Space-Varying Restoration of Optical Images

James G. Nagy*
V. Paul Pauca and **Robert J. Plemmons**†
Todd C. Torgersen‡

August 22, 1996

Abstract

The improvement in optical image quality is now generally attempted in two stages. The first stage involves techniques in adaptive-optics and occurs as the observed image is initially formed. The second stage of enhancing the quality of optical images generally occurs off-line, and consists of the postprocessing step of *image restoration*. Image restoration is an ill-posed inverse problem which involves the removal or minimization of degradations caused by noise and blur in an image, resulting from, in this case, imaging through a medium. Our work here concerns a new space-varying regularization approach, and associated techniques for accelerating the convergence of iterative image postprocessing computations. Denoising methods, including total variation minimization, followed by segmentation-based preconditioning methods for minimum residual conjugate gradient iterations, are investigated. Regularization is accomplished by segmenting the image into (smooth) segments and varying the preconditioners across the segments. The method appears to work especially well on images that are piecewise smooth. Our algorithm has computational complexity of only $O(\ell n^2 \log n)$, where n^2 is the number of pixels in the image and ℓ is the number of segments used. Also, parallelization is straightforward. Numerical tests are reported on both simulated and actual atmospheric imaging problems. Comparisons are made with the case where segmentation is not used. It is found that our approach is especially attractive for restoring images with low signal to noise ratios, and that magnification of noise is effectively suppressed in the iterations, leading to a numerically efficient and robust regularized iterative restoration algorithm.

*Department of Mathematics, Southern Methodist University, Dallas, TX 75275-0156. Research supported in part by a National Science Foundation Postdoctoral Research Fellowship.

†Department of Mathematics and Computer Science, Wake Forest University, Winston-Salem, NC 27109. Research by these authors sponsored by the Air Force Office of Scientific Research and by the National Science Foundation.

‡Department of Mathematics and Computer Science, Wake Forest University, Winston-Salem, NC 27109. Research sponsored by the National Science Foundation.

1 Introduction

Researchers are actively seeking to overcome the degradation of optical image quality caused by the effects of imaging through a medium such as air, liquids, and other image degradation media. For example, atmospheric turbulence has especially frustrated astronomers since telescopes were invented. The optical effects are in part due to the mixing of warm and cold atmospheric layers, resulting in nonuniformities in the density and refractive index of air. These variations cause parts of the light waveforms from an object to be slowed by different amounts, distorting the image. The resulting twinkling of the stars and other effects are the main limitations of imaging through the atmosphere [7].

Modern methods for the improvement in optical image quality is often attempted in two steps. The first step occurs as the observed image is initially formed. Specially designed *deformable mirrors* operating in a closed-loop adaptive-optics system can partially compensate for the effects of medium turbulence. Optical systems detect the distortions. A wavefront sensor measures the turbulence-induced wavefront distortions using light from either a natural guide star or a guide star artificially generated using range-gated laser backscatter. These distortions are at least partially nullified by adjusting the surface shape of the deformable mirror. To be effective, these corrections have to be performed at real-time speed. Adaptive-optics control systems form the subject of considerable recent investigation, see, e.g., Ellerbroek et. al. [18, 19] and Hardy [27].

The second stage of compensating for the effects of imaging through a medium generally occurs off-line, and consists of the postprocessing step of *image reconstruction* or *restoration*. Here, large-scale computations, using either a simultaneous image of a natural guide star or a large ensemble of images corresponding to different atmospheric realizations, are used to deconvolve the blurring effects of atmospheric turbulence, e. g., Lagendijk and Biedmond [29] and Nagy, Plemmons and Torgersen [32, 33]. Our concern in this paper is the development of robust regularized acceleration techniques for large-scale image postprocessing iterations. Adaptive linear and nonlinear methods are considered.

There are a number of applications for optical image postprocessing:

- Ground-Based Atmospheric Imaging. Here, restoration procedures are used to deconvolve the blurring effects of atmospheric turbulence.
- Endoscopy. Images are obtained through aberration-inducing fluids, and thus optical image restoration can often improve the image quality.
- Microscopy. This involves confocal, fluorescence, and scanning microscopy, sometimes in three dimensions, producing images which often must be restored by postprocessing.
- Underwater Imaging. Again, this technology is known to be limited by turbulence and turbidity.

The outline of our paper follows. A review of relevant background material on image denoising and image deblurring methods is given in §2. A space-varying regularized iterative

restoration method based on image segmentation is proposed in §3. Numerical tests on simulated and real atmospheric imaging data are reported in §4. Some observations and directions for future work are discussed in §5.

2 Image Postprocessing

Linear image postprocessing enhancement methods, especially *direct* FFT-based deconvolution and related approaches, are currently used in many practical applications because of their simplicity and speed. This is especially true in medical imaging, where even the use of iterative methods [5, 11, 15] has not yet been widely accepted in practice, e.g., Biedmond et. al [5]. Convolution backprojection and related direct filtering schemes are often used as standard approaches to image reconstruction and restoration. Two main problems with standard linear methods are oscillations and smoothing. Traditional enhancement methods prevent oscillations in the restored image by smooth regularization techniques [23, 37]. Here, discontinuous or singular image features and oscillatory textures sometimes cannot be restored. We focus on overcoming these restrictions by using space-varying regularization methods. We also provide methods for reducing the computational requirements of certain iterative image enhancement schemes, in order to make their use more practical for applications where both enhancement quality and speed are important considerations.

If we let $\mathcal{H}$ denote the blurring operator and η the noise process, then the image restoration problem with additive noise can be expressed as a linear operator equation

$$g = \mathcal{H}f + \eta, \tag{1}$$

where g and the unknown f denote functions containing the information of the recorded and original images, respectively. Note that when $\mathcal{H} = \mathcal{I}$, the identity operator, the image restoration problem means to extract the image f from a noisy image g. This problem is usually referred to as the *denoising* problem.

Let u and v denote two-dimensional variables. If $\mathcal{H}$ is a convolution operator, as is often the case in optical imaging, then the operator acts uniformly (i.e., in a spatially invariant manner) on f. Here, (1) can be written as

$$\mathcal{H}f(u) = \int_\Omega h(u-v)f(v)dv. \tag{2}$$

The problem is to both deconvolve and denoise the recorded image during the reconstruction process, and we refer to this as the *denoising and deblurring* problem. In optical imaging, the kernel h in (2) is called the convolution *point spread function* (PSF). After discretization of (1), the spatial operator $\mathcal{H}$ defined by h in (2) is a matrix that we denote by H. Here, in the spatially invariant case, H is a block Toeplitz matrix with Toeplitz blocks [29].

A classical approach employed for solving (1) is that of penalized least squares, which is also called Tikhonov regularization in the inverse problems literature. This requires minimization of the expression

$$\|\mathcal{H}f - g\|^2 + \alpha J(f), \tag{3}$$

where $\|\cdot\|$ denotes the norm on $L^2(\Omega)$, α is a positive (regularization) parameter and the functional $J(f)$ serves the purpose of stabilizing the least squares problem and penalizing certain undesirable artifacts like spurious oscillations in the computed f. Various choices of $J(f)$ can be made, including $\|Sf\|^2$, where S is some smoothing differential operator, or the identity. This model leads to fast linear methods for computing f, and is often the method of choice by practitioners [29]. However the use of other norms such as the $\mathcal{L}_1$ norm, lead to nonlinear minimization methods [2] which sometimes result in superior enhancement of blocky, noisy images, but with added computational cost. Such approaches are reviewed in the next subsection.

2.1 Variational Methods

Osher and Rudin [36] have suggested an image enhancement method based on solving a nonlinear PDE constrained minimization problem where the function being minimized is the Total Variation (TV) of the image $f = f(x, y)$. Numerous papers on various TV approaches for image denoising or restoration have been written in the past five years. A partial list is provided in the references [1, 2, 10, 12, 17, 36, 39, 40, 41]. By now the principles of TV methods are well-understood, and recent papers have concentrated on speeding up the computations. Significant computational requirements are often necessary for these nonlinear TV-based methods.

The variational method of Rudin, Osher and Fatemi [39, 40] considers the following constrained minimization problem:

$$\min_f \int_\Omega |\nabla f| du \qquad \text{subject to} \quad \|\mathcal{H}f - g\|^2 = \sigma^2, \tag{4}$$

where ∇f denotes the gradient of f, and σ is the noise level. At a point $u = (x, y)$ in the image domain, $f(u) = f(x, y)$, and so

$$|\nabla f(u)| = \sqrt{f_x^2 + f_y^2}. \tag{5}$$

The quantity

$$\int_\Omega |\nabla f| du = \int_\Omega \sqrt{f_x^2(x,y) + f_y^2(x,y)} dx dy \tag{6}$$

is called the *total variation* norm of f. The minimization in (4) is a form of regularization, a step necessary in solving most ill-posed inverse problems. The TV method [39, 40] is especially effective for recovering a blocky, discontinuous, function from noisy data [17].

Several authors [1, 2, 10, 12, 17, 41] have considered the following closely-related Tikhonov regularization problem (3), where

$$\alpha J(f) = \alpha \int_\Omega |\nabla f| du, \tag{7}$$

see, e.g., Acar and Vogel [1] and Vogal and Oman [41]. Here α is a positive regularization parameter which measures the trade-off between a good fit and an oscillatory solution. This

method corresponds to the use of the $\mathcal{L}_1$ norm in the discrete case. At a stationary point of (7), the gradient of (3) vanishes, giving:

$$z(f) \equiv \mathcal{H}^*(\mathcal{H}f - g) - \alpha \nabla \cdot \left(\frac{\nabla f}{|\nabla f|} \right) = 0, \qquad u = (x, y) \in \Omega. \tag{8}$$

Due to the term $1/|\nabla f|$, (8) is a degenerate nonlinear second order diffusion equation. The degeneracy can be alleviated by modifying the diffusion coefficient; see Vogel and Oman [41]. More precisely, let $\hat{f}$ be an approximation to f given by

$$\kappa_\beta(f) = \frac{1}{\sqrt{|\nabla f|^2 + \beta}} \qquad \beta > 0, \tag{9}$$

$$\mathcal{L}_\beta(f)\hat{f} = -\nabla \cdot (\kappa_\beta(f) \nabla \hat{f}), \tag{10}$$

and define

$$\mathcal{R}_\beta(f)\hat{f} = (\mathcal{H}^*\mathcal{H} + \alpha \mathcal{L}_\beta(f))\hat{f}. \tag{11}$$

Then (8) becomes the following non-degenerate system

$$\mathcal{R}_\beta(f)\hat{f} = \mathcal{H}^* g, \quad u = (x, y) \in \Omega, \quad \text{with} \quad \frac{\partial f}{\partial n} = 0, \quad u = (x, y) \in \partial\Omega. \tag{12}$$

A recent survey of related approaches to image enhancement can be found in Alvarez and Morel [2].

Various numerical schemes have been devised to obtain the minimizer of the functional (7). For example in Rudin and Osher [39] and in Rudin et. al. [40], an explicit time marching scheme is used. However, the time step must be chosen small. Thus the number of iterations to optimal convergence can be quite large [2]. Vogel and Oman [41] have introduced a *lagged diffusivity fixed point iteration* approach, which we denote by FP, to solve (12). If $R_\beta(f^k)$, H and L_β denote the discretization matrices of $\mathcal{R}_\beta(f^k)$, $\mathcal{H}$ and $\mathcal{L}_\beta$, respectively, then the FP iteration will produce a sequence of approximations $\{f^k\}$ to the solution f and can be expressed as a sequence of systems of linear equations:

$$R_\beta(f^k) f^{k+1} \equiv (H^*H + \alpha L_\beta(f^k)) f^{k+1} = H^* g, \qquad k = 0, 1, \ldots \tag{13}$$

In the denoising case, numerical experiments cited in Vogel and Oman [41] indicated that the FP iteration method often gives a faster convergence rate than the time marching method, with overall greater speed for the entire process. Note that in (13), obtaining f^{k+1} from f^k requires one to solve a large linear system with coefficient matrix $H^*H + \alpha L_\beta(f^k)$. For deconvolution, H is block Toeplitz with Toeplitz blocks. In any case, the matrix $\alpha L_\beta(f^k)$ is a 2-D nonconstant Laplacian with five bands. Its spectrum can vary widely over the outer iterations, i.e. with the index k [10].

A disadvantage of experimentally obtained data sets representing the PSFs (obtained, for example, using guide stars) is that they are also subject to degradations caused by noise during

the image formation process. Thus, removal of such degradations may be necessary *prior to* any computations using the PSF in the image restoration postprocessing step. In this paper we consider, among other approaches, TV-based denoising of the PSF whose model is formulated as

$$s = h + \eta, \tag{14}$$

where h is the blur produced by the atmospheric turbulence on a single point, and s is the actual measurement or observed PSF. For the denoising problem, $H^*H = I$, so that the coefficient matrix in (13) is a sum of the identity matrix and $L_\beta(f^k)$.

The question arises here that if the computationally expensive TV-based minimization is used for denoising, then why not use it as well to solve the complete image restoration problem for atmospheric imaging. The reason for this is that TV denoising of the PSF can be accomplished with little extra computation since point spread functions h for optical imaging can usually be treated as having small extent [7]. Thus, the cost of preprocessing the PSF is much less than that of deblurring f by TV methods. Our approach to the deblurring step is considered in the next section, where a numerically efficient and stable preconditioned iterative regularization method is proposed.

2.2 Preconditioned Conjugate Gradients

In this subsection we consider numerical methods for approximating the solution to the linear restoration problem in discretized (matrix) form obtained from the operator equation given in (1), namely:

$$g = Hf + \eta. \tag{15}$$

Here, for simplicity we assume that the image in question is square, and accordingly H is $n^2 \times n^2$, where n is the number of pixels in each direction. The vectors $f, g \in \mathcal{R}^{n^2}$ represent the observed and true (unknown) image pixel values, respectively, unstacked by rows.

There are many ways to handle this inverse problem of recovering f, which include both direct methods such as the singular value decomposition, and linear iterative methods such as conjugate gradients, and/or nonlinear TV methods as defined earlier by (13). Direct methods can sometimes be impractical for problems having matrices of large dimensions, as is the case in image restoration. Recently, though, Fish, Grochmalicki and Pike [20] proposed a "scanning SVD" which computes an approximate singular value decomposition. Here, we consider iterative methods. The fact that (15) is an ill-posed inverse problem must also be considered. By ill-posed it is meant that computed solutions to $Hf = g$ are likely to be extremely corrupted by noise.

Regularization methods attempt to alleviate sensitivity to the noise by "filtering" out eigen-components of the solution belonging to the noise subspace. For some iterative methods, it has been established that early termination of the iterations accomplishes this regularization effect. That is, the eigen-components of the signal subspace are reconstructed in the first (possibly many) iterations and, after reaching a certain approximate restoration, the components in the noise subspace begin to be reconstructed. It is at this point, where the noise begins to contaminate the reconstruction, that the iterations are halted, see, e.g., Hanke and Hansen [23].

The linear conjugate gradient method (CG) has the previously mentioned regularizing properties [23, 34]. The classical CG method solves symmetric positive definite systems of equations [21]. The blurring operator (PSF) associated with atmospheric turbulence is generally modeled by a symmetric Gaussian function [29]. In optical image restoration, though, the matrix H is typically constructed from observations (such as a bright guide star) and, hence, is generally not quite symmetric in practice. One way to overcome this difficulty, and still work with the classical CG method, is to apply CG to the factored form of the normal equations $H^*(g - Hf) = 0$. (A variant of this method is known as CGLS [6]). However, there are two main disadvantages with this approach:

- The ill-conditioning of the (implicit) normal equations matrix H^*H is much worse than that of H. Thus the convergence rate may be slow, and full reconstruction of the components in the signal subspace may require many iterations.

- Each CGLS iteration requires matrix-vector multiplications with both H and H^*.

The first of these difficulties can be partially addressed through *preconditioning*; the second difficulty is considered in the following subsection. It is known that convergence of CGLS will be fast if the singular values of H^*H are clustered around one [4]. Preconditioning amounts to obtaining a matrix C such that $(C^*C)^{-1}H^*H$ has a more favorable spectrum. Here, $Q \equiv (C^*C)^{-1}$ is an an *approximate inverse preconditioner* for H^*H. With preconditioned CGLS, linear systems with C and C^* need to be solved at each iteration.

Furthermore, preconditioning of ill-posed inverse problems is a delicate matter. Although clustering of the spectrum will provide fast convergence rates, it also can have the disadvantage of mixing the signal and noise subspaces [25]. To summarize, we observe that an effective preconditioner for ill-posed inverse problems should satisfy the following two properties:

- The spectrum of the preconditioned system must cluster the components belonging to the signal subspace. In this way reconstructing an approximation to f can be done quickly, without contaminating the early iterations with noise.

- It must be relatively easy to solve systems of equations with C and C^*.

For spatially invariant point spread functions h in image restoration, both of these goals can be achieved. Briefly, since the discretized h, given as H, is a block Toeplitz matrix with Toeplitz blocks, it can be approximated by a block circulant with circulant blocks matrix C [13]. Then C can be easily diagonalized by the 2-D unitary fast Fourier transform (FFT) [16]; that is, $C = F^*\Lambda F$, where Λ is a diagonal and generally complex matrix. Using FFTs, solutions of linear systems with C and C^* can be computed using only $O(n^2 \log n)$ arithmetic operations, as can matrix-vector multiplications with H and H^*.

Since the eigenvalues of C are readily available, those with small absolute values, which are associated with the noise subspace, can be replaced by ones (A method for separating the signal and noise subspaces is summarized in the following section). In this way, C clusters the singular values of H associated with the signal subspace, and acts like the identity on the noise subspace; see Hanke et. al. [25] for further details.

We point out that conjugate gradients has a regularizing effect on ill-posed inverse problems such as (15). In particular, the regularizing method now takes the form

$$\min \|\mathcal{H}f - g\|^2 \quad \text{subject to} \quad f \in \mathcal{S}_k, \tag{16}$$

where $\mathcal{S}_k$ is a certain Krylov-like subspace of dimension k, and the number of iterations k is a regularization parameter; see Hanke and Hansen [23].

2.3 A Preconditioned Iterative Regularization Method

We next address the question of how one can avoid matrix multiplications with H^*, as well as solving systems with C^* when using preconditioned iterative regularization with C. As previously mentioned, often the theoretical model of the blurring operator is symmetric (*e.g.,* a Gaussian PSF in atmospheric blurring). In our optical imaging applications, H is often obtained from observations of the PSF which are degraded by noise. Hence H is not likely to retain a symmetric structure. However, since the theoretical blurring model is symmetric, we would expect H to be nearly symmetric, and that one could obtain a good symmetric approximation to this operator. This is the symmetrization approach taken by Hanke and Nagy [24] for spatially invariant blurs in astronomical imaging and adopted here as well. Thus, without loss of generality, in the following work we assume that the blurring operator H is symmetric (if not, it is symmetrized using the techniques suggested by Hanke and Nagy [24]).

However, due to noise and approximation errors, we cannot expect H to be positive definite. Therefore, the classical CG method is still not applicable. CGLS could be used, but it does not take full advantage of the new symmetric property imposed in H (each iteration still requires two multiplications with H). Although there are other variants of CG that one might attempt [4], the question of whether they still retain a regularization effect through early termination has, until very recently, been unanswered. An important result in this direction was established by Hanke [22] for a *minimum residual* (MR) type variant of CG, which he refers to as MR-II. The MR-II method is a mathematically, but not computationally, equivalent formulation of a method first proposed by Paige and Saunders [38] whose iterates differ in general from those of MR, since the residuals are minimized in different Krylov subspaces.

Hanke and Nagy [24], have used the MR-II method for restoration of atmospherically blurred images in combination with a preconditioning approach, known as preconditioned regularization, similar to the one described in the previous subsection. In addition to the advantage that each iteration needs only one matrix-vector multiplication and one solution to a system involving the preconditioner (as opposed to two each for PCGLS), it was observed by these authors that convergence for (P)MR-II can be substantially faster than (P)CGLS.

The preconditioned iterative regularization scheme was first introduced by Hanke, Nagy, and Plemmons [25]. Given a block circulant with circulant blocks matrix

$$C = F^* \Lambda F \tag{17}$$

described earlier, and a a truncation parameter τ as defined in Hanke et. al. [25], which effectively associates the spectrum of C with the signal and noise subspaces, a *regularized preconditioner*

C_τ is defined as

$$C_\tau = F^* \Lambda(\tau) F, \tag{18}$$

where the diagonal matrix $\Lambda(\tau)$ is defined by

$$\Lambda(\tau)_{jj} = \begin{cases} |\Lambda_{jj}| \text{ if } |\Lambda_{jj}| > \tau \\ 1 \text{ otherwise} \end{cases} . \tag{19}$$

The *approximate inverse regularized preconditioner* is then given by

$$Q_\tau \equiv C_\tau^{-1} = F^* \Lambda(\tau)^{-1} F, \tag{20}$$

so that Q_τ is an approximate inverse of H. Then $Hf = g$ becomes the approximate inverse regularized preconditioned system

$$Q_\tau H f = Q_\tau g. \tag{21}$$

In our application, H is symmetric, but generally indefinite, while C_τ is symmetric positive definite. Moreover, only a single column of H and of C_τ are needed in order to perform the computations.

We next show how the parameter τ is found. In Hanke and Nagy [24], the L-curve criterion was used in order to chose the τ that selects the noise and signal subspaces, as it would be used in determining a truncation index for the truncated singular value decomposition. Specifically, given $|\lambda_1| \geq \tau \geq |\lambda_{n^2}|$ where $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_{n^2}$ are the eigenvalues in Λ, let

$$z_\tau = \sum_{|\lambda_j| > \tau} \frac{v_j^* g}{\lambda_j} v_j, \quad \text{and} \quad r_\tau = g - C z_\tau \tag{22}$$

where v_j is the jth column of F^*. Then the L-curve consists of a log-log plot of

$$||z_\tau||^2 \quad \text{verses} \quad ||r_\tau||^2. \tag{23}$$

For larger τ the solutions z_τ have norm typically on the same order as f while the residual tends to decrease. Thus this part of the plot remains relatively flat. As τ is decreased, the following situation usually occurs:

- more components of the noise subspace begin to corrupt z_τ, so that $||z_\tau||$ begins to grow large, and
- $||r_\tau||$ begins to approach the noise level where it essentially remains constant as τ is further decreased.

Therefore, the plot begins to rise at some point, and thus has a distinct L-shaped appearance. The value of τ corresponding to the corner of the L typically provides a reasonable estimate of the separation of the signal and noise subspaces. That is, the eigen-components corresponding to those λ_j satisfying $|\lambda_j| > \tau$ are associated with the signal subspace, and the remaining with the noise subspace. For further details on the properties of the L-curve, we refer to Hansen and O'Leary [26] and the references therein. The preconditioned MR-II method is given as Algorithm 1 in Figure 1 and an example of an L-curve for actual USAF satellite image data presented in §4, is given in Figure 2.

Algorithm 1: MR-II **with approximate inverse preconditioning by** Q_τ

Input

g – observed data

H – symmetrized point spread function operator

τ – truncation parameter

Q_τ – approximate inverse preconditioner defined by (20)

f_0 – initial guess

Output

$f_k \approx f$

Method

$q_1 = g - Hf_0$
$f_1 = f_0$
$r_1 = Q_\tau q_1$
$v_{-1} = 0$
$v_0 = Q_\tau H r_1$
$u_{-1} = 0$
$u_0 = Hv_0$
$w_{-1} = 0$
$w_0 = Q_\tau u_0$
$\beta = \langle w_0, u_0 \rangle^{1/2}$
$v_0 = \frac{1}{\beta} v_0$
$w_0 = \frac{1}{\beta} w_0$
$u_0 = \frac{1}{\beta} u_0$
$k = 1$
while (not stop) do
  $\rho = \langle r_k, u_{k-1} \rangle$
  $f_{k+1} = f_k + \rho v_{k-1}$
  $r_{k+1} = r_k - \rho w_{k-1}$
  $\alpha = \langle w_{k-1}, Hw_{k-1} \rangle$
  $v_k = w_{k-1} - \alpha v_{k-1} - \beta v_{k-2}$
  $u_k = Hw_{k-1} - \alpha u_{k-1} - \beta u_{k-2}$
  $w_k = Q_\tau H w_{k-1} - \alpha w_{k-1} - \beta w_{k-2}$
  $\beta = \langle w_k, u_k \rangle^{1/2}$
  $v_k = \frac{1}{\beta} v_k$
  $w_k = \frac{1}{\beta} w_k$
  $u_k = \frac{1}{\beta} u_k$
  $k = k + 1$
end while

Figure 1: Algorithm 1

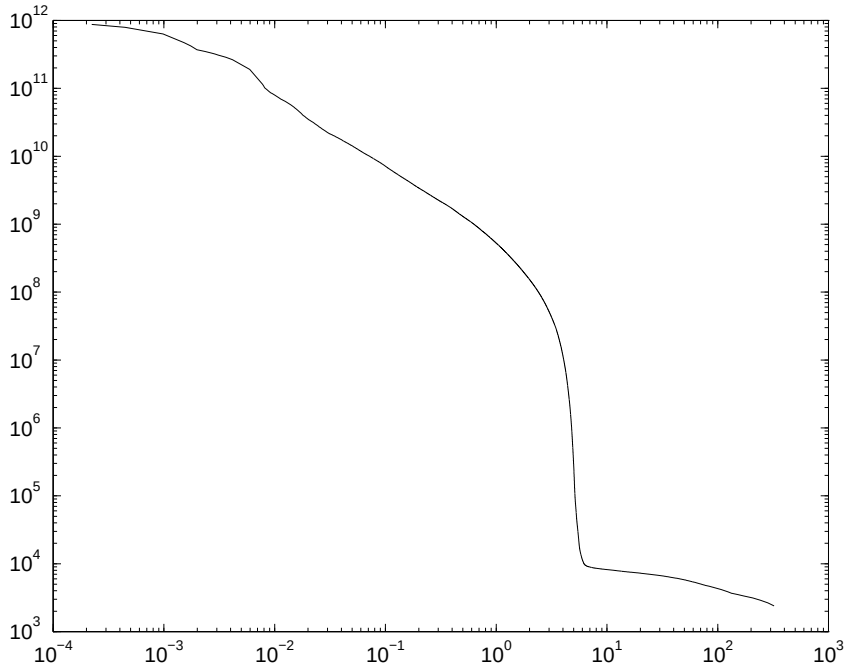

Figure 2: L-curve or log–log plot of $||f_\tau||^2$ vs. $||r_\tau||^2$ for an actual satellite image computation.

3 Space–Varying Restoration (SVR)

In this section we introduce our space-varying restoration (SVR) approach for the purpose of deblurring optical images. SVR is a variant of the preconditioned regularization MR-II method discussed in §2.3. It provides a space-varying regularized iterative restoration scheme based on image segmentation. Recently, other image restoration methods based on image segmentation have been investigated [3, 14, 31].

The motivation for our work is based on the observation that the preconditioned regularization MR-II method of Hanke and Nagy [24] applies the same level of regularization and preconditioning to the entire image, regardless of how much information or noise is present in various portions of the image. Note that in regions containing lower signal to noise ratios (such as background areas) one would expect fast convergence to amplify noise; that is, there may be too much preconditioning in these regions. Therefore, we propose to apply regularized preconditioning in a space-varying manner. Specifically, our approach considers a multi-level preconditioning scheme in order to vary the level of restoration depending on the amount of activity in various regions of the image. These activity regions are identified by a segmentation scheme described in the next section. Once these regions are defined, we use the L-curve scheme [26] on each region to identify the signal and noise subspaces, and hence the level of restoration. Thus, we expect that more regularized preconditioning will be used in regions of high activity and less regularized preconditioning in regions of low activity. The method should work especially well on images that are piecewise smooth.

For the image data presented in §4, the boundaries of the activity regions tend to coincide with steep edges in the image, thus the segmentation scheme also results in essentially smooth segments. The net effect is to suppress the amplification of the noise signal by adaptively accelerating convergence in areas of high signal to noise ratio. By regularizing the computation on independent smooth segments, a stable, space-varying restoration method is achieved.

3.1 Image Segmentation

Image segmentation has been widely studied and a number of approaches have been investigated (see, e.g., Castleman [9]). Heuristic segmentation schemes for adaptive restoration methods are often based on comparing the image intensity, the gradient of the image intensity, the Laplacian of the image intensity, or some combination of these measures to a set of threshold values. Other approaches include watershed type methods, region growing and numerous edge operators, e.g., Sobel [9].

Nevertheless, segmentation remains a problem dependent technique [35]. A successful method for one class of problems may yield poor solutions for other problems. In the case of restoration of optical images, segmentation is especially troublesome: "good" segmentations in the presence of noise are considerably difficult to obtain. Furthermore, segmentation, as a preprocessing step in space-varying restoration, is required to be computationally inexpensive.

A simple approach to obtaining a fast segmentation for our optical image restoration purposes is the technique of multiple thresholding [9]. Let the matrix G denote an observed $n \times n$ image of pixel values and let g denote the n^2-vector obtained by unstacking G row-wise. The image G is segmented into ℓ non-overlapping regions $G_1, G_2, ..., G_\ell$ using a monotone increasing sequence of threshold values $g_{min} = t_0, t_1, t_2, ..., t_\ell = g_{max}$ and the following rule: pixel g_i is in segment G_m, for $1 \leq m \leq \ell$, if and only if $t_{m-1} < g_i \leq t_m$.

It was determined that for the data used in our tests, some form of smoothing or denoising is necessary to improve the effectiveness of the segmentation procedure. Here, median filtering was chosen to smooth the images and reduce noise prior to segmentation by multiple thresholding. Similarly, effectiveness can be improved by enhancing high activity features using the gradient or the Laplacian, or by transforming the intensity pixel distribution to a more favorable one. While gradient-based methods seem to be particularly sensitive to noise in the observed image, the transformation of the intensity pixel distribution using the standard technique of histogram equalization [9] to improve contrast and facilitate spacing of threshold boundaries turns out to be more promising for the purpose of restoration of atmospheric images.

For all of the image data presented in §4, median filtering is used to smooth the observed image and reduce noise prior to segmentation. For the restorations shown, threshold values $t_0, ..., t_\ell$ are chosen by inspection. The threshold values chosen by inspection were found to be consistent with histogram based techniques [9] which separate multi-modalities in a histogram of the pixel values. Figure 3 illustrates the segmentation scheme described above as applied to some real data (see Example 3 in §4).

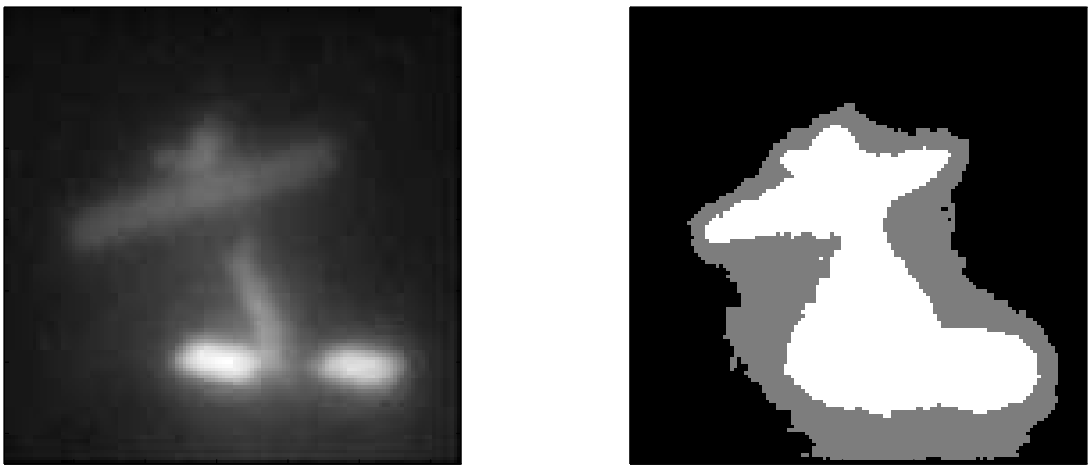

Figure 3: Segmentation applied to satellite data.

3.2 An SVR Algorithm

We now describe our scheme for extending and improving upon the method of Hanke and Nagy by incorporating segmentation and space-varying restoration in different regions according to activity level. As previously mentioned, the basic idea is to use a slightly different preconditioner over each of the regions so that convergence is accelerated in regions of high activity.

Consider again matrix C, a block circulant with circulant blocks approximation to H, and recall that $C = F^*\Lambda F$, where Λ is a diagonal matrix containing the eigenvalues of C, and F is the 2-D unitary Fourier transform matrix. Given the non-overlapping segmented regions $G_1, \ldots, G_\ell$ of the $n \times n$ observed image G, the L-curve criterion suggested in §2.3 can be used to choose a truncation tolerance τ_i for the signal and noise subspaces of each region. Let g be the vector obtained from G by unstacking the image pixels row-wise. Then the segmentation of G can be equivalently defined as

$$g = D_1 g + D_2 g + \cdots + D_\ell g, \tag{24}$$

where D_i is an $n^2 \times n^2$ diagonal masking matrix given by

$$[D_i]_{kk} = \begin{cases} 1 & \text{if pixel } g_k \text{ belongs to segment } G_{(i)} \\ 0 & \text{otherwise.} \end{cases} \tag{25}$$

Then an approximate inverse preconditioner for region $G_{(i)}$ is defined to be Q_{τ_i} given in (20), replacing τ by τ_i.

With the notation in place, we can now introduce our space-varying restoration scheme. The SVR method proposed here is shown as Algorithm 2 in Figure 4. SVR is a minimum residual type method that belongs to a class of Krylov subspace methods [22] and, more specifically, is closely related to PMR-II. Using ℓ segments, SVR solves a problem $Hf = g$ with n^2 unknowns

Algorithm 2: SVR

Input

g – observed data

ℓ – number of segments

$I = D_1 + D_2 + \ldots + D_\ell$ – masking matrices for segmentation

$\{\tau_i\}_{i=1}^{\ell}$ – truncation parameters

$\{Q_{\tau_i}\}_{i=1}^{\ell}$ – approximate inverse preconditioners

H – symmetrized point spread function operator

f_0 – initial guess

Output

$f_k \approx f$

Method

$q_1 = g - Hf_0$

$f_1 = f_0$

$r_1 = \sum_{i=1}^{\ell} D_i r_1^{(i)}$, $r_1^{(i)} = Q_{\tau_i} q_1$ for $i = 1, \ldots, \ell$

$v_{-1} = 0$, $v_{-1}^{(i)} = 0$ for $i = 1, \ldots, \ell$

$v_0 = \sum_{i=1}^{\ell} D_i v_0^{(i)}$, $v_0^{(i)} = Q_{\tau_i} H r_1^{(i)}$ for $i = 1, \ldots, \ell$

$u_{-1} = 0$, $u_{-1}^{(i)} = 0$ for $i = 1, \ldots, \ell$

$u_0 = \sum_{i=1}^{\ell} D_i u_0^{(i)}$, $u_0^{(i)} = H v_0^{(i)}$ for $i = 1, \ldots, \ell$

$w_{-1} = 0$, $w_{-1}^{(i)} = 0$ for $i = 1, \ldots, \ell$

$w_0 = \sum_{i=1}^{\ell} D_i w_0^{(i)}$, $w_0^{(i)} = Q_{\tau_i} u_0^{(i)}$ for $i = 1, \ldots, \ell$

$\beta = \langle w_0, u_0 \rangle^{1/2}$

$v_0 = \frac{1}{\beta} v_0$, $w_0 = \frac{1}{\beta} w_0$, $u_0 = \frac{1}{\beta} u_0$

$w_0^{(i)} = \frac{1}{\beta} w_0^{(i)}$ for $i = 1, \ldots, \ell$

$k = 1$

while (not stop) do

  $\rho = \langle r_k, u_{k-1} \rangle$

  $f_{k+1} = f_k + \rho v_{k-1}$

  $r_{k+1} = r_k - \rho w_{k-1}$

  $\alpha = \langle w_{k-1}, \sum_{i=1}^{\ell} D_i H w_{k-1}^{(i)} \rangle$

  $v_k = w_{k-1} - \alpha v_{k-1} - \beta v_{k-2}$

  $u_k = \sum_{i=1}^{\ell} D_i H w_{k-1}^{(i)} - \alpha u_{k-1} - \beta u_{k-2}$

  $w_k = \sum_{i=1}^{\ell} D_i w_k^{(i)}$, $w_k^{(i)} = Q_{\tau_i} H w_{k-1}^{(i)}$ for $i = 1, \ldots, \ell$

  $w_k = w_k - \alpha w_{k-1} - \beta w_{k-2}$

  $\beta = \langle w_k, u_k \rangle^{1/2}$

  $v_k = \frac{1}{\beta} v_k$, $w_k = \frac{1}{\beta} w_k$, $u_k = \frac{1}{\beta} u_k$

  $w_k^{(i)} = \frac{1}{\beta} w_k^{(i)}$ for $i = 1, \ldots, \ell$

  $k = k + 1$

end while

Figure 4: Algorithm 2

by computing a projection of a sequence of implicit iterates of dimension ℓn^2 produced by (modified) PMR-II applied to a larger, related problem. To understand the relation between SVR and PMR-II, consider the block diagonal approximate inverse preconditioned system

$$QTx = Qy \tag{26}$$

where the unknown $x \in \mathcal{R}^{\ell n^2}$ and

$$\begin{aligned} Q &= \mathrm{diag}(Q_{\tau_1}, Q_{\tau_2}, \ldots, Q_{\tau_\ell}), \\ T &= \mathrm{diag}(H, H, \ldots, H), \\ y &= [g^T, g^T, \ldots, g^T]^T. \end{aligned} \tag{27}$$

Note that the block diagonal matrix Q is symmetric positive definite and (26) is equivalent to the system

$$Q^{1/2}TQ^{1/2}z = Q^{1/2}y, \qquad x = Q^{1/2}z. \tag{28}$$

A key observation in the understanding of SVR in Algorithm 2 is that the quantities $w_k^{(i)}$ in the main loop of the algorithm are of the form $(QT)^k w_0$ where w_0 is formed by stacking the $w_0^{(i)}$ for $1 \le i \le \ell$ into a vector of dimension ℓn^2.

Further, let D be an $n^2 \times \ell n^2$ matrix representing a projection D: $\mathcal{R}^{\ell n^2} \to \mathcal{R}^{n^2}$ given by

$$D = \left[\begin{array}{cccc} D_1, & D_2, & \ldots, & D_\ell \end{array} \right] \tag{29}$$

where D_i are the diagonal matrices defined in (25) for $1 \le i \le \ell$. Observe that an approximate restoration f_k computed by SVR is a projection under D of a vector x_k belonging to a Krylov subspace given by

$$\mathcal{K}_k(r_1; QT) = \mathrm{span}(r_1, (QT)r_1, (QT)^2 r_1, \ldots, (QT)^k r_1) \tag{30}$$

where $r_1 = Q(y - Tx_0)$ is the initial (preconditioned) residual and x_0 is formed from ℓ repetitions of the initial guess f_0. The careful reader may observe that SVR can be derived by substituting Q, T, x_0, and y into PMR-II and making the following two modifications.

First, observe the scalar quantities $\rho, \alpha,$ and β in SVR are computed similarly to the corresponding scalar quantities in PMR-II, except the inner product used in SVR is defined relative to the projection D. More precisely, let $\langle \cdot, \cdot \rangle_D$ denote a semi-definite inner product on a, $b \in \mathcal{R}^{\ell n^2}$ given by

$$\langle a, b \rangle_D = \langle Da, Db \rangle. \tag{31}$$

See Hanke [22] for a discussion of choice of norms and the corresponding derivation of various conjugate gradient type methods. It is possible that for general a and b in $\mathcal{R}^{\ell n^2}$ that the parameter $\beta = \langle w_k, u_k \rangle_D = \langle Dw_k, Du_k \rangle$ may become zero. This can be verified by constructing a vector b with a zero in every position where the diagonal matrix $D^T D$ in $\mathcal{R}^{\ell n^2 \times n^2}$ has a one. In this case $\langle a, b \rangle_D = 0$ for any a. Although at this writing it is not resolved whether β could become zero for our particular applications, we emphasize that this situation seems unlikely because of the repetitive way in which our vectors w_k and u_k of size ln^2 are formed.

Second, the computation and storage of uninteresting quantities are suppressed, e.g., an approximate solution (restoration) f_k requires only n^2 values. The approximated solution f_k is obtained by projecting x_k at each iteration onto $\mathcal{R}^{n^2}$ using D; that is,

$$f_k = Dx_k. \tag{32}$$

Note also that SVR reduces to PMR-II, that is, SVR computes exactly the same sequence of vector iterates as PMR-II, when $\tau_1 = \tau_2 = \ldots = \tau_\ell$.

A number of variants of SVR are possible based on different choices of which variables in PMR-II are projected, and which are fully represented in a space of dimension ℓn^2. At one extreme, all variables in PMR-II could be fully represented in $\mathcal{R}^{\ell n^2}$. Here, this variant of SVR is equivalent to applying MR-II to the preconditioned system $QTx = Qy$, and then combining the desired portions of the independently computed solutions using the diagonal matrices D_i. At the other extreme, all of the variables in PMR-II could be represented by their projections, i.e., a solution to the system $DQTx = DQy$ is sought. The SVR variant presented as Algorithm 2 in Figure 4 takes a position in between these two extremes. Projected values are used for the purposes of computing the scalars ρ, α, and β.

Development of SVR variants remains a topic for future investigation. Formal convergence analysis of SVR will also be presented in future work. At this point, it seems that our restoration scheme as presented above could be slightly modified to deal with non-uniform or spatially varying blur. Also, we note that SVR could be used for multiple-frame image restoration of an object. By multiple-frame restoration is meant the restoration of an image from a *set* of observations of both the PSF and the observed image. This particular application is of interest because of the inherent sensitivity of optical image restoration algorithms to changes in the observations of the PSF.

Our restoration method has computational complexity of only $O(\ell n^2 \log n)$, where n^2 is the number of pixels in the image and ℓ is the number of segments used. Moreover, parallelization of Algorithm 2 seems a natural step to take. The idea is that every image segment can be processed concurrently in ℓ different nodes. Using the master/slave paradigm, a master node determines when the inner preconditioned linear systems are solved, and when other tasks, such as $H * w_{k-1}^{(i)}$, are performed. By simulating a distributed data structure, where each slave node i maintains local variables associated with segment i, communication can be minimized to mostly *sending* and *receiving* instructions. The only time when "heavy" communication occurs and synchronicity is required, is when the master node forms the solution of an outer preconditioned system. We have a parallelized version of Algorithm 2 and it is implemented using a Parallel Toolbox for MATLAB [28]. It is not the purpose of the present paper to give an analysis of such a parallel algorithm. However, in our experiments, we note that the computing time of the parallelized version of SVR using 3 to 5 segments was no more than that of MR-II with preconditioning applied to a single segment.

4 Numerical Tests

In this section we present numerical tests of three data samples to illustrate the effectiveness of the space-varying restoration (SVR) method, in comparison to the unpreconditioned MR-II and preconditioned-based regularized MR-II methods in Hanke and Nagy [24]. The first data sample consists of true and degraded images of a simple known block object. The second data sample consists of true and degraded images of simulated satellite data as it would be captured by a ground-based telescope using an adaptive-optics system. The third and last data sample includes actual ground-based telescope images obtained at the USAF Phillips Laboratory Starfire Optical Range using an adaptive-optics imaging procedure. Here, total variation minimization is used for the preprocessing step of denoising the guide star. For the SVR method, the degraded images in each sample were partitioned into three segments using the segmentation scheme described in §3. All the computations were done in MATLAB 4.2c.

Example 1: Regular Boundaries. This example consists of a simple 256×256 image (with 65,536 unknown pixel values) containing two blocks having different gray scale intensities with a common boundary, as shown in Figure 5a. To obtain a blurred image, we used a Gaussian-type point spread function shown in Figure 5b (and described in further detail in Example 2) and convolved it with the original image. The degraded image was then obtained by adding 10% normally distributed random noise, with mean 0 and variance 1.0. This degraded image is shown in Figure 5c.

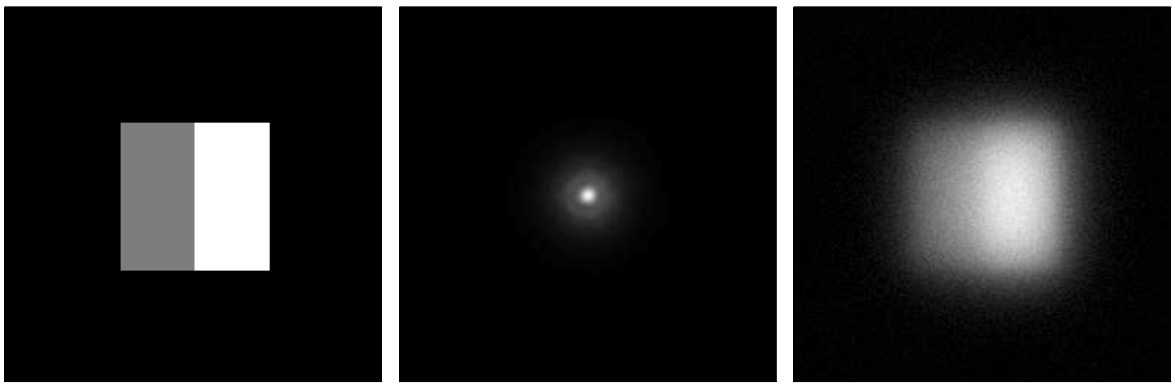

a. True image b. PSF c. Degraded image

Figure 5: Block image data.

The L-curve method was used to select the tolerance or truncation parameters τ_i for the preconditioners Q_{τ_i} associated with each segment. This procedure is described in §2.3. Recall that the truncation parameter τ_i separates the spectrum into signal and noise subspaces, thereby enabling the space-varying preconditioner to cluster the singular values of H associated with the signal subspace, and act like the identity on the noise subspace. For the regularized preconditioned MR-II method (SVR with one segment), the truncation parameter we obtained

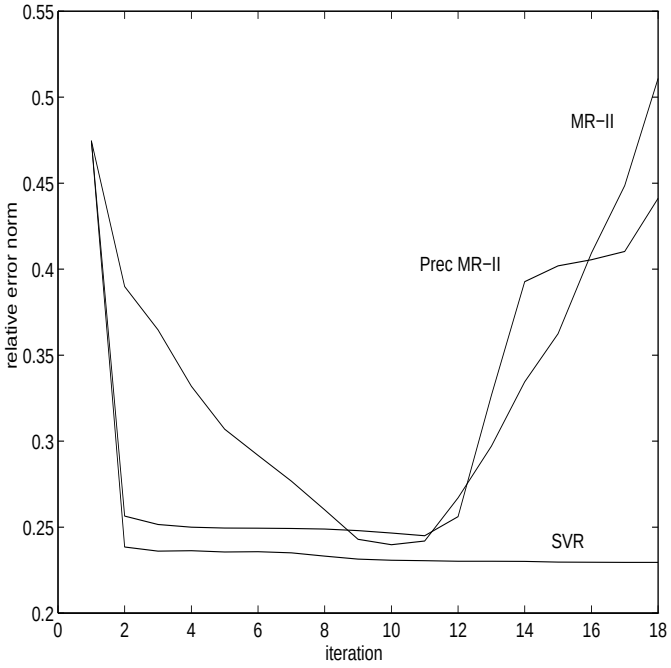

Figure 6: Relative errors for block image data.

resulted in a signal subspace involving 350 eigenvalues; that is, 350 eigenvalues were greater than the truncation parameter (following the definitions in (19) and (18)). For the SVR method, the truncation parameters we obtained for each of the three regions resulted in signal subspaces involving 100 eigenvalues for the region associated with the background, 350 eigenvalues for the region associated with the border regions, and 450 eigenvalues for the region containing most of the object. (following the discussion of §3.2).

Because we have the true solution, for each method we are able to quantify solutions in terms of their relative errors. A plot of the relative error norms is shown in Figure 6. As can be seen from the plot, the SVR method allows us to obtain smaller relative errors, and in that sense, better solutions. Moreover, SVR appears to stabilize the iterations in that the relative errors are reduced quickly, but then remain at about the same level for many iterations. For instance, comparing the relative errors at the 15th iteration, we see that SVR still provides a good restoration, whereas restorations by the other methods are corrupted with noise. This is a distinct advantage since it is difficult, e.g., Hanke and Hansen [23], to predict the "optimal" stopping point in practice, and hence being a few iterations further away from the minimum relative error solution will not have much effect on the computed restoration. Thus, robustness is incorporated into our SVR scheme, leading to more stability in the iterations. Figure 7 shows the computed restorations for each method, which further illustrates the stable behavior of SVR.

Example 2: Simulated Ground-Based Telescope Data. The image boundaries in Example 1 are very regular, and hence the segmentation used to construct the SVR preconditioners is relatively simple to obtain. In the next example, consider a 256×256 image with irregular

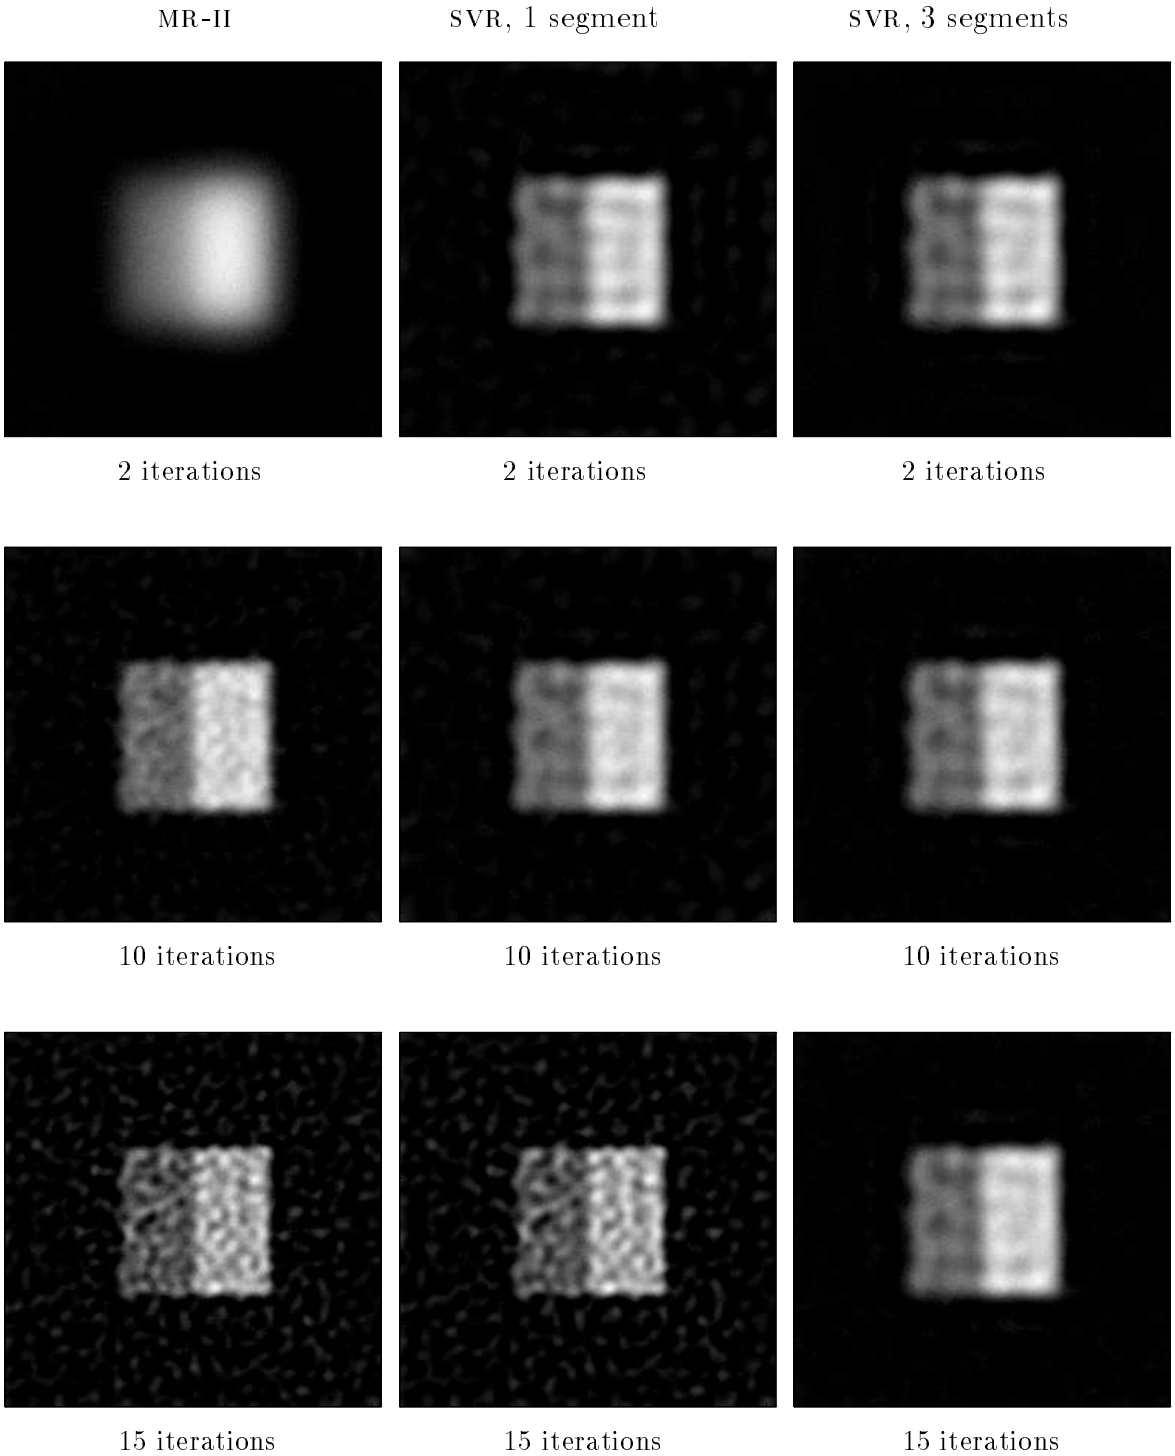

Figure 7: Computed restorations of the block image.

boundaries. This data was obtained from the US Air Force Phillips Laboratory, Lasers and Imaging Directorate, Kirtland Air Force Base, New Mexico. Specifically, the true object is an ocean reconnaissance satellite, which is shown in Figure 8a. A computer simulation algorithm at Phillips Laboratory was used to produce an image of the satellite, shown in Figure 8c, as would be observed from a ground-based telescope using adaptive-optics compensation [18]. The satellite was modeled as being 12 meters in length and in an orbit 500 kilometers above the surface of the earth. The simulated charge-coupled device (CCD) for forming the image was a 65,536 pixel square array. CCD root-mean-square read-out noise variance was fixed at 15 microns per pixel to reflect a realistic state-of-the-art detector.

In actual field experiments, several hundred measurement are averaged to reduce the effects of noise. The PSF and degraded image used in this data, which are shown in Figure 8b, were created to simulate this situation [8, 30].

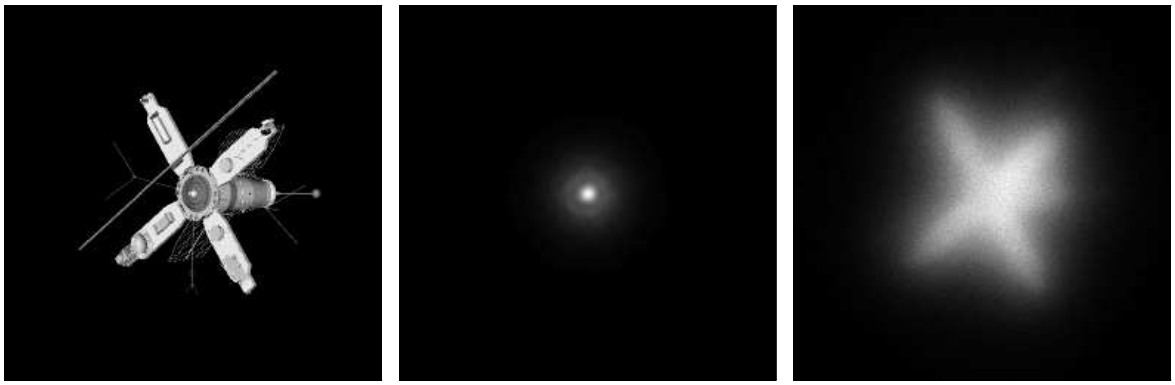

a. True image b. PSF c. Degraded image

Figure 8: Simulated ground-based telescope image data.

Again we use the L-curve method to select the tolerance parameters for the preconditioners. When preconditioning on only one segment, following (19), this corresponded to a signal subspace involving 1335 eigenvalues. With SVR, and following (19), the L-curve suggested using 600, 1335, and 1347 eigenvalues for the signal subspaces of each of the three respective regions. A plot of the relative errors is shown in Figure 9. Observe that the SVR scheme does not quite obtain as small a relative error as the non-preconditioned MR-II method. We attribute this observation to the low noise levels in the data, which essentially reduces the need for regularization in this particular problem. However, the SVR algorithm exhibits a robust and stable behavior in the sense that the relative errors decrease quickly and remain at a low level for many iterations. Computed restorations at various iterations are shown in Figure 10.

Example 3: Ground-Based Telescope Data. We consider here a real ground-based telescope image taken by a USAF facility of a satellite circling the earth, and a corresponding PSF image of a nearby natural guide star, Deneb. The guide star was chosen in the near field-of-view for the satellite, in order to enforce spatial invariance of the PSF for the operator equation given in

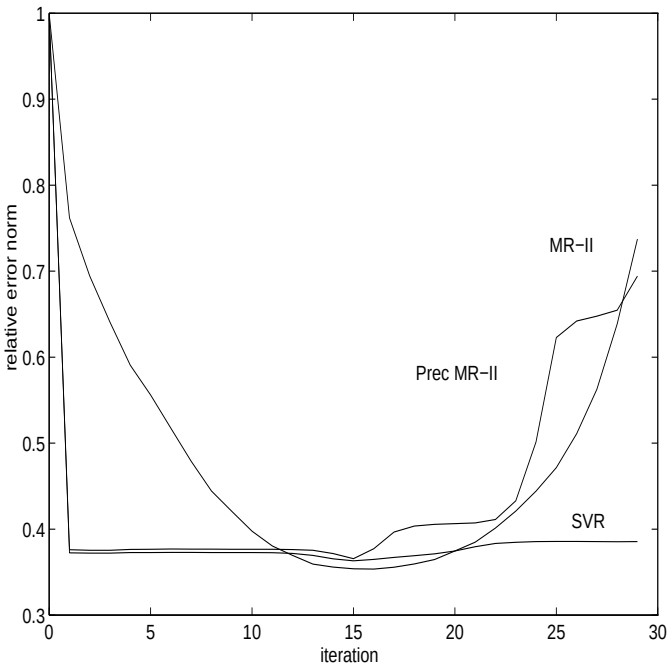

Figure 9: Relative errors for simulated satellite image data.

(1). Deneb is Arabic for "the tail" and this star is the tail of the Cygnus constellation, a great swan flying southward down the Milky Way. It is a brilliant white star with visual magnitude 1.4. Deneb and wing stars of the swan form an asterism called, for obvious reasons, the Northern Cross. Figure 11 shows the observed degraded image and an image of the guide star Deneb. We remark that the effects of atmospheric turbulence have been partially compensated by an adaptive-optics system [27] as the satellite and guide star were imaged through the telescope. The adaptive-optics system used was controlled by a procedure similar to that developed by Ellerbroek, Van Loan, Pitsianis, and Plemmons [18, 19].

Figure 12a shows a surface plot of the observed PSF image, and it is not difficult to see that the image contains a substantial amount of noise. A preprocessing step to reduce this noise is necessary to obtain a better restoration. In our experiments we attempted to restore the image both with and without denoising the PSF, and found that significantly better restorations were obtained when denoising. We tried various techniques for removing noise from the PSF, such as local averaging, median filtering and total variation minimization. Best results were obtained with the total variation (TV) method, as described in §2.1. Here, the sharp spike associated with the PSF is preserved very nicely, and the computationally intensive TV method is feasible in this case since the PSF has small support. Figure 12b shows a surface plot of the TV denoised PSF. We mention that because the TV method requires solving a nonlinear minimization problem, in general it can be expensive to implement. However the extent of the PSF is generally very small compared to the dimensions of the image, and so the cost of denoising the PSF is reasonably inexpensive compared to the cost of restoration.

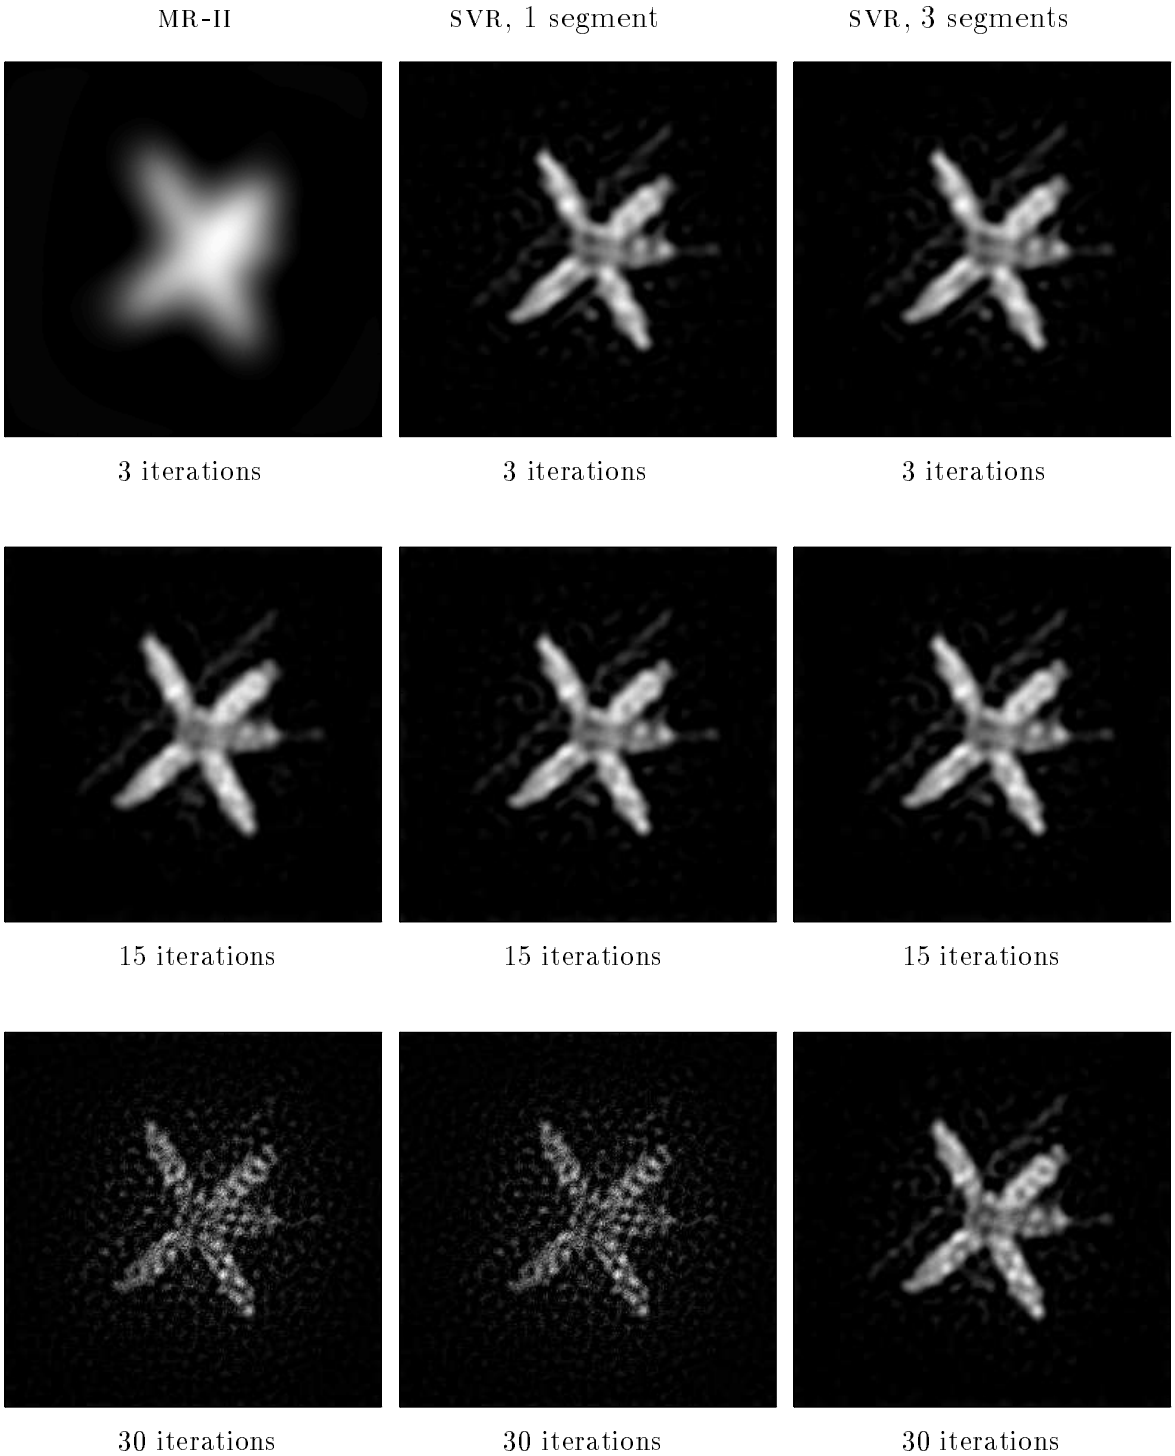

Figure 10: Computed restorations of the simulated satellite image.

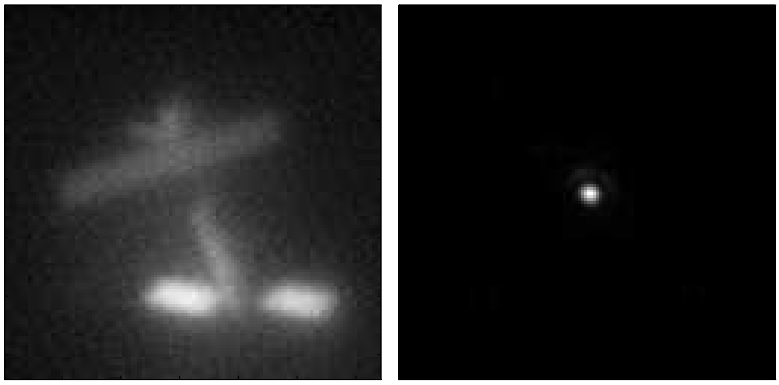

a. Observed image b. Observed PSF (Deneb)

Figure 11: Real ground-based telescope data.

It should be emphasized that this data is real, and that a true image of the object is not available. Therefore, we cannot use a plot of the relative errors to determine a stopping criteria. Many methods have been proposed for terminating the iterations, including the discrepancy principle, generalized cross validation, and the L-curve method [23]. There are advantages and disadvantages to each of these techniques, and we do not want to advocate any one method or go into any detailed discussion of them. However, a significant observation from our previous tests is the stable behavior of the SVR iterates. This suggests that examining the norm of the difference in successive iterates may determine when the computed restoration can no longer be improved significantly. Figure 13 shows computed restorations for each of the images at various iterations. Once again we observe the stability of the restorations using the SVR method.

5 Conclusions

In this paper we have presented a new segmentation-based preconditioning method, which we refer to as SVR, leading to space-varying regularized iterative restoration of optical images. Denoising methods, including total variation minimization, followed by segmentation-based preconditioning methods for minimum residual conjugate gradient iterations, were investigated. Numerical tests were made on both simulated and practical atmospheric imaging problems. The SVR preconditioning method appears to have certain advantages:

- Even if a stopping criteria does not halt the algorithm precisely at the appropriate solution (either a few iterations before or after), it appears that the SVR method will still provide a restoration that is close to optimal. We have observed that the SVR method rapidly reduces the relative error, and then stabilizes at a level close to that of a near optimal solution for several iterations. Examining the norms of the difference of successive iterations, in conjunction with one of the standard stopping criteria, leads to a robust iterative method.

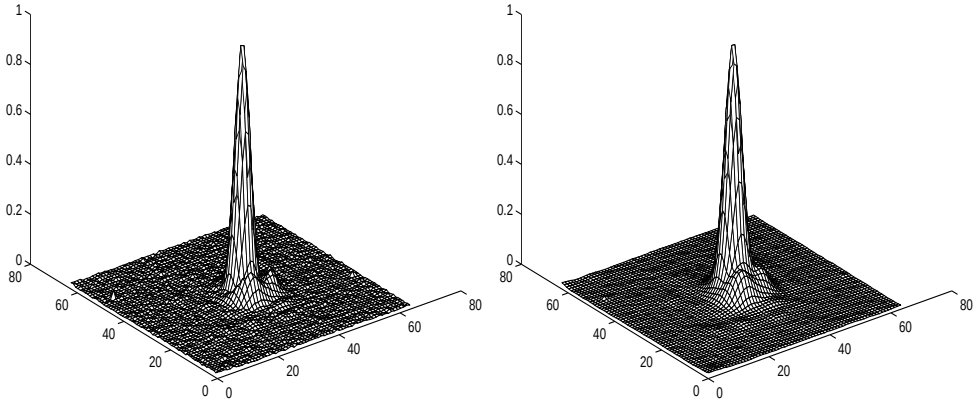

a. Observed PSF b. TV denoised PSF

Figure 12: Preprocessing of Point Spread Function (PSF).

- As described in §3, the SVR preconditioning scheme can easily be implemented in parallel. In this case, the total computational cost of our scheme is essentially the same as for the fast preconditioned MR-II restoration method, which has computational complexity of only $O(n^2 \log n)$, where n^2 is the number of pixels in the image.
- The appropriate segmentation procedure used, as well as the number of segments chosen, is problem dependent. More studies and comparisons are needed in this regard. Analyzing the image in a well-chosen, piecewise smooth form appears to yield better overall separation of the signal and noise subspaces for the purpose of choosing the space-varying preconditioners. The resulting iterative restoration using SVR can effectively suppress noise contamination of the iterates. Over smoothing the image is avoided, and singular image features such as sharp edges and oscillatory textures can be restored without resorting to a computationally expensive nonlinear optimization method such as total variation restoration.

Acknowledgements. The authors wish to thank R. Carreras, B. Ellerbroek and C. Matson from the US Air Force Phillips Laboratory, Lasers and Imaging Directorate, Kirtland Air Force Base, New Mexico for supplying much of the simulated and real data used in the numerical tests. We also thank R. Chan from the Chinese University of Hong Kong for supplying a section of the total variation code used to denoise the guide star Deneb.

References

[1] R. Acar and C. Vogel, *Analysis of Bounded Variation Penalty Methods*, Inverse Problems, 10 (1994), pp. 1217–1229.

[2] L. Alvarez and J. Morel, *Formalization and Computational Aspects of Image Analysis*, Acta Numerica (1994), pp. 1–59.

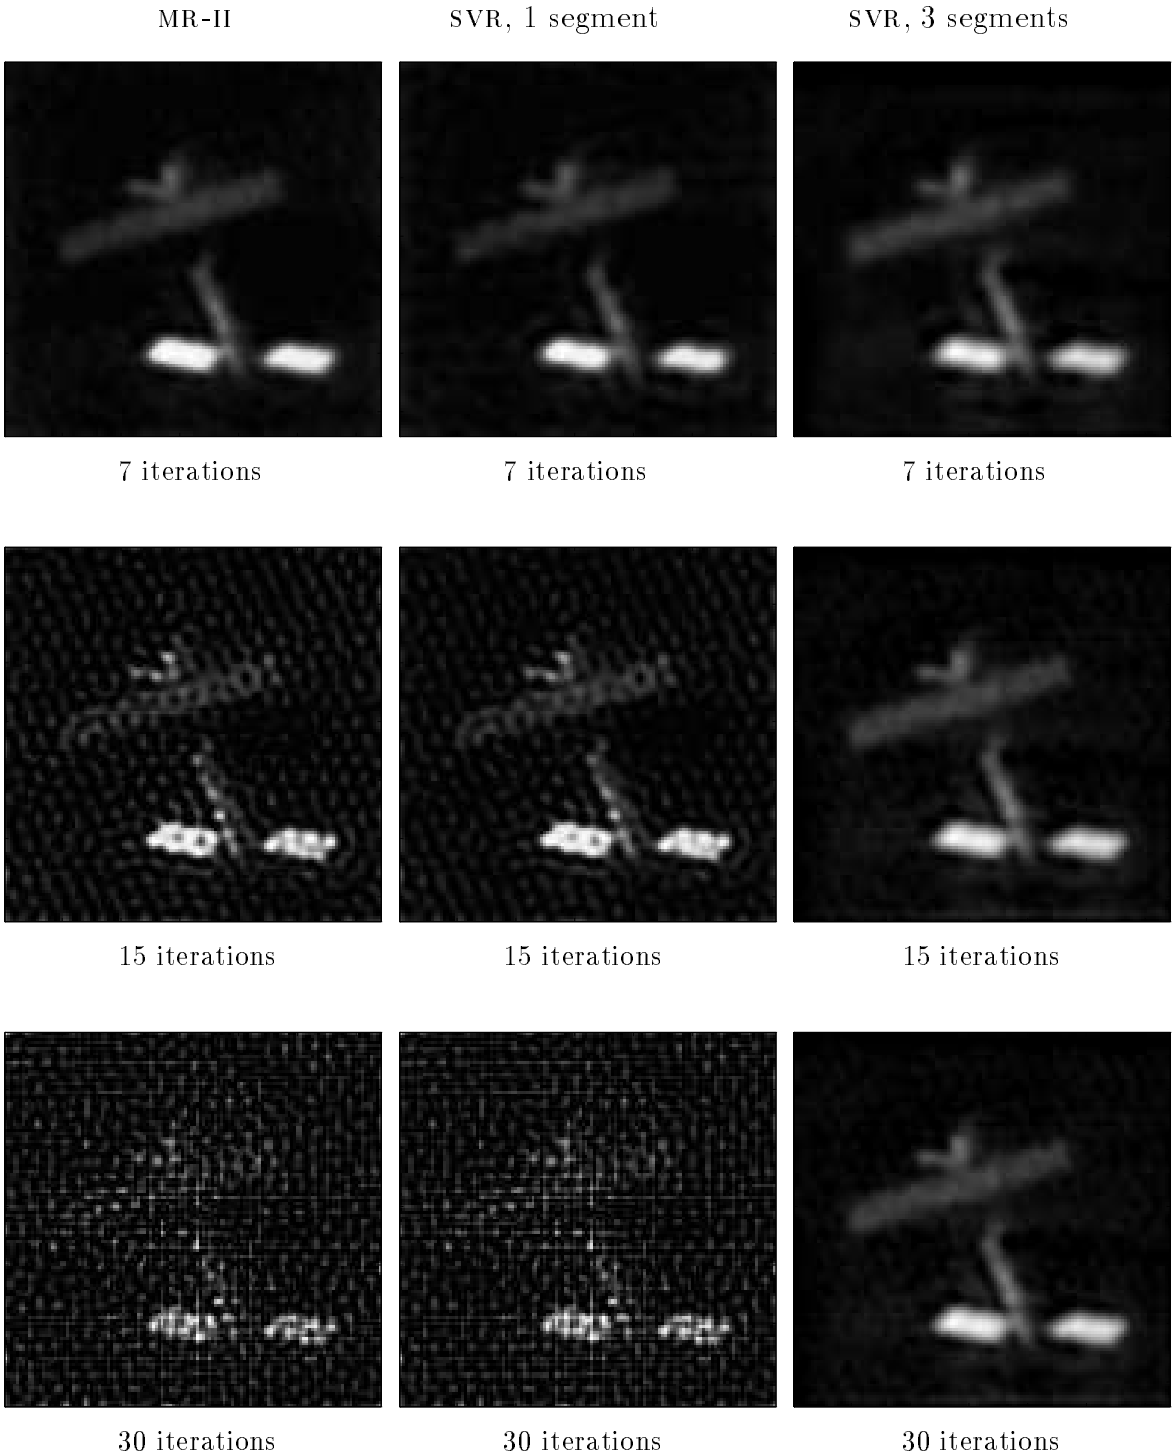

Figure 13: Computed restorations of the real satellite image.

[3] C. Avila, *An Adaptive Regularized Method for Deconvolution of Signals with Edges by Convex Projections*, IEEE Trans. on Signal Processing, 42 (1994), pp. 1849–1851.

[4] O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, Cambridge, 1994.

[5] J. Biedmond, R. Lagendijk and R. Mersereau, *Iterative Methods for Image Deblurring*, Proc. of the IEEE, 78 (1990), pp. 856–883.

[6] Å. Björck, *Numerical Methods for Least Squares Problems*, SIAM, Philadelphia, PA, 1996.

[7] J. Brown, *Overview of Topical Issues on Inverse Problems in Astronomy*, Inverse Problems, 11 (1995), pp. 635–638.

[8] R. Carreras, *Personal Correspondence*, US Air Force Phillips Laboratory, Lasers and Imaging Directorate, Kirtland Air Force Base, New Mexico.

[9] K. Castleman, *Digital Image Processing*, Prentice–Hall, NJ, 1996.

[10] R. Chan, T. Chan and C. Wong, *Cosine Transform Based Preconditioners for Total Variation Minimization Problems in Image Processing*, Iterative Methods in Linear Algebra II, V3, IMACS Series in Computational and Applied Mathematics, Proceedings of the Second IMACS International Symposium on Iterative Methods in Linear Algebra, Bulgaria, June, 1995, pp. 311–329, eds. S. Margenov and P. Vassilevski.

[11] R. Chan, J. Nagy and R. Plemmons, *FFT-based Preconditioners for Toeplitz–Block Least Squares Problems*, SIAM J. Numer. Anal., 30 (1993), pp. 1740–1768.

[12] T. Chan, G. Golub and P. Mulet, *A Nonlinear Primal–Dual Method for Total Variation Minimization Problems in Image Processing*, UCLA Computational and Applied Mathematics Technical Report CAM 95-43, September, 1995.

[13] T. F. Chan and J. A. Olkin, *Preconditioners for Toeplitz–block Systems*, Numer. Algor., 6 (1993), pp. 89–101.

[14] S. Choy, Y. K. Chan and W. C. Siu, *A New Adaptive Iterative Image Restoration Algorithm*, Proc. IEEE International Conf. on Image Processing, II (1994), pp. 670–674.

[15] N. Clinthorne, et al, *Preconditioning Methods for Improved Convergence Rates in Iterative Reconstructions*, IEEE Trans. on Medical Imaging, 12 (1993), pp. 78–83.

[16] P. J. Davis, *Circulant Matrices*, Wiley, New York, 1979.

[17] D. Dobson and F. Santosa, *Recovery of Blocky Images from Noisy and Blurred Data*, SIAM J. Appl. Math., 56 (1996), pp. 1181–1198.

[18] B. Ellerbroek, C. Van Loan, N. Pitsianis, and R. Plemmons, *Optimizing Closed–loop Adaptive-Optics Performance using Multiple Control Bandwidths*, J. Optical Soc. Amer. A, 11 (1994), pp. 2871–2886.

[19] B. Ellerbroek, C. Van Loan, N. Pitsianis, and R. Plemmons, *Multiple Control Bandwidth Computations in Adaptive-Optics*, Preprint, 1996.

[20] D. Fish, J. Grochmalicki and E. Pike, *Scanning Singular-Value-Decomposition Method for Restoration of Images with Space-Variant Blur*, J. Optical Soc. Amer. A, 13 (1996), pp. 1–6.

[21] G. Golub and C. Van Loan, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, MD, 2nd Edition, 1989.

[22] M. Hanke, *Conjugate Gradient Type Methods for Ill–posed Problems*, Pitman Research Notes in Mathematics, Longman Scientific & Technical, Harlow, Essex, 1995.

[23] M. Hanke and P. Hansen, *Regularization Methods for Large–scale Problems*, Surveys Math. Indust., 3 (1993), pp. 253–315.

[24] M. Hanke and J. Nagy, *Restoration of Atmospherically Blurred Images by Symmetric Indefinite Conjugate Gradient Techniques*, Inverse Problems, 12 (1996), pp. 157–173.

[25] M. Hanke, J. Nagy and R. Plemmons, *Preconditioned Iterative Regularization for Ill–posed Problems*, Numerical Linear Algebra and Scientific Computing, Eds. L. Reichel, A. Ruttan, R. Varga, Walter de Gruyter Press, Berlin, pp. 141–163.

[26] P. Hansen and D. P. O'Leary, *The use of the L–Curve in the Regularization of Discrete Ill–posed Problems*, SIAM J. Sci. Comp., 14 (1993), pp. 1487–1503.

[27] J. W. Hardy, *Adaptive-Optics*, Scientific American, 270, no. 6 (1994), pp. 60–65.

[28] J. Hollingsworth, K. Lui and P. Pauca, *Parallel Toolbox for Matlab*, User's Manual, Developed in the Department of Mathematics and Computer Science, Wake Forest University, 1995.

[29] R. Lagendijk and J. Biedmond, *Iterative Identification and Restoration of Images*, Kluwer Press, Boston, 1991.

[30] C. Matson and M. Roggemann, *Noise Reduction in Adaptive-Optics Imagery with the use of Support Constraints*, Applied Optics, 34 (1995), pp. 767–780.

[31] J. Nagy and D. P. O'Leary, *Restoring Images Degraded by Space Variant Blur*, Computer Science Department Report CS-TR-3426, Institute for Advanced Computer Studies Report UMIACS-95-26, University of Maryland, February 1995, to appear in SIAM J. Sci. Comput.

[32] J. Nagy, R. Plemmons and T. Torgersen, *Fast Restoration of Atmospherically Blurred Images*, in Advanced Signal Processing Algorithms, Architectures, and Implementations IV, Franklin T. Luk, Editor, 2295 (1994), pp. 542–553.

[33] J. Nagy, R. Plemmons and T. Torgersen, *Iterative Image Restoration using Approximate Inverse Preconditioning*, IEEE Trans. on Image Processing, 5 (1996), pp. 1151-1162.

[34] A. S. Nemirovskii, *The Regularizing Properties of the Adjoint Gradient Method in Ill–posed Problems*, USSR Comput. Math. Math. Phys., 26 (1986), pp. 7–16.

[35] J. M. Morel and S. Solimini, *Variational Methods in Image Segmentation*, Birkhäuser Press, Boston, 1995.

[36] S. Osher and L. Rudin, *Feature–oriented Image Enhancement Using Shock Filters*, SIAM J. Numer. Anal., 27 (1990), pp. 919–940.

[37] S. Reaves, *Optimal Space–Varying Regularization in Iterative Image Restoration*, IEEE Trans. on Image Processing, 3 (1994), pp. 319–324.

[38] M. Paige and M. A. Saunders, *Solution of Sparse Indefinite Systems of Linear Equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.

[39] L. Rudin and S. Osher, *Total Variation Based Image Restoration with Free Local Constraints*, Proc. IEEE International Conf. on Image Processing, II (1994), pp. 31–35.

[40] L. Rudin, S. Osher and E. Fatemi, *Nonlinear Total Variation Based Noise Removal Algorithms*, Physica D., 60 (1992), pp. 259–268.

[41] C. Vogel and M. Oman, *Iterative Methods for Total Variation Denoising*, SIAM J. Sci. Comput., 17 (1996), pp. 227–238.

[42] C. Vogel and M. Oman, *A Fast, Robust Algorithm for Total Variation Based Reconstruction of Noisy Blurred Images*, submitted, IEEE Trans. Image Processing, 1995.