

A computational method for the restoration of images with an unknown, spatially-varying blur

Johnathan Bardsley

*Department of Mathematical Sciences
University of Montana, Missoula, MT 59812.*

Research supported in part by University of Montana Research Grant MRA682.

bardsleyj@mso.umt.edu

Stuart Jefferies

*Institute for Astronomy
University of Hawaii, Kula, HI 96790*

Research supported in part by the Air Force Office of Scientific Research under grant F49620-02-1-0107.

stuartj@maile.ifa.hawaii.edu

James Nagy

*Department of Mathematics and Computer Science
Emory University, Atlanta, GA 30322.*

Research supported in part by the NSF under grant DMS-05-11454 and by an Emory University Research Council grant.

nagy@mathcs.emory.edu

Robert Plemmons

*Departments of Computer Science and Mathematics
Wake Forest University, Winston-Salem, NC 27109.*

Research supported in part by the Air Force Office of Scientific Research under grant F49620-02-1-0107, by the Army Research Office under grants DAAD19-00-1-0540 and W911NF-05-1-0402, and by the DTO under Contract No. 2364-AR03-A1.

plemmons@wfu.edu

Abstract: In this paper, we present an algorithm for the restoration of images with an unknown, spatially-varying blur. Existing computational methods for image restoration require the assumption that the blur is known and/or spatially-invariant. Our algorithm uses a combination of techniques. First, we section the image, and then treat the sections as a sequence of frames whose unknown PSFs are correlated and approximately spatially-invariant. To estimate the PSFs in each section, phase diversity is used. With the PSF estimates in hand, we then use a technique by Nagy and O'Leary for the restoration of images with a known, spatially-varying blur to restore the image globally. Test results on star cluster data are presented.

© 2006 Optical Society of America

OCIS codes: (100.3020) Image reconstruction-restoration; (010.1330) Atmospheric turbulence; (100.3190) Inverse problems; (100.5070) Phase retrieval

References and links

1. J. Biretta, "WFPC and WFPC 2 instrumental characteristics," in *The Restoration of HST Images and Spectra II*, R. J. Hanisch and R. L. White, eds., pp. 224–235 (Space Telescope Science Institute, Baltimore, MD, 1994).
2. H. J. Trussell and S. Fogel, "Identification and restoration of spatially variant motion blurs in sequential images," *IEEE Trans. Image Proc.* **1**, 123–126 (1992).
3. R. G. Paxman, B. J. Thelen, and J. H. Seldin, "Phase-diversity correction of turbulence-induced space-variant blur," *Opt. Lett.* **19**(16), 1231–1233 (1994).
4. M. C. Roggemann and B. Welsh, *Imaging Through Turbulence* (CRC Press, Boca Raton, FL, 1996).
5. M. Faisal, A. D. Lanterman, D. L. Snyder, and R. L. White, "Implementation of a Modified Richardson-Lucy Method for Image Restoration on a Massively Parallel Computer to Compensate for Space-Variant Point Spread Function of a Charge-Coupled Device Camera," *J. Opt. Soc. Am. A* **12**, 2593–2603 (1995).
6. A. F. Boden, D. C. Redding, R. J. Hanisch, and J. Mo, "Massively Parallel Spatially-Variant Maximum Likelihood Restoration of Hubble Space Telescope Imagery," *J. Opt. Soc. Am. A* **13**, 1537–1545 (1996).
7. J. G. Nagy and D. P. O'Leary, "Restoring images degraded by spatially-variant blur," *SIAM J. Sci. Comput.* **19**, 1063–1082 (1998).
8. R. A. Gonsalves, "Phase diversity in adaptive optics," *Opt. Eng.* **21**, 829–832 (1982).
9. C. R. Vogel, T. Chan, and R. J. Plemmons, "Fast algorithms for phase diversity-based blind deconvolution," in *Adaptive Optical System Technologies*, vol. 3353 (SPIE, 1998).
10. R. G. Paxman, T. Schulz, and J. Fienup, "Joint estimation of object and aberrations by using phase diversity," *J. Opt. Soc. Am. A* **9**, 1072–1085 (1992).
11. L. Gilles, C. R. Vogel, and J. M. Bardsley, "Computational Methods for a Large-Scale Inverse Problem Arising in Atmospheric Optics," *Inverse Probl.* **18**, 237–252 (2002).
12. J. Nocedal and S. J. Wright, *Numerical Optimization* (Springer-Verlag, New York, 1999).
13. H.-M. Adorf, "Towards HST restoration with space-variant PSF, cosmic rays and other missing data," in *The Restoration of HST Images and Spectra II*, R. J. Hanisch and R. L. White, eds., pp. 72–78 (1994).
14. D. A. Fish, J. Grochmalicki, and E. R. Pike, "Scanning singular-value-decomposition method for restoration of images with space-variant blur," *J. Opt. Soc. Am. A* **13**, 1–6 (1996).
15. H. J. Trussell and B. R. Hunt, "Image Restoration of Space-Variant Blurs by Sectional Methods," *IEEE Trans. Acoust. Speech, Signal Processing* **26**, 608–609 (1978).
16. J. G. Nagy and D. P. O'Leary, "Fast iterative image restoration with a spatially varying PSF," in *Advanced Signal Processing Algorithms, Architectures, and Implementations VII*, F. T. Luk, ed., vol. 3162, pp. 388–399 (SPIE, 1997).
17. H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems* (Kluwer Academic Publishers, Dordrecht, 2000).
18. P. C. Hansen, *Rank-deficient and discrete ill-posed problems* (SIAM, Philadelphia, PA, 1997).
19. C. R. Vogel, *Computational Methods for Inverse Problems* (SIAM, Philadelphia, PA, 2002).

1. Introduction

The mathematical model for image formation is given by the linear operator equation

$$d = Sf + \eta, \quad (1.1)$$

where d is the blurred, noisy image, f is the unknown true image, η is additive noise, and S is the blurring operator. In the case of spatially invariant blurs, Sf can be written as a convolution of the associated point spread function (PSF) s and the object f ; that is,

$$Sf(u, v) = s \star f(u, v) := \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s(u - \xi, v - \eta) f(\xi, \eta) d\xi d\eta.$$

However, if the blur is spatially variant, Sf cannot be written as a simple convolution operation. Space variant blur can occur from distortions due to telescope optics, such as was the case for the original Hubble Space Telescope Wide-Field/Planetary Camera, which had a large amount of spatial variation because of errors in the shaping mirrors [1]. Other important examples include objects moving at different velocities [2], and turbulence outside the telescope pupil [3]. Accurately modeling the spatially variant PSF in these applications can lead to substantially

better reconstructions of the image. On the other hand, allowing for a fully spatially-varying PSF results in a computationally intractable problem [4]. Thus one typically makes the assumption that s is spatially-invariant on subregions of the image. That is, that the image d can be broken up into regions $\{\Omega_n\}_{n=1}^p$ such that the blur in each region is spatially-invariant. Then, if we let c_n be the indicator function on Ω_n , i.e. $c_n(u, v) = 1$ for $(u, v) \in \Omega_n$ and 0 otherwise, and we define s_n to be the (approximately) spatially-invariant PSF on Ω_n , the PSF s will be given by

$$s(u, v; \xi, \eta) = \sum_{n=1}^p s_n(u - \xi, v - \eta) c_n(\xi, \eta). \quad (1.2)$$

Substituting (1.2) into (1.1), we obtain

$$d = \sum_{n=1}^p s_n \star f_n + \eta, \quad (1.3)$$

where $f_n(\xi, \eta) := c_n(\xi, \eta) f(\xi, \eta)$.

Computational methods for the restoration of images arising from (1.3) have been sought by several researchers, though only in the case that the s_n 's are known. Faisal et al. [5] apply the Richardson-Lucy algorithm, and discuss a parallel implementation. Boden et al. [6] also describe a parallel implementation of the Richardson-Lucy algorithm, and consider a model that allows for smooth transitions in the PSF between regions. Nagy and O'Leary [7] use a conjugate gradient (CG) algorithm with piecewise constant and linear interpolation, and also suggest a preconditioning scheme (for both interpolation methods) that can substantially improve rates of convergence.

In this paper, we present an algorithm for restoring images arising from model (1.3) when the s_n 's are unknown. In our approach, the PSF in each region is estimated using the technique of *phase diversity* [8]. Once the PSF estimates have been obtained, we use a global restoration scheme of Nagy and O'Leary mentioned in the previous paragraph to obtain the restored image.

The paper is organized as follows. We begin in Section 2 with a description of the technique of phase diversity. Our computational method is then presented in Section 3. Results on some tests using simulated star field images taken through atmospheric turbulence are provided in Section 4, while conclusions and comments on the spatially-varying blur removal problem are given in Section 5.

2. PSF estimation and phase diversity

As discussed in, e.g. [9], with atmospheric turbulence the phase varies both with time and position in space. Adaptive optics systems use deformable mirrors to correct for these phase variations. Errors in this correction process arise from a variety of sources, e.g., errors in the measurement of phase, inability of the mirror to conform exactly to the phase shape, and lag time between phase measurement and mirror deformation. Thus, a spatially variant model can potentially produce better restorations.

In this section we describe a phase diversity-based scheme to approximate the PSF associated with each section of a segmented image. For a sufficiently fine segmentation, the PSF can be assumed to be essentially spatially invariant, and thus a blind deconvolution scheme can be applied. The mathematics of this phase recovery process was first described by Gonsalves [8], and has been applied extensively for imaging through atmospheric turbulence [4].

Assuming that light emanating from the object is incoherent, the dependence of the PSF on the phase is given by

$$s[\phi] = |\mathcal{F}^{-1}(pe^{i\phi})|^2, \quad (2.4)$$

where p denotes the pupil, or aperture, function, $\iota = \sqrt{-1}$, and \mathcal{F} denotes the 2-D Fourier transform,

$$(\mathcal{F}h)(\mathbf{y}) = \int \int_{\mathbb{R}^2} h(\mathbf{x}) e^{-i2\pi\mathbf{x}\cdot\mathbf{y}} d\mathbf{x}, \quad \mathbf{y} \in \mathbb{R}^2. \quad (2.5)$$

The pupil function $p = p(x_1, x_2)$ is determined by the extent of the telescope's primary mirror. In phase diversity-based blind deconvolution, the k^{th} diversity image is given by

$$d_k = s[\phi + \theta_k] \star f + \eta_k, \quad k = 1, \dots, K, \quad (2.6)$$

where η_k represents noise in the data, f is the unknown object, s is the point spread function (PSF), ϕ is the unknown phase function, θ_k is the k^{th} phase diversity function.

In atmospheric optics [4], the phase $\phi(x_1, x_2)$ quantifies the deviation of the wave front from a reference planar wave front. This deviation is caused by variations in the index of refraction (wave speed) along light ray paths, and is strongly dependent on air temperature. Because of turbulence, the phase varies with time and position in space and is often modelled as a stochastic process.

Additional changes in the phase ϕ can occur after the light is collected by the primary mirror, e.g., when adaptive optics are applied. This involves mechanical corrections obtained with a deformable mirror to restore ϕ to planarity. By placing beam splitters in the light path and modifying the phase differently in each of the resulting paths, one can obtain more independent data. The phase diversity functions θ_k represent these deliberate phase modifications applied *after* light is collected by the primary mirror. The easiest to implement is *defocus blur*, modelled by a quadratic

$$\theta_k(x_1, x_2) = b_k (x_1^2 + x_2^2), \quad (2.7)$$

where the parameters b_k are determined by defocus lengths. In practice, the number of diversity images is often quite small, e.g., $K = 2$ in the numerical simulations to follow. In addition, one of the images, which we will denote using index $k = 1$, is obtained with no deliberate phase distortion, i.e., $\theta_1 = 0$ in (2.6).

2.1. The minimization problem

To estimate the phase ϕ , and the object f from data (2.6), we consider the least squares fit-to-data functional

$$J_{data}[\phi, f] = \frac{1}{2K} \sum_{k=1}^K \|s[\phi + \theta_k] \star f - d_{k,t}\|^2. \quad (2.8)$$

Here $\|\cdot\|$ denotes the standard L^2 norm. By the convolution theorem and the fact that Fourier transforms preserve the L^2 norm, one can express this in terms of Fourier transforms, $F = \mathcal{F}\{f\}$, $D_k = \mathcal{F}\{d_k\}$, and $S_k[\phi] = \mathcal{F}\{s[\phi + \theta_k]\}$,

$$J_{data}[\phi, F] = \frac{1}{2K} \sum_{k=1}^K \|S_k[\phi]F - D_k\|^2. \quad (2.9)$$

Since deconvolution and phase retrieval are both ill-posed problems, any minimizer of J_{data} is unstable with respect to noise in the data. Hence we add regularization terms to obtain the *full cost functional*,

$$J_{full}[\phi, f] = J_{data}[\phi, f] + \gamma J_{object}[f] + \alpha J_{phase}[\phi]. \quad (2.10)$$

Here the regularization parameters γ and α are positive real numbers, and the regularization functionals J_{object} and J_{phase} provide stability and incorporate prior information.

Because of atmospheric turbulence, variations in the refractive index, and hence the phase itself, can be modelled as a random process [4]. We apply the von Karman turbulence model,

which assumes this process is second order, wide sense stationary, and isotropic with zero mean; see, e.g., [4]. It can be characterized by its power spectral density,

$$\Phi(\omega) = \frac{C_1}{(C_2 + |\omega|^2)^{11/6}}, \quad (2.11)$$

where $\omega = (\omega_x, \omega_y)$ represents spatial frequency. Corresponding to this stochastic model for phase, we take the phase regularization functional

$$J_{phase}[\phi] = \frac{1}{2} \langle \Phi^{-1} \mathcal{F}\{\phi\}, \mathcal{F}\{\phi\} \rangle, \quad (2.12)$$

where $\langle f, g \rangle = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\omega) g^*(\omega) d\omega$, and the superscript * denotes complex conjugate. For regularization of the object, we take the “minimal information prior”

$$J_{object}[f] = \frac{1}{2} \|f\|^2 = \frac{1}{2} \|F\|^2. \quad (2.13)$$

Note that the object regularization functional (2.13) is quadratic and the dependence of the fit-to-data functional (2.8) on the object f is quadratic. Moreover, the Hessian with respect to the object of the full cost functional (2.10) is symmetric and positive definite with eigenvalues bounded below by γ . By setting the gradient with respect to the object equal to zero, one obtains a linear equation whose solution yields the object at a minimizer for J_{full} [10]. From (2.9)-(2.10) and (2.13) one obtains the Fourier representation for the minimizing object,

$$F = \frac{P[\phi]^*}{Q[\phi]}, \quad (2.14)$$

where

$$P[\phi] = \sum_k D_k^* S_k[\phi], \quad Q[\phi] = \gamma + \sum_k |S_k[\phi]|^2. \quad (2.15)$$

Substituting (2.14) back in (2.9)-(2.10), one obtains the cost functional that we will minimize,

$$J[\phi] = J_{reduced\ data}[\phi] + \alpha J_{phase}[\phi], \quad (2.16)$$

where

$$J_{reduced\ data}[\phi] = \sum_k \|D_k\|^2 - \left\langle \frac{P[\phi]}{Q[\phi]}, P[\phi] \right\rangle. \quad (2.17)$$

See Appendix B of [9] for a detailed derivation.

Our objective is to minimize (2.16) in order to obtain a reconstruction of the phase ϕ . In order to do that, we need an effective computational method. Such a method is discussed in the next section.

2.2. The minimization algorithm

To minimize J in (2.16) we will use, as in [11], a quasi-Newton, or secant, method known as limited memory BFGS (L-BFGS) [12]. A generic quasi-Newton algorithm with line search globalization is presented below. We denote the gradient of J at ϕ by $\text{grad}J(\phi)$ and the Hessian of J at ϕ by $\text{Hess}J(\phi)$.

Quasi-Newton / Line Search Algorithm

```

v := 0;
phi_0 := initial guess;
begin quasi-Newton iterations
  g_v := gradJ(phi_v);      % compute gradient
  B_v := SPD approximation to HessJ(phi_v);
  d_v := -B_v^{-1} g_v;     % compute quasi-Newton step
  tau_v := argmin_{tau>0} J(phi_v + tau d_v); % line search
  phi_{v+1} := phi_v + tau_v d_v; % update approximate solution
  v := v + 1;
end quasi-Newton iterations

```

In practice, the line search subproblem is solved inexactly [12].

The BFGS method [12] is one of the most popular quasi-Newton techniques. Given B_v , this method generates a new Hessian approximation B_{v+1} in terms of the differences in the successive approximates to the solution and its gradient,

$$s_v \stackrel{\text{def}}{=} \phi_{v+1} - \phi_v, \quad (2.18)$$

$$y_v \stackrel{\text{def}}{=} \text{grad}J(\phi_{v+1}) - \text{grad}J(\phi_v). \quad (2.19)$$

L-BFGS is based on a recursion for the inverse of the B_v 's,

$$B_{v+1}^{-1} = \left(I - \frac{y_v s_v^T}{y_v^T s_v} \right) B_v^{-1} \left(I - \frac{s_v y_v^T}{y_v^T s_v} \right) + \frac{s_v s_v^T}{y_v^T s_v}, \quad v = 0, 1, \dots \quad (2.20)$$

Given $v \in \mathbb{R}^n$, computation of $B_{v+1}^{-1}v$ requires a sequence of inner products involving v and the s_v 's and y_v 's, together with the application of B_0^{-1} . If B_0 is symmetric positive definite (SPD) and the ‘‘curvature condition’’ $y_v^T s_v > 0$ holds for each v , then each of the B_v 's is also SPD, thereby guaranteeing that $-B_v^{-1} \text{grad}J(\phi_v)$ is a descent direction. The curvature condition can be maintained by implementing the line search correctly [12].

‘‘Limited memory’’ means that at most N vector pairs $\{(s_v, y_v), \dots, (s_{v-N+1}, y_{v-N+1})\}$ are stored and at most N steps of the recursion are taken, i.e., if $v \geq N$, apply the recursion (2.20) for $v, v-1, \dots, v-N$, and set B_{v-N}^{-1} equal to an SPD matrix M_v^{-1} . We will refer to M_v as the *preconditioning matrix*. In standard implementations, M_v is taken to be a multiple of the identity [12]. For our application we choose an M_v which has the operator representation on an operand ψ , given by

$$M_v \psi = \mathcal{F}^{-1} (\Phi^{-1} \mathcal{F}(\psi)),$$

where Φ is defined in (2.11).

Under mild assumptions, the local convergence rate for the BFGS method is superlinear [12]. In the limited memory case, this rate can become linear.

Before continuing, we make the observation that the above computational approach provides estimates of both the phase ϕ and object f (via (2.14)). In our application of removing spatially-varying blur, the main interest is in using this approach for obtaining the PSF as a function of the phase, as given in (2.4). The object f is then reconstructed in a second stage, which we now discuss.

3. Sectioning the image and applying a global restoration scheme

For many spatially variant blurs, in small regions of the image the blur can be well approximated by a spatially invariant PSF. This property has motivated several types of *sectioning methods*

[2, 13, 14, 15] that partition the image, restoring each local region using its corresponding spatially invariant PSF. The results are then sewn together to obtain the restored image. To reduce blocking artifacts at the region boundaries, larger, overlapping regions are used, and then the restored sections are extracted from their centers. Trussell and Hunt [2] proposed using the Landweber iteration for the local deblurring, and suggested a complicated stopping criteria based on a combination of local and global convergence constraints. Fish, Grochmalicki and Pike [14] use a truncated singular value decomposition (TSVD) to obtain the local restorations.

An alternative scheme can be derived by partitioning the image into subregions on which the blur is assumed to be spatially invariant; but, rather than deblurring the individual subregions locally and then sewing the individual results together, we first sew (interpolate) the individual PSFs, and restore the image globally. A global reconstruction avoids the problem of boundary artifacts that can occur in traditional sectioning methods [2, 13, 14, 15]. In algebraic terms, the blurring matrix S , given in (1.1), can be written as

$$S = \sum_{i=1}^p \sum_{j=1}^p D_{ij} S_{ij}, \quad (3.21)$$

where S_{ij} a matrix representing the spatially invariant PSF in region (i, j) , and D_{ij} is a nonnegative diagonal matrix satisfying $\sum \sum D_{ij} = I$. For example, if piecewise constant interpolation is used, then the l th diagonal entry of D_{ij} is 1 if the l th point is in region (i, j) , and 0 otherwise. In principle, any partitioning scheme can be used. The most obvious approach is to partition the image into rectangular subregions, or to use concentric annular regions in the case that the PSFs vary radially from the on-axis PSF.

Faisal et al. [5] use this formulation of the spatially variant PSF, apply the Richardson-Lucy algorithm with piecewise constant interpolation of the PSFs, and discuss a parallel implementation. Boden et al. [6] also describe a parallel implementation of the Richardson-Lucy algorithm, and consider piecewise constant as well as piecewise linear interpolation. Nagy and O'Leary [7] use a conjugate gradient algorithm with piecewise constant and linear interpolation, and also suggest a preconditioning scheme (for both interpolation methods) that can substantially improve the rate of convergence. Furthermore, it is shown in [16] that by generalizing overlap-add and overlap-save convolution techniques, very efficient FFT-based algorithms can be obtained when the image is partitioned into rectangular subregions. Although other partitioning schemes can be used, the computational cost will increase. If the PSFs vary smoothly, then a rectangular partitioning should provide a good approximation of the spatially variant blur. For these reasons, we assume throughout the paper that a rectangular partitioning of the image is used.

Note that, using (3.21), the image formation model can be written as:

$$d = \sum_{i=1}^p \sum_{j=1}^p D_{ij} (s_{ij} \star f) + n, \quad (3.22)$$

where s_{ij} is the PSF for region (i, j) . This model assumes the PSFs are known, so the question is how to use this for blind deconvolution. The algorithm proceeds follows:

1. First determine a region (i, j) for which a good estimate of the PSF can be obtained. Use the L-BFGS method discussed in Section 2.2 on an extended region of (i, j) to obtain reconstruction of the phase ϕ in that region. From ϕ an approximation of the spatially invariant PSF for that region is obtained via (2.4). A reconstruction of the image contained in that region, if desired, could be computed by applying the inverse Fourier transform to F given in (2.14).

2. Now, assuming that the overall space varying PSF varies slowly across the image, the PSFs for the regions neighboring region (i, j) should be similar to the one in region (i, j) . That is, PSF s_{ij} should be a good estimate of the PSFs $s_{i,j+1}$, $s_{i,j-1}$, $s_{i+1,j}$ and $s_{i-1,j}$. Thus, one can use ϕ_{ij} for the initial guess ϕ_0 in the L-BFGS iterations, obtaining reconstruction for the phase, and image if desired. Once again, via (2.4), a reconstruction of the PSF is obtained.
3. Execute steps 1 and 2 for each region. When done, one hopefully has a set of good PSFs, and restored regions of the image. A global image restoration algorithm, with the individual (good) PSFs, can be used to restore the blurred image (as done by Nagy and O’Leary [7]). Note that the restored images in the regions can be pieced together, and used as an initial guess for the global restoration scheme.

4. Numerical examples

In this section we describe some numerical experiments to illustrate the potential advantages of the space variant approach discussed in this paper. All tests were done on a simulated star field image, shown in Fig. 1, which was obtained from the Space Telescope Science Institute (www.stsci.edu).



Fig. 1. Simulated image of a star field.

To simulate an image blurred by a spatially variant point spread function, we begin by generating 1024 different PSFs. The PSFs were created by generating a single pupil function and a moving phase screen, each on a 32×32 grid. Figure 2 shows the pupil function, two selected phase screens, and their corresponding PSFs.

The blurred image, shown in Fig. 3, was then created by convolving each of the 1024 PSFs with successive 32×32 pixels of the true image. By overlapping these regions, we obtain successive 8×8 pixel regions of the blurred image with virtually no blocking (boundary) artifacts; see Fig. 3.

4.1. Advantages of space variant model

To illustrate the potential advantages of using a space variant model, we construct certain average PSFs as follows:

- By averaging all of the 1024 true PSFs, we obtain a single PSF which represents a spatially invariant approximation of the spatially variant blurring operator.
- Next we section the image into 4 equally sized regions (that is, we use a 2×2 partitioning of the image). Note that because of the way we constructed the blurred image, there are





Fig. 2. Pupil function, two sample phases and their corresponding PSFs.




Fig. 3. Simulated space variant blur of star field image; a linear scale is used to display the image on the left, and a log scale is used to display the image on the right.

$\frac{1024}{4} = 256$ true PSFs corresponding to each region. We obtain a single approximate PSF for each region by averaging the 256 true PSFs in their respective regions. We then use these 4 average PSFs to construct an approximation of the space variant blur, as described in Section 3.

- Using an analogous approach on a 4×4 partitioning of the image, we construct 16 PSFs to approximate the space variant blurring operator, each obtained by averaging the $\frac{1024}{16} = 64$ true PSFs corresponding to each region.
- Finally, we use a 16×16 partitioning of the image. In this case $\frac{1024}{256} = 4$ of the true PSFs are averaged to obtain a single PSF for each region.

The conjugate gradient (CG) method was used to reconstruct the image with the various approximations of the PSF as described above. Efficient implementation of the matrix vector multiplications for the space variant cases (4 PSFs, 16 PSFs, and 64 PSFs) was done using the approach outlined in Section 3. Goodness of fit is measured using the relative error,

$$\frac{\|f_{\text{true}} - f_k\|_2}{\|f_{\text{true}}\|_2}$$

where f_{true} is the (diffraction limited) true image, and f_k is the (diffraction limited) solution at the k th CG iteration.

A plot of the relative errors for the special case of noise free data is shown in Fig. 4. Note that when more PSFs are used, a better approximation of the space variant blur is obtained. This is confirmed by the results in Fig. 4. The computed reconstructions at iterations corresponding to the minimum relative errors for each case are shown in Fig. 5.

Similar results occur if noise is added to the blurred image, however the ill-conditioning of the problem means that the reconstructions will be more contaminated with noise. For example, with Poisson and Gaussian noise (mean 0, standard deviation 2) added to the blurred image, the relative errors for CG are shown in Fig. 6 and the corresponding reconstructions are shown in Fig. 7.




Fig. 4. Relative errors at each iteration of the conjugate gradient method, using increasingly more accurate approximations of the spatially variant blur. These results were computed using a noise free blurred image.

4.2. Practical results

The results from the previous section clearly show that if we are able to get accurate approximations of the PSFs, then the space variant model will provide better reconstructions than by




Fig. 5. Reconstructions computed by the conjugate gradient algorithm, using accurate approximations of the PSF, and with a noise free blurred image. For comparison, we also show the true image and the true image convolved with the diffraction-limited PSF.




Fig. 6. Relative errors at each iteration of the conjugate gradient method, using increasingly more accurate approximations of the spatially variant blur. These results were computed using a noisy (Poisson and Gaussian) blurred image.




Fig. 7. Reconstructions computed by the conjugate gradient algorithm, using accurate approximations of the PSF, with Poisson and Gaussian noise added to the blurred image. For comparison, we also show the true image and the true image convolved with the diffraction-limited PSF.

using a space invariant model. In this section we show results when the PSFs are reconstructed using a blind deconvolution algorithm. In particular, we use phase diversity, with one frame and two channels. The blind deconvolution (BD) algorithm we use is described in Section (2.2). Our regularization parameters were taken to be $\alpha = 1 \times 10^{-1}$ and $\gamma = 1 \times 10^{-4}$. We remark that the problem of choosing appropriate regularization parameters is one of the most difficult issues in the numerical solution of ill-posed problems. There are techniques for estimating parameters for simple noise models; see, for example [17, 18, 19]. In many cases good heuristics based on computational experience with the algorithm and the particular application can be useful. The parameters we use are similar to those reported in the literature; see [11].

In the phase estimation step, regularization is applied before computing a minimum of the least squares functional. Thus it is reasonable to seek estimates of the phase and object that yield a gradient norm near zero. We therefore stop L-BFGS iterations once the norm of the gradient had been reduced by six orders of magnitude in all cases. Note that no noise amplification occurs provided the regularization parameters are properly chosen, and, as previously mentioned, the choice of these parameters is dependent on the data.

Once the PSFs are computed from the BD algorithm, we then use the CG algorithm to reconstruct the image, as was done in the previous subsection. Figure 9 shows plots of the relative

errors when no noise is added to the data. At first glance these appear to be disappointing results; 4 PSFs produce lower relative errors than a single PSF, but additional PSFs actually increase the relative error. It is important to note, however, that to obtain more PSFs, the blind deconvolution algorithm must be implemented on small regions of the image. If the small regions do not contain significant object information, then we cannot hope to obtain good approximations of the PSFs. Consequently, given the fact that the object that we are using in our experiment has significant black background, it is not surprising that several such small regions occur. Thus the poor results when using 64 PSFs.

On the other hand, some of the small regions may contain enough significant object information so that the blind deconvolution algorithm can reconstruct good PSFs. This motivates us to consider an approach where we refine the partitioning in each subregion only if further refinement produces good reconstructions of the PSFs. For example:

- Suppose we have the partitioning shown on the left in Fig. 8, and that we have obtained reconstructed PSFs on each of the four regions.
- Refine the partitioning of the image as shown in the middle of Fig. 8, and reconstruct PSFs for each of the (now smaller) subregions.
- Determine if each of these reconstructed PSFs are sufficiently accurate.
 - If a reconstructed PSF is not sufficiently accurate, then reject it and use a PSF from the previous step for this subregion. In this case, further refinement and reconstruction of PSFs for these regions is not needed.
 - If a reconstructed PSF is sufficiently accurate, then accept this PSF.
- Repartition all of the subregions corresponding to accepted PSFs, and repeat the above process. The right plot in Fig. 8 shows a possible repartitioning of selected subregions.

Thus, this adaptive approach uses a “mix” of PSFs computed by the various partitionings of the image.




Fig. 8. Example of adaptive refinement of region partitions. The left plot shows an initial partitioning of an image, the middle shows a refinement of the partitioning, and the right plot shows what might happen in an adaptive approach where only certain subregions are refined further.

To automate the process of determining when a PSF is reconstructed sufficiently accurately, an appropriate measure must be used. Such a measure will be dependent on the application, and the type of images being reconstructed. In order to test the potential of this adaptive scheme, we used a systematic approach for choosing a good PSF by comparing the computed PSFs with the average PSFs described in the previous subsection. Of course in a realistic problem the average PSFs will not be available, and some other mechanism must be used to decide on the quality of the computed PSFs. In any case, our systematic approach shows that if the best computed PSFs

can be found, then substantially better reconstructions can be obtained; see the bottom curve in Fig. 9. The reconstructed images are shown in Fig. 10.

Similar results are obtained with noisy data. For example, with Poisson and Gaussian noise (mean 0, standard deviation 2) added to the blurred image, the relative errors using the CG iterative method are shown in Fig. 11. The corresponding reconstructions are shown in Fig. 12, where we also include arrows to indicate regions in which the space variant approach produces significantly better reconstructions (compared to the diffraction limited true image) than the space invariant (1 PSF) model.




Fig. 9. Relative errors at each iteration of the conjugate gradient method, using increasingly more accurate approximations of the spatially variant blur. The PSFs used to approximate the space variant blur were computed using phase diversity-based blind deconvolution. These results were computed using noise free blurred image data.

5. Conclusions

We have presented a computational approach for solving image restoration problems in which the blur in the image is both unknown and spatially varying. The approach has three stages. The first stage involves sectioning the image into regions in which the blur is believed to be approximately spatially invariant. In the second stage, phase diversity-based blind deconvolution via the L-BFGS optimization algorithm is implemented in order to obtain an estimate of the phase in each region. From these reconstructed phases, the corresponding PSF in each region can be computed. In the final stage, with these PSFs in hand, the object can be reconstructed globally via the algorithm of Nagy and O’Leary.

Our numerical experiments show, first, that using a spatially varying model when a spatially varying blur is present does indeed provide more accurate results. Secondly, we find that in regions with little object information, the phase, and hence, PSF, reconstructions can be inaccurate. This motivates a “PSF mixing” scheme, in which the object is divided into further subregions only in areas in which there is enough object information. A conclusion of particular importance that follows from our numerical experiments is that in the presence of an unknown, spatially varying blur, our approach is much more effective than is standard one-PSF phase-diversity.

We remark that, in principle, our space variant approach should be applicable with other blind deconvolution methods, but as evidenced from our numerical results, the quality of the reconstructed PSFs is important. Thus we would expect that additional diversity information should improve the results. Similarly, we would expect a multi-frame scheme, coupled with our space variant technique, to outperform its single frame counterpart.




Fig. 10. Reconstructions computed by the conjugate gradient algorithm, using PSFs computed from a phase diversity blind deconvolution algorithm. The blurred image data in this case is noise free.




Fig. 11. Relative errors at each iteration of the conjugate gradient method, using increasingly more accurate approximations of the spatially variant blur, and with a noisy (Poisson and Gaussian) blurred image.




Fig. 12. Reconstructions computed by the conjugate gradient algorithm, using PSFs computed from a phase diversity-based blind deconvolution algorithm, and with a noisy (Poisson and Gaussian) blurred image.

The problem of optimal sectioning of the image, and efficient implementation details, is worth further exploration. One might base the partitioning on a-priori knowledge of how the PSF varies; for example, one could partition into concentric annular regions for a radially varying PSF. Another approach is to base the partitioning on the amount of information (e.g., signal to noise ratio) in various regions. In this case, one could use an approach where regions are repartitioned into smaller and smaller subregions until, for example, the mean intensity of the subregion falls below a predetermined threshold (which might be obtained through extensive simulations).