Stabilizing and Deblurring Atmospheric Turbulence

Xiang Zhu and Peyman Milanfar*

Electrical Engineering Department, University of California, Santa Cruz, CA 95064

xzhu@soe.ucsc.edu, milanfar@ee.ucsc.edu

Abstract

A new approach is proposed to correct geometric distortion and reduce space and time-variant blur in videos that suffer from atmospheric turbulence. We first register the frames to suppress geometric deformation using a B-spline based non-rigid registration method. Next, a fusion 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. Experiments using real data illustrate that this approach is capable of alleviating blur and geometric deformation caused by turbulence, recovering details of the scene and significantly improving visual quality.

1. Introduction

Atmospheric turbulence caused by variation of refractive index along the optical transmission path can strongly affect the performance of a long-distance imaging system [14, 16, 13, 18, 19]. It produces geometric distortion, space and time-variant defocus blur, and motion blur (if the exposure time is not sufficiently short). Several approaches have been proposed to solve this problem [21, 22, 18, 19, 23]. All these works are based on videos or image sequences, under the assumption that the scene is static. The imaging process can be modeled as [18]:

$$\mathbf{g}_k = \mathbf{H}_k \mathbf{F}_k \mathbf{f} + \mathbf{n}_k$$

where $\mathbf{f}$ denotes the ideal image, $\mathbf{F}_k$ and $\mathbf{H}_k$ represent the geometric deformation matrix and the blurring matrix respectively, $\mathbf{n}_k$ denotes additive noise, and $\mathbf{g}_k$ is the k-th observed frame.

Existing restoration algorithms for this problem can generally be divided into two main categories. One category is based on a multi-frame reconstruction framework [18, 23]. These approaches first require a non-rigid image registration technique to register each observed frame with respect to a turbulence-free grid, and use the registration parameters to estimate the corresponding deformation matrix $\mathbf{F}_k$. Then, a sharp image is reconstructed through a Bayesian reconstruction method (see diagram in Fig. 1 (a)). In [18] Shimizu *et al.* chose a B-spline based registration technique to estimate $\mathbf{F}_k$. They introduced a stabilization term into the registration cost function, and showed that it is capable of improving estimation accuracy in the regions containing less texture. This stabilization term changes the value of deformation parameters according to local image gradients. Meanwhile, atmospheric turbulence is in fact independent from image content; hence their approach may result in unnecessary bias being added to the estimate. In [23] another registration method incorporating a symmetry constraint was introduced. Experiments illustrate that this method outperforms both classic B-spline based registration technique and the one proposed in [18]. However, the main problem for such multi-frame reconstruction algorithms is that in general they can hardly estimate the actual point spread function (PSF), which is spatially and temporally changing. Both [18] and [23] employ a fixed Gaussian model to approximate the PSF, which strongly limits their performance.

Another class of approaches called "lucky exposure" employ image selection and fusion methods to reduce the blurring effects caused by turbulence [15, 6, 22, 19]. 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 [15, 6]. This method is based on the observation that for short-exposure images, turbulence creates "mutations" in image quality, and randomly makes some sufficiently sharp images (see examples in Fig. 2) [6, 1, 8, 2]. 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 – a small angle that can be viewed as containing space-invariant blur [19]. In [2] Vorontsov *et al.* proposed a "lucky region" restoration approach for *anisoplanatic* scenarios. Small lucky (isoplanatic) regions are detected by a local sharpness metric, and are fused to produce a large high quality
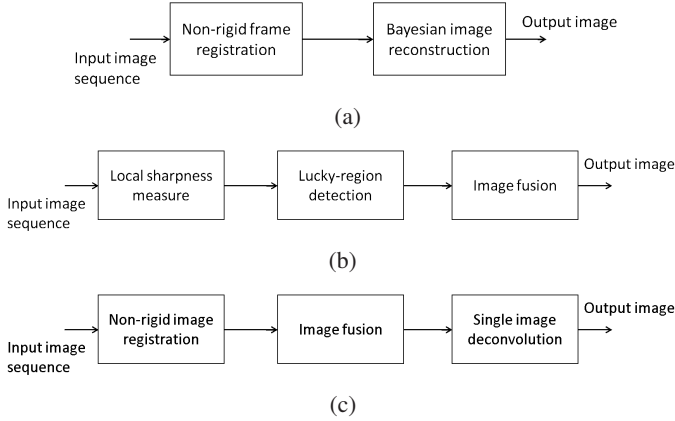

(a)

(b)

(c)

Figure 1. Block diagrams: (a) Multi-frame reconstruction approach [18, 23]; (b) Lucky-region fusion approach [2]; (c) Proposed approach.

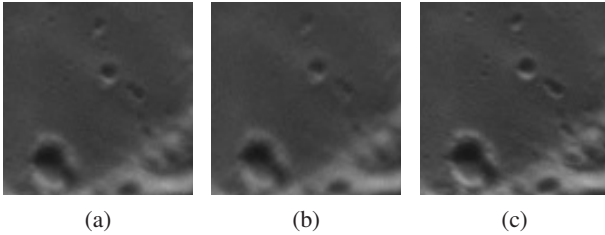(a)   (b)   (c)

Figure 2. Isoplanatic patches from multiple short exposure images of the moon's surface, taken through a ground based telescope, where a variation of blur can be observed. (c) illustrates a lucky frame example, which is sharper than the others.

image (see the diagram in Fig. 1 (b)). In another similar method developed by Joshi *et al.* [9] 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 balance between noise reduction and sharpness preservation. A dehazing process is also used to enhance the visual quality. One shortcoming of this method is that even though turbulence caused blur is strongly alleviated through the fusion process, the output still suffers from the blur caused by the diffraction-limited PSF [19, 11].

In this paper, a new framework is proposed for restoring a single image from an image sequence acquired in aniso-planatic scenarios. The 3-D physical scene is assumed to be static, while the air between the scene and sensor is affected by atmospheric turbulence. Our approach is designed to reduce the space and time-variant deblurring problem to a shift invariant one. It focuses on the observed regions convolved by near-diffraction-limited PSFs, which can be viewed as space and time-invariant, and thus can be estimated along with the latent sharp image through a blind deconvolution algorithm. Experiments show that this frame-
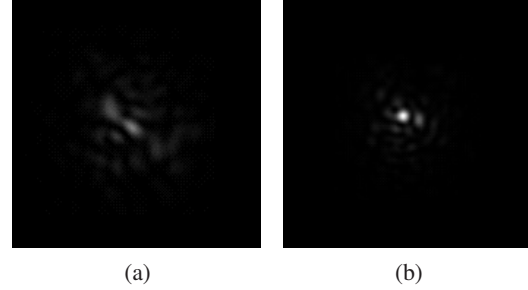

(a)   (b)

Figure 3. PSF examples from [19]: (a) an example of PSF caused by air turbulence; (b) an example of near-diffraction-limited PSF with the highest 0.1% of Strehl ratios selected from several thousand samples.

work is capable of alleviating both geometric deformation and blur, and significantly improving the visual quality.

2. Restoration Framework

Consider the imaging model described in Section 1:

$$\mathbf{g}_k = \mathbf{H}_k \mathbf{F}_k \mathbf{f} + \mathbf{n}_k, \tag{1}$$

where additive noise $\mathbf{n}_k$ is defined as zero-mean: $\mathbf{E}(\mathbf{n}_k) = \mathbf{0}$ and i.i.d. with its covariance matrix $\mathbf{cov}(\mathbf{n}_k) = \sigma_\mathbf{n}^2 \mathbf{I}$. Here $\mathbf{I}$ denotes the identity matrix. In some articles specific models (e.g. Gaussian) are used to describe the shape of local PSFs caused by air turbulence, which do not fit practical cases very well [18, 23]. In the real world, while turbulence-caused PSFs can look quite arbitrary, they still have some common characteristics. For example, usually they are restricted within a small region, and contain a "core", which is the peak intensity of a PSF (see Fig. 3). In the above model, since geometric distortion is separated from the blurring matrix, the core of each local PSF (also the entry with highest value in each row) is typically located on the diagonal of $\mathbf{H}_k$.

To estimate the latent image $\mathbf{f}$ a new restoration framework is proposed, which contains three main steps (see the diagram in Fig. 1 (c)):

1. Non-rigid image registration;

2. Fusion-based near-diffraction-limited image restoration;

3. Single image blind deconvolution.

The first step registers each frame onto a fixed reference grid, removing geometric distortion from the observed data. The second step fuses the registered sequence to produce a single image convolved by near-diffraction-limited PSFs, which have short support and can be approximately viewed

as spatially invariant. The fused image is finally deconvolved through a blind deblurring algorithm based on a natural image prior. Details about each step are given in the following subsections.

2.1. Non-rigid Image Registration

In [18] an experiment is provided to illustrate that the local turbulent motion has a zero-mean distribution, and thus the geometric distortion can be removed by simply averaging the observed frames. Such averaged image, though even more blurred than the observed data, can serve as a reference frame to register each observed frame. The B-spline based non-rigid registration method in [23] is directly implemented in our approach. This method incorporates a symmetry constraint that can effectively improve the estimation accuracy. We refer interested readers to [23] for details.

In this paper, we are interested in how the registration operator physically changes local PSFs. Assume that the warping matrix $\mathbf{F}_k$ is non-singular. The registration operator can then be denoted as $\mathbf{F}_k^{-1}$. Multiplying both sides by $\mathbf{F}_k^{-1}$ we have

$$\mathbf{q}_k = \widetilde{\mathbf{H}}_k \mathbf{f} + \tilde{\mathbf{n}}_k \tag{2}$$

where

$$\widetilde{\mathbf{H}}_k = \mathbf{F}_k^{-1} \mathbf{H}_k \mathbf{F}_k, \quad \tilde{\mathbf{n}}_k = \mathbf{F}_k^{-1} \mathbf{n}_k. \tag{3}$$

We denote the registered image as $\mathbf{q}_k = \mathbf{F}_k^{-1} \mathbf{g}_k$.

The blurring matrices $\widetilde{\mathbf{H}}_k$ and $\mathbf{H}_k$ are related by a similarity transformation. Since a similarity transformation preserves the spectrum (eigenvalues) of $\mathbf{H}_k$, the transformation should not significantly change the physical shape of PSFs. In particular, if the motion is pure translational and applies only integer movement (pixel to pixel), then $\mathbf{F}_k$ is a block-circulant matrix with each row and column containing precisely a single 1 and with 0s everywhere else. That is, we can approximate $\mathbf{F}_k$ by a permutation matrix, which means that $\mathbf{F}_k^T = \mathbf{F}_k^{-1}$. For such matrices, the corresponding similarity transformation circulates the rows and columns of $\mathbf{H}_k$. In other words, the registration operator only moves the locations of local PSFs without changing their physical shape. Of course the overall motion caused by turbulence is not translational. However, the local motion inside an isoplanatic region can be well-approximated as translational [20, 19, 9]. If the support of any local PSF is also small enough, then inside an isoplanatic region the transformation $\mathbf{F}_k^{-1} \mathbf{H}_k \mathbf{F}_k$ preserves the physical shape of local PSF [20] (see examples in Fig. 4).[1] Meanwhile, it can be shown

[1]If the PSF is spatially invariant, then $\mathbf{H}_k$ and $\mathbf{F}_k$ commute since both of them are block-circulant [4], and we have $\widetilde{\mathbf{H}}_k = \mathbf{H}_k$. For example, ignoring the PSF variation caused by air turbulence, if we assume that there only exist the diffraction-limited blur denoted as $\mathbf{H}$, then $\mathbf{F}_k^{-1} \mathbf{H} \mathbf{F}_k = \mathbf{H}$.
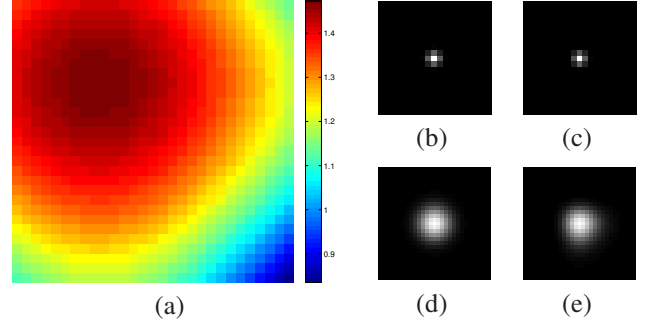

Figure 4. (a) Magnitude of a simulated motion field which can be viewed as locally translational. (b)-(c) illustrate a small size PSF before and after the transformation $\mathbf{F}_k^{-1} \mathbf{H}_k \mathbf{F}_k$, where the shape is preserved. (d)-(e) illustrate a big size PSF before and after the transformation, where we can see the change of its shape.

that the diagonal entries of $\mathbf{H}_k$ still remain on the diagonal of $\widetilde{\mathbf{H}}_k$ with their locations permutated.

Furthermore, since $\mathbf{F}_k$ can be approximated as a permutation matrix (pixel to pixel movement) we can also write:

$$\mathbf{cov}(\tilde{\mathbf{n}}_k) = \mathbf{F}_k^{-1} \mathbf{cov}(\mathbf{n}_k) \mathbf{F}_k \approx \sigma_{\mathbf{n}}^2 \mathbf{I}. \tag{4}$$

2.2. Near-diffraction Limited Image Restoration

By decomposing the overall blur $\widetilde{\mathbf{H}}_k = \mathbf{H} + \Delta \widetilde{\mathbf{H}}_k$ equation (2) can be rewritten as:

$$\begin{aligned} \mathbf{q}_k &= \left(\mathbf{H} + \Delta \widetilde{\mathbf{H}}_k \right) \mathbf{f} + \tilde{\mathbf{n}}_k \\ &= \mathbf{z} + \mathbf{e}_k + \tilde{\mathbf{n}}_k \end{aligned} \tag{5}$$

where $\mathbf{H}$ is diffraction-limited blur and $\mathbf{z} = \mathbf{H}\mathbf{f}$ denotes the diffraction-limited image. $\Delta \widetilde{\mathbf{H}}_k$ represents turbulence-caused blur, and $\mathbf{e}_k = \Delta \widetilde{\mathbf{H}}_k \mathbf{f}$ is the corresponding blurring artifact.

The diffraction-limited blur is unknown but space and time *invariant*, and there are numerous blind deconvolution algorithms available for removing it [5, 17, 10]. Unfortunately, estimating $\Delta \widetilde{\mathbf{H}}_k$ is not trivial. However, if we treat $\mathbf{e}_k$ as an additive noise, then it is still possible to estimate the diffraction-limited image $\mathbf{z}$. This requires an analysis of the statistical properties of $\mathbf{e}_k$, which we present next.

Both $\widetilde{\mathbf{H}}_k$ and $\mathbf{H}$ can be viewed as row-stochastic[2], which means $\tilde{\mathbf{h}}_i^T \mathbf{1} = \mathbf{h}_i^T \mathbf{1} = 1$, and thus

$$\Delta \tilde{\mathbf{h}}_i^T \mathbf{1} = (\tilde{\mathbf{h}}_i - \mathbf{h}_i)^T \mathbf{1} = 0. \tag{6}$$

Here $\Delta \tilde{\mathbf{h}}_i^T$, $\tilde{\mathbf{h}}_i^T$ and $\mathbf{h}_i^T$ represent the i-th rows of matrices $\Delta \widetilde{\mathbf{H}}_k$, $\widetilde{\mathbf{H}}_k$ and $\mathbf{H}$ respectively. Because the support of the PSFs $\tilde{\mathbf{h}}_i$ and $\mathbf{h}_i$ are restricted to a small isoplanatic region

[2]Row-stochastic means that each row of a matrix consists of nonnegative real numbers, which sum up to 1.

around the i-th pixel denoted as w_i, all the non-zeros values of $\Delta\tilde{\mathbf{h}}_i$ should also be located inside w_i:

$$\Delta\tilde{h}_{ij} = 0 \quad \forall j \notin w_i \tag{7}$$

where $\Delta\tilde{h}_{ij}$ represents the j-th element of vector $\Delta\tilde{\mathbf{h}}_i$. By defining a diagonal matrix $\mathbf{M}_i$ denoting a mask having 1s in the locations inside w_i and 0s everywhere else, we have $\mathbf{M}_i\Delta\tilde{\mathbf{h}}_i = \Delta\tilde{\mathbf{h}}_i$.

We model the latent sharp image $\mathbf{f}$ as a random field, whose mean vector is $\mathbf{m_f} = [m_{\mathbf{f}1}, m_{\mathbf{f}2}, \ldots]^T$, and its covariance matrix is assumed to be diagonal $\mathbf{C_f} = \mathrm{diag}[\sigma_{\mathbf{f}1}^2, \sigma_{\mathbf{f}2}^2, \ldots]$. The statistics of $\mathbf{f}$ are considered piecewise stationary:

$$m_{\mathbf{f}j} \approx m_{\mathbf{f}i}, \quad \sigma_{\mathbf{f}j} \approx \sigma_{\mathbf{f}i} \quad \forall j \in w_i, \tag{8}$$

and thus

$$\mathbf{M}_i\mathbf{m_f} \approx m_{\mathbf{f}i}\mathbf{M}_i\mathbf{1} \tag{9}$$

So considering (6)-(9) we have:

$$\begin{aligned}
\Delta\tilde{\mathbf{h}}_i^T\mathbf{m_f} &= (\mathbf{M}_i\Delta\tilde{\mathbf{h}}_i)^T\mathbf{m_f} \\
&= m_{\mathbf{f}i}\Delta\tilde{\mathbf{h}}_i^T\mathbf{M}_i\mathbf{1} \\
&= 0,
\end{aligned} \tag{10}$$

which indicates that the mean of $\mathbf{e}_k$ is zero. Now we have zero-mean additive noise $\mathbf{e}_k$, whose covariance matrix is $\mathbf{C}_{\mathbf{e}_k} = \Delta\widetilde{\mathbf{H}}_k\mathbf{C_f}\Delta\widetilde{\mathbf{H}}_k^T$ with diagonal entries

$$\sigma_{\mathbf{e}_k i}^2 = \sum_j \sigma_{\mathbf{f}j}^2\Delta\tilde{h}_{ij}^2 \tag{11}$$

Again, by applying (7) and (8) the above entries can be written as:

$$\sigma_{\mathbf{e}_k i}^2 = \sum_{j \in w_i} \sigma_{\mathbf{f}j}^2\Delta\tilde{h}_{ij}^2 = \sigma_{\mathbf{f}i}^2 \sum_{j \in w_i}\Delta\tilde{h}_{ij}^2 = \sigma_{\mathbf{f}i}^2\|\Delta\tilde{\mathbf{h}}_i\|^2. \tag{12}$$

Combining the additive noise $\tilde{\mathbf{n}}_k$ and $\mathbf{e}_k$ together we define the total additive noise

$$\varepsilon_k = \tilde{\mathbf{n}}_k + \mathbf{e}_k. \tag{13}$$

Since $\tilde{\mathbf{n}}_k$ and $\mathbf{e}_k$ are independent, the covariance matrix of ε_k becomes $\mathbf{C}_k = \mathbf{C}_{\mathbf{e}_k} + \sigma_{\mathbf{n}}^2\mathbf{I}$.

Then the diffraction-limited image $\mathbf{z}$ can be estimated through:

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \sum_k (\mathbf{q}_k - \mathbf{z})^T \mathbf{C}_k^{-1}(\mathbf{q}_k - \mathbf{z}); \tag{14}$$

For simplicity, we only take the diagonal part of $\mathbf{C}_k$ denoted as $\mathbf{U}_k = \mathrm{diag}[u_{k1}, u_{k2}, \ldots]$, whose i-th diagonal entry can be written as:

$$u_{ki} = \sigma_{\mathbf{n}}^2 + \sigma_{\mathbf{f}i}^2\|\Delta\tilde{\mathbf{h}}_i\|^2 \tag{15}$$

Now the estimation problem in (14) can be simplified as:

$$\begin{aligned}
\hat{\mathbf{z}} &= \arg\min_{\mathbf{z}} \sum_k (\mathbf{q}_k - \mathbf{z})^T \mathbf{U}_k^{-1}(\mathbf{q}_k - \mathbf{z}) \\
&= \left(\sum_k \mathbf{U}_k^{-1}\right)^{-1} \sum_k \mathbf{U}_k^{-1}\mathbf{q}_k
\end{aligned} \tag{16}$$

Or in pixel-wise form:

$$\hat{z}_i = \frac{\sum_k u_{ki}^{-1}q_{ki}}{\sum_k u_{ki}^{-1}} \tag{17}$$

which is nothing but an *image fusion* process, where the weight value u_{ki}^{-1} is determined by $\sigma_{\mathbf{n}}^2$ and $\sigma_{\mathbf{f}i}^2\|\Delta\tilde{\mathbf{h}}_i\|^2$. If $\sigma_{\mathbf{n}}^2$ is temporally constant, the higher $\sigma_{\mathbf{f}i}^2\|\Delta\tilde{\mathbf{h}}_i\|^2$ is (more blurry), the lower u_{ki}^{-1} becomes (lower weight). In practice, noise variance $\sigma_{\mathbf{n}}^2$ is viewed as spatially invariant and can be estimated using, for example, median absolute deviation (MAD) method [7].

Once sufficient observations are collected, diffraction-limited image patches (also the sharpest isoplanatic patches) that occasionally appear due to the turbulence variation can be detected through sharpness metrics such as Strehl ratio, image gradient, or local intensity variance [15, 3, 12]. These diffraction-limited patches suffer from noise and cannot be used directly as outputs since noise will be magnified in the following deconvolution step. However, they can be utilized as references for calculating the local variances of $\mathbf{e}_k$ and the weights (15) for the fusion. Let us consider a given isoplanatic region w_i centered at the i-th pixel across all the registered frames. Assume that the sharpest patch of w_i, which can be approximated as a diffraction-limited patch, is detected in the k'-th frame:

$$\mathbf{M}_i\mathbf{q}_{k'} = \mathbf{M}_i\mathbf{H}\mathbf{f} + \mathbf{M}_i\mathbf{n}_{k'}. \tag{18}$$

Then, given the k-th frame we can write the patch difference (see Fig. 5):

$$\mathbf{M}_i(\mathbf{q}_k - \mathbf{q}_{k'}) = \mathbf{M}_i\Delta\widetilde{\mathbf{H}}_k\mathbf{f} - \mathbf{M}_i\mathbf{n}_{k'} + \mathbf{M}_i\mathbf{n}_k \tag{19}$$

Because of the isoplanaticism of local PSF, and of the piecewise constancy of image variance, it can be deduced that:

$$\begin{aligned}
\mathbf{var}[\mathbf{M}_i(\mathbf{q}_k - \mathbf{q}_{k'})] &= \frac{1}{|w_i|}\mathrm{tr}\left(\mathbf{M}_i\Delta\widetilde{\mathbf{H}}_k\mathbf{C_f}\Delta\widetilde{\mathbf{H}}_k^T\mathbf{M}_i \right. \\
&\quad \left. +2\sigma_{\mathbf{n}}^2\mathbf{M}_i^2\right) \\
&= \frac{1}{|w_i|}\mathrm{tr}\left(\mathbf{C_f}\Delta\widetilde{\mathbf{H}}_k^T\mathbf{M}_i\Delta\widetilde{\mathbf{H}}_k\right) + 2\sigma_{\mathbf{n}}^2 \\
&\approx \sigma_{\mathbf{f}i}^2\|\Delta\tilde{\mathbf{h}}_i\|^2 + 2\sigma_{\mathbf{n}}^2,
\end{aligned} \tag{20}$$

where $|w_i|$ is the cardinality of the set. Thus the weight-related value u_{ki} in (15) can be approximated by

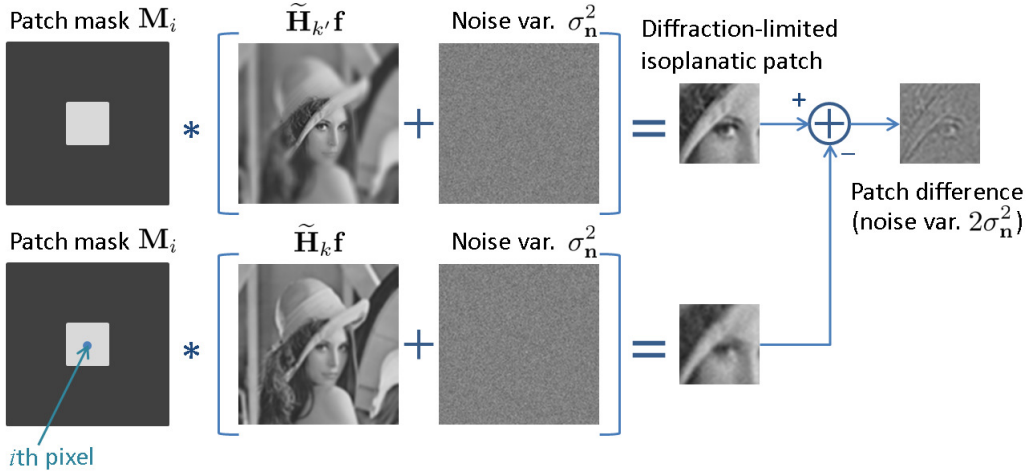$$u_{ki} = \mathbf{var}[\mathbf{M}_i(\mathbf{q}_k - \mathbf{q}_{k'})] - \sigma_{\mathbf{n}}^2. \tag{21}$$

Figure 5. Illustration of calculating the patch difference between the k'-th and k-th frames. Diffraction-limited isoplanatic patch around the i-th pixel is detected in the k'-th frame. This patch difference can be used to produce the weight for the i-th pixel in the k-th frame.

It is worth noting that due to the covariance matrix simplification in (16) and the limited registration accuracy, the PSF of the fused image is, in practice, not the pure diffraction-limited PSF. Namely, it includes diffraction-limited blur, residual turbulence blur, and blur caused by registration error. So we call such PSF *near-diffraction limited*. Because the fusion process strongly reduces the PSF variation, such PSF can be approximately viewed as spatially invariant.

A concise description of the algorithm for the near-diffraction-limited image restoration step is provided below:

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

---

1. Given a registered frame sequence $\{\mathbf{q}_k\}$, divide each frame into $N \times N$ overlapping patches centered at each pixel, and calculate the variance of each patch as a local sharpness measure.

2. For patches centered at the i-th position across all the frames, find the sharpest patch, say in frame k', as the diffraction-limited reference.

3. Set $u_{k'i} = \sigma_{\mathbf{n}}^2$, and for the remaining patches in frame $k \neq k'$ use (21) to calculate u_{ki}.

4. Restore the i-th pixel according to the regression form (17).

5. Go to the $(i+1)$-th pixel and return to step 2.

2.3. Blind Deconvolution

Finally, the unknown image $\mathbf{f}$ and the near-diffraction-limited PSF $\mathbf{h}$ (which is the vector form of $\mathbf{H}$) can be estimated using a Bayesian image deconvolution algorithm de-scribed as:

$$< \hat{\mathbf{f}}, \hat{\mathbf{h}} >= \arg\min_{\mathbf{f}, \mathbf{h}}(\mathbf{z} - \mathbf{H}\mathbf{f}) + \lambda_1 R_f(\mathbf{f}) + \lambda_2 R_h(\mathbf{h}). \quad (22)$$

where R_f and R_h are the regularization terms based on prior knowledge about the latent sharp image $\mathbf{f}$ and the PSF $\mathbf{h}$. Recent research on natural image statistics has shown that image gradients obey *heavy-tailed* distributions that have most of the mass on small values but give significantly more probability to large values than Gaussian distributions [5]. Based on these studies, several *sparsity*-based regularization methods have been introduced and have achieved great success in solving the blind deconvolution problem. One example is Shan *et al.*'s method [17], which is directly employed in the proposed framework as the final step. In what follows, we used the default parameter settings as described in the authors' project page [17] 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 [17] for details.

3. Experiments

Some real video data are tested to illustrate the performance of the proposed restoration framework. In implementing the fusion step in Section 2.2, we set the patch size $N = 9$. The lucky-region algorithm from [2] and the multi-frame reconstruction approach from [23] are also implemented as comparison. The first set of images ($410 \times 380 \times 80$) show the Moon surface imaged from a ground-based telescope (see Fig. 6 (a)). From (b) we can see that though the output image of [2] looks slightly sharper than one of the observed frames, it is still quite blurry probably due to the diffraction-limit blur and the limited number of

frames. [23] provides a better result but with some details (small craters) vanished (Fig. 6 (c)). The proposed method gives a significant improvement in visual quality (Fig. 6 (d)). It successfully removed blur and meanwhile recovered many small craters on the surface (Fig. 6 (h)) that can hardly be seen from either original frame (Fig. 6 (e)), or the outputs of other two methods (Fig. 6 (f), (g)).

The second video stream is an example of long distance imaging and horizontal atmospheric turbulence. It contains a water tower located above ground, and imaged using a stationary camera 2.4 km away. The video is quite noisy and highly blurred (due to long exposure time). 80 frames were taken from the video to produce the result image (300×220) demonstrated in Fig.7. Again, the lucky region method did not provide much improvement in the result and slightly changed the object shape (Fig. 7 (b)). This is probably because the method uses simple translational motion compensation to remove turbulence deformation, which is not quite accurate. The reconstruction method in [23] slightly increased the sharpness (Fig. 7 (c)). The output of the proposed method looks much sharper and clean (Fig. 7 (d)).

4. Conclusion

In this paper we proposed a new framework for restoring a single image from an image sequence distorted by air turbulence. The proposed algorithm first registers the frames to suppress geometric deformation using a B-spline based non-rigid image registration method. Next, a fusion process is carried out to produce an image convolved with near-diffraction-limited PSFs, which can be viewed as space-invariant. Finally, a blind deconvolution algorithm is implemented to remove diffraction-limited blur from the fused image to generate the final output. Experiments using real data illustrate that this framework is capable of alleviating geometric deformation and space-and-time-varying blur caused by turbulence, recovering details of the scene and significantly improving the visual quality.

5. Acknowledgements

References

[1] M. Aubailly, M. A. Vorontsov, G. W. Carhart, and M. T. Valley. Image enhancement by local information fusion with pre-processing and composed metric. *Proceedings of SPIE*, 7090, 2008. 1

[2] 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. *Proceedings of SPIE*, 7463, 2009. 1, 2, 5, 7, 8

[3] J. C. Christou, K. J. Mighell, and R. B. Makidon. Strehl ratio and image sharpness for adaptive optics. *Advances in Adaptive Optics II. Proceedings of the SPIE*, 6272:62721Y, 2006. 4

[4] M. Elad and Y. Hel-Or. A fast super-resolution reconstruction algorithm for pure translational motion and common space-invariant blur. *IEEE Transactions on Image Processing*, 10(8):1187–1193, August 2001. 3

[5] R. Fergus, B. Singh, A. Hertsmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single image. *ACM Transactions on Graphics (SIGGRAPH)*, 2006. 3, 5

[6] D. L. Fried. Probability of getting a lucky short-exposure image through turbulence. *Optical Society of America, Journal*, 68:1651–1658, 1978. 1

[7] F. R. Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69:383–393, 1974. 4

[8] S. John and M. A. Vorontsov. Multiframe selective information fusion from robust error estimation theory. *IEEE Transactions on Image Processing*, 14(5):577–584, May 2005. 1

[9] N. Joshi and M. Cohen. Seeing Mt. Rainier: Lucky imaging for multi-image denoising, sharpening, and haze removal. *IEEE ICCP 2010*, March 2010. 2, 3

[10] N. Joshi, R. Szeliski, and D. Kriegman. PSF estimation using sharp edge prediction. *CVPR*, 2008. 3

[11] N. M. Law. Lucky imaging: Diffraction-limited astronomy from the ground in the visible. *Ph. D. Thesis, Cambridge University*, May 2003. 2

[12] N. M. Law, C. D. Mackay, and J. E. Baldwin. Lucky imaging: High angular resolution imaging in the visible from the ground. *Astron. Astrophys.*, 446:739–745, 2006. 4

[13] D. Li, R. M. Mersereau, and S. Simske. Atmospheric turbulence-degraded image restoration using principal components analysis. *IEEE Geoscience and Remote Sensing Letters*, 4(3):340–344, July 2007. 1

[14] W. E. K. Middleton. *Vision Through The Atmosphere*. University of Toronto Press, Canada, 1958. 1

[15] 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 Engineering*, 33(10):3254–3264, October 1994. 1, 4

[16] M. C. Roggemann and B. M. Welsh. Imaging through turbulence. *CRC Press, Boca Raton, Fla.*, 1996. 1

[17] Q. Shan, J. Jia, and A. Agarwala. High-quality motion deblurring from a single image. *ACM Transactions on Graphics (SIGGRAPH)*, 2008. 3, 5

[18] M. Shimizu, S. Yoshimura, M. Tanaka, and M. Okutomi. Super-resolution from image sequence under influence of hot-air optical turbulence. *CVPR 2008*, June 2008. 1, 2, 3

[19] R. N. Tubbs. Lucky exposures: Diffraction limited astronomical imaging through the atmosphere. *Ph. D. Thesis, Cambridge University*, September 2003. 1, 2, 3
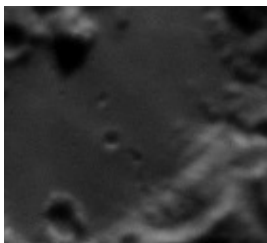
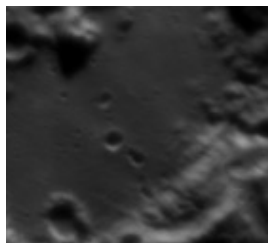(a) One observed frame

(b) Lucky region fusion [2]

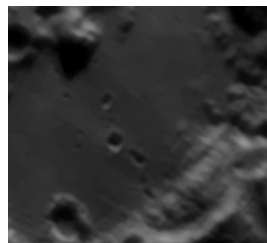(c) Multi-frame reconstruction [23]

(d) Proposed approach

(e) Zoomed part of (a)

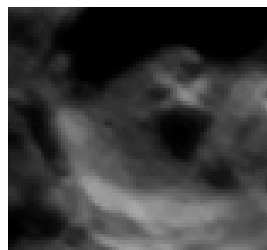(f) Zoomed part of (b)

(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)

Figure 6. Image reconstruction result using 80 frames taken from the video *Moon Surface* distorted by real atmospheric turbulence.
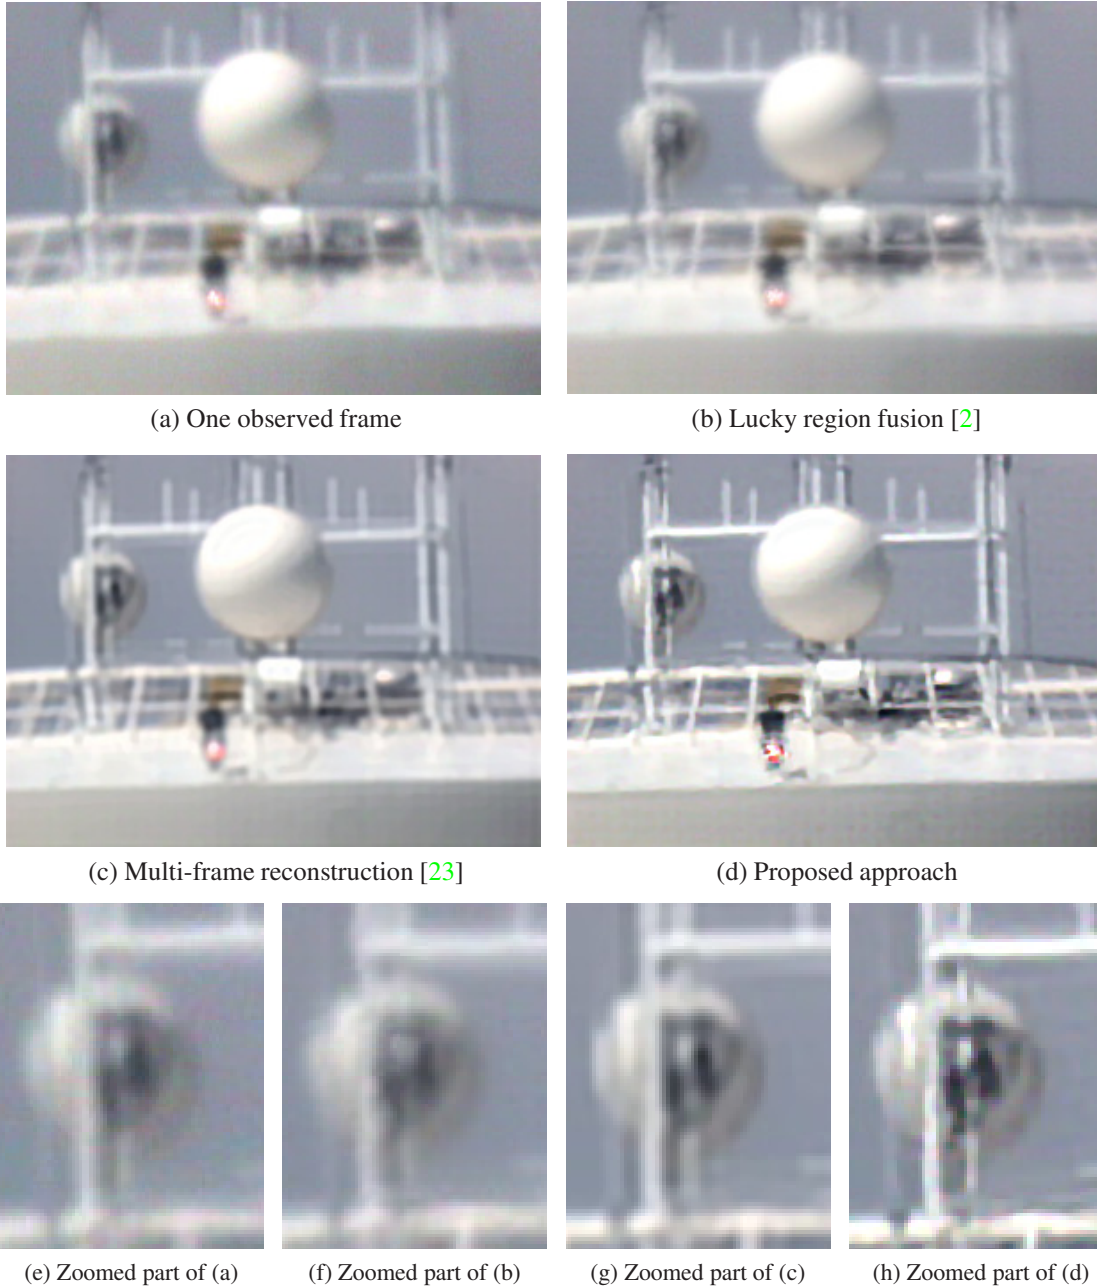
(a) One observed frame   (b) Lucky region fusion [2]

(c) Multi-frame reconstruction [23]   (d) Proposed approach

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

Figure 7. Image reconstruction result using 80 frames taken from the video *Water Tower* distorted by real atmospheric turbulence.

[20] R. K. Tyson. *Principles of Adaptive Optics*. Academic Press, San Diego, USA, 1998. 3

[21] M. A. Vorontsov. Parallel image processing based on an evolution equation with anisotropic gain: integrated optoelectronic architectures. *J. Opt. Soc. Am. A*, 16:1623–1637, 1999. 1

[22] 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. Opt. Soc.* *Am. A*, 18(6):1312–1324, June 2001. 1

[23] X. Zhu and P. Milanfar. image reconstruction from videos distorted by atmospheric turbulence. In *SPIE Electronic Imaging, Conference 7543 on Visual Information Processing and Communication*, San Jose, CA, January 2010. 1, 2, 3, 5, 6, 7, 8