problem to a shift invariant one. It first registers each frame to suppress geometric deformation through B-spline-based nonrigid registration. Next, a temporal regression process is carried out to produce an image from the registered frames, which can be viewed as being convolved with a space *invariant* near-diffraction-limited blur. Finally, a blind deconvolution algorithm is implemented to deblur the fused image, generating a final output. Experiments using real data illustrate that this approach can effectively alleviate blur and distortions, recover details of the scene, and significantly improve visual quality.

Index Terms—Image restoration, atmospheric turbulence, nonrigid image registration, point spread function, sharpness metric

✦

---

1 INTRODUCTION

A long-distance imaging system can be strongly affected by atmospheric turbulence, which randomly changes the refractive index along the optical transmission path, generating geometric distortion (motion), space and time-varying blur, and sometimes even motion blur if the exposure time is not sufficiently short [1], [2], [3], [4], [5]. Aside from hardware-based adaptive optics approaches [6], several signal processing approaches have been proposed to solve this problem [7], [8], [4], [5], [9], [10]. These approaches attempt to restore a single high-quality image from an observed frame sequence distorted by air turbulence. As with these other works based on videos or image sequences, we work under the assumption that the scene and the image sensor are both static and that observed motions are due to the air turbulence alone. The imaging process can be modeled as [11], [10]

$$G_k[\mathbf{x}] = (F \otimes h_{k,\mathbf{x}} \otimes h)[\mathbf{x}] + N_k[\mathbf{x}], \tag{1}$$

where $\otimes$ represents a 2D convolution operator. F, G_k, and N_k denote the ideal image, the kth observed frame, and sensor noise, respectively.[1] The vector $\mathbf{x} = (x, y)^T$ denotes a 2D spatial location. $h_{k,\mathbf{x}}$ represents the space-varying (air turbulence-caused) point spread function (PSF) for the position $\mathbf{x}$ in the kth frame, which includes both blur and motion effects [10]. h is the space-invariant diffraction-limited PSF, typically due to sensor optics and settings.

---

1. In this paper, we use capital letters (e.g., F, G, N) to denote image-wise models, and lowercase (e.g., f, g, n) to represent patch-wise models.

- *The authors are with the Department of Electrical Engineering, University of California, Santa Cruz, 1156 High St., Santa Cruz, CA 95064. E-mail: {xzhu, milanfar}@ee.ucsc.edu.*

Because the unknown $h_{k,\mathbf{x}}$ is *spatially and temporally changing* in (1), restoring the high-quality image F is not trivial. Some multiframe reconstruction approaches [4], [9] first employ a nonrigid image registration technique to register each observed frame with respect to a fixed reference grid, and use the registration parameters to estimate the corresponding motion field for each frame. Then, a sharp image is formed through a Bayesian reconstruction method. The main problem for such algorithms is that they do not estimate the actual PSF $h_{k,\mathbf{x}}$. Both [4] and [9] employ a fixed Gaussian model to approximate the PSF, which strongly limits their performance.

Recently, a method called Efficient Filter Flow (EFF) was introduced by Hirsch et al. for space-varying blind deconvolution and has been applied for astronomical imaging to alleviate turbulence distortion and also for magnetic resonance imaging (MRI) to reduce blur caused by object motion [10]. This method first divides each frame into overlapping patches. Because the size of these patches is small, they can be viewed as *isoplanatic* regions—small regions containing space-*invariant* blur [6], [5]—and thus can be processed through a multiframe blind deconvolution algorithm [12]. Given an isoplanatic patch sequence of a local region extracted from the input video, the blind deconvolution algorithm estimates the PSF *separately for each patch*. Final output is then generated by a nonblind image estimation step. Though capable of suppressing turbulence effects, the EFF method cannot remove diffraction-limited blur since the blind deconvolution step does not utilize much prior knowledge of the ideal sharp image except for a nonnegativity constraint and a simple Tikhonov regularizer. Besides, due to the limited accuracy of the local PSF estimation and probably also due to the sensor noise effect, the results can contain strong deblurring artifacts, such as ringing [12].

Another class of approaches called "lucky imaging" employs image selection and fusion methods to reduce the blurring effects caused by turbulence [13], [14], [8], [5]. The image selection technique attempts to find frames of the best quality (lucky frames) from a short-exposure video stream. The output image is produced by fusing these lucky frames together [13], [14]. This method is based on the

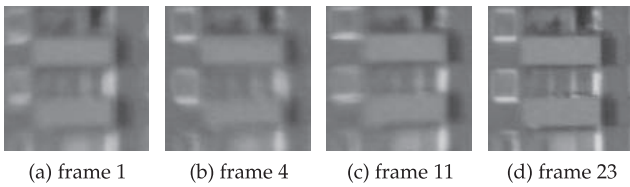| (a) frame 1 | (b) frame 4 | (c) frame 11 | (d) frame 23 |

Fig. 1. Isoplanatic patches from the short exposure image sequence *Building* taken through hot air, where a variation of blur can be observed. (d) illustrates a lucky frame example which is much sharper than the others.

observation that for short-exposure images, turbulence creates "mutations" in image quality and randomly makes some images sufficiently sharp (see examples in Fig. 1) [14], [15], [16], [17]. This strategy is favored in many astronomical imaging applications where the image of the object of interest (e.g., a star) is usually taken inside an *isoplanatic* angle. In [17], Vorontsov et al. proposed a "lucky region" restoration approach for *anisoplanatic* scenarios. Small lucky (isoplanatic) regions are detected by a local sharpness metric (which can be viewed as being blurred only by the diffraction-limited PSF) and are fused to produce a large high-quality image. In another similar method developed by Joshi and Cohen [18], a local block-based image alignment is first carried out to reduce geometric distortion caused by turbulence; then a lucky imaging-based weighting scheme is employed to generate a single image that is sharp everywhere. The weighting scheme is carefully designed to balance between noise reduction and sharpness preservation. A dehazing process is finally used to enhance the visual quality. One difficulty with such a method is that even though turbulence-caused blur is strongly alleviated through the lucky imaging process, the output still suffers from the blur caused by the diffraction-limited PSF [5], [11]. Besides, the blur in lucky regions would inevitably be increased once temporal averaging is used for noise suppression [18].

Due to the isoplanatism of turbulence effects (including blurring and geometric deformation) in small local regions, lucky imaging focuses on patch-wise restoration. However, limited patch size also limits the restoration performance. As Levin et al. pointed out in [19], PSF estimation can be much improved as the size of the observed patch increases. In other words, if PSF can be estimated from a large region (e.g., the whole image), then we may be able to achieve better deconvolution performance. However, to achieve this goal the PSF spatial variation needs to be reduced.

Another important factor that may affect restoration is sensor noise. High shutter speed is favored in long-distance imaging systems in general to avoid motion blur caused by turbulence, and this would inevitably increase the noise level. Noise can strongly affect the performance of deblurring as it distorts image details and produces ringing artifacts [20]. Fortunately, the frames within a sequence are highly correlated with each other, and such a correlation can be used to advantage denoising in the temporal domain.

In this paper, a new framework is proposed for restoring a single image from an image sequence acquired in general anisoplanatic scenarios. The 3D physical scene is assumed to be static, as is the imaging sensor, while the air between the scene and sensor is affected by atmospheric turbulence. Our approach is designed to reduce the spatial variation of PSFs over the whole image space so that the blur can be approximately treated as *spatially invariant* and the latent image content can be estimated globally instead of by local deconvolution. By doing this, we can improve the estimation accuracy and, more importantly, natural image statistics can be introduced so that diffraction-limited blur can be effectively removed. An image with reduced PSF variation (which we call the *near-diffraction-limited (NDL) image*) is generated through a fusion process. To avoid noise effects in the subsequent deconvolution, temporal kernel regression is employed in our approach for denoising.

Note that this restoration framework can be viewed as an extended version of the one in our earlier paper in ICCP 2011 [21]. However, the fundamental ideas behind the two are quite different. The framework in [21] is based on a simplified least-square estimation derived from the imaging model with several assumptions on image statistics and approximations. The key point of the present paper is reducing the variance of spatially and temporally changing PSFs through *nonparametric kernel regression*. Compared with [21], the idea in this paper is much cleaner and technically sounder. It also requires far fewer a priori assumptions than [21].

The paper is organized as follows: Section 2 describes the restoration framework in detail. Experimental results are given in Section 3 to show the performance as compared with other methods, and finally we conclude and discuss directions of future research in Section 4.

2 RESTORATION ALGORITHM

The proposed restoration framework contains three main steps (see the diagram in Fig. 2):

1. nonrigid image registration,
2. near-diffraction-limited image reconstruction,
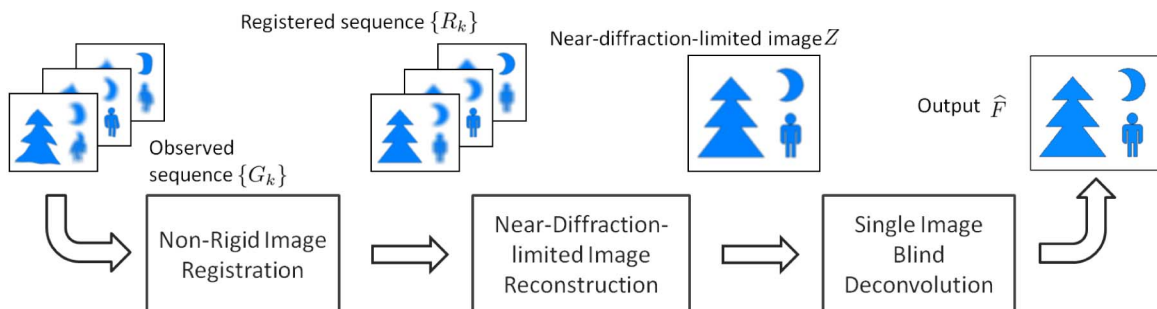3. single image blind deconvolution.

Fig. 2. Block diagram for the proposed restoration framework.

Given an observed sequence $\{G_k\}$, step 1 of our proposed approach registers each frame onto a fixed reference grid, generating a new sequence $\{R_k\}$ without geometric deformation. This process makes sure that for a given pixel located at a position, say $\mathbf{x}$, the most correlated pixels in each frame across time are aligned at this same position. This step improves the performance of the subsequent temporal filtering.

Step 2 restores a single image Z from the registered $\{R_k\}$. For each local region, the sharpest patch (convolved by a diffraction-limited PSF) can be detected from the patch sequence. Next, patch-wise temporal kernel regression is carried out to reduce the noise level. An image Z is then generated by fusing all the denoised sharp patches together. This image is still blurred by a diffraction-limited PSF, which can be approximately viewed as *spatially invariant*.

Finally, a single image blind deconvolution algorithm based on natural image statistics is implemented on Z to further remove the diffraction-limited blur and to enhance image quality. Details of each step are given in the following sections.

2.1 Nonrigid Image Registration

In [9], we introduced a B-spline based registration algorithm to estimate the motion field in each observed frame. This method can be implemented in the proposed framework to remove geometric deformation. In this section, we give a review of this registration algorithm. Assume G represents a given observed distorted image and R denotes a reference image without turbulent deformation (which can be obtained by averaging the frame sequence [4]). A nonrigid deformation model can be employed to describe the geometric distortion between these two images. In this model, the complex motion is represented by the movement of m control points, whose initial positions $\hat{\mathbf{x}}_{0i} = (\hat{x}_{0i}, \hat{y}_{0i})^T$, $i = 1, \ldots, m$, are equally spaced on the reference image R. The displacement of all the control points on the given image G is denoted as the deformation vector: $\overrightarrow{\mathbf{p}} = [\Delta \hat{x}_1, \ldots, \Delta \hat{x}_m, \Delta \hat{y}_1, \ldots, \Delta \hat{y}_m]^T$, where $\Delta \hat{x}$ and $\Delta \hat{y}$ (also called deformation parameters) are the horizontal and vertical displacement from the initial position [22], [4]. The deformed position of any given pixel located at $\mathbf{x} = (x, y)^T$ from image R can then be described as

$$\mathbf{W}(\mathbf{x}; \overrightarrow{\mathbf{p}}) = \mathbf{x} + \mathbf{A}(\mathbf{x})\overrightarrow{\mathbf{p}}, \tag{2}$$

where $\mathbf{A}(\mathbf{x})$ denotes the basis function matrix for $\mathbf{x}$:

$$\mathbf{A}(\mathbf{x}) = \begin{bmatrix} c_1 & \ldots & c_m & 0 & \ldots & 0 \\ 0 & \ldots & 0 & c_1 & \ldots & c_m \end{bmatrix}, \tag{3}$$

and $\mathbf{A}(\mathbf{x})\overrightarrow{\mathbf{p}}$ is the motion vector, which is a linear combination of the movements of all control points. The weight (or spline basis) c_i is determined by the distance between $\mathbf{x}$ and $\hat{\mathbf{x}}_{0i}$ using B-spline functions:

$$c_i = \beta\left(\frac{x - \hat{x}_{0i}}{\epsilon_x}\right)\beta\left(\frac{y - \hat{y}_{0i}}{\epsilon_y}\right), \tag{4}$$

$$\beta(\kappa) = \begin{cases} 2/3 - (1 - |\kappa|/2)\kappa^2, & \text{if } 0 \leq |\kappa| \leq 1 \\ (2 - |\kappa|)^3/6, & \text{if } 1 < |\kappa| < 2 \\ 0, & \text{otherwise,} \end{cases} \tag{5}$$
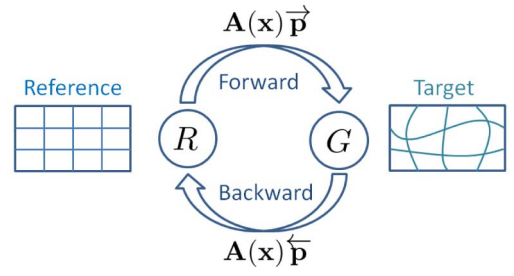

Fig. 3. Symmetry constraint, where forward motion vector and backward motion vector are the inverse of each other.

where ϵ_x and ϵ_y are the horizontal and vertical intervals between neighboring control points. This model guarantees local smoothness of the turbulence-caused motion field.

In the classic B-spline-based registration approach, the deformation vector is estimated by minimizing the following cost function [22]:

$$C(\overrightarrow{\mathbf{p}}) = \sum_{\mathbf{x}} |G(\mathbf{W}(\mathbf{x}; \overrightarrow{\mathbf{p}})) - R(\mathbf{x})|^2. \tag{6}$$

However, because the above formulation lacks prior constraint, the resulting estimate is unstable and easily gets stuck in local minima. To improve this method, a stabilization constraint is proposed in [4], which makes the estimated deformation parameters remain small in the regions that have less texture (low image gradient). However, in the present application, we know that the deformation caused by atmospheric turbulence is independent of image content. In other words, such stabilization constraint would bring unnecessary bias into the estimation.

Instead, to accurately estimate the deformation (motion) vectors from image grid R to G, we introduce a more natural symmetry constraint [23], [24] into the B-spline registration algorithm. This constraint is based on the important property that the registration should be symmetric or inverse consistent [23]. Let $\overrightarrow{\mathbf{p}}$ denote the deformation vector that transforms reference grid R into G (forward deformation) and let $\overleftarrow{\mathbf{p}}$ denote the inverse vector that transforms G into R (backward deformation). Then, approximately,[2] we should have: $\overrightarrow{\mathbf{p}} = -\overleftarrow{\mathbf{p}}$ (see Fig. 3). Combining the two vectors into one, $\mathbf{p}^T = [\overrightarrow{\mathbf{p}}^T, \overleftarrow{\mathbf{p}}^T]$, the proposed cost function to minimize becomes

$$C(\mathbf{p}) = \sum_{\mathbf{x}} |G(\mathbf{W}(\mathbf{x}; \overrightarrow{\mathbf{p}})) - R(\mathbf{x})|^2 + \sum_{\mathbf{x}} |R(\mathbf{W}(\mathbf{x}; \overleftarrow{\mathbf{p}})) - G(\mathbf{x})|^2 + \gamma(\overrightarrow{\mathbf{p}} + \overleftarrow{\mathbf{p}})^T(\overrightarrow{\mathbf{p}} + \overleftarrow{\mathbf{p}}), \tag{7}$$

where scalar γ controls the effect of the soft symmetry constraint.

The Gauss-Newton method is used to minimize $C(\mathbf{p})$, and the update of the parameter set $\mathbf{p}$ can be derived as follows:

$$\mathbf{p}^{l+1} = \mathbf{p}^l - \mathbf{E}^{-1}\mathbf{b}, \tag{8}$$

2. The forward and backward motion vectors for each point in the physical world should be the exact opposite of each other. However, in the registration model the motion field is indexed by pixel locations instead of physical points. Thus, the motion vectors for the same position in the reference and the target image are not exactly the opposite of each other since they belong to slightly different points. So do the deformation vectors.
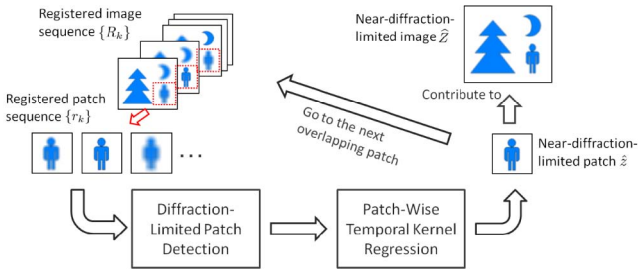
Fig. 4. Block diagram for the near-diffraction-limited image reconstruction step.
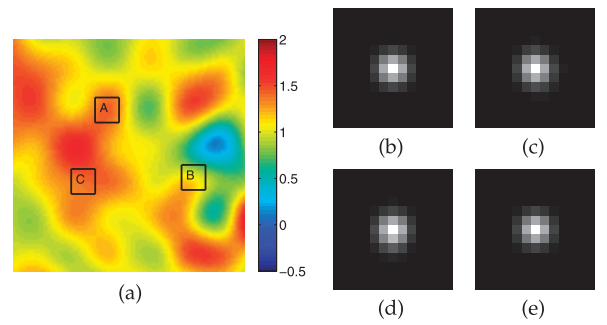

Fig. 5. PSF distortion after the registration step. (a) Magnitude of a motion field estimated from a real frame distorted by air turbulence (video *Moon Surface* in Fig. 12) using the proposed registration method. (b) A Gaussian shaped diffraction-limited PSF before alignment. (c)-(e) PSFs after the alignment using the local motion fields given in the squares A, B, and C, respectively, from (a), where it can be observed that since the local movement could be treated as translational, the shapes of the PSFs are preserved.

where

$$\mathbf{E} = \begin{bmatrix} \overrightarrow{\mathbf{E}} + \gamma\mathbf{I} & \gamma\mathbf{I} \\ \gamma\mathbf{I} & \overleftarrow{\mathbf{E}} + \gamma\mathbf{I} \end{bmatrix},$$

$$\overrightarrow{\mathbf{E}} = \sum_{\mathbf{x}} \overrightarrow{\mathbf{d}}(\mathbf{x})\overrightarrow{\mathbf{d}}(\mathbf{x})^T, \quad \overleftarrow{\mathbf{E}} = \sum_{\mathbf{x}} \overleftarrow{\mathbf{d}}(\mathbf{x})\overleftarrow{\mathbf{d}}(\mathbf{x})^T,$$

$$\overrightarrow{\mathbf{d}}(\mathbf{x})^T = \frac{\partial G(\mathbf{W}(\mathbf{x}; \overrightarrow{\mathbf{p}}^l))}{\partial \mathbf{W}} \mathbf{A}(\mathbf{x}),$$

$$\overleftarrow{\mathbf{d}}(\mathbf{x})^T = \frac{\partial R(\mathbf{W}(\mathbf{x}; \overleftarrow{\mathbf{p}}^l))}{\partial \mathbf{W}} \mathbf{A}(\mathbf{x}),$$

and

$$\mathbf{b}^T = \left[\left(\overrightarrow{\mathbf{b}} + \overrightarrow{\mathbf{p}}^l + \overleftarrow{\mathbf{p}}^l\right)^T, \left(\overleftarrow{\mathbf{b}} + \overrightarrow{\mathbf{p}}^l + \overleftarrow{\mathbf{p}}^l\right)^T \right],$$

$$\overrightarrow{\mathbf{b}} = \sum_{\mathbf{x}} \overrightarrow{\mathbf{d}}(\mathbf{x})\left[G(\mathbf{W}(\mathbf{x}; \overrightarrow{\mathbf{p}}^l)) - R(\mathbf{x}) \right],$$

$$\overleftarrow{\mathbf{b}} = \sum_{\mathbf{x}} \overleftarrow{\mathbf{d}}(\mathbf{x})\left[R(\mathbf{W}(\mathbf{x}; \overleftarrow{\mathbf{p}}^l)) - G(\mathbf{x}) \right].$$

In the above algorithm, the computational cost for calculating the matrices $\overrightarrow{\mathbf{E}}$ and $\overleftarrow{\mathbf{E}}$ is tremendous if it is implemented directly. Instead, we designed a fast implementation method to alleviate this problem. Details can be found in [9]. Once the deformation vector $\mathbf{p}$ is estimated for each frame, a registered sequence $\{R_k\}$ is generated through bilinear interpolation.

2.2 Near-Diffraction-Limited Image Reconstruction

In this section, we estimate a *diffraction-limited* image Z from $\{R_k\}$, which can be globally deconvolved. A concise description of the reconstruction procedure is provided in Algorithm 1 (see the diagram in Fig. 4).

To better understand this procedure, let us consider a *patch-wise* imaging model under the isoplanatic scenario (see Fig. 4):

$$\begin{aligned} g_k &= f \otimes h_k \otimes h + n_k \\ &= f \otimes h \otimes h_k + n_k \\ &= z \otimes h_k + n_k, \end{aligned} \quad (9)$$

where g_k can denote any local patch of size $L \times L$ extracted from frame G_k, and its corresponding latent patch is denoted by f. Differently from (1), turbulence-caused PSF h_k is now patch-wise constant and temporally changing. We call $z = f \otimes h$ the *diffraction-limited* patch, which is convolved by the space and time invariant h, and thus can be accurately deconvolved in a larger scene [19] (such as the whole image). n_k represents local noise, which is assumed to be zero-mean and IID.

The motion field estimated from the nonrigid image registration in Section 2.1 is smooth, and the registration

process can be viewed approximately as patch-wise constant translational movement (see an example given in Fig. 5). So the relationship between an observed local patch g_k and the corresponding registered one r_k extracted from R_k can be described as

$$r_k = g_k \otimes \delta_{\Delta\mathbf{x}}, \quad (10)$$

where $\delta_{\Delta\mathbf{x}}$ represents a 2D Kronecker Delta function shifted by the local registration vector $\Delta\mathbf{x}$. Hence, (9) and (10) can be combined as follows:

$$\begin{aligned} r_k &= z \otimes h_k \otimes \delta_{\Delta\mathbf{x}} + n_k \otimes \delta_{\Delta\mathbf{x}} \\ &= z \otimes \tilde{h}_k + \tilde{n}_k \\ &= q_k + \tilde{n}_k. \end{aligned} \quad (11)$$

The PSF $\tilde{h}_k = h_k \otimes \delta_{\Delta\mathbf{x}}$ is simply a shifted version of h_k, which means the registration process preserves the shape (and spectrum) of local PSFs. Similarly, the shifted noise patch $\tilde{n}_k$ also has the same statistical properties as n_k. As mentioned before, the registration process increases the correlation among the pixels in the same position from different patches across time. Such a correlation can be utilized for estimating the noise-free sequence $\{q_k\}$, as we will describe in Section 2.2.2.

Algorithm 1. Procedure for Restoring A Near-Diffraction-Limited Image from Registered Frames

1. Given a registered sequence $\{R_k\}$, divide each frame into $L \times L$ overlapping patches centered at each pixel, and calculate the intensity variance of each patch as a local sharpness measure.
2. For a patch sequence $\{r_k\}$ centered at location $\mathbf{x}$, detect the sharpest one r_{k^*} by maximizing local sharpness measure (outliers need to be detected and excluded).
3. Set r_{k^*} as a reference patch, and restore its center pixel value using temporal kernel regression. Assign this value to the pixel $Z[\mathbf{x}]$.
4. Go to the next pixel and return to step 2.

2.2.1 Diffraction-Limited Patch Detection

As mentioned in Section 1, once sufficient observations are collected, relatively sharp image patches that occasionally
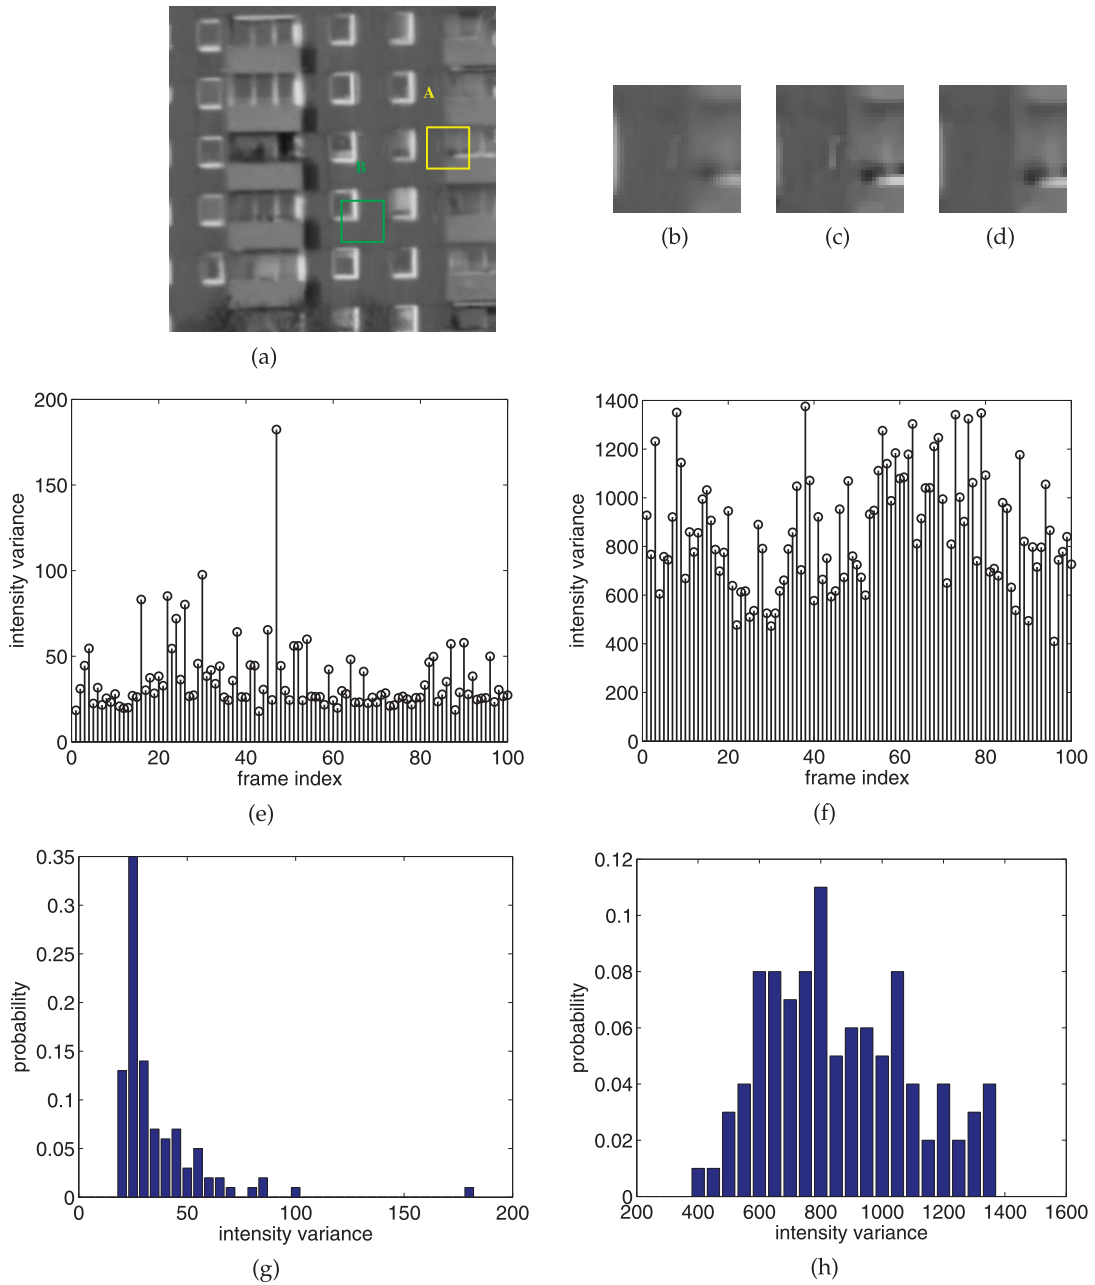
Fig. 6. Outlier identification example. (a) A fused image without outlier removal, where patch A contains outlier pixels and patch B does not contain any outlier, (b) zoomed patch A from (a), (c) patch A after deblurring, where outlier effect is magnified, (d) deblurred patch A with outlier removal process, (e) intensity variance of patch A over 100 frames, (f) intensity variance of patch B over 100 frames, (g) histogram of (e), (h) histogram of (f).

appear due to the turbulence variation can be found. Suppose a k^*th patch is a diffraction-limited one: $r_{k^*} \approx z + \tilde{n}_{k^*}$, then simply denoising r_{k^*} provides a good estimation of the patch z.

In practice, diffraction-limited patches can be detected through various local sharpness measures. For example, in astronomical imaging, Strehl ratio is widely used, which is the ratio of the peak intensity in the aberrated PSF to the diffraction-limited PSF [6]. This can be easily measured if there exists a reference star inside the patch. However, in general cases Strehl ratio is difficult to measure, and alternative metrics based on local image intensity are preferred [13], [25], [26]. In [26], it was shown that image intensity variance has a unique relationship with the Strehl

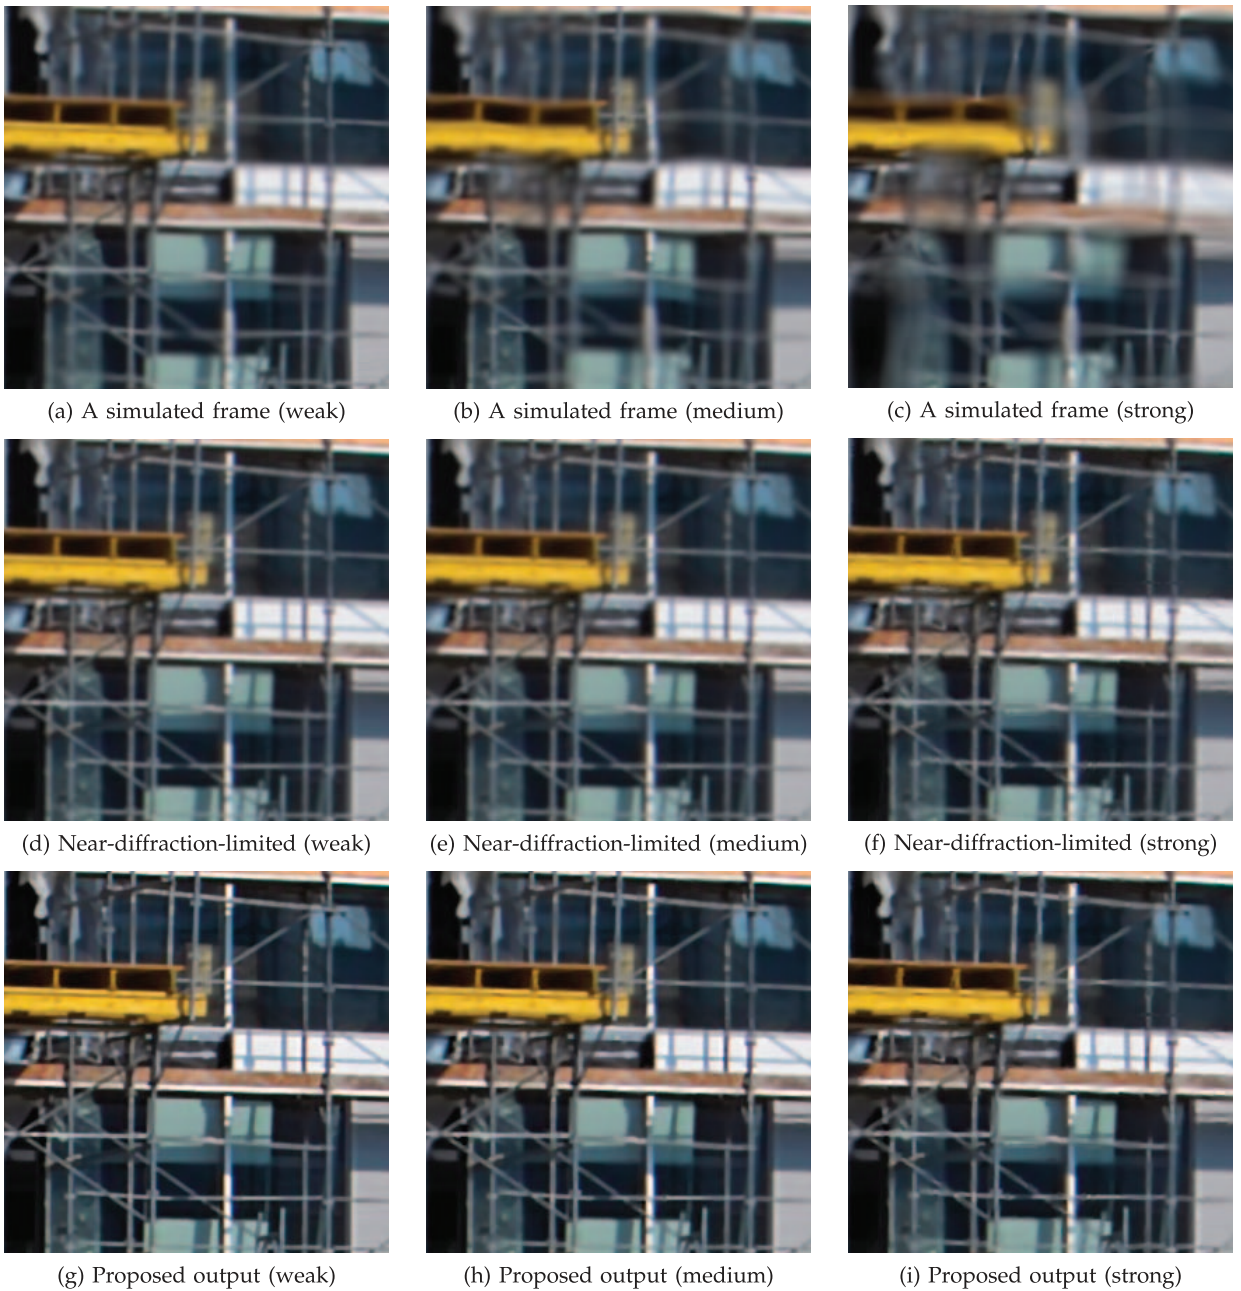Fig. 7. Latent sharp image used for simulation.

(a) A simulated frame (weak)   (b) A simulated frame (medium)   (c) A simulated frame (strong)

(d) Near-diffraction-limited (weak)   (e) Near-diffraction-limited (medium)   (f) Near-diffraction-limited (strong)

(g) Proposed output (weak)   (h) Proposed output (medium)   (i) Proposed output (strong)

Fig. 8. Simulated experiments. (a)-(c) Simulated frames with three turbulence levels (weak, medium, strong). (d)-(f) Near-diffraction-limited images from the three test sequences with noise variance $\sigma_n^2 = 1$ and turbulence level weak, medium, and strong, respectively. (g)-(i) Final outputs deblurred from (d)-(f).

ratio and can be used for patch selection. The intensity variance of an $L \times L$ patch r_k is defined as

$$s_k = \frac{1}{L^2 - 1} \sum_{\mathbf{x}} (r_k[\mathbf{x}] - \bar{r}_k)^2, \qquad (12)$$

where $\bar{r}_k$ represents the mean value of patch r_k. This metric is employed in this step.

Sharp patch selection by maximizing the above metric can be sensitive to outliers, such as those caused by objects incorrectly registered or moving objects occasionally appearing in the scene (e.g., a bird flying through the field of view). Examples are illustrated in Fig. 6a, where some artifacts can be observed in flat regions (see patch A). These artifacts are

caused by registration error and can be strongly magnified (see Fig. 6c) after the consequential deblurring process (Section 2.3). In Fig. 6e, we can observe extremely high values among the sharpness measures for patch A. Patch B gives another example (see Fig. 6f) where no outlier exists.[3] Histograms of these two patches are given in Figs. 6g and 6h, where the sharpness values for the outlier-free patch have a roughly symmetric distribution, while outlier-contaminated distributions contain a long tail corresponding to high values.

---

3. Patch B contains strong structure (high SNR), and thus its sharpness measures have values much higher than patch A (see Figs. 6e and 6f). This also explains why the registration accuracy is lower in patch A compared with patch B.

TABLE 1
Performance of the Restoration Approaches Evaluated in PSNR Values (dB)

Method	$\sigma_n^2 = 1$ weak / medium / strong	$\sigma_n^2 = 9$ weak / medium / strong	$\sigma_n^2 = 25$ weak / medium / strong
Proposed	**23.52** / **23.17** / **22.79**	**23.47** / **23.11** / **22.77**	**23.35** / **23.10** / **22.60**
Proposed NDL	22.22 / 21.67 / 21.64	22.24 / 21.70 / 21.63	22.29 / 21.74 / 21.62
Lucky Region [17]	22.44 / 21.32 / 21.28	22.40 / 21.29 / 21.27	22.33 / 21.23 / 21.21
Method [9]	21.80 / 20.67 / 18.81	21.77 / 20.59 / 18.80	21.67 / 20.52 / 18.77
EFF [10]	21.29 / 20.15 / 18.89	21.85 / 20.23 / 18.86	21.71 / 20.24 / 18.85
Averaged Input	20.67 / 19.33 / 18.06	20.61 / 19.28 / 18.03	20.49 / 19.19 / 17.96

(a) Observed frame 1   (b) Observed frame 2   (c) Observed frame 3

(d) Registered frame 1   (e) Registered frame 2   (f) Registered frame 3
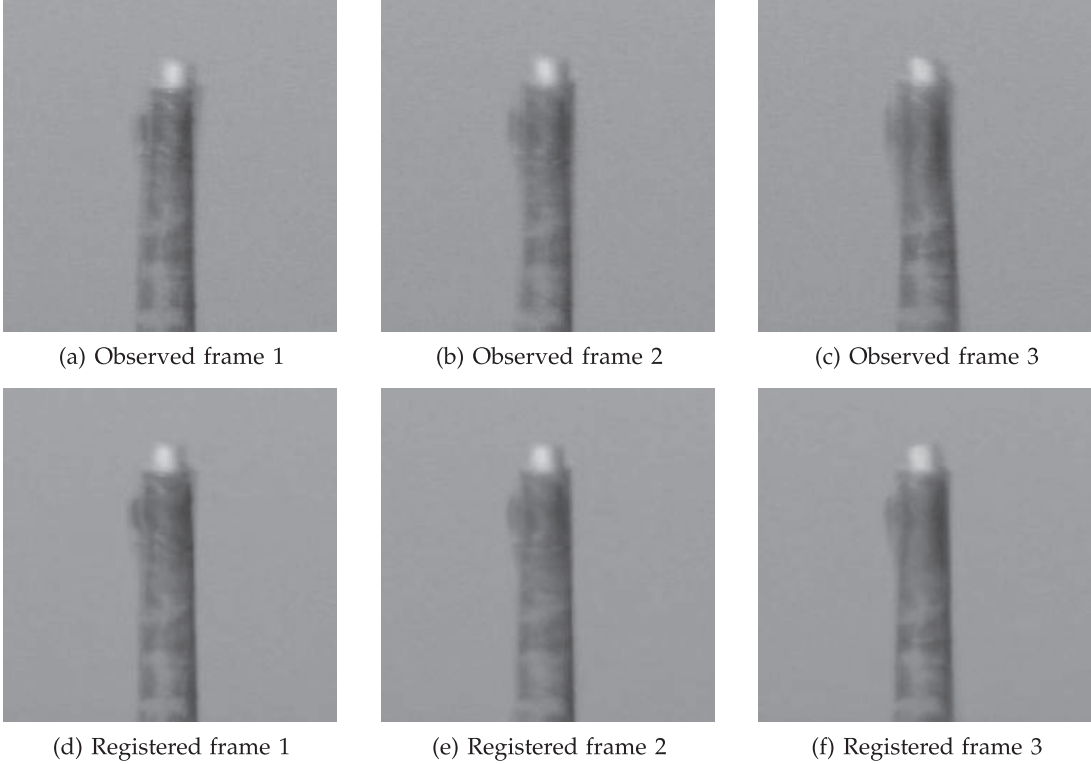
Fig. 9. Image registration results using 100 frames taken from the video *Chimney* distorted by real atmospheric turbulence.

To alleviate the outlier problem, Hampel's outlier identification method is employed [27]. We use a Gaussian model $\mathcal{N}(\eta, \sigma)$ to approximate the distribution of outlier-free samples (which are sharpness measurements in our case). A given sample s_k is identified as an outlier if

$$|s_k - \eta| > \varrho_{1-\alpha/2}\sigma, \qquad (13)$$

where ϱ_x is the x quantile of the $\mathcal{N}(0,1)$ distribution and α denotes the confidence level. Hampel suggested using the median and median absolute deviation (MAD) method [28] to robustly estimate the data mean η and the standard deviation σ, respectively [28]. In our experiment, we set $\varrho_{1-\alpha/2} = 6$ to strongly suppress the type I error. Any sample s_k that is above the threshold $\tau = 6\hat{\sigma}$ is detected as outlier and would be ignored when maximizing the local sharpness metric to select the diffraction-limited patch. In Fig. 6d we can see that with outlier removal the artifacts caused by incorrect registration are successfully eliminated from the fused image.

2.2.2 Patch-Wise Temporal Kernel Regression

To avoid possible artifacts that may appear in the subsequent deconvolution step, noise in the selected diffraction-limited patches needs to be suppressed. We formulate this denoising problem under a general temporal regression framework, where the imaging model is

$$r_k = q_k + \tilde{n}_k. \qquad (14)$$

The value of a pixel at $\mathbf{x}$ in the lth frame ($q_l[\mathbf{x}]$) can be estimated through zeroth order kernel regression [29], [30]:

$$\hat{q}_l[\mathbf{x}] = \arg\min_{q_l[\mathbf{x}]} \sum_k (r_k[\mathbf{x}] - q_l[\mathbf{x}])^2 U(\mathbf{x}; k, l), \qquad (15)$$

where the weight $U(\mathbf{x}; k, l)$ is a positive function measuring the "similarity" between $q_l[\mathbf{x}]$ and $q_k[\mathbf{x}]$. One simple but effective way of defining $U(\cdot)$ is based on patch-wise photometric distance and a Gaussian kernel function:

$$U(\mathbf{x}; k, l) = \exp\left(\frac{-\|q_k - q_l\|^2}{L^2\mu^2}\right). \qquad (16)$$

Here, L^2 is the total number of pixels in the patch and the scalar μ is called the *smoothing parameter* [29]. The noise-free

(a) Averaged image   (b) Ground truth   (c) [10]

(d) [9]   (e) Near-diffraction-limited   (f) Proposed approach

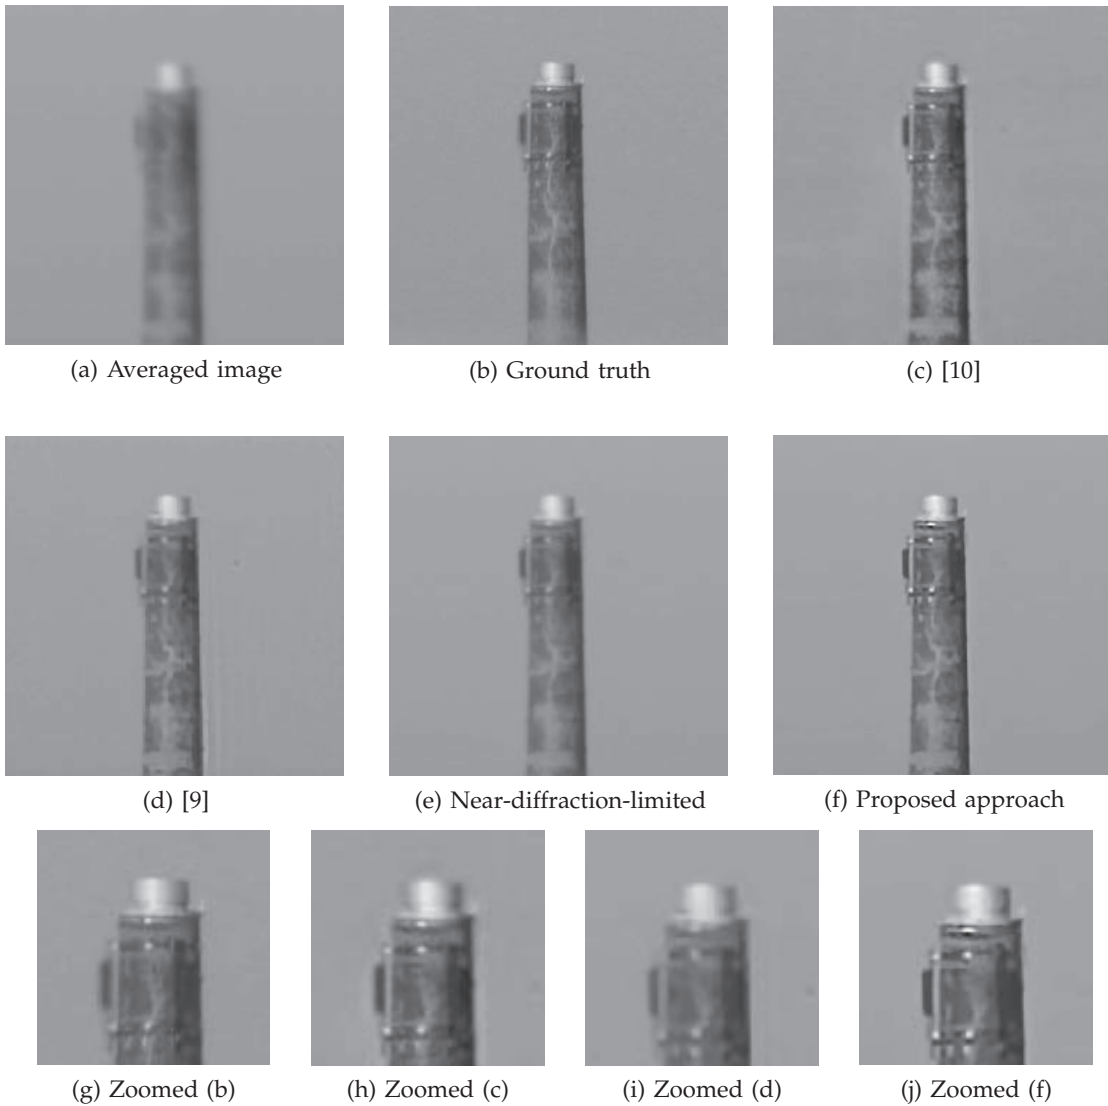(g) Zoomed (b)   (h) Zoomed (c)   (i) Zoomed (d)   (j) Zoomed (f)

Fig. 10. Image reconstruction results using 100 frames taken from the video *Chimney* distorted by real atmospheric turbulence.

photometric distance $\frac{1}{L^2}\|q_k - q_l\|^2$ can be unbiasedly estimated using the following function [31]:

$$\frac{\|q_k - q_l\|^2}{L^2} \approx \frac{\|r_k - r_l\|^2}{L^2} - 2\sigma_n^2, \qquad (17)$$

where the noise variance σ_n^2 can be estimated using, for example, the MAD method [28].

The solution to (15) is

$$\hat{q}_l[\mathbf{x}] = \frac{\sum_k U(\mathbf{x}; k, l) r_k[\mathbf{x}]}{\sum_k U(\mathbf{x}; k, l)}, \qquad (18)$$

which is nothing but a filter generated by normalizing the kernel $U(\cdot)$. According to (16), the kernel value *within a patch* is independent of the pixel position and space-invariant, which means the spatial constancy of the PSF in the estimated patch $\hat{q}_l$ is preserved after the regression.

Of course, in this step we only estimate the center pixel in the selected k^*th patch, which will be assigned to the output image Z according to the reconstruction procedure given in Algorithm 1. One important issue is that in the proposed

regression, all the weights are positive, which means that when restoring the diffraction-limited value, the output is always more blurry than the observed $r_{k^*}[\mathbf{x}]$. That is why we call the restored image *near-diffraction-limited*.

2.3 Single Image Deconvolution

Finally, a single image deblurring algorithm is required as a postprocess to deconvolve the near-diffraction-limited image Z. The degradation model is

$$Z = F \otimes h + \varepsilon, \qquad (19)$$

where ε represents error caused by the process generating the estimate of Z. Such a blind deconvolution algorithm can be described generally using the following:

$$<\widehat{F}, \hat{h}> = \arg\min_{F,h} \|Z - h \otimes F\|^2 + \lambda_1 R_f(F) + \lambda_2 R_h(h),$$

$$\qquad (20)$$

where R_f and R_h are the regularization terms based on prior knowledge about the latent sharp image F and the PSF h. Recent research on natural image statistics has shown that

(a) Ground truth   (b) One observed frame   (c) [10]

(d) [9]   (e) Proposed approach

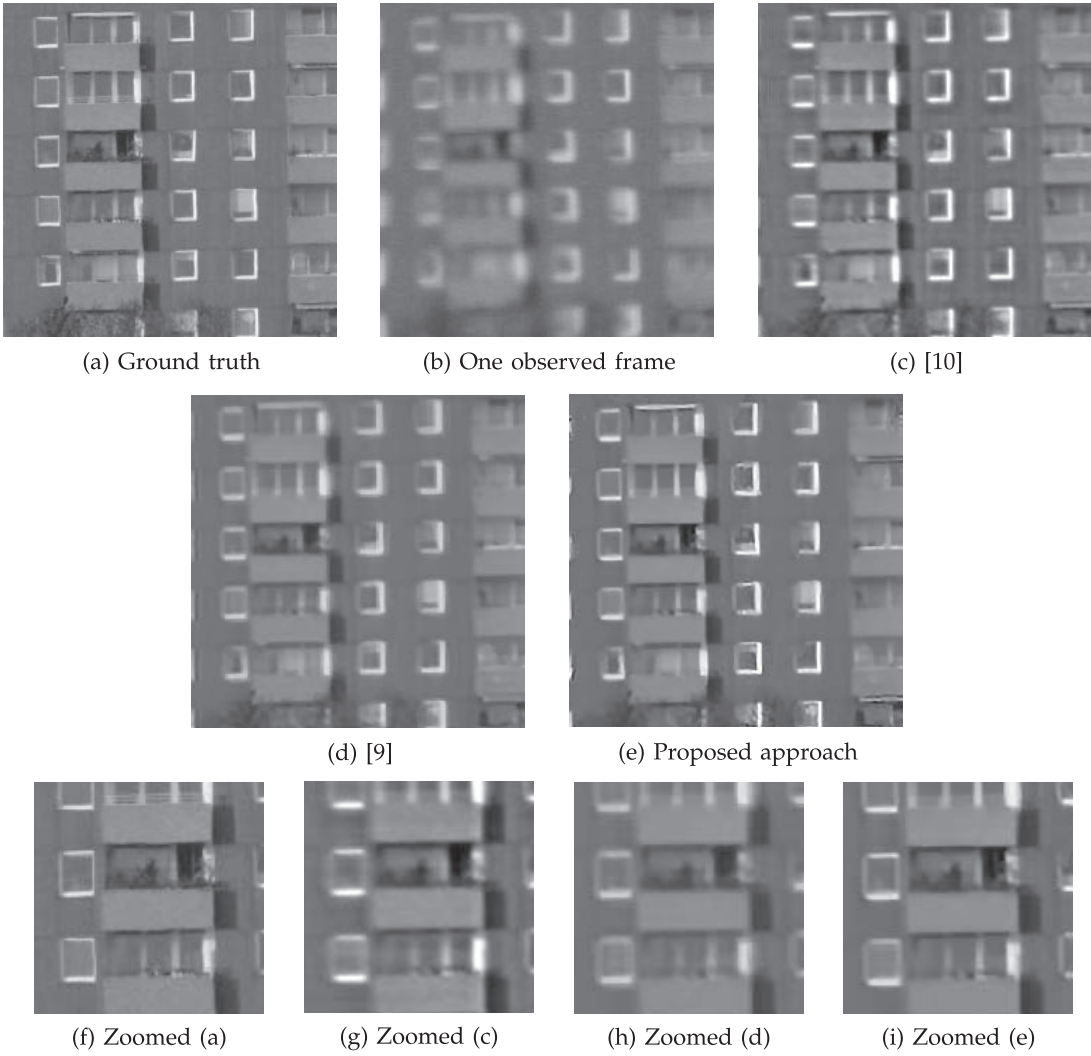(f) Zoomed (a)   (g) Zoomed (c)   (h) Zoomed (d)   (i) Zoomed (e)

Fig. 11. Image reconstruction results using 100 frames taken from the video *Building* distorted by real atmospheric turbulence.

image gradients obey *heavy-tailed* distributions that have most of their mass on small values but give significantly more probability to large values than Gaussian distributions [32]. Based on these studies, several *sparse* regularization methods have been introduced and have achieved great success in solving the blind deconvolution problem [32], [33], [20]. One example is the method proposed by Shan et al. [20], which is directly implemented in this step to calculate a final output. This method uses basically the same estimation form as (20) with the following sparse regularization term:

$$R_f(F) = \|\rho(F_x) + \rho(F_y)\|_1, \tag{21}$$

where F_x and F_y denote the derivatives of F in horizontal and vertical directions, respectively, and

$$\rho(\kappa) = \begin{cases} -\theta_1|\kappa| & \kappa \le l_t \\ -(\theta_2\kappa^2 + \theta_3) & \kappa > l_t. \end{cases} \tag{22}$$

Here, l_t, θ_1, θ_2, and θ_3 are all fixed parameters [20]. Sparsity is also utilized in regularizing h [20]:

$$R_h(h) = \|h\|_1. \tag{23}$$

The cost function (20) is optimized by alternating the estimation of f and h [33], [20].

In what follows, we used the default parameter settings as described in the authors' project page[4] for [20] except the noise level parameter "noiseStr," which is chosen in the range $[0.01, 0.05]$ according to the actual noise level observed in the given data. We refer interested readers to [20] for details.

3 EXPERIMENTS

In this section, we will illustrate the performance of the proposed approach using both simulated and real image sequences. Throughout all the experiments, the intervals of the control points in the registration step are set as $\epsilon_x = \epsilon_y = 16$ pixels and the symmetry constraint parameter: $\gamma = 5,000$. In implementing the restoration step in Section 2.2, we set the patch size $L = 9$. Results of the lucky-region algorithm from [17], the multiframe reconstruction approach from [9], and the EFF approach from [10] are also shown for comparison.[5]

4. http://www.cse.cuhk.edu.hk/~leojia/projects/motion_deblurring/index.html.

5. In this section, the outputs of method [9] are generated using the original code. The EFF outputs are directly from the authors of [10]. The outputs of method [17] are produced by our own implementation.

3.1   Simulated Experiments

To quantitatively evaluate the algorithm performance, a set of image sequences with different degrees of turbulence and noise are generated. The latent sharp image[6] (300×300) is shown in Fig. 7. The sequences are produced using the imaging model described in (1) and (9), where motion fields, spatially variant PSFs, and spatially invariant diffraction-limited PSF are required. To produce the motion fields, we first randomly generate a set of deformation vectors through a Gaussian distribution, and then calculate a motion field for each frame through the B-spline model (2)-(5). The turbulence strength is determined by the variance of the Gaussian distribution. The spatially variant PSFs are simulated also using a Gaussian function, where the variance of a local PSF is proportional to the magnitude of the corresponding local motion vector. The diffraction-limited PSF is generated using a disc function. Three degrees (weak, medium, and strong) of turbulence are produced (see Figs. 8a, 8b, and 8c). Also, different levels of white Gaussian noise (with variance $\sigma_n^2 = 1, 9, 25$) are added into the sequences. Each sequence contains 100 frames.

Due to the space limits, we only provide results of three sets of sequences with noise variance $\sigma_n^2 = 1$ in Fig. 8. It is clear that the proposed approach significantly improved the visual quality, and recovered many high-frequency details of the image content.

Table 1 gives the Peak Signal-to-Noise Ratio (PSNR) values for all the outputs with four different restoration algorithms and the averaged PSNR values of each input sequence. It can be seen that the proposed approach outperforms in all test sequences in terms of PSNR. Results of the near-diffraction-limited images are also provided in this table, from which we can tell that the final blind-deconvolution step generally increases the PSNR by 1-1.5 dB.

3.2   Real Video Experiments

Several real videos are tested to illustrate the performance of the proposed restoration framework in practical situations.[7] We first show results on data taken under controlled but real conditions. The two sets of videos (see *Chimney* and *Building* in Figs. 9, 10, and 11) of size 237×237 were captured through hot air exhausted by a building's vent, which could be closed to take "ideal" images of the same scene [10]. Each sequence consists of 100 frames degraded by strong turbulence effects (see Figs. 9a, 9b, 9c, and 11b). In the *Chimney* experiment, some registration results are also given in Figs. 9d, 9e, and 9f, where we can see geometric deformation with respect to the averaged image Fig. 10a has effectively been removed. From the output images in Figs. 10c, 10d, and 10f we can observe that EFF (c) and the proposed method (f) provide the best restoration results and faithfully recover details of the object. The proposed result (f) looks even sharper, while the EFF output (c) contains some halo artifacts in the edge regions. Similar restoration results can be observed in the *Building* experiment shown in Fig. 11. The near-diffraction-limited image generated from the *Chimney* sequence is also provided in Fig. 10e.

Additional sets of video data taken from long-distance imaging systems are processed next. The first set of images ($410 \times 380 \times 80$) shows the moon's surface taken from a ground-based telescope (see Fig. 12a). From Fig. 12b we can see that though the output image of [17] looks slightly sharper than one of the observed frames, it is still quite blurry, probably due to the diffraction-limited blur and the limited number of frames. The method in [9] provides a better result but with some details (like small craters) vanished (Fig. 12c). The proposed method gives a significant improvement in visual quality (Fig. 12d). It successfully removed blur and meanwhile recovered many small craters on the surface (Fig. 12h) that can hardly be seen from either the original frame (Fig. 12e) or the outputs of the other two methods (Figs. 12f and 12g).

The scene of the next video stream consists of a water tower located above the ground, imaged at a (horizontal) distance of 2.4 kilometers. The (300×220) video is quite noisy and highly blurred (due to long exposure time). Eighty frames were taken from the video to produce the result shown in Fig. 13. Again, the lucky region method did not provide much improvement in the result and slightly changed the object shape (Fig. 13b). The reconstruction method in [9] slightly increased the sharpness (Fig. 13c). The output of the proposed method looks much sharper and clean, with many details well restored (Fig. 13d).

4   DISCUSSION AND CONCLUSIONS

In this paper, we proposed a new approach for restoring a single high-quality image from an image sequence distorted by air turbulence. The proposed algorithm first registers the frames to suppress geometric deformation using B-spline-based nonrigid image registration, which incorporates a symmetry constraint to effectively improve the estimation accuracy. Next, a regression-based process is carried out to produce an image convolved with a near-diffraction-limited PSF, which can be viewed as spatially invariant. Finally, a blind deconvolution algorithm is implemented to remove diffraction-limited blur from the fused image to generate the final output. Experiments using controlled and real data illustrate that this approach is capable of alleviating geometric deformation and space-time varying blur caused by turbulence, recovering unprecedented details of the scene and significantly improving visual quality.

Another technique which may be related to the present work is *seeing through water* [34], [35]. Compared with air turbulence, in most cases the geometric warping effect induced by water is much stronger, but the blur effect is relatively milder. In the following experiments, we replace the image registration step in Section 2.1 with a robust iterative registration algorithm from a seeing through water approach [34]. The sequence *Chimney* with severe turbulence motion is tested. Results are given in Fig. 14, where (a) shows the direct output of method [34], which is highly blurry since it is just the temporal mean of the registered sequence. We then take the registered sequence to generate a near-diffraction-limited image (b) using the temporal regression step described in Section 2.2, this intermediate result is much sharper than (a). The final deblurred output is given in (c), which is visually very close to the proposed result given in Fig. 10f except for a mild shape change at the

---

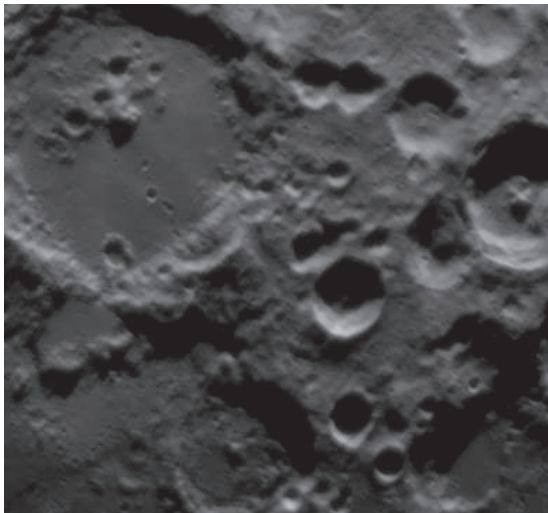6. The pixel intensity range here is $[0, 255]$.

7. Videos and experimental results are given on the webpage: http://users.soe.ucsc.edu/~xzhu/doc/turbulence.html.
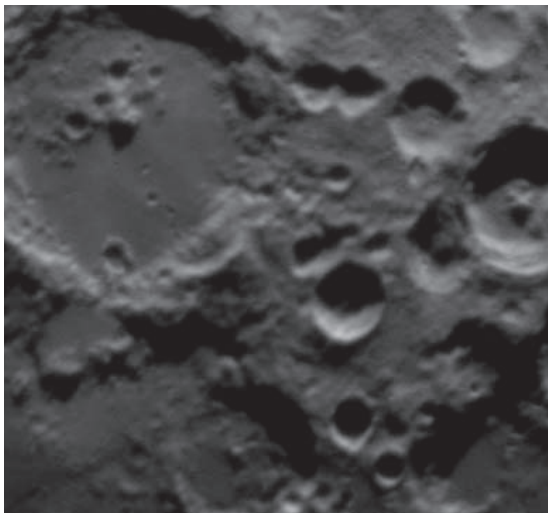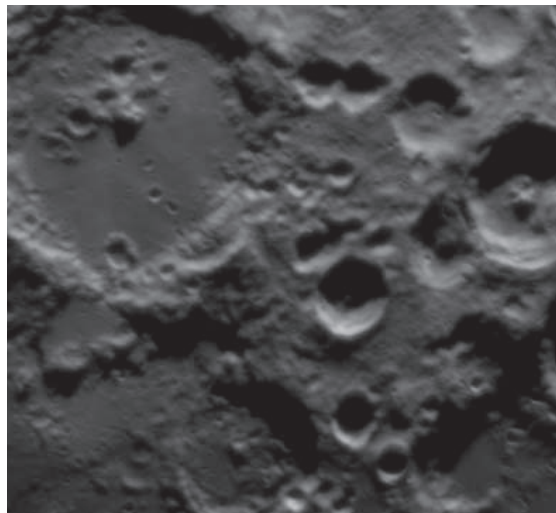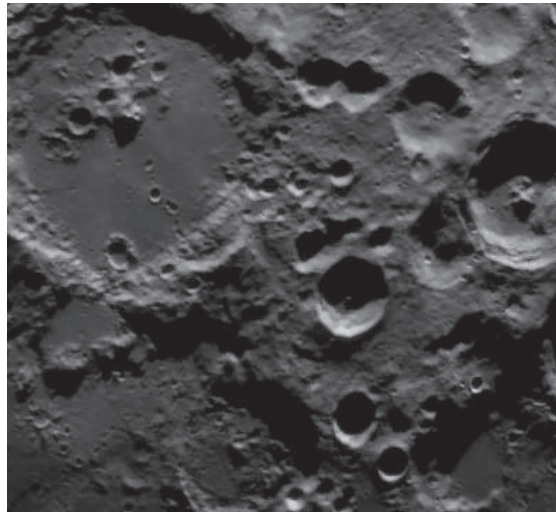
(a) One observed frame

(b) [17]

(c) [9]

(d) Proposed approach

(e) Zoomed part of (a)

(f) Zoomed part of (b)

(g) Zoomed part of (c)

(h) Zoomed part of (d)

(i) Zoomed part of (a)

(j) Zoomed part of (b)

(k) Zoomed part of (c)

(l) Zoomed part of (d)

Fig. 12. Image reconstruction result using 80 frames taken from the video *Moon Surface* distorted by real atmospheric turbulence.
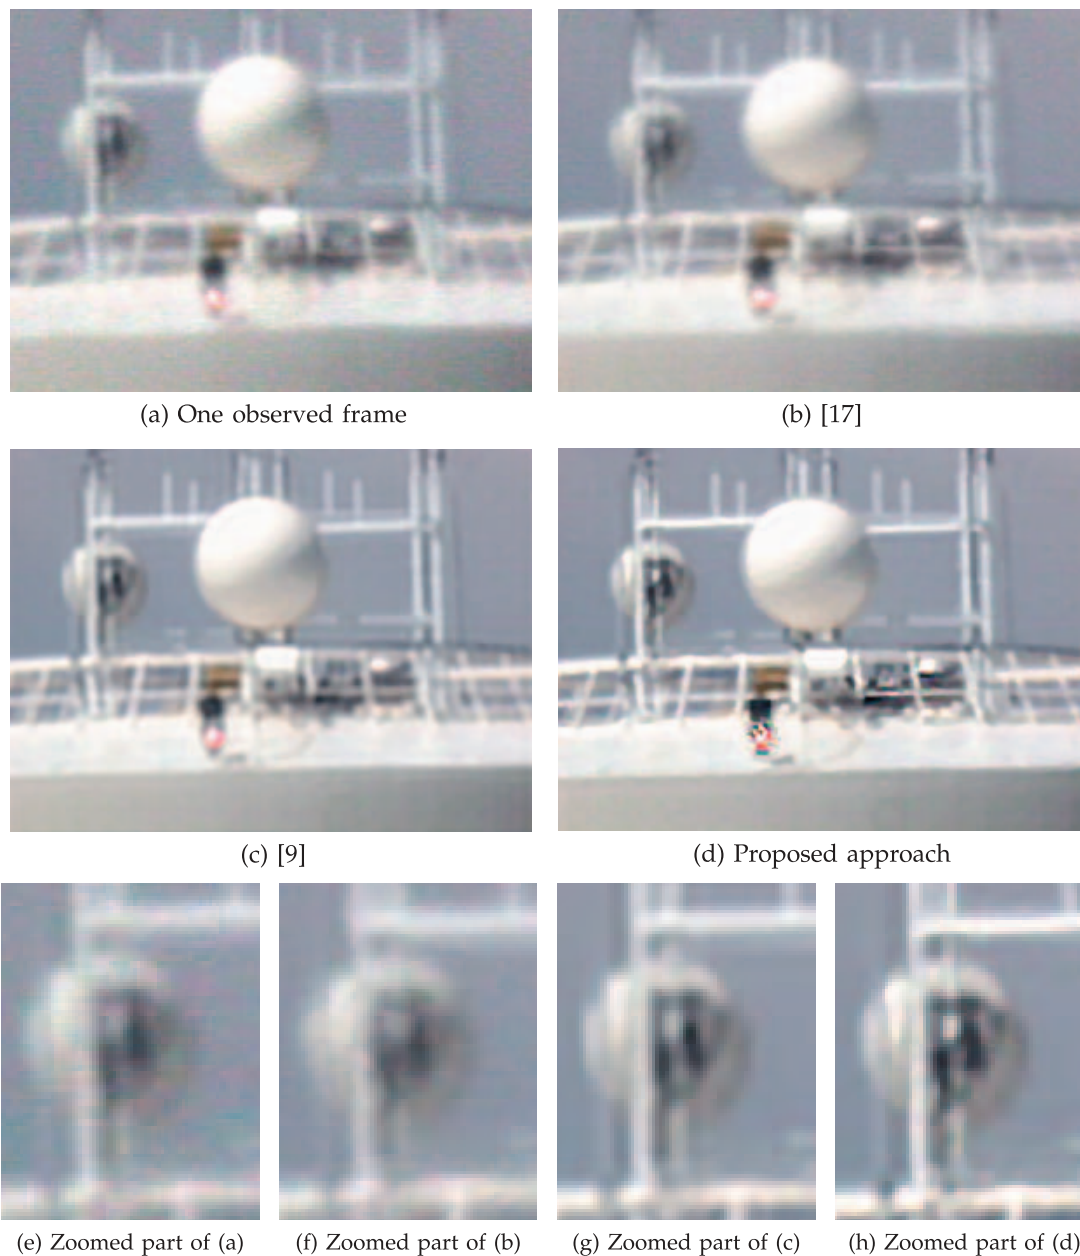
(a) One observed frame

(b) [17]

(c) [9]

(d) Proposed approach

(e) Zoomed part of (a)   (f) Zoomed part of (b)   (g) Zoomed part of (c)   (h) Zoomed part of (d)

Fig. 13. Image reconstruction result using 80 frames taken from the video *Water Tower* distorted by real atmospheric turbulence.

top part of the Chimney. Similar experiments using test video *Building* are illustrated in Fig. 15. Probably the robust registration algorithm addressing water distortions is capable of correcting strong geometric deformations quite well, but it seems that the registration step in Section 2.1 is sufficient for handling most air induced motion.

One limitation of the proposed framework is that the near-diffraction-limited patch reconstruction step requires the existence of a diffraction-limited isoplanatic patch for all the local regions. This requirement may be hard to achieve if we do not have sufficiently many frames. However, in that case it is still possible to deconvolve the diffraction-limited patch z. We can consider the local imaging model:

$$g_k = z \otimes h_k + n_k,$$

as a single-input multioutput (SIMO) system [36]. Such systems have been thoroughly studied in the past few years

and one group of algorithms restoring z is multichannel deconvolution filters [36], [37]. These are defined as

$$\hat{z} = \sum_k t_k \otimes g_k, \qquad (24)$$

where t_k is a frame-wise filter of size $N \times N$. In [12], the EFF approach attempts to estimate z and $\{h_k\}$ simultaneously given $\{g_k\}$, which can be viewed as a special case of multichannel deconvolution filtering. Assuming that $\{h_k\}$ can be exactly estimated in the EFF process, then the EFF reconstruction can be interpreted as calculating each t_k as the inverse filter of h_k and averaging all the filtered frames to get the final output. In this case, t_k needs a large support N in general. Giannakis and Heath showed that if we have sufficient observations, it is also possible (in the absence of noise) to perfectly reconstruct z given $\{t_k\}$ of very small
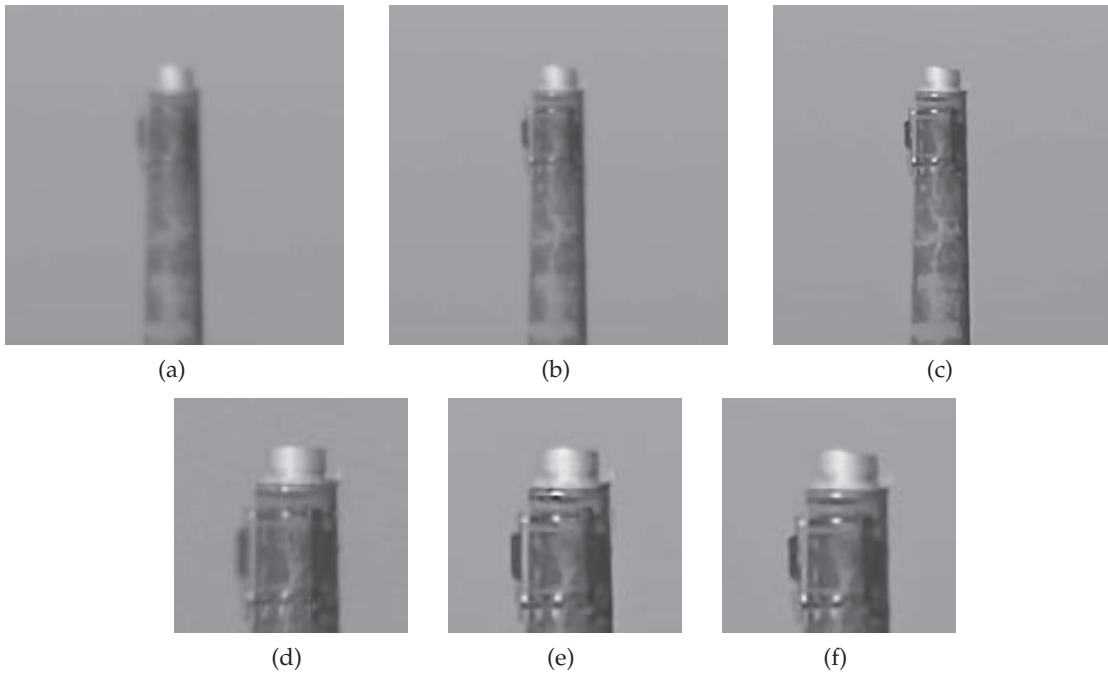
Fig. 14. Experiments using the seeing through water method [34]. Video *Chimney* distorted by real atmospheric turbulence is tested. (a) Output of method [34]. (b) Near-diffraction-limited image generated using the sequence registered by [34]. (c) Image deblurred from (b). (d) Zoomed ground truth. (e) Zoomed proposed result (same as Fig. 10j). (f) Zoomed (c).
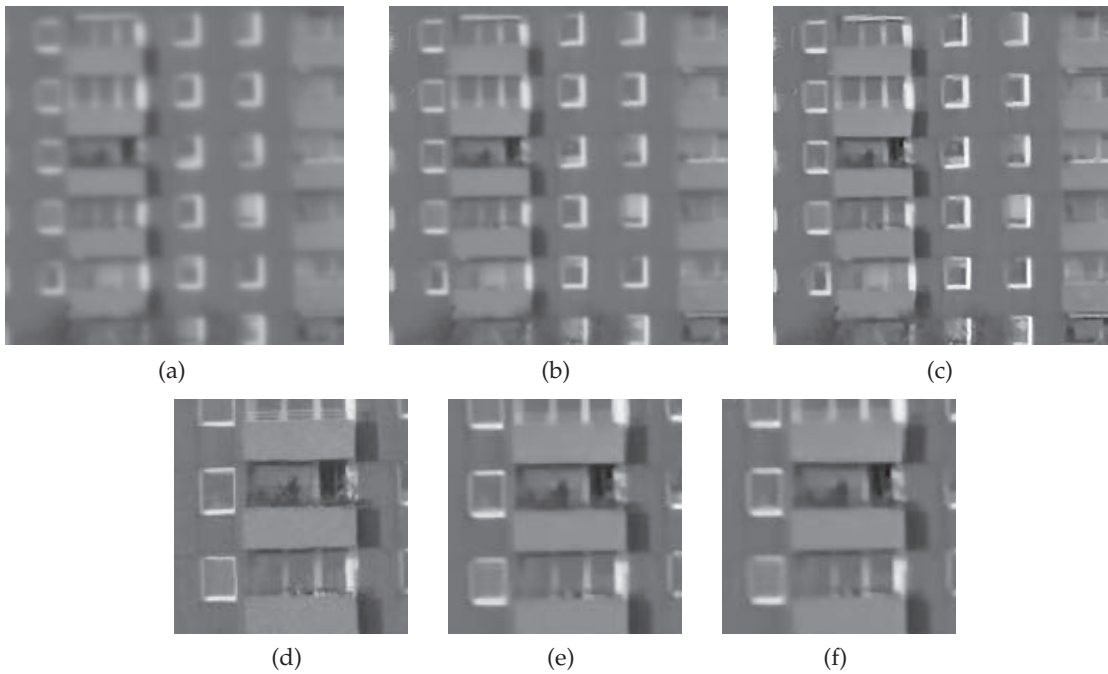

Fig. 15. Experiments using the seeing through water method [34]. Video *Building* distorted by real atmospheric turbulence is tested. (a) Output of method [34]. (b) Near-diffraction-limited image generated using the sequence registered by Oreifej et al. [34]. (c) Image deblurred from (b). (d) Zoomed ground truth. (e) Zoomed proposed result (same as Fig. 11i). (f) Zoomed (c).

support [36]. In other words, we can "mildly" deblur each frame and get a perfectly restored image through fusion. Furthermore, if the filter size N is reduced to 1, the filtering process would become a pixel-wise fusion, which is very similar to what we presented in Section 2.2. An accurate estimation of z under a more general scenario requires an accurate PSF estimation procedure, which is difficult to achieve, especially in noisy cases. Further study should be carried out in this direction.

REFERENCES

[1] W.E.K. Middleton, *Vision through the Atmosphere.* Univ. of Toronto Press, 1958.

[2] M.C. Roggemann and B.M. Welsh, *Imaging through Turbulence.* CRC Press, 1996.

[3] D. Li, R.M. Mersereau, and S. Simske, "Atmospheric Turbulence-Degraded Image Restoration Using Principal Components Analysis," *IEEE Geoscience and Remote Sensing Letters,* vol. 4, no. 3, pp. 340-344, July 2007.

[4] M. Shimizu, S. Yoshimura, M. Tanaka, and M. Okutomi, "Super-Resolution from Image Sequence under Influence of Hot-Air Optical Turbulence," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* June 2008.

[5] R.N. Tubbs, "Lucky Exposures: Diffraction Limited Astronomical Imaging through the Atmosphere," PhD thesis, Cambridge Univ., Sept. 2003.

[6] R.K. Tyson, *Principles of Adaptive Optics.* Academic Press, 1998.

[7] M.A. Vorontsov, "Parallel Image Processing Based on an Evolution Equation with Anisotropic Gain: Integrated Optoelectronic Architectures," *J. Optical Soc. Am. A,* vol. 16, pp. 1623-1637, 1999.

[8] M.A. Vorontsov and G.W. Carhart, "Anisoplanatic Imaging Through Turbulent Media: Image Recovery by Local Information Fusion from a Set of Short-Exposure Images," *J. Optical Soc. Am. A,* vol. 18, no. 6, pp. 1312-1324, June 2001.

[9] X. Zhu and P. Milanfar, "Image Reconstruction from Videos Distorted by Atmospheric Turbulence," *Proc. SPIE Electronic Imaging Conf. Visual Information Processing and Comm.,* Jan. 2010.

[10] M. Hirsch, S. Sra, B. Schölkopf, and S. Harmeling, "Efficient Filter Flow for Space-Variant Multiframe Blind Deconvolution," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* pp. 607-614, June 2010.

[11] N.M. Law, "Lucky Imaging: Diffraction-Limited Astronomy from the Ground in the Visible," PhD thesis, Cambridge Univ., May 2003.

[12] S. Harmeling, M. Hirsch, S. Sra, and B. Schölkopf, "Online Blind Deconvolution for Astronomical Imaging," *Proc. IEEE Int'l Conf. Computational Photography,* Apr. 2009.

[13] M.C. Roggemann, C.A. Stoudt, and B.M. Welsh, "Image-Spectrum Signal-to-Noise-Ratio Improvements by Statistical Frame Selection for Adaptive-Optics Imaging through Atmospheric Turbulence," *Optical Eng.,* vol. 33, no. 10, pp. 3254-3264, Oct. 1994.

[14] D.L. Fried, "Probability of Getting a Lucky Short-Exposure Image through Turbulence," *Optical Soc. Am., J.,* vol. 68, pp. 1651-1658, 1978.

[15] M. Aubailly, M.A. Vorontsov, G.W. Carhart, and M.T. Valley, "Image Enhancement by Local Information Fusion with Pre-Processing and Composed Metric," *Proc. SPIE,* vol. 7090, 2008.

[16] S. John and M.A. Vorontsov, "Multiframe Selective Information Fusion from Robust Error Estimation Theory," *IEEE Trans. Image Processing,* vol. 14, no. 5, pp. 577-584, May 2005.

[17] M. Aubailly, M.A. Vorontsov, G.W. Carhat, and M.T. Valley, "Automated Video Enhancement from a Stream of Atmospherically-Distorted Images: The Lucky-Region Fusion Approach," *Proc. SPIE,* vol. 7463, 2009.

[18] N. Joshi and M. Cohen, "Seeing Mt. Rainier: Lucky Imaging for Multi-Image Denoising, Sharpening, and Haze Removal," *Proc. IEEE Int'l Conf. Computational Photography,* Mar. 2010.

[19] A. Levin, Y. Weiss, F. Durand, and W.T. Freeman, "Understanding and Evaluating Blind Deconvolution Algorithms," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2009.

[20] Q. Shan, J. Jia, and A. Agarwala, "High-Quality Motion Deblurring from a Single Image," *ACM Trans. Graphics,* vol. 27, article 73, 2008.

[21] X. Zhu and P. Milanfar, "Stabilizing and Deblurring Atmospheric Turbulence," *Proc. IEEE Int'l Conf. Computational Photography,* Apr. 2011.

[22] R. Szeliski and J. Coughlan, "Spline-Based Image Registration," *Int'l J. Computer Vision,* vol. 22, no. 93, pp. 199-218, 1997.

[23] S. Farsiu, M. Elad, and P. Milanfar, "Constrained, Globally Optimal, Multi-Frame Motion Estimation," *Proc. IEEE/SP 13th Workshop Statistical Signal Processing,* pp. 1396-1401, July 2005.

[24] M.F. Beg and A. Khan, "Symmetric Data Attachment Terms for Large Deformation Image Registration," *IEEE Trans. Medical Imaging,* vol. 26, no. 9, pp. 1179-1189, Sept. 2007.

[25] R.F. Dantowitz, S.W. Teare, and M.J. Kozubal, "Ground-Based High-Resolution Imaging of Mercury," *The Astronomical J.,* vol. 119, pp. 2455-2457, May 2000.

[26] J.C. Christou, K.J. Mighell, and R.B. Makidon, "Strehl Ratio and Image Sharpness for Adaptive Optics," *Proc. SPIE Advances in Adaptive Optics II,* vol. 6272, p. 62721Y, 2006.

[27] L. Davies and U. Gather, "The Identification of Multiple Outliers," *J. Am. Statistical Assoc.,* vol. 88, pp. 782-792, Sept. 1993.

[28] F.R. Hampel, "The Influence Curve and Its Role in Robust Estimation," *J. Am. Statistical Assoc.,* vol. 69, pp. 383-393, 1974.

[29] M.P. Wand and M.C. Jones, *Kernel Smoothing,* series Monographs on Statistics and Applied Probability. Chapman and Hall, 1995.

[30] H. Takeda, S. Farsiu, and P. Milanfar, "Kernel Regression for Image Processing and Reconstruction," *IEEE Trans. Image Processing,* vol. 16, no. 2, pp. 349-366, Feb. 2007.

[31] A. Buades, B. Coll, and J.M. Morel, "A Review of Image Denoising Algorithms, with a New One," *Multiscale Modeling and Simulation (SIAM Interdisciplinary J.),* vol. 4, no. 2, pp. 490-530, 2005.

[32] R. Fergus, B. Singh, A. Hertsmann, S.T. Roweis, and W.T. Freeman, "Removing Camera Shake from a Single Image," *Proc. ACM Siggraph,* 2006.

[33] A. Levin, R. Fergus, F. Durand, and W.T. Freeman, "Image and Depth from a Conventional Camera with a Coded Aperture," *ACM Trans. Graphics,* vol. 26, article 70, 2007.

[34] O. Oreifej, G. Shu, T. Pace, and M. Shah, "A Two-Stage Reconstruction Approach for Seeing through Water Clearly," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* 2011.

[35] Y. Tian and S. Narasimhan, "Seeing through Water: Image Restoration Using Model-Based Tracking," *Proc. 12th IEEE Int'l Conf. Computer Vision,* 2009.

[36] G.B. Giannakis and R.W. Heath, "Blind Identification of Multi-channel FIR Blurs and Perfect Image Restoration," *IEEE Trans. Image Processing,* vol. 9, no. 11, pp. 1877-1896, Nov. 2000.

[37] W. Souidene, K. Abed-Meraim, and A. Beghdadi, "A New Look to Multichannel Blind Image Deconvolution," *IEEE Trans. Image Processing,* vol. 18, no. 7, pp. 1487-1500, July 2009.

Xiang Zhu received the BS and MS degrees in electrical engineering from Nanjing University, Nanjing, China, in 2005 and 2008, respectively. He is currently working toward the PhD degree in electrical engineering at the University of California, Santa Cruz. His research interests are in the domain of image processing (denoising, deblurring, super-resolution, and image quality assessment) and computer vision. He is a student member of the IEEE.

Peyman Milanfar received the BS degree in electrical engineering and mathematics from the University of California, Berkeley, and the MS and PhD degrees in electrical engineering from the Massachusetts Institute of Technology, respectively. Until 1999, he was a senior research engineer at SRI International, Menlo Park, California. He is a professor of electrical engineering and an associate dean for research in the Baskin School of Engineering at the University of California, Santa Cruz. From 1998 to 2000, he was a consulting assistant professor of computer science at Stanford University, where he was also a visiting associate professor in 2002. His research interests include statistical signal, image, and video processing, and computational photography and vision. He won a US National Science Foundation (NSF) CAREER award in 2000, and the best paper award from the IEEE Signal Processing Society in 2010. He is a member of the Signal Processing Society's Awards Board, and the Image, Video and Multidimensional Signal Processing technical committee. He is a fellow of the IEEE.