Krylov subspace iterative methods for nonsymmetric discrete ill-posed problems in image restoration

D. Calvetti [a], B. Lewis [b] and L. Reichel [c]

[a]Department of Mathematics, Case Western Reserve University, Cleveland, OH 44106

[b]Rocketcalc, LLC, Richfield, OH 44286

[c]Department of Mathematics, Kent State University, Kent, OH 44242

ABSTRACT

The BiCG and QMR methods are well-known Krylov subspace iterative methods for the solution of linear systems of equations with a large nonsymmetric, nonsingular matrix. However, little is known of the performance of these methods when they are applied to the computation of approximate solutions of linear systems of equations with a matrix of ill-determined rank. Such linear systems are known as linear discrete ill-posed problems. We describe an application of the BiCG and QMR methods to the solution of linear discrete ill-posed problems that arise in image restoration, and compare these methods to the conjugate gradient method applied to the associated normal equations and to total variation-penalized Tikhonov regularization.

Keywords: BiCG, QMR, regularization, ill-posed problem, iterative method, image restoration

1. INTRODUCTION

The blurring of an image by spatially invariant blur can be modeled by a convolution

$$g(\xi,\eta)=\int_{\mathbf{R}^2} h(\xi-\xi',\eta-\eta')f(\xi',\eta')d\xi' d\eta', \qquad (\xi,\eta)\in\mathbf{R}^2, \tag{1}$$

where the functions f and g represent the original and blurred images, respectively. The smooth kernel h is referred to as a *point spread function*. The problem of determining the function f, given the smooth kernel h and the function g, by solving the integral equation (1) is ill-posed, because the solution f, when it exists, does not depend continuously on g. The original and blurred images, i.e., the functions f and g, are assumed to be defined in a compact set with a piecewise smooth boundary (e.g., the unit square in $\mathbf{R}^2$), and equation (1) is discretized to yield a finite system of linear equations

$$Ax=b, \qquad A\in\mathbf{R}^{m^2\times m^2}, \qquad x,b\in\mathbf{R}^{m^2}. \tag{2}$$

In the present paper, the vector x contains nonnegative pixel values, ordered row-wise, of the $m\times m$ discrete image that corresponds to f in equation (1). The matrix A is the blurring matrix obtained by discretizing the integral operator (1), and the vector b contains the pixel values, ordered row-wise, of the blurred discretized image associated with x. Spatial invariance of the blur can be exploited to give a blurring matrix A of block Toeplitz form with Toeplitz blocks. We refer to matrices with this structure as BTTB matrices. The matrix A has many singular values of different orders of magnitude close to the origin; thus, A is of ill-determined rank. Following Hansen,[1] we refer to linear systems of equations (2) with a matrix of ill-determined rank as linear discrete ill-posed problems. A recent review of many numerical methods for the solution of linear discrete ill-posed problems is provided by Hansen.[1] We remark that in many image restoration problems, a discrete point spread function is determined empirically and is often nonsymmetric, i.e., $h(\xi,\eta)\not\equiv h(\eta,\xi)$. This situation generally leads to a nonsymmetric blurring matrix A.

In many applications of image restoration, the vector b is not available. Instead, the vector

$$\tilde{b}=b+d \tag{3}$$

E-mail: D.C.: dxc57@po.cwru.edu, B.L.: blewis@rocketcalc.com, L.R.: reichel@math.kent.edu.

is known, where the entries of the vector $d \in \mathbf{R}^n$ contain measurement and transmission errors. We refer to d as "noise." Thus, the vector $\tilde{b}$ represents an available discretized image, that is contaminated by blur and noise. In the present paper, we assume that the norm of the noise

$$\delta = \|d\| \tag{4}$$

is explicitly known, but that the noise vector d is not. Throughout this paper $\|\cdot\|$ denotes the Euclidean vector norm.

The desired, but unknown, discretized image x solves the linear system of equations (2) with unknown right-hand side vector b. We seek to determine an approximation of x by computing an approximate solution of the linear system of equations

$$A\tilde{x} = \tilde{b}, \qquad \tilde{x}, \tilde{b} \in \mathbf{R}^{m^2}, \tag{5}$$

with known right-hand side (3). To simplify the discussion, we assume for the remainder of the paper that the matrix A is nonsingular. Thus, the linear system (5) has a unique solution, which is given by $\tilde{x} = A^{-1}\tilde{b}$. Note that since the matrix A has many singular values close to the origin, it is severely ill-conditioned. Due to the ill-conditioning of A and the error d in $\tilde{b}$, the solution $\tilde{x}$ of (5) typically is a poor approximation of the solution x of (2). Therefore, instead of computing the solution $\tilde{x}$ of (5), we approximate the matrix A^{-1} by a matrix $R_\alpha \in \mathbf{R}^{m^2 \times m^2}$ that is less ill-conditioned than A and compute the approximation

$$\tilde{x}_\alpha = R_\alpha \tilde{b} \tag{6}$$

of the solution x of (2). The approximation of the matrix A^{-1} by R_α is referred to as *regularization*, the matrix R_α as a regularized inverse and the parameter α as the regularization parameter. This parameter determines how well R_α approximates A^{-1}. The vector $\tilde{x}_\alpha$ defined by (6) is referred to as a regularized approximate solution of (5). We would like to choose R_α so that $\tilde{x}_\alpha$ is a meaningful approximation of the solution x of equation (2). When a bound δ of the norm of the noise d in the right-hand side $\tilde{b}$ is known, the matrix R_α is often chosen so that the regularized approximate solution $\tilde{x}_\alpha$ of (5) determined by (6) satisfies

$$\|b - A\tilde{x}_\alpha\| \approx \delta. \tag{7}$$

One of the most popular regularization methods is Tikhonov regularization, where A^{-1} is approximated by

$$R_\alpha = (A^T A + \alpha L^T L)^{-1} A^T, \qquad L \in \mathbf{R}^{m^2 \times m^2}, \tag{8}$$

cf. Groetsch.[2] The matrix L is referred to as the regularization operator. Thus, Tikhonov regularization replaces the solution of the linear system (5) by the solution of the system of equations

$$(A^T A + \alpha L^T L)\tilde{x}_\alpha = A^T \tilde{b}. \tag{9}$$

The solution $\tilde{x}_\alpha$ of (9) satisfies the minimization problem

$$\min_{x \in \mathbf{R}^{m^2}} \{\|\tilde{b} - Ax\|^2 + \alpha \|Lx\|^2\}. \tag{10}$$

The term $\alpha\|Lx\|^2$ in (10) penalizes solutions whose image under L have large norm. In image restoration problems involving images with "sharp" edges, L is popularly chosen so that $\|Lx\|^2$ approximates a total variation (TV) functional of x; see, e.g., Vogel and Oman.[3] This regularization operator depends on x and makes (10) a nonlinear minization problem, which we refer to as TV-penalized Tikhonov regularization.

The computation of the regularized solution $\tilde{x}_\alpha$ that solves the TV-penalized minimization problem (10) can be expensive. It is the purpose of the present paper to compare this approach to image restoration with regularization by the application of a few steps of Krylov subspace iterative methods applied directly to the solution of (5), a computationally less demanding approach to regularization. In particular, we consider application of the bi-conjugate gradient (BiCG) and the quasi-minimal residual (QMR) methods. These are well-known Krylov subspace iterative methods for the solution of linear systems of equations with a nonsingular and nonsymmetric matrix; see, e.g., Freund and Nachtigal[4] and Saad.[5] However, the application of the BiCG and QMR methods to the computation of regularized approximate solutions of nonsymmetric linear discrete ill-posed problems (5) has so far not been considered in the literature.

When applying a few, say j, steps of the BiCG or QMR methods to the linear system (5), the regularized inverse R_α is defined implicitly. The number of steps plays the role of the regularization parameter α in (6). In order to emphasize the dependence on j, we denote the regularized inverse defined by taking j steps of a Krylov subspace iterative method by R_j, and the computed solution of (6) with initial approximate solution $\tilde{x}_0 = 0$ by $\tilde{x}_j$, i.e.,

$$\tilde{x}_j = R_j \tilde{b}. \tag{11}$$

Thus, the matrix R_j depends on the iterative method, the number of steps j and the vector $\tilde{b}$.

The conjugate gradient (CG) method is another widely used Krylov subspace iterative method. The CG method is an attractive iterative method for the solution of linear systems of equations with a symmetric positive definite matrix; it also can be applied to the solution of consistent linear systems of equations with a symmetric positive semidefinite matrix. Recently, the computation of an approximate solution of (5) by carrying out a few iterations with the CG method applied to the normal equations associated with (5),

$$A^T A\tilde{x} = A^T \tilde{b}, \tag{12}$$

has received considerable attention. An analysis of this method is given by Hanke[6] and successful applications to the restoration of images degraded by noise and spatially variant blur are reported by Nagy and O'Leary[7,8]; see also Calvetti et al.,[9] Hanke and Nagy,[10] and Hanke and Hansen[11] for related discussions.

Since the singular values of the matrix $A^T A$ are the square of the singular values of A, it follows that if A has many singular values of different orders of magnitude close to the origin, then so does $A^T A$. Hence, the matrix $A^T A$ is of ill-determined rank when A is, and the linear system of equations (12) is a linear discrete ill-posed problem whenever the linear system (5) is.

The QMR and BiCG methods with initial approximate solution $\tilde{x}_0 = 0$ determine approximate solutions to the linear system (5) in the Krylov subspaces

$$\mathcal{K}_j(A, \tilde{b}) = \text{span}\{\tilde{b}, A\tilde{b}, A^2\tilde{b}, \ldots, A^{j-1}\tilde{b}\}, \qquad j = 1, 2, \ldots\ . \tag{13}$$

The CG method applied to the normal equations (12) with the same initial approximate solution computes approximate solutions of (5) that lie in the Krylov subspaces $\mathcal{K}_j(A^T A, A^T\tilde{b})$, $j = 1, 2, \ldots$; thus any approximate solution to (5) computed by the CG method applied to the normal equations lies in the range of A^T. We are interested in image restoration problems in which the blur- and noise-free image contains "sharp" edges; that is, x represents the discretization of a non-smooth function f. It has recently been demonstrated that iterative methods that do not require that iterates lie in the ranges of A or A^T can produce approximate solutions of better quality than methods that restrict iterates to the ranges of A or A^T in this situation; see Calvetti et al..[9]

We compare the performance of the QMR and BiCG methods with that of TV-penalized Tikhonov regularization for restoration of an image with "sharp" edges that has been degraded by noise and blur. We also examine the performance of the CG method applied to the iterative solution of the normal equations (12). We use the implementation CGLS by Björck[12] which does not require the symmetric positive definite (or semidefinite) matrix $A^T A$ to be formed. We terminate the iterations with each iterative method as soon as an approximate solution $\tilde{x}_j$ has been determined, such that the associated residual error $A\tilde{x}_j - \tilde{b}$ is of norm smaller than or equal to δ, where δ is the norm of the noise in the right-hand side vector $\tilde{b}$; cf. (3) and (4). Thus, we terminate the computations as soon as an approximate solution $\tilde{x}_j$ has been computed, such that

$$\|A\tilde{x}_j - \tilde{b}\| \leq \delta, \qquad \|A\tilde{x}_{j-1} - \tilde{b}\| > \delta. \tag{14}$$

The residual error $A\tilde{x}_j - \tilde{b}$ is sometimes referred to as the discrepancy. The stopping criterion (14) is based on the observation that in general it is not meaningful to determine an iterate $\tilde{x}_j$ such that $\|A\tilde{x}_j - \tilde{b}\|$ is much smaller than the norm of the noise in the right-hand side $\tilde{b}$. This observation is often referred to as the discrepancy principle; see Hansen[1] or Morozov[13] for discussions. An analysis of the regularizing property of the BiCG method terminated by the stopping criterion (14) is provided by Calvetti et al..[14]

We remark that iterates $\tilde{x}_j$ determined by the BiCG method can exhibit erratic convergence. Indeed, instability of the BiCG method provided incentive for the development of the QMR method; see Freund and Nachtigal.[4] Poor convergence can sometimes be improved by the use of a preconditioner. In fact, it may be possible to reduce the number of iterations required by all of the methods considered in this paper by preconditioning. However, the determination of suitable preconditioners for linear discrete ill-posed problems is generally not straightforward. In order to keep the comparison simple, we do not precondition the linear systems of equations in the computed example reported in Section 3.

2. THE BICG AND QMR METHODS

The BiCG and QMR methods are derived from the two-sided Lanczos process, stated as Algorithm 2.1 below. We use this algorithm to define the regularized inverses produced by the BiCG and QMR methods, cf. (6). For implementations of the BiCG and QMR methods, further discussion of their properties and derivation, and references; see, e.g., Freund and Nachtigal[4] and Saad.[5] We let e_j denote the jth column of the identity matrix of appropriate order.

ALGORITHM 2.1. *The two-sided Lanczos process*

> *Input:* A, $\tilde{b}$, $\ell \in \mathbf{N}$;
> *Output:* $k \leq \ell$, *nontrivial entries of tridiagonal matrix* $T_k = [t_{ij}]_{i,j=1}^{k} \in \mathbf{R}^{k\times k}$, $t_{k+1,k} \in \mathbf{R}$,
> $\{u_j\}_{j=1}^{k+1}$, $\{w_j\}_{j=1}^{k+1}$, $\{\theta_j\}_{j=1}^{k+1}$;
>
> $k := 1$; $u_0 := 0$; $w_0 := 0$; $\hat{u}_1 := \tilde{b}$; $\hat{w}_1 := \tilde{b}$; $\theta_0 := 1$;
> $\alpha_1 := \|\hat{u}_1\|$; $\beta_1 := \|\hat{w}_1\|$; $u_1 := \hat{u}_1/\alpha$; $w_1 := \hat{w}_1/\beta$; $\theta_1 := w_1^T u_1$;
> *while* $k \leq \ell$ *do*
>
> $\quad t_{k,k} := w_k^T A u_k/\theta_k$;
> $\quad t_{k-1,k} := \beta_k\theta_k/\theta_{k-1}$;
> $\quad \hat{u}_{k+1} := Au_k - t_{k,k}u_k - t_{k-1,k}u_{k-1}$;
> $\quad \hat{w}_{k+1} := A^T w_k - t_{k,k}w_k - \frac{t_{k-1,k}\alpha_k}{\beta_k}w_{k-1}$;
> $\quad \alpha_{k+1} := \|\hat{u}_{k+1}^{\delta}\|$; $\beta_{k+1} := \|\hat{w}_{k+1}^{\delta}\|$;
> $\quad t_{k+1,k} := \alpha_{k+1}$;
> $\quad$ *if* $\alpha_{k+1} = 0$ *or* $\beta_{k+1} = 0$ *then*
>
> $\qquad u_{k+1} := 0$; $w_{k+1} := 0$; $\theta_{k+1} = 0$; *exit;*
>
> $\quad$ *endif*
> $\quad u_{k+1} := \hat{u}_{k+1}/\alpha_{k+1}$; $w_{k+1} := \hat{w}_{k+1}/\beta_{k+1}$;
> $\quad \theta_{k+1} := w_{k+1}^T u_{k+1}$;
> $\quad$ *if* $\theta_{k+1} = 0$ *then exit;*
> $\quad k = k + 1$;
>
> *end while*

Algorithm 2.1 may break down because u_k or w_k vanish, or because $w_k^T u_k$ vanishes even if u_k and w_k do not. Assume that Algorithm 2.1 applied to $\{A, \tilde{b}\}$ does not break down before step ℓ and that the resulting matrix T_ℓ is nonsingular. Then the sets of vectors $\{u_k\}_{k=1}^{\ell}$ and $\{w_k\}_{k=1}^{\ell}$ are biorthogonal bases of the Krylov subspaces $\mathcal{K}_\ell(A, \tilde{b})$ and $\mathcal{K}_\ell(A^T, \tilde{b})$, respectively; they satisfy

$$w_k^T u_j = \begin{cases} 0, & j \neq k, \\ \theta_k, & j = k, \end{cases} \tag{15}$$

where the θ_k, $1 \leq k \leq \ell$, are nonvanishing.

Introduce the $m^2 \times \ell$ matrices $U_\ell = [u_1, \ldots, u_\ell]$ and $W_\ell = [w_1, \ldots, w_\ell]$, and the $m^2 \times (\ell+1)$ matrices $U_{\ell+1} = [u_1, \ldots, u_{\ell+1}]$ and $W_{\ell+1} = [w_1, \ldots, w_{\ell+1}]$. Let D_ℓ be the $\ell \times \ell$ diagonal matrix with diagonal entries given by $\{\theta_1, \theta_2, \ldots, \theta_\ell\}$ and define $D_{\ell+1}$ analogously. The recursion relations of Algorithm 2.1 can be written in the form

$$AU_\ell = U_\ell T_\ell + \hat{u}_{\ell+1} e_\ell^T = U_{\ell+1} \bar{T}_\ell, \tag{16}$$

where $T_\ell \in \mathbf{R}^{\ell \times \ell}$ is the tridiagonal matrix determined by Algorithm 2.1 and $\bar{T}_\ell$ is the $(\ell+1) \times \ell$ matrix obtained by appending the row $t_{\ell+1,\ell} e_\ell^T$ to the matrix T_ℓ. It follows from (15) and (16) that

$$W_\ell^T A U_\ell = D_\ell T_\ell. \tag{17}$$

The iterate $\tilde{x}_\ell$ determined by the BiCG method applied to (5) with initial approximate solution $\tilde{x}_0 = 0$ is characterized by the Petrov-Galerkin condition that the associated residual $\tilde{b} - A\tilde{x}_\ell$ be orthogonal to $\mathcal{K}_\ell(A^T, \tilde{b})$. Under the assumptions stated above on the sets $\{u_j\}_{j=1}^\ell$ and $\{w_j\}_{j=1}^\ell$ and on the matrix T_ℓ determined by Algorithm 2.1, the Petrov-Galerkin condition yields the following expressions for the regularized inverse and regularized approximate solution determined by the BiCG method,

$$R_\ell = U_\ell T_\ell^{-1} D_\ell^{-1} W_\ell^T, \qquad \tilde{x}_\ell = U_\ell T_\ell^{-1} \|\tilde{b}\| e_1. \tag{18}$$

We now turn to the QMR method. Let $x = U_\ell y$, where $y \in \mathbf{R}^\ell$. Then $x \in \mathcal{K}_\ell(A, \tilde{b})$ and it follows from the definitions of $U_\ell, U_{\ell+1}$ and $\bar{T}_\ell$ that

$$\tilde{b} - Ax = \tilde{b} - AU_\ell y = U_{\ell+1}(\|\tilde{b}\| e_1 - \bar{T}_\ell y). \tag{19}$$

The iterate $\tilde{x}_\ell$ determined by the QMR method applied to (5) with initial approximate solution $\tilde{x}_0 = 0$ satisfies $\tilde{x}_\ell = U_\ell y_\ell$, where y_ℓ solves the minimization problem

$$\min_{y \in \mathbf{R}^\ell} \left\| \|\tilde{b}\| e_1 - \bar{T}_\ell y \right\|. \tag{20}$$

We remark that if the basis $\{u_1, u_2, \ldots, u_{\ell+1}\}$ produced by Algorithm 2.1 is orthonormal, or if $\{u_1, u_2, \ldots, u_\ell\}$ forms an orthonormal set and $u_{\ell+1} = 0$, then it follows from (19) and (20) that $\tilde{x}_\ell$ solves the minimization problem

$$\min_{x \in \mathcal{K}_\ell(A, \tilde{b})} \|\tilde{b} - Ax\|, \tag{21}$$

a desirable property. However, the basis $\{u_1, u_2, \ldots, u_{\ell+1}\}$ generally is not orthonormal, and therefore the iterate $\tilde{x}_\ell$ generally does not satisfy (21).

The minimization (20) gives the following expressions for the regularized inverse and regularized approximate solution defined by the QMR method,

$$R_\ell = U_\ell \bar{T}_\ell^\dagger D_{\ell+1}^\dagger W_{\ell+1}^T, \qquad \tilde{x}_\ell = U_\ell \bar{T}_\ell^\dagger \|\tilde{b}\| e_1, \tag{22}$$

where $\bar{T}_\ell^\dagger$ denotes the Moore-Penrose pseudoinverse of the matrix $\bar{T}_\ell$.

We remark that breakdown of Algorithm 2.1 before the BiCG or QMR methods determine a suitable regularized approximate solution to (5) can in general be circumvented by modifying the recursion relations for the u_k and w_k in Algorithm 2.1 by utilizing look-ahead; see, e.g., Freund and Nachtigal.[4]

Algorithm 2.1 and the BiCG and QMR methods require two matrix-vector product evaluations per iteration; one with the matrix A and one with A^T. When A is a BTTB matrix, these matrix-vector products can be computed efficiently by using 2D discrete fast Fourier transforms; see, e.g., Jain.[15]

3. NUMERICAL EXPERIMENTS

We consider the restoration of a noisy and blurred image. The noise- and blur-free discrete model image is from the Regularization Tools package by Hansen[16] and is represented by 256×256 pixels. Thus, $m = 256$ in (2) and (5). This model image displays a few simple shapes that exhibit "sharp" edges. The noisy and blurred image to be restored is constructed as follows. Let the entries of the vector x contain the pixels of the noise- and blur-free model image ordered row-wise and let A represent the blurring matrix. Then the vector $b = Ax$ represents the blurred, but noise-free, image associated with the model image represented by x. We construct the noise vector d by first determining a vector in $\mathbf{R}^{m^2}$ with uniformly distributed random entries in the interval $[0, 1]$ and then scale this vector to have norm $0.05\|Ax\|$. The vector so obtained is the noise vector d. We have $\delta = \|d\| = 11.1$. The nonsymmetric blurring matrix A is constructed from an empirically determined discrete point spread function that models blur induced by the original imaging optics of the Hubble space telescope. The point spread function can be obtained by ftp from the address *ftp.stsci.edu* in the directory */pub/software/stsdas/testdata/restore/sims/qso*. The model image and the blurred and noisy image are displayed in Figure 3. The blurring matrix obtained is a BTTB matrix.

We compare the BiCG and QMR methods applied to the linear system (5) with TV-penalized Tikhonov regularization (9) and with the CGLS method applied to the normal equations (12). We used the initial approximate solution $\tilde{x}_0 = 0$ for the QMR, BiCG, and CGLS methods.

All computations were carried out on an Intel Pentium workstation with about 16 significant decimal digits. Computations for the TV-penalized Tikhonov regularization used the Image Deblurring package for Matlab by Hamilton and Vogel, available on the web at *http://www.math.montana.edu/˜vogel/Software/deconv*. This package uses a quasi-Newton method to solve the minimization problem (10) for fixed α. Each step of the quasi-Newton method requires the solution of a sparse, symmetric positive definite linear system of equations with m^2 unknowns. The CG method without preconditioning is used to solve this linear system. Computations for the BiCG, QMR and CGLS methods were carried out with Octave 2.0.16. We used the program fv, see Blackburn,[17] to convert the point spread function stored in fits format to floating-point arrays for input to Matlab and Octave. The program fv is part of the ftools package available at *http://heasarc.gsfc.nasa.gov/ftools*. We converted the floating-point representations of the images to PostScript with xv.

The iterations with the BiCG, QMR and CGLS methods were terminated as soon as an iterate that satisfied the stopping criterion (14) was computed. The regularization parameter α in (10) was chosen manually to satisfy (7). Since the model image x is available, we can evaluate the relative error norm $\|x - \tilde{x}_j\|/\|x\|$ for the iterates determined by the BiCG, QMR and CGLS methods, and for the relative error norm $\|x - \tilde{x}_\alpha\|/\|x\|$ for the TV-penalized Tikhonov regularization method. We compare the number of matrix-vector products with the matrices A or A^T required for each method. The evaluation of these matrix-vector products constitutes the bulk of the floating-point operations performed by each of the methods, and therefore the count of matrix-vector product evaluations provides a good indication of the total amount of arithmetic work required.

Figure 1 shows the restored images determined by the methods. The BiCG method satisfied the stopping criterion (14) after 3 steps, as did the QMR method. The CGLS method satisfied the discrepancy principle after 6 steps. The TV-penalized Tikhonov regularization method required three quasi-Newton iterations with a total of 140 steps of the CG method, which requires about 280 matrix-vector product evaluations; each step of the CG method requires the evaluation of one matrix-vector product with the matrix A and one with A^T. Table 1 summarizes key numerical quantities. The restored image produced by the BiCG and QMR methods terminated by the discrepancy principle are noticeably "sharper" than those produced by the other methods.

We remark that for the BiCG and QMR methods, the regularized approximate solutions obtained when the methods are terminated by the discrepancy principle (14) are also the iterates that exhibit the lowest relative error norms. This is not the case for the CGLS method. Figure 2 displays the restored image produced by the CGLS method that exhibits the lowest relative error norm. This image is obtained after 15 iterations.

Similarly, for TV-penalized Tikhonov regularization, the choice of regularization parameter α that satisfies (7) does not produce the "best" image as measured by its relative error norm. A nearly "best" restored image determined by TV-penalized Tikhonov regularization is shown in Figure 2. Its computation required a total of 99 steps of the CG method, i.e., the evaluation of about 198 matrix-vector products with the matrices A or A^T. Table 1 shows associated numerical quantities. Both the restored images in Figure 2 are of markedly higher quality than the restored images determined by the same methods with a stopping rule based on the discrepancy principle. However, the visual quality of the images in Figure 2 in terms of "sharpness" is similar to the restored images produced by the BiCG and QMR methods displayed in Figure 1. As shown by Table 1, the images determined by the BiCG and QMR methods were by far the least expensive to compute.

We note that the BiCG, QMR, CGLS iterative methods and the TV-penalized Tikhonov regularization method were found to perform analogously as for the example shown in this section for a large number of image restoration problems for images with "sharp" edges, i.e., the BiCG and QMR methods determined restored images of higher or the same quality than the CGLS method and TV-penalized Tikhonov regularization with less computational work.

We finally remark that a related investigation of the performance of another popular Krylov subspace iterative method, the GMRES method, when applied to the solution of linear nonsymmetric discrete ill-posed problems, is reported by Calvetti et al..[18,19] The iterates $\tilde{x}_\ell$ computed by the GMRES method satisfy the minimization problem (21) for $\ell = 1, 2, \ldots$. A comparison with the CGLS method showed the GMRES method to be competitive; see Calvetti et al.[18,19] for details.

	BiCG	QMR	CGLS		TV	
Number of iterations	3	3	6	15	-	-
Matrix-vector products	6	6	12	30	280	198
Residual norm	$9.62 \cdot 10^0$	$8.25 \cdot 10^0$	$9.23 \cdot 10^0$	$5.25 \cdot 10^0$	$1.11 \cdot 10^1$	$5.44 \cdot 10^0$
Relative error norm	$2.25 \cdot 10^{-1}$	$2.28 \cdot 10^{-1}$	$2.48 \cdot 10^{-1}$	$1.89 \cdot 10^{-1}$	$2.96 \cdot 10^{-1}$	$2.18 \cdot 10^{-1}$

Table 1. Summary of numerical results: Number of iterations is counted only for the BiCG, QMR and CGLS methods. Matrix-vector products shows the total number of matrix-vector product evaluations with the matrices A and A^T. The residual norm is defined as $\|\tilde{b} - A\tilde{x}_\alpha\|$, where $\tilde{x}_\alpha$ is the regularized approximate solution produced by the respective method. The relative error norm is defined by $\|x - \tilde{x}_\alpha\|/\|x\|$. Stopping criterion (14) is used for the BiCG, QMR and CGLS methods. Results for the latter are reported in left-hand side column for the CGLS method. The right-hand side column for the CGLS method reports the case when the iterations are terminated when the relative error norm is minimal. The regularization parameter α for TV-penalized Tikhonov regularization is chosen so that either (7) is satisfied (reported in the left-hand side column) or so that the relative error norm is minimized (reported in the right-hand side column).

4. ACKNOWLEDGEMENTS

This work was supported in part by NSF grants DMS-9806702 and DMS-9806413.

REFERENCES

1. P. C. Hansen, Rank-Deficient and Discrete Ill-Posed Problems, SIAM, Philadelphia, 1998.
2. C. W. Groetsch, The Theory of Tikhonov Regularization for Fredholm Equations of the First Kind, Pitman, Boston, 1984.
3. C. R. Vogel and M. E. Oman, Fast, robust total variation-based reconstruction of noisy, blurred images, IEEE Trans. Image Processing, 7 (1998), pp. 813–824.
4. R. W. Freund and N. M. Nachtigal, QMR: a quasi-minimal residual method for non-Hermitian linear systems, Numer. Math., 60 (1991), pp. 315–339.
5. Y. Saad, Iterative Methods for Sparse Linear Systems, PWS, Boston, 1996.
6. M. Hanke, Conjugate Gradient Type Methods for Ill-Posed Problems, Longman, Harlow, 1995.
7. J. G. Nagy and D. P. O'Leary, Restoring images degraded by spatially-variant blur, SIAM J. Sci. Comput., 19 (1998), pp. 1063–1082.
8. J. G. Nagy and D. P. O'Leary, Fast iterative image restoration with a space-varying PSF, in Advanced Signal Processing Algorithms, Architectures, and Implementations IV, ed. F. T. Luk, Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE), vol. 3162, The International Society for Optical Engineering, Bellingham, 1997, pp. 388–399.
9. D. Calvetti, B. Lewis and L. Reichel, On the choice of subspace for iterative methods for linear discrete ill-posed problems, Int. J. Appl. Math. Comput. Sci., to appear.

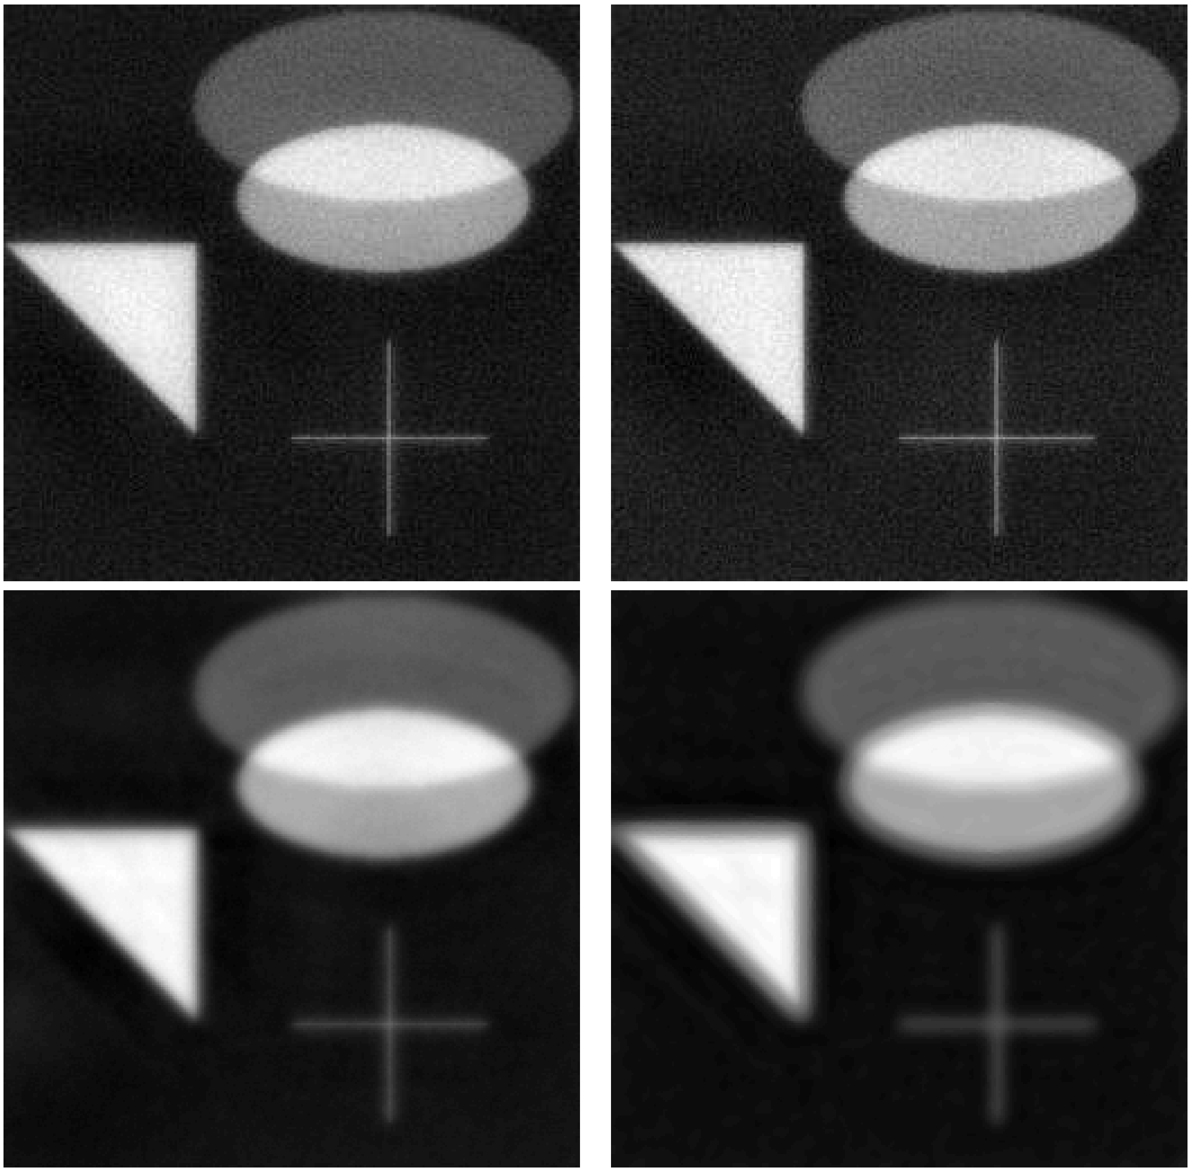

Figure 1. Restored image produced by the QMR method after 3 steps (upper left), restored image produced by the BiCG method after 3 steps (upper right), restored image produced by the CGLS method after 6 steps (lower left), restored image produced by the TV-penalized Tikhonov regularization (lower right).

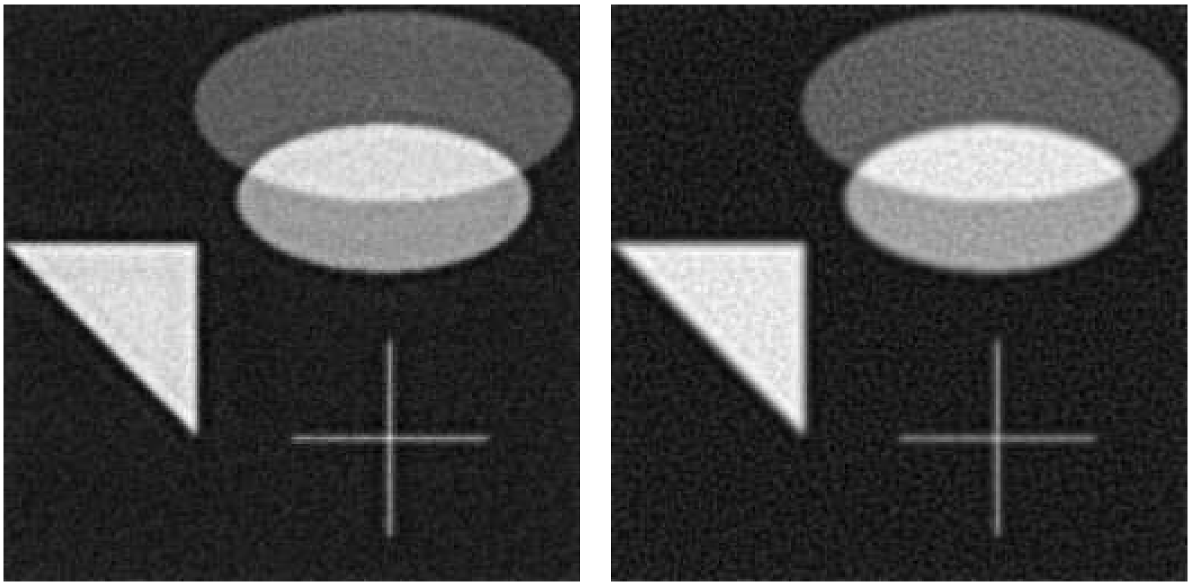

Figure 2. The "best" restored image produced by the CGLS method, as measured by the relative error $\|x - \tilde{x}_j^{CGLS}\|/\|x\|$, after $j = 15$ steps (left), and the "best" restored image produced by the TV-penalized Tikhonov regularization (right).

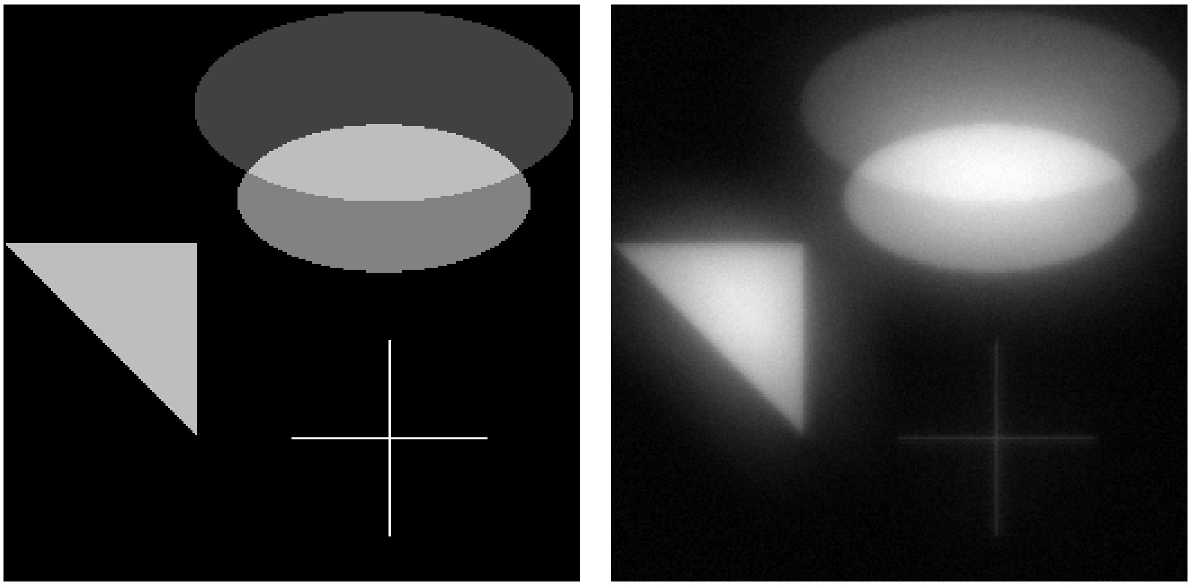

Figure 3. Blur- and noise-free image (left), and the blurred and noisy test image (right)

10. M. Hanke and J. G. Nagy, Restoration of atmospherically blurred images by symmetric indefinite conjugate gradient techniques, Inverse Problems, 12 (1996), pp. 157–173.
11. M. Hanke and P. C. Hansen, Regularization methods for large-scale problems, Surv. Math. Ind., 3 (1993), pp. 253–315.
12. Å. Björck, Numerical Methods for Least Squares Problems, SIAM, Philadelphia, 1996.
13. V. A. Morozov, On the solution of functional equations by the method of regularization, Soviet Math. Dokl., 7 (1966), pp. 414–417.
14. D. Calvetti, B. Lewis and L. Reichel, Regularizing properties of the BiCG method, manuscript.
15. A. K. Jain, Fundamentals of Digital Image Processing, Prentice-Hall, Englewood Cliffs, 1989.
16. P. C. Hansen, Regularization tools: A Matlab package for analysis and solution of discrete ill-posed problems, Numer. Algorithms, 6 (1994), pp. 1–35.
17. J. K. Blackburn, in Astronomical Data Analysis Software and Systems IV, ed. R. A. Shaw, H. E. Payne, and J. J. E. Hayes, ASP Conf. Ser., vol. 77, San Francisco, 1995, p. 367.
18. D. Calvetti, B. Lewis and L. Reichel, Restoration of images with spatially variant blur by the GMRES method, in Advanced Signal Processing Algorithms, Architectures, and Implementations X, ed. F. T. Luk, Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE), vol. 4116, The International Society for Optical Engineering, Bellingham, 2000, pp. 364–374.
19. D. Calvetti, B. Lewis and L. Reichel, On the regularizing properties of the GMRES method, Numer. Math., to appear.