Received: Added at production | Revised: Added at production | Accepted: Added at production

DOI: xxx/xxxx

ARTICLE TYPE

On the block Lanczos and block Golub–Kahan reduction methods applied to discrete ill-posed problems[†]

A. Alqahtani[1,2] | S. Gazzola[3] | L. Reichel[2] | G. Rodriguez*[4]

[1]Department of Mathematics, King Khalid University, Saudi Arabia
[2]Department of Mathematical Sciences, Kent State University, OH, USA
[3]Department of Mathematical Sciences, University of Bath, United Kingdom
[4]Department of Mathematics and Computer Science, University of Cagliari, Italy

Correspondence
*Giuseppe Rodriguez, via Ospedale 72, 09124 Cagliari, Italy. Email: rodriguez@unica.it

Present Address
Present address

Abstract

The reduction of a large-scale symmetric linear discrete ill-posed problem with multiple right-hand sides to a smaller problem with a symmetric block tridiagonal matrix can easily be carried out by the application of a small number of steps of the symmetric block Lanczos method. We show that the subdiagonal blocks of the reduced problem converge to zero fairly rapidly with increasing block number. This quick convergence indicates that there is little advantage in expressing the solutions of discrete ill-posed problems in terms of eigenvectors of the coefficient matrix when compared to using a basis of block Lanczos vectors, which are simpler and cheaper to compute. Similarly, for nonsymmetric linear discrete ill-posed problems with multiple right-hand sides, we show that the solution subspace defined by a few steps of the block Golub–Kahan bidiagonalization method usually can be applied instead of the solution subspace determined by the singular value decomposition of the coefficient matrix without significant, if any, reduction of the quality of the computed solution.

KEYWORDS:
large-scale discrete ill-posed problem; symmetric Lanczos block tridiagonalization; Golub–Kahan block bidiagonalization; Tikhonov regularization.

1 | INTRODUCTION

Consider the minimization problem

$$\min_{X \in \mathbb{R}^{n\times p}} \|AX - B\|_F. \tag{1}$$

with a large matrix $A \in \mathbb{R}^{\ell\times n}$, whose singular values gradually approach zero without significant gap. Thus, A is very ill-conditioned and may be rank deficient. The data matrix $B \in \mathbb{R}^{\ell\times p}$ with $1 < p \ll \ell$ is a "block vector" with many more rows than columns. The Frobenius norm $\|M\|_F$ of a matrix M is defined as follows. For two matrices $M_1, M_2 \in \mathbb{R}^{n\times p}$, we introduce the inner product

$$\langle M_1, M_2\rangle_F = \mathrm{trace}(M_1^T M_2),$$

where the superscript ^{T} denotes transposition and $\mathrm{trace}(\cdot)$ stands for the trace of a square matrix. Then

$$\|M\|_F = \sqrt{\langle M, M\rangle_F}\,.$$

The usual inner product of elements $u, v \in \mathbb{R}^n$ is denoted by $\langle u, v\rangle_2 = u^T v$ and the Euclidean norm by $\|u\|_2 = \sqrt{\langle u, u\rangle}$. Finally, in the following, $\mathcal{R}(M)$ stands for the range of the matrix M.

[†]This is an example for title footnote.

Minimization problems like the one appearing in equation (1) with a matrix with the properties described are commonly referred to as discrete ill-posed problems; see, e.g.,[11] and the references therein. They arise, for instance, from the discretization of linear ill-posed problems, such as Fredholm integral equations of the first kind. Applications include color and hyperspectral image restoration; see, e.g.,[2, 13].

In discrete ill-posed problems of the form (1) that arise in applications in science and engineering, the matrix B typically represents measured data that are contaminated by an error $E \in \mathbb{R}^{\ell \times p}$. Thus,

$$B = B_{\text{true}} + E, \tag{2}$$

where $B_{\text{true}} \in \mathbb{R}^{\ell \times p}$ represents the (unknown) noise-free block vector associated with B. We would ideally like to compute an approximation of the solution $X_{\text{true}} \in \mathbb{R}^{n \times p}$ of minimal Frobenius norm of the minimization problem

$$\min_{X \in \mathbb{R}^{n \times p}} \|AX - B_{\text{true}}\|_F.$$

Let $A^\dagger$ denote the Moore–Penrose pseudoinverse of the matrix A. Then, $X_{\text{true}} = A^\dagger B_{\text{true}}$. Note that the solution of (1), given by

$$X := A^\dagger B = A^\dagger (B_{\text{true}} + E) = X_{\text{true}} + A^\dagger E,$$

is not a useful approximation of X_{true} because, generally, $\|A^\dagger E\|_F \gg \|X_{\text{true}}\|_F$ due to the presence of tiny positive singular values of A.

The computation of a meaningful approximation of X_{true} from (1) requires that the system be *regularized* before solution, i.e., the system (1) has to be modified so that its solution is less sensitive to the error E in B than the solution of (1). We regularize the system (1) in two steps: first A is projected to a generally fairly small block tridiagonal or block bidiagonal matrix by application of a few iterations of the block Lanczos tridiagonalization (BLT) algorithm to A (when A is symmetric) or of the block Golub–Kahan bidiagonalization (BGKB) algorithm (when A is non-symmetric), respectively; then the reduced problem so obtained is solved by Tikhonov regularization. Discussions of these block algorithms for discrete inverse problems of the form (1) can be found in[1] and[10, Section 10.3.6] (for BLT), and[2] (for BGKB), as well as in Section 2. Also, recent advances in understanding the convergence behavior of block Krylov methods based on the Arnoldi algorithm can be found in[16], where ways of constructing matrices and right-hand sides producing any admissible convergence behavior are presented. The point of view adopted in this paper is fundamentally different, as the derivations presented here are targeted at problems of the kind (1). Indeed, it is the purpose of this paper to discuss the structure and properties of the block tridiagonal and block bidiagonal matrices determined by the BLT or BGKB algorithms, respectively, and to show the performance of Tikhonov regularization used jointly with these decompositions.

This paper is organized as follows. Section 2 reviews some background material, namely: first, summaries are given about the BLT method for symmetric matrices A and the BGKB algorithm for non-symmetric, possibly rectangular, matrices A; then, a description is added about how Tikhonov regularization can be applied to solve the reduced problems obtained by such algorithms. Section 3 presents new theoretical bounds for the diagonal and subdiagonal BLT and BGKB blocks when A stems from the discretization of a linear ill-posed problem. A few computed examples are presented in Section 4. Finally, Section 5 contains concluding remarks.

2 | BLOCK ALGORITHMS AND TIKHONOV REGULARIZATION

Summaries of the BLT and BGKB algorithms are given in Sections 2.1 and 2.2, respectively. A solution method based on Tikhonov regularization applied to the projected problems associated to BLT and BGKB is described in Section 2.3.

2.1 | Block Lanczos tridiagonalization (BLT)

Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix and let $B = X_1 S_1$ be a compact QR factorization of $B \in \mathbb{R}^{n \times p}$ with $1 \leq p \ll n$, where $X_1 \in \mathbb{R}^{n \times p}$ has orthonormal columns and $S_1 \in \mathbb{R}^{p \times p}$ is upper triangular. Then, application of $m \ll n/p$ steps of the block Lanczos method to A with initial block vector X_1 yields a decomposition of the form

$$AQ_m = Q_{m+1} T_{m+1,m}, \tag{3}$$

where the block columns of $Q_m = [X_1, \dots, X_m] \in \mathbb{R}^{n\times pm}$ and $Q_{m+1} = [Q_m, X_{m+1}] \in \mathbb{R}^{n\times p(m+1)}$ are such that $X_i \in \mathbb{R}^{n\times p}$ and

$$X_i^T X_j = \begin{cases} I_p, & i = j, \\ O_p, & i \neq j, \end{cases}$$

for $i, j = 1, 2, \dots, m+1$. Here $I_p \in \mathbb{R}^{p\times p}$ denotes the identity matrix and $O_p \in \mathbb{R}^{p\times p}$ the zero matrix. Moreover, the matrix

$$T_{m+1,m} = \begin{bmatrix} M_1 & S_2^T & & & & \\ S_2 & M_2 & S_3^T & & & \\ & S_3 & M_3 & \ddots & & \\ & & \ddots & \ddots & S_{m-1}^T & \\ & & & S_{m-1} & M_{m-1} & S_m^T \\ & & & & S_m & M_m \\ & & & & & S_{m+1} \end{bmatrix} \in \mathbb{R}^{p(m+1)\times pm}, \tag{4}$$

is block tridiagonal with a leading symmetric $pm \times pm$ submatrix, which we denote by $T_{m,m}$. The diagonal blocks $M_i \in \mathbb{R}^{p\times p}$, $i = 1, 2, \dots, m$, are symmetric, and the subdiagonal blocks $S_j \in \mathbb{R}^{p\times p}$, $j = 2, 3, \dots, m+1$, are upper triangular. We tacitly assume that m is small enough so that the decomposition (3) exists and can be computed as summarized in Algorithm 1.

Algorithm 1 Block Lanczos tridiagonalization (BLT).

Input: A, B, m.
1. Compute the compact QR factorization $B = X_1 S_1$.
2. $M_1 = X_1^T A X_1$.
3. $B_2 = A X_1 - X_1 M_1$.
4. Compute the compact QR factorization $B_2 = X_2 S_2$.
5. For $j = 2, \dots, m$
  (a) $M_j = X_j^T A X_j$.
  (b) $B_{j+1} = A X_j - X_j M_j - X_{j-1} S_j^T$.
  (c) Compute the compact QR factorization $B_{j+1} = X_{j+1} S_{j+1}$.
6. EndFor
Output: Block Lanczos decomposition (3)

The block columns X_i, $i = 1, 2, \dots, m$, of the matrix Q_m form an orthonormal basis for the block Krylov subspace

$$\mathbb{K}_m(A, X_1) = \text{span}\{X_1, AX_1, A^2X_1, \dots, A^{m-1}X_1\}, \qquad m \geqslant 1.$$

An approximate solution of (1) can be computed by the truncated block Lanczos tridiagonalization method as follows: compute the solution $Y_m \in \mathbb{R}^{pm\times p}$ of the small minimization problem on the right-hand side of

$$\min_{X\in\mathbb{K}_m(A,X_1)} \|AX - B\|_F = \min_{Y\in\mathbb{R}^{pm\times p}} \|T_{m+1,m}Y - E_1 S_1\|_F, \tag{5}$$

where $E_1 = [I_p, O_p, \dots, O_p]^T \in \mathbb{R}^{p(m+1)\times p}$. Then, $\widehat{X}_m := Q_m Y_m$ is the solution of the large minimization problem on the left-hand side of (5), as well as an approximate solution of (1). By choosing m suitably small, we can ensure that the matrix $T_{m+1,m}$ is of full rank and that the effect of the error E in B on the computed solution $\widehat{X}_m$ is smaller than if we attempt to solve the original problem (1). The latter is a consequence of the fact that the condition number of $T_{m+1,m}$, given by

$$\kappa(T_{m+1,m}) := \|T_{m+1,m}\|_2 \|T_{m+1,m}^\dagger\|_2,$$

is an increasing function of m. Here and below $\|\cdot\|_2$ denotes the spectral norm of a matrix. A large condition number indicates that the solution Y_m of the problem on the right-hand side of (5) is very sensitive to errors in the data as well as to round-off errors introduced during the computations. In Section 3 we will discuss properties of the block tridiagonal matrix $T_{m+1,m}$, and in Section 2.3 how to stabilize the solution process by Tikhonov regularization; the solution method so obtained does not require the matrix $T_{m+1,m}$ to be of full rank.

2.2 | Block Golub–Kahan bidiagonalization (BGKB)

A large nonsymmetric and possibly rectangular matrix $A \in \mathbb{R}^{\ell \times n}$ can be reduced to a small lower block bidiagonal matrix by application of a few steps of the block Golub–Kahan bidiagonalization (BGKB) algorithm. This reduction method can be used to determine an approximate solution of the minimization problem

$$\min_{X \in \mathbb{R}^{n \times p}} \|AX - B\|_F, \tag{6}$$

where the block vector $B \in \mathbb{R}^{\ell \times p}$ is error-contaminated and can be written as (2). We assume for notational simplicity that $1 \leqslant n \leqslant \ell$. Introduce the compact QR factorization $B = P_1 R_1$, where $P_1 \in \mathbb{R}^{\ell \times p}$ has orthonormal columns and $R_1 \in \mathbb{R}^{p \times p}$ is upper triangular. Then, $m \ll n/p$ steps of the BGKB algorithm applied to A with initial block vector P_1 give the decompositions

$$AW_m = U_{m+1} C_{m+1,m}, \qquad A^T U_m = W_m C_{m,m}^T, \tag{7}$$

where the matrices $U_{m+1} = [P_1, \ldots, P_{m+1}] \in \mathbb{R}^{\ell \times p(m+1)}$ and $W_m = [Z_1, \ldots, Z_m] \in \mathbb{R}^{n \times pm}$ have orthonormal columns and

$$C_{m+1,m} = \begin{bmatrix} L_1 & & & \\ R_2 & L_2 & & \\ & \ddots & \ddots & \\ & & R_m & L_m \\ & & & R_{m+1} \end{bmatrix} \in \mathbb{R}^{p(m+1) \times pm}$$

is lower block bidiagonal with lower triangular diagonal blocks $L_j \in \mathbb{R}^{p \times p}$ and upper triangular subdiagonal blocks $R_j \in \mathbb{R}^{p \times p}$. The matrix U_m consists of the m first block columns of U_{m+1}, and $C_{m,m}$ is the $pm \times pm$ leading principal submatrix of $C_{m+1,m}$. We assume that the number of steps, m, is small enough so that the decompositions (7) with the stated properties exists. The main steps required to compute these decompositions are summarized in Algorithm 2.

Algorithm 2 Block Golub–Kahan bidiagonalization (BGKB).

Input: A, B, m
1. Compute the compact QR factorization $B = P_1 R_1$
2. $F_1 = A^T P_1$
3. Compute the compact QR factorization $F_1 = Z_1 L_1^T$
4. For $j = 1, \ldots, m$
  (a) $H_j = AZ_j - P_j L_j$
  (b) Compute the compact QR factorization $H_j = P_{j+1} R_{j+1}$
  (c) If $j < m$
    i. $F_{j+1} := A^T P_{j+1} - Z_j R_{j+1}^T$
    ii. Compute the compact QR factorization $F_{j+1} = Z_{j+1} L_{j+1}^T$
  (c) EndIf
4. EndFor
Output: Block Golub–Kahan decompositions (7)

We will use the connection between the BGKB of A and the BLT of $A^T A$ in our analysis of the decompositions (3) and (7). Multiplying the left-hand side decomposition of (7) by A^T from the left-hand side gives

$$A^T A W_m = A^T U_{m+1} C_{m+1,m} = W_{m+1} \underbrace{C_{m+1,m+1}^T C_{m+1,m}}_{=: T_{m+1,m}}. \tag{8}$$

Thus, this decomposition is analogous to (3). In particular, the matrix $T_{m+1,m}$ is block tridiagonal with block size $p \times p$ and its leading $pm \times pm$ submatrix is symmetric. We conclude that (8) is a block Lanczos tridiagonalization of $A^T A$ with initial block vector Z_1. Since $T_{m+1,m}$ is block tridiagonal, equation (8) shows that the block columns Z_j of W_m satisfy a three-term recurrence relation. Moreover, the block columns $Z_1, Z_2, \ldots, Z_m$ form an orthonormal basis for the block Krylov subspace

$$\mathbb{K}_m(A^T A, A^T P_1) = \text{span}\{A^T P_1, (A^T A) A^T P_1, \ldots, (A^T A)^{m-1} A^T P_1\}, \qquad m \geqslant 1.$$

The block LSQR method applied to the solution of (6) solves at step m the minimization problem

$$\min_{X \in \mathbb{K}_m(A^T A, A^T P_1)} \|AX - B\|_F = \min_{Y \in \mathbb{R}^{pm \times p}} \|C_{m+1,m} Y - E_1 R_1\|_F, \tag{9}$$

where the right-hand side is obtained by substituting decomposition (7) into the left-hand side. Assume for the moment that the matrix $C_{m+1,m}$ is of full rank, and denote the solution of the problem appearing on the right-hand side of (9) by Y_m. Then, $\widehat{X}_m = W_m Y_m$ is the solution of the problem appearing on the left-hand side of (9), which is an approximate solution of (6).

2.3 | Tikhonov regularization

The block tridiagonal or lower block bidiagonal matrices in the reduced problems (5) and (9), respectively, might be numerically rank deficient. This often is the case when these matrices are large, because the singular values of the matrix A "cluster" at the origin. It follows that the reduced problems may have to be regularized before solution. We will apply Tikhonov regularization to the reduced problems obtained by BLT and BGKB; we provide details for the former only (i.e., for the problem appearing on the right-hand side of (5)). Tikhonov regularization applied to this setting gives a minimization problem of the form

$$\min_{Y \in \mathbb{R}^{pm \times p}} \left\{ \|T_{m+1,m} Y - E_1 S_1\|_F^2 + \mu \|Y\|_F^2 \right\}. \tag{10}$$

For a given value of the regularization parameter $\mu > 0$, the solution of (10) can be expressed as

$$Y_\mu = \left(T_{m+1,m}^T T_{m+1,m} + \mu I \right)^{-1} T_{m+1,m}^T E_1 S_1. \tag{11}$$

There are several techniques for determining a suitable value of μ, including the discrepancy principle, generalized cross validation, and the L-curve criterion; see[4, 5, 7, 9, 14, 15, 18] for discussions on these and other methods. In the computed examples of this paper, we will use the discrepancy principle, which was first discussed by Morozov in[17]. This approach to determine μ requires that a bound for the error E in B, cf. (2), be known,

$$\|E\|_F \leq \rho,$$

and prescribes that $\mu > 0$ be chosen so that the solution (11) of (10) satisfies

$$\|T_{m+1,m} Y_\mu - E_1 S_1\|_F = \tau \rho, \tag{12}$$

where $\tau > 1$ is a user-chosen constant that is independent of ρ; when the available estimate of $\|E\|_F$ is deemed accurate, the parameter τ is generally chosen to be close to unity. We note that there is a $\mu > 0$ that satisfies (12) only if the number of steps, m, is large enough. It follows from (10) that $\mu \to \|T_{m+1,m} Y_\mu - E_1 S_1\|_F$ is an increasing function of $\mu \geq 0$. In our examples, we choose m as small as possible so that $\|T_{m+1,m} Y_\mu - E_1 S_1\|_F < \tau \rho$. Then a zero-finder is applied to solve (12) for $\mu > 0$; see[2] for further details. Thus, the discrepancy principle is used both to determine the number of steps m and $\mu > 0$.

Similar derivations and analogous expressions for the norm of the residual errors can be obtained when applying the BGKB algorithm.

3 | BLT AND BGKB APPLIED TO LINEAR DISCRETE ILL-POSED PROBLEMS

This section first discusses the convergence of the subdiagonal and diagonal block entries of the matrices $T_{m+1,m}$ and $C_{m+1,m}$ in (3) and (7), respectively, with increasing block index. Unless otherwise stated, and with a slight abuse of notation, here and in the following, we will denote by A the square matrix of order $n + q$, whose leading principal submatrix of order n is the coefficient matrix A appearing in (3), padded with $0 \leq q < p$ rows and columns of zeros, where p is the block size used in the block Lanczos or block Golub–Kahan algorithms. For A symmetric, the following proofs use the spectral factorization

$$A = \mathcal{W} \Lambda \mathcal{W}^T, \tag{13}$$

where the matrix $\mathcal{W} = [w_1, w_2, \ldots, w_{n+q}] \in \mathbb{R}^{(n+q)\times(n+q)}$ is orthogonal and

$$\Lambda = \mathrm{diag}[\lambda_1, \lambda_2, \ldots, \lambda_{n+q}] \in \mathbb{R}^{(n+q)\times(n+q)}, \qquad |\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_{n+q}| \geq 0. \tag{14}$$

We will use the notation

$$E_i = [O_p, \ldots, O_p, I_p, O_p, \ldots, O_p]^T \in \mathbb{R}^{(n+q)\times p}, \quad i = 1, 2, \ldots, r,$$

where I_p is the ith block, and where $0 < r := (n+q)/p$ with q being the smallest non-negative integer such that $r \in \mathbb{N}$.

Theorem 1. Assume that the block Lanczos method applied to the symmetric and positive semidefinite matrix $A \in \mathbb{R}^{(n+q)\times(n+q)}$ with initial block matrix $X_1 \in \mathbb{R}^{(n+q)\times p}$ with orthonormal columns does not break down, i.e., that $r := (n+q)/p$ steps of the method can be carried out. Let the eigenvalues of A be ordered according to (14), and let $S_2, S_3, \ldots, S_{m+1}$, $m \leq r$, be the subdiagonal blocks of the matrix $T_{m+1,m}$ determined by m steps of the block Lanczos methods; cf. (4). Define $S_{r+1} = O_p$. Then,

$$\|S_{m+1}S_m \cdots S_2\|_2 \leq \prod_{j=1}^{m} \lambda_j, \qquad m = 1, 2, \ldots, r. \tag{15}$$

Proof. Introduce the monic polynomial $p_m(t) = \prod_{j=1}^{m}(t - \lambda_j)$ defined by the m largest eigenvalues of A. Using the spectral factorization (13), we obtain

$$\|p_m(A)\|_2 = \|p_m(\Lambda)\|_2 = \max_{m+1\leq j\leq n+q} |p_m(\lambda_j)| \leq |p_m(0)| = \prod_{j=1}^{m} \lambda_j,$$

where the inequality follows from the fact that all λ_j are nonnegative. Hence,

$$\|p_m(A)X_1\|_2 \leq \|p_m(A)\|_2 \cdot \|X_1\|_2 = \|p_m(A)\|_2 \leq \prod_{j=1}^{m} \lambda_j. \tag{16}$$

Application of r steps of the block Lanczos method gives the decomposition $A = Q_r T_{r,r} Q_r^T$, where $T_{r,r} \in \mathbb{R}^{(n+q)\times(n+q)}$ is symmetric block tridiagonal. We have

$$p_m(A)X_1 = p_m(Q_r T_{r,r} Q_r^T)X_1 = Q_r p_m(T_{r,r})\, Q_r^T X_1 = Q_r p_m(T_{r,r})E_1.$$

Thus,

$$\|p_m(A)X_1\|_2 = \|p_m(T_{r,r})E_1\|_2 \geq \|E_{m+1}^T p_m(T_{r,r})E_1\|_2. \tag{17}$$

The above inequality follows by direct computations. We are going to show by induction over m that

$$E_{m+1}^T p_m(T_{r,r})E_1 = S_{m+1}S_m \cdots S_2, \tag{18}$$

for any $m < r$. When $m = 1$, equation (18) becomes

$$E_2^T(T_{r,r} - \lambda_1 I_{n+q})E_1 = \left[S_2,\ (M_2 - \lambda_1 I_p),\ S_3^T,\ O_p,\ \cdots,\ O_p\right] E_1 = S_2.$$

When $m = 2$, let us consider

$$E_3^T(T_{r,r} - \lambda_2 I_{n+q})(T_{r,r} - \lambda_1 I_{n+q})E_1. \tag{19}$$

The first two factors in (19) are

$$E_3^T(T_{r,r} - \lambda_2 I_{n+q}) = \left[O_p,\ S_3,\ (M_3 - \lambda_2 I_p),\ S_4^T,\ O_p,\ \cdots,\ O_p\right] = S_3 E_2^T + (M_3 - \lambda_2 I_p)E_3^T + S_4^T E_4^T,$$

while the remaining two factors are

$$(T_{r,r} - \lambda_1 I_{n+q})E_1 = \begin{bmatrix} M_1 - \lambda_1 I_p \\ S_2 \\ O_p \\ \vdots \\ O_p \end{bmatrix} = (M_1 - \lambda_1 I_p)E_1 + S_2 E_2.$$

It follows that the expression (19) can be written as

$$\left(S_3 E_2^T + (M_3 - \lambda_2 I_p)E_3^T + S_4^T E_4^T\right)\left((M_1 - \lambda_1 I_p)E_1 + S_2 E_2\right) = S_3 S_2.$$

More generally, by induction, assume that (18) is valid for $2 \leq m < r - 1$. This means

$$E_{m+1}^T p_m(T_{r,r})E_1 = S_{m+1}S_m \cdots S_2. \tag{20}$$

We would like to show that (18) is valid for $2 \leq m + 1 < r$. From

$$\begin{aligned} E_{m+2}^T p_{m+1}(T_{r,r})E_1 &= E_{m+2}^T(T_{r,r} - \lambda_{m+1} I_{n+q})p_m(T_{r,r})E_1 \\ &= \left(S_{m+2}E_{m+1}^T + (M_{m+2} - \lambda_{m+1} I_p)E_{m+2}^T + S_{m+3}^T E_{m+3}^T\right) p_m(T_{r,r})E_1 \\ &= S_{m+2}E_{m+1}^T p_m(T_{r,r})E_1 \qquad \text{(since } p_m(T_{r,r}) \text{ is } (2m+1)\text{-block-banded)} \\ &= S_{m+2}S_{m+1}S_m \cdots S_2, \qquad \text{(by (20))} \end{aligned}$$

it follows that

$$E_{m+2}^T p_{m+1}(T_{r,r})E_1 = S_{m+2}S_{m+1}\cdots S_2.$$

Hence, (18) is valid. Combining (16), (17) and (18) proves the theorem. □

We are interested in problems for which the right-hand side of (15) decreases to zero an m increases. This holds for all symmetric linear discrete ill-posed problems that we have encountered in various applications. It depends on that the eigenvalues of A "cluster" at the origin. In fact, we found the bound (15) to be quite sharp. We give a numerical illustration of the latter in Section 4.

Corollary 1. Let the matrix $A \in \mathbb{R}^{(n+q)\times(n+q)}$ be symmetric and positive semidefinite. Assume that the eigenvalues of A "cluster" at the origin so that the right-hand side of (15) decreases to zero as m increases, and that the block Lanczos method applied to A with initial block vector X_1 with orthonormal columns does not break down. Further, assume that, for all $j > s$,

$$\frac{\|S_j\|_2}{\min\limits_{1\leqslant i\leqslant s}\|S_i\|_2} \leqslant C \tag{21}$$

for some constant C independent of j and s. Then, both the diagonal and subdiagonal block matrices M_j and S_j of the block tridiagonal Lanczos matrix $T_{r,r}$, defined by (3), converge to O_p as j increases.

Proof. We first remark that when we let the index j increase in (21), we also may have to increase m in (3). By Theorem 1, we have the bound

$$\|S_{m+1}S_m\cdots S_2\|_2 \leq \prod_{j=1}^{m}\lambda_j.$$

The product $\prod_{j=1}^{m}\lambda_j$ converges to zero as m increases. It follows that $\|S_{m+1}S_m\cdots S_2\|_2$ converges to zero as m increases. Therefore,

$$S_{m+1}S_m\cdots S_2 \to O_p, \qquad \text{as} \qquad m\to\infty.$$

In view of (21), the subdiagonal blocks S_j of $T_{r,r}$ approach O_p as j increases.

We turn to the block diagonal entries M_j of the matrix $T_{r,r}$. Let $\delta > 0$ be arbitrarily small. Since $A = Q_r T_{r,r} Q_r^T$, the matrices A and $T_{r,r}$ are similar. Therefore, the eigenvalues of matrix $T_{r,r}$ "cluster" at the origin, which is the only cluster point. Split the matrix $T_{r,r} = \widetilde{T}_{j,j} + \widetilde{E}_j$, where

$$\widetilde{T}_{j,j} = \begin{bmatrix} M_1 & S_2^T & & & & & & & \\ S_2 & M_2 & S_3^T & & & & & & \\ & \ddots & \ddots & \ddots & & & & & \\ & & S_{j-1} & M_{j-1} & S_j^T & & & & \\ & & & S_j & M_j & O_p & & & \\ & & & & O_p & \ddots & \ddots & & \\ & & & & & \ddots & & O_p & \\ & & & & & & O_p & M_{r-1} & O_p \\ & & & & & & & O_p & M_r \end{bmatrix}, \qquad \widetilde{E}_j = \begin{bmatrix} \mathrm{O} & & & & & \\ & & & & & \\ & & & & & \\ & O_p & S_{j+1}^T & & & \\ & S_{j+1} & O_p & \ddots & & \\ & & \ddots & \ddots & & \\ & & & & O_p & S_r^T \\ & & & & S_r & O_p \end{bmatrix},$$

and j is chosen so that $\|\widetilde{E}_j\|_2 \leq \delta$. Thus, $\widetilde{T}_{j,j}$ is a symmetric block tridiagonal matrix, which is obtained by setting the subdiagonal blocks of $T_{r,r}$ in the block rows $j+1, j+2, \ldots, r$ to zero (the corresponding superdiagonal blocks also are set to zero). The matrix $\widetilde{E}_j$ contains the blocks set to zero in $T_{r,r}$.

Since the eigenvalues of $T_{r,r}$ "cluster" at the origin, it follows from the Bauer–Fike theorem that the eigenvalues of the matrix $\widetilde{T}_{j,j}$ "cluster" in the interval $[-\delta, \delta]$. For some $\eta > 0$ arbitrarily small, there is an index s, depending on η, such that all eigenvalues of the blocks M_k are in the interval $[-\delta-\eta, \delta+\eta]$ for all $k \geq s$. Hence, $\|M_k\|_2 \leq \delta+\eta$ for all $k \geq s$. Since δ and η can be chosen arbitrarily small, this shows that the diagonal blocks M_j converge to the zero matrix as j increases. □

Corollary 1 is stated in[8] for block size $p = 1$ without the condition (21). Numerical experiments with a large number of discrete ill-posed problems indicate that this condition does not have to be imposed. We conjecture that this is the case.

The following example illustrates that condition (21) is required if the matrix A is not a discretization of an ill-posed operator equation.

Example. Let $p = 1$ and consider the symmetric tridiagonal matrix T_{2m} with subdiagonal entries $S_{2j} = 1$, $j = 1, 2, \ldots, m$, and $S_{2j+1} = 10^{-j}$, $j = 1, 2, \ldots, m-1$, and diagonal entries M_j equal to the sum of the subdiagonal and superdiagonal entries in the same row. Then, T_{2m} satisfies the conditions of Corollary 1 except for (21). Its eigenvalues cluster at the origin and at 2. Since the eigenvalues of T_{2m} cluster at two points, the matrix is not a discretization of a linear operator of an ill-posed problem. Neither the diagonal nor subdiagonal entries of T_{2m} converge to zero for increasing index number as m increases. $\square$

We observe that the decrease of the subdiagonal blocks S_j of $T_{r,r}$ to the zero matrix follows from the clustering of the eigenvalues of A. It is not necessary that they cluster at the origin. This can be seen by replacing the matrix A in Corollary 1 by $A + cI_n$ for some constant $c \in \mathbb{R}$.

We turn to symmetric indefinite matrices.

Theorem 2. Let the eigenvalues $\{\lambda_j\}_{j=1}^{n+q}$ of the symmetric matrix $A \in \mathbb{R}^{(n+q)\times(n+q)}$ be ordered according to (14). Assume that the block Lanczos method applied to A with initial matrix X_1 does not break down. Then

$$\|S_{m+1}S_m \cdots S_2\|_2 \leq \prod_{j=1}^{m}(\,|\lambda_{m+1}| + |\lambda_j|), \qquad m = 1, 2, \ldots, r-1. \tag{22}$$

Proof. Let $p_m(t)$ be the monic polynomial of the proof of Theorem 1. Then, just like in that proof

$$\|p_m(A)\|_2 = \|p_m(\Lambda)\|_2 = \max_{m+1\leq j\leq n+q} |p_m(\lambda_j)|.$$

Due to the ordering (14), it follows that the eigenvalues $\lambda_{m+1}, \lambda_{m+2}, \ldots, \lambda_{n+q}$ are contained in the interval $[-|\lambda_{m+1}|, |\lambda_{m+1}|]$. Thus,

$$\begin{aligned}\|p_m(A)\|_2 &= \max_{m+1\leq j\leq n+q} |p_m(\lambda_j)| \leqslant \max_{-|\lambda_{m+1}|\leq t\leq|\lambda_{m+1}|} |p_m(t)| = \max_{-|\lambda_{m+1}|\leq t\leq|\lambda_{m+1}|} \prod_{k=1}^{m} |t - \lambda_k| \\ &\leqslant \max_{-|\lambda_{m+1}|\leq t\leq|\lambda_{m+1}|} \prod_{k=1}^{m}(|t| + |\lambda_k|) = \prod_{k=1}^{m}(|\lambda_{m+1}| + |\lambda_k|).\end{aligned}$$

Therefore,

$$\|p_m(A)X_1\|_2 \leq \|p_m(A)\|_2 \cdot \|X_1\|_2 \leqslant \prod_{k=1}^{m}(\,|\lambda_{m+1}| + |\lambda_k|\,).$$

Also, we have shown in (18) that

$$\|p_m(A)X_1\|_2 \geq \|E_{m+1}^T p_m(T_{r,r})E_1\|_2 = \|S_{m+1}S_m \cdots S_2\|_2.$$

Hence,

$$\prod_{k=1}^{m}(\,|\lambda_{m+1}| + |\lambda_k|\,) \geqslant \|p_m(A)X_1\|_2 \geqslant \|S_{m+1}S_m \cdots S_2\|_2.$$

$\square$

Assume that the eigenvalues of A cluster at the origin. Then, Theorem 2 shows that the quantity $\|S_{m+1}S_m \cdots S_2\|_2$ decreases to zero, because the factors $|\lambda_{m+1}| + |\lambda_k|$ decrease to zero as m and k increase, with $1 \leqslant k \leqslant m$. Moreover, the more block Lanczos steps are taken, the tighter is the bound for the norm of the product of the subdiagonal block matrices of the matrix $T_{r,r}$.

We can obtain sharper bounds if more information about the spectrum of A is available. For instance, if all but a few eigenvalues of A are known to be nonnegative, then only factors with negative eigenvalues have to be modified as in Theorem 2, resulting in improved bounds for $\|S_{m+1}S_m \cdots S_2\|_2$. In the next corollary, we derive a simpler, but cruder, bound than (22).

Corollary 2. Let the eigenvalues $\{\lambda_j\}_{j=1}^{n+q}$ of the symmetric matrix $A \in \mathbb{R}^{(n+q)\times(n+q)}$ be ordered according to (14). Assume that the block Lanczos method applied to A with initial block vector X_1 with orthonormal columns does not break down. Then

$$\|S_{m+1}S_m \cdots S_2\|_2 \leq \prod_{k=1}^{m}(2|\lambda_k|), \qquad m = 1, 2, \ldots, r-1.$$

Proof. By Theorem 2, since $|\lambda_{m+1}| \leq |\lambda_k|$ for $1 \leq k \leq m$, we have

$$\|S_{m+1}S_m \cdots S_2\|_2 \leq \prod_{k=1}^{m}(|\lambda_{m+1}| + |\lambda_k|) \leq \prod_{k=1}^{m}(|\lambda_k| + |\lambda_k|) = \prod_{k=1}^{m}(2|\lambda_k|).$$

$\square$

A block vector $X \in \mathbb{R}^{(n+q)\times p}$ is said to be invariant under the matrix $A \in \mathbb{R}^{(n+q)\times(n+q)}$ if $AX \subset \mathcal{R}(X)$, where $\mathcal{R}(X)$ is the range space of X. Thus, there is a matrix $M \in \mathbb{R}^{p\times p}$ such that $AX = XM$. Let $\epsilon > 0$. We say that the block vector X with orthonormal columns is ϵ-invariant under A if there is a matrix $M \in \mathbb{R}^{p\times p}$ such that

$$\|AX - XM\|_2 \leq \epsilon.$$

Theorem 3. Let the conditions of Corollary 1 hold, and let $\epsilon > 0$. Then, for j sufficiently large, the block vectors $X_j \in \mathbb{R}^{(n+q)\times p}$ determined by the block Lanczos algorithm are ϵ-invariant under A with $M = O_p$.

Proof. We have

$$\|AX_j\|_2 = \|AQ_m E_j\|_2 = \|Q_{m+1}T_{m+1,m}E_j\|_2 = \|X_{j-1}S_j^T + X_jM_j + X_{j+1}S_{j+1}\|_2.$$

Because M_j and S_j approach O_p as j increases, we can conclude that the Lanczos block vectors X_j are ϵ-invariant under A with $M = O_p$ for j large. □

We conclude this subsection by deriving an estimate that is analogous to the one in Theorem 1 for nonsymmetric and potentially rectangular matrices that require the use of the BGKB partial factorization. This estimate involves both the diagonal and lower diagonal blocks of the matrix $C_{m+1,m}$ in (7), and leverages the fact that (8) is a block Lanczos decomposition of A^TA, analogous to (3). Here we assume that $A \in \mathbb{R}^{\ell\times(n+q)}$, with $\ell \geq (n+q)$, is a matrix whose first n columns contain the coefficient matrix A appearing in (7), padded with $q \geq 0$ columns of zeros; q is the smallest non-negative integer such that $r := (n+q)/p \in \mathbb{N}$.

Corollary 3. Let $A \in \mathbb{R}^{\ell\times(n+q)}$ have the singular values $\sigma_1 \geqslant \sigma_2 \geqslant \cdots \geqslant \sigma_{n+q} \geqslant 0$, and assume that the BGKB algorithm applied to A with initial block vector P_1 does not break down, i.e., that $r := (n+q)/p$ steps of the method can be carried out. Then

$$\|L_{m+1}^TR_{m+1}L_m^TR_m \cdots L_2^TR_2\|_2 \leq \prod_{j=1}^{m} \sigma_j^2, \qquad m = 1, 2, \ldots, r. \tag{23}$$

Proof. The block entries of the block tridiagonal matrix in (8) can be expressed as

$$T_{m+1,m} = C_{m+1,m+1}^T C_{m+1,m} = \begin{bmatrix} D_1 & R_2^TL_2 & & & \\ L_2^TR_2 & D_2 & R_3^TL_3 & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & R_m^TL_m \\ & & & L_m^TR_m & D_m \\ & & & & L_{m+1}^TR_{m+1} \end{bmatrix},$$

with $D_j = L_j^TL_j + R_{j+1}^TR_{j+1}$. Since the block subdiagonal entries of the matrix $T_{m+1,m}$ are $L_j^TR_j$, and the eigenvalues of A^TA are σ_j^2, the result follows from Theorem 1. □

Note that, since the singular value decomposition of A can be characterized in terms of the eigendecompositions of A^TA and AA^T, the bound (23) can only be given in terms of both the diagonal and lower diagonal blocks of the matrix $C_{m+1,m}$.

We are interested in problems for which the right-hand side of (23) decreases to zero an m increases. This holds for all linear discrete ill-posed problems that we have come across in many applications.

4 | COMPUTED EXAMPLES

To illustrate the properties discussed in the previous sections, we applied the symmetric block Lanczos method and the block Golub–Kahan bidiagonalization method to a set of test matrices that stem from the discretization of ill-posed problems. The numerical experiments were carried out using MATLAB R2017a with about 15 significant decimal digits, on a Xeon E-2244G computer (8 cores, 16 threads) with 16 Gbyte RAM. Although the results of many numerical tests validating the use of BLT and BGKB are already provided in[1] and[2], respectively, the ones presented in this section illustrate the theoretical results in this paper, include new comparisons with other direct and sequential solution methods, and span a wider set of test problems.

A first set of experiments uses fairly small square test matrices of order 200 (unless otherwise stated), so that computing the eigendecomposition and the SVD is computationally affordable; indeed, the properties discussed in this paper can be observed already for quite small matrices. The symmetric test matrices are listed in the upper part of Table 1, and the nonsymmetric ones in the bottom part of the same table. Among the symmetric matrices, `gravity` is positive definite, `deriv2` is negative definite, and `phillips` is indefinite. In the case of the nonsymmetric test matrix `tomo`, we set the size to 400×400, because of the very slow decay of its singular values. All matrices but one (i.e., `lotkin` from MATLAB's `gallery`) in this set of experiments are from the REGULARIZATION TOOLS package[12]. More precisely, we use the test problems from[12] to define matrices A, the first column of B_{true}, and the associated error-free solution x_0 in the first column of the block vector solution X_{true}. The other columns of X_{true} are obtained by setting $x_i = x_{i-1} + \frac{y}{2}$ for $i = 1, 2, \dots, p-1$, where y is a vector obtained by discretizing a function of the form $\frac{1}{2}\cos\frac{t}{3} + \frac{1}{4}$ at equidistant points on the interval $-6 \leqslant t \leqslant 6$. Consequently, the other columns of B_{true} are obtained by taking $b_i = Ax_i$ for $i = 1, 2, \dots, p-1$. The solution of the `lotkin` example is the same as for the `phillips` example. The contaminated data block vector is given by (2) with

$$E = \breve{E}\|B_{\text{true}}\|_F \delta,$$

where the random block vector $\breve{E} \in \mathbb{R}^{n\times p}$ models Gaussian white noise with mean zero and variance one, and δ is a chosen noise level. In our experiments, we let $\delta \in \{10^{-6}, 10^{-4}, 10^{-2}\}$. Unless otherwise stated, the blocksize is $p = 5$. As prescribed by Algorithms 1 and 2, the BLT and BGKB algorithms are initiated with the block vector B. One reorthogonalization step is carried out; the process is repeated if needed. The computed results do not agree with the theory developed in the previous section when no reorthogonalization is carried out. The quantity $K \leq r$ denotes the number of BLT or BGKB steps performed for each test problem. The last experiment models image deblurring of a color image, and uses some of the functionalities available within the IR TOOLS package[6].

In the first set of experiments we use an error-free initial block, both in the BLT and BGKB algorithms, that is, we set $B = B_{\text{true}}$. We verified that the graphs in Figures 1–5 do not change significantly for noise levels up to 10^{-2}.

We first illustrate the properties derived in Section 3. Figure 1 displays, in logarithmic scale, the values taken by the left-hand side and right-hand side in the inequalities (15), (22), and (23), as functions of the number of iterations. Iterations were carried out until breakdown, that is, $m = 1, 2, \dots, K \leq r$. The graphs show that for symmetric discrete ill-posed problems the decay of the subdiagonal blocks of $T_{m+1,m}$ to zero may be much faster than suggested by the bounds (15) and (22). It follows that the ability of the Lanczos block vectors to approximate the space spanned by the principal eigenvectors often is stronger than indicated by the bounds (15) and (22). The same holds true for the BGKB method. We also remark that round-off errors introduced during the computation of the eigenvalues and subdiagonal blocks of the matrices $T_{m,m}$, $m = 1, 2, \dots$, may affect the graphs. In any case, when m is large, the matrix $T_{m,m}$ has eigenvalues of "tiny" absolute value.

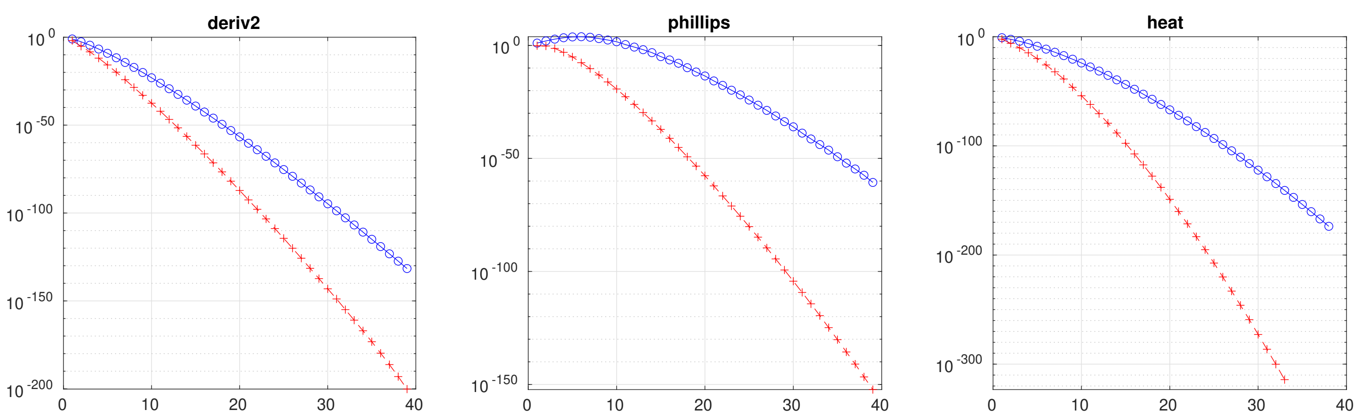

FIGURE 1 Behavior of the bounds (15) (left), (22) (center), and (23) (right), as functions of the iteration number m. The test matrices are (from left to right) symmetric positive definite, symmetric indefinite, nonsymmetric. The left-hand side of each inequality is represented by crosses, and the right-hand side by circles. The sign of the test matrix `deriv2`, which is negative definite, has been inverted.

We next illustrate that the subspaces $\mathcal{R}(Q_k)$ generated by the block Lanczos method (3) essentially contain subspaces of eigenvectors of A associated with the eigenvalues of largest absolute value. In addition, we show the convergence of the largest eigenvalues (in absolute value) of the matrices $T_{k,k}$ in (3) to the largest eigenvalues (in absolute value) of A as k increases. Here, $T_{k,k} \in \mathbb{R}^{pk\times pk}$ denotes the matrix obtained by neglecting the last block row of the matrix $T_{k+1,k} \in \mathbb{R}^{p(k+1)\times pk}$ in (3), with m replaced by k. The block Lanczos method is applied until breakdown occurs. For each k, consider the spectral factorization $T_{k,k} = \breve{W}_k \breve{\Lambda}_k \breve{W}_k^T$, where

$$\breve{\Lambda}_k = \mathrm{diag}[\breve{\lambda}_1^{(k)}, \breve{\lambda}_2^{(k)}, \ldots, \breve{\lambda}_{pk}^{(k)}], \qquad |\breve{\lambda}_1^{(k)}| \geq \cdots \geq |\breve{\lambda}_{pk}^{(k)}|, \quad \text{and} \quad \breve{W}_k = [\breve{w}_1^{(k)}, \breve{w}_2^{(k)}, \ldots, \breve{w}_{pk}^{(k)}].$$

The eigenvalues $\{\breve{\lambda}_i^{(k)}\}_{i=1}^{pk}$ are commonly referred to as Ritz values of A. We compare the Ritz values of largest absolute value to the corresponding eigenvalues of the matrix A. For each step k of the block Lanczos algorithm, we compute the relative difference

$$R_k^{\lambda} := \max_{i=1,2,\ldots,\lceil pk/3\rceil} \frac{|\breve{\lambda}_i^{(k)} - \lambda_i|}{|\lambda_i|}, \qquad k = 1, 2, \ldots, K,$$

i.e., we compare the $\lceil pk/3\rceil$ eigenvalues of largest absolute value of $T_{k,k}$ and A, where $\lceil \alpha \rceil$ denotes the integer closest to $\alpha \geq 0$. Figure 2 shows excellent agreement between the first $\lceil pk/3\rceil$ Ritz values of A and the corresponding eigenvalues already for small k.

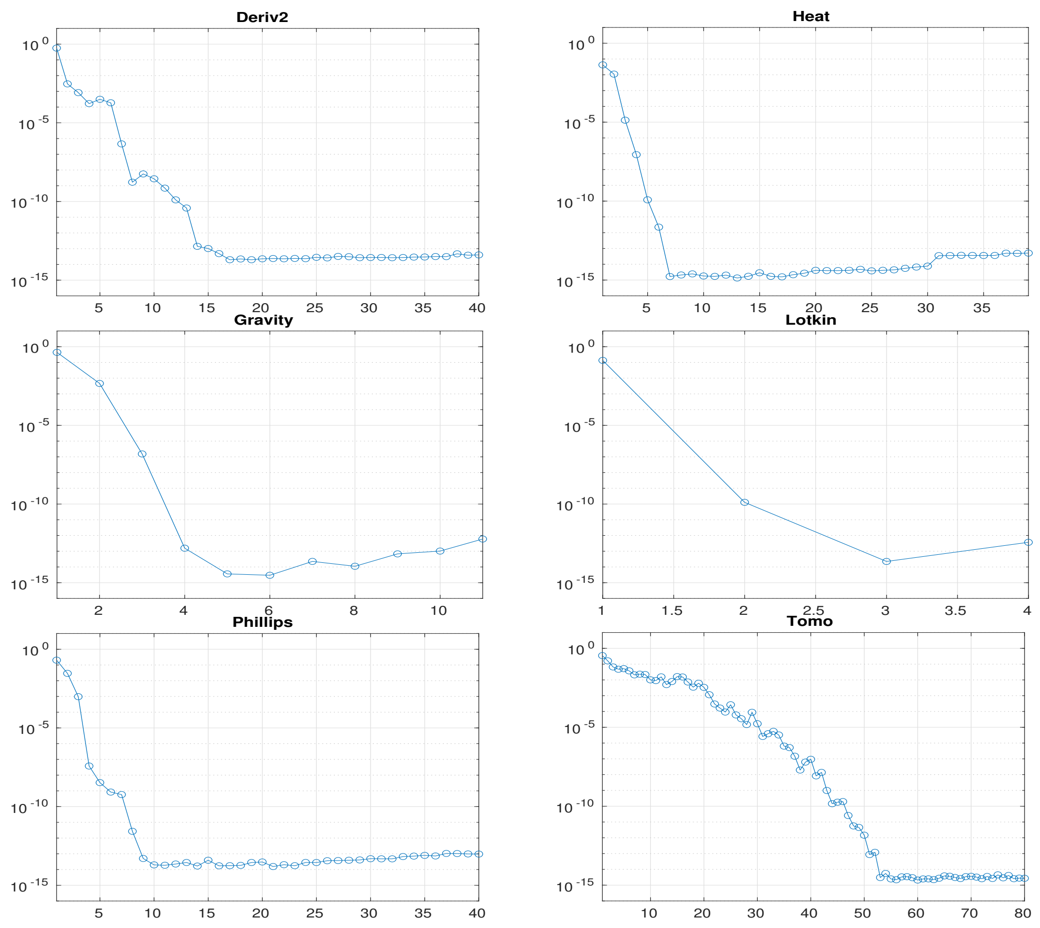

FIGURE 2 The graphs in the left-hand side column display the relative difference R_k^{λ} versus k between the $\lceil pk/3\rceil$ eigenvalues of largest absolute value of the symmetric test matrices and the corresponding Ritz values. The right-hand side column shows the behavior of R_k^{σ} versus k for nonsymmetric problems.

We turn to a comparison of subspaces determined by the span of the Lanczos block vectors of A associated with the Ritz values of largest absolute value. For each k, consider $Q_k \in \mathbb{R}^{n\times pk}$ made up of the first k block columns of the matrix Q_m in (3). Partition the eigenvector matrix of A, cf. (13), according to $\mathcal{W} = [\mathcal{W}_i^{(1)}\ \mathcal{W}_{n-i}^{(2)}]$, where the columns w_j $(j = 1,\dots,i)$ of $\mathcal{W}_i^{(1)} \in \mathbb{R}^{n\times i}$ are the first i eigenvectors, and let the columns $\mathcal{W}_{n-i}^{(2)} \in \mathbb{R}^{n\times(n-i)}$ be the remaining eigenvectors. The columns of $\mathcal{W}_i^{(1)}$ and $\mathcal{W}_{n-i}^{(2)}$ span orthogonal subspaces.

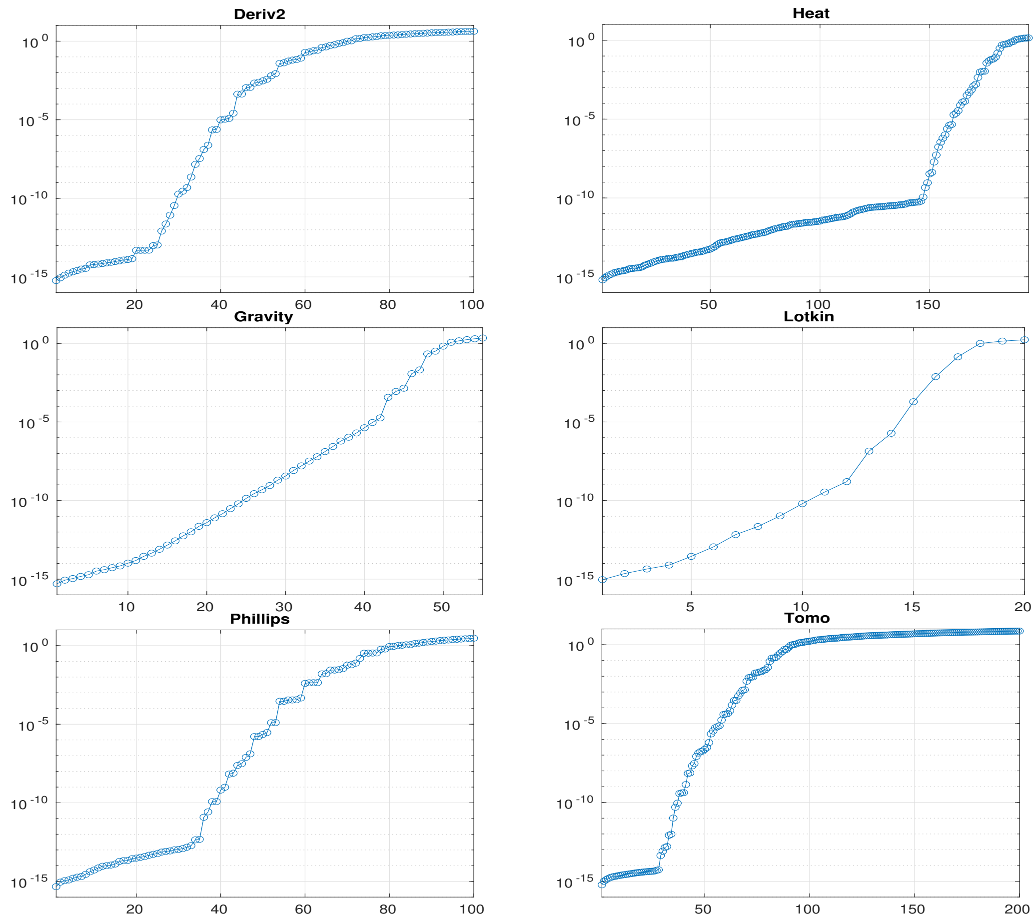

FIGURE 3 Distances $R^{w}_{k,i}$ (resp. $R^{(u,v)}_{k,i}$), versus $i = 1,2,\dots,pk$, between the subspaces spanned by the first i eigenvectors (resp. singular vectors) of the symmetric (resp. nonsymmetric) test matrices, and the subspaces spanned by the corresponding k Lanczos (resp. Golub–Kahan) block vectors. Here, $k = \lceil n/(2p)\rceil$, unless a breakdown occurred.

Let $\mathcal{Q}_k = I_n - Q_k Q_k^T$ be the orthogonal projector onto $\mathcal{R}(Q_k)^{\perp}$, the subspace orthogonal to the range of Q_k. We consider the quantities

$$R^{w}_{k,i} := \|\mathcal{Q}_k \mathcal{W}_i^{(1)}\|_F\,, \qquad k = 1,2,\dots,K, \quad i = 1,2,\dots,pk\,.$$

The value of $R^{w}_{k,i}$ is small when $\mathrm{span}\{w_j\}_{j=1}^{i}$ is approximately contained in $\mathrm{span}\{q_j\}_{j=1}^{pk}$, that is, when the solution subspace generated by the block Lanczos vectors essentially contains the space generated by the first i eigenvectors. The graphs in the left-hand side column of Figure 3 depict $R^{w}_{k,i}$ for $k = \lceil n/(2p)\rceil$ ($k = K$ if a breakdown occurred) and $i = 1,2,\dots,pk$, for the symmetric test matrices. They show that, for a fixed k, only a fraction of the eigenvectors are well approximated by pk Lanczos vectors.

The left-hand side column of Figure 4 displays the values of $R^{w}_{k,\lceil pk/3\rceil}$ ($k = 1,2,\dots,K$, if a breakdown occurred, and $pK < \lceil n/2\rceil$), while the right-hand side column of the same figure represents the behavior of $R^{w}_{k,\lceil pk/2\rceil}$, ($k = 1,2,\dots,K$, if a breakdown occurred, and $pK < \lceil n/2\rceil$).

A few comments on the graphs of Figure 4 are in order. The left-hand side graphs show that the span of the first $\lceil pk/3\rceil$ eigenvectors of A is numerically contained in the span of the first pk Lanczos vectors already for quite small values of k. We

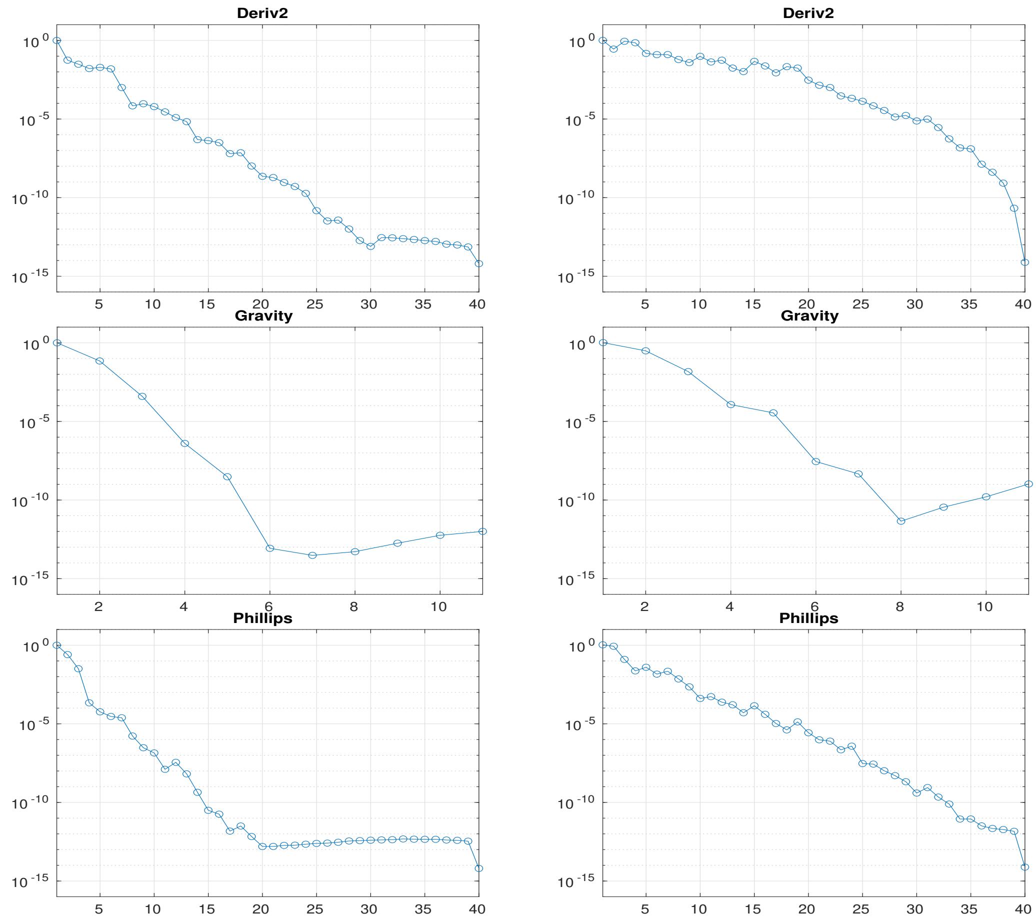

FIGURE 4 The graphs in the left-hand side column display the distances $R^w_{k,\lceil pk/3\rceil}$, versus $k = 1, 2, \dots, \lceil n/p\rceil$, between the space spanned by the $\lceil pk/3\rceil$ principal eigenvectors of the symmetric test matrices and the space spanned by the first k Lanczos block vectors. The right-hand side column shows the behavior of $R^w_{k,\lceil pk/2\rceil}$.

remark that this is not true if we compare the spaces spanned by the first pk eigenvectors of A and by the first k Lanczos block vectors. Graphs in the right-hand side column, that compare the span of the first $\lceil pk/2\rceil$ eigenvectors of A with the span of the first pk Lanczos vectors, look similar to the graphs in the left-hand side column, but display slower convergence.

We turn to nonsymmetric matrices A. Introduce the singular value decomposition

$$A = U\Sigma V^T. \tag{24}$$

Thus, $U \in \mathbb{R}^{\ell\times\ell}$ and $V \in \mathbb{R}^{n\times n}$ are orthogonal matrices, and

$$\Sigma = \mathrm{diag}[\sigma_1, \sigma_2, \dots, \sigma_n] \in \mathbb{R}^{\ell\times n}, \qquad \sigma_1 \geqslant \sigma_2 \geqslant \cdots \geqslant \sigma_r > \sigma_{r+1} = \cdots = \sigma_n = 0,$$

where r is the rank of A. The block Lanczos method in the above experiments is replaced by the block Golub–Kahan method (7). The latter method is applied until iteration $pK \leq r$, when breakdown occurs. The graphs in the right-hand side column of Figure 2 show the relative differences

$$R^\sigma_k := \max_{i=1,2,\dots,\lceil pk/3\rceil} \frac{|\breve{\sigma}^{(k)}_i - \sigma_i|}{|\sigma_i|}, \qquad k = 1, 2, \dots, K,$$

between the singular values $\breve{\sigma}^{(k)}_i$ of $C_{k+1,k}$ and the corresponding singular values of A.

Let U and V be the orthogonal matrices in the singular value decomposition (24) of A, let $\ell = n$, and partition these matrices similarly to what was done for the eigenvector matrix for symmetric matrices A, that is, we let $U = [U^{(1)}_i, U^{(2)}_{n-i}]$ and $V = [V^{(1)}_i, V^{(2)}_{n-i}]$, where the submatrices $U^{(1)}_i$ and $V^{(1)}_i$ contain the first i left and right singular vectors, and the submatrices $U^{(2)}_{n-i}$ and $V^{(2)}_{n-i}$ contain the remaining $n-i$ left and right singular vectors, respectively.

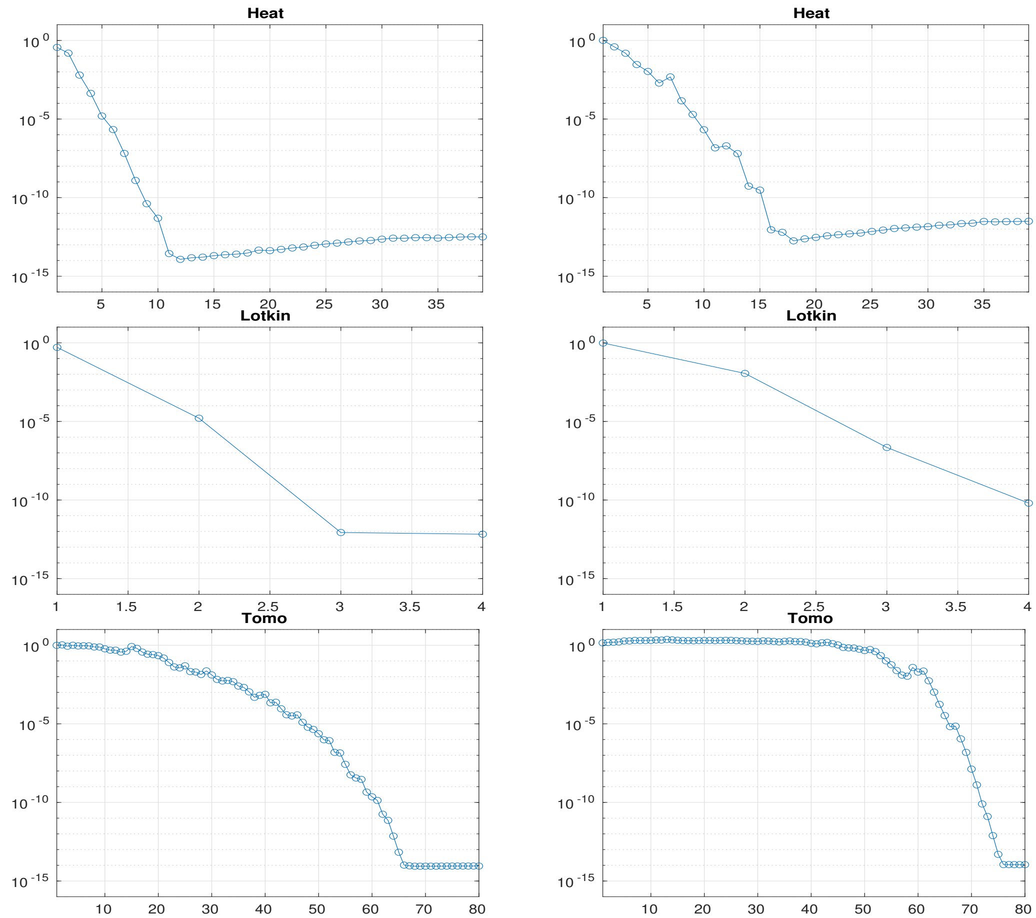

FIGURE 5 The graphs in the left-hand side column display the distance $R^{u,v}_{k,\lceil pk/3\rceil}$, versus $k = 1, \ldots, \lceil n/p\rceil$, between the space spanned by the first $\lceil pk/3\rceil$ singular vectors of the nonsymmetric test matrices and the first k Golub–Kahan block vectors. The right-hand side column shows the behavior of $R^{u,v}_{k,\lceil pk/2\rceil}$.

To investigate the convergence of subspaces, we introduce the orthogonal projectors

$$\mathcal{P}^L_k = I_n - U_k U_k^T, \qquad \mathcal{P}^R_k = I_n - W_k W_k^T,$$

where the matrices U_k and W_k contain the first k block columns of the matrices U_m and W_m, respectively, in the decompositions (7). To measure the distance between the spaces spanned by the singular vectors of A and those spanned by vectors computed with the BGKB method, we define the following merit index

$$R^{(u,v)}_{k,i} := \max\{\|\mathcal{P}^L_k U^{(1)}_i\|_F, \|\mathcal{P}^R_k V^{(1)}_i\|_F\}, \qquad k = 1, 2, \ldots, K, \quad i = 1, 2, \ldots, pk\,.$$

The quantities $R^{(u,v)}_{k,i}$ are displayed, for $k = \lceil n/(2p)\rceil$ ($k = K$ in case a breakdown occurred) and $i = 1, 2, \ldots, pk$, in the right-hand side column of Figure 3. The figures illustrate that the subspaces spanned by the first few columns determined by the block Lanczos and block Golub–Kahan algorithms are close to the subspace spanned by the first few eigenvectors and singular vectors, respectively, of the matrix A. Figure 5 depicts graphs for the quantities $R^{(u,v)}_{k,\lceil pk/3\rceil}$ and $R^{(u,v)}_{k,\lceil pk/2\rceil}$, for $k = 1, 2, \ldots, \lceil n/p\rceil$ ($k = 1, 2, \ldots, K$, if a breakdown occurred).

We finally illustrate the performances of the BLT and BGKB algorithms when applied to the solution of discrete ill-posed problems. Since we assume that the desired solution X_{true} is known, we first use it to elucidate that the solution subspaces determined by these algorithms can give approximations of X_{true} of as high quality as the solution subspaces defined by the truncated eigenvalue or singular value decompositions of A. Subsequently, we present examples that regularize the discrete ill-posed problem by Tikhonov regularization as described in Section 2.3. The latter examples do not require knowledge of X_{true} and show how applications to real-world discrete ill-posed problems can be carried out.

In the experiments reported in Table 1,the test matrices are of size 200×200 (400×400 for the `tomo` matrix), and the block size is $p = 5$. We measure the accuracy of the approximations of X_{true} determined by each regularization method by the relative error

$$E_{\text{method}} = \frac{\|X_{k_{\text{method}}} - X_{\text{true}}\|_F}{\|X_{\text{true}}\|_F} = \min_{k=1,2,\ldots,m} \frac{\|X_k - X_{\text{true}}\|_F}{\|X_{\text{true}}\|_F}, \tag{25}$$

which is obtained by choosing the value $k = k_{\text{method}}$ that minimizes the error in the computed solution. We remark that this approach to choosing k is not practical, but it shows the smallest possible error that can be determined by using the computed solution subspaces.

TABLE 1 Solution of symmetric linear systems: the errors E_{BLT} and E_{TEIG} are optimal for truncated block Lanczos iteration and truncated eigenvalue decomposition, the errors E_{BGKB} and E_{TSVD} are optimal for BGKB and truncated singular value decomposition (TSVD). The corresponding truncation parameters are denoted by k_{BLT}, k_{TEIG}, k_{BGKB}, and k_{TSVD}. The Tikhonov regularization parameter μ is presented in the 4th column. Three noise levels are considered; m denotes the number of iterations performed. The test matrices are of size 200×200 (400×400 for the `tomo` matrix).

Noise level	Matrix	m	μ	E_{BLT}	k_{BLT}	E_{TEIG}	k_{TEIG}
10^{-6}	`deriv2`	39	2.63×10^{-5}	3.48×10^{-3}	8	4.19×10^{-3}	81
	`gravity`	10	1.02×10^{-3}	1.24×10^{-3}	5	1.24×10^{-3}	15
	`phillips`	39	3.18×10^{-3}	4.17×10^{-4}	6	3.61×10^{-4}	29
10^{-4}	`deriv2`	39	3.54×10^{-4}	8.40×10^{-3}	5	9.29×10^{-3}	19
	`gravity`	10	1.96×10^{-2}	5.39×10^{-3}	5	4.96×10^{-3}	11
	`phillips`	39	2.87×10^{-2}	2.25×10^{-3}	4	1.69×10^{-3}	12
10^{-2}	`deriv2`	39	3.06×10^{-3}	2.58×10^{-2}	5	2.58×10^{-2}	5
	`gravity`	10	2.15×10^{-1}	2.59×10^{-2}	4	2.59×10^{-2}	7
	`phillips`	39	2.40×10^{-1}	9.66×10^{-3}	3	9.79×10^{-3}	7
Noise level	Matrix	m	μ	E_{BGKB}	k_{BGKB}	E_{TSVD}	k_{TSVD}
10^{-6}	`heat`	39	3.43×10^{-5}	1.87×10^{-2}	18	1.84×10^{-2}	79
	`lotkin`	4	2.99×10^{-5}	2.47×10^{-1}	3	2.29×10^{-1}	10
	`tomo`	79	2.46×10^{-6}	5.74×10^{-2}	79	3.16×10^{-2}	398
10^{-4}	`heat`	39	8.01×10^{-4}	2.83×10^{-2}	11	2.79×10^{-2}	37
	`lotkin`	4	4.17×10^{-3}	3.05×10^{-1}	2	3.03×10^{-1}	7
	`tomo`	79	3.48×10^{-2}	5.95×10^{-2}	79	3.79×10^{-2}	397
10^{-2}	`heat`	39	1.01×10^{-2}	8.49×10^{-2}	5	8.93×10^{-2}	15
	`lotkin`	4	2.13×10^{-1}	3.68×10^{-1}	3	3.71×10^{-1}	3
	`tomo`	79	9.22×10^{-1}	2.36×10^{-1}	60	2.29×10^{-1}	362

The upper part of Table 1 reports the approximate solutions obtained by truncated block Lanczos decomposition (5) and truncated eigenvalue decomposition, for test problems with symmetric matrices. The minimal error (25) obtained by applying the block Lanczos method and the truncated eigenvalue decomposition method, denoted by E_{BLT} and E_{TEIG}, respectively, are reported in the fifth and seventh columns. The truncation parameter values that produce the minimal errors are listed in the sixth and eighth columns. The third column shows how many block Lanczos iterations were executed; an entry smaller than 40 indicates that breakdown occurred. The results in Table 1 suggest that, for the test problems considered, the truncated block Lanczos projection method is able to produce solutions of essentially the same quality as truncated eigenvalue decomposition. We remark that the application of BLT is much cheaper than the evaluation of the truncated eigenvalue decomposition. We also remark that, since the best approximation of A of rank k is furnished by the k largest singular triplets of A, we may require more vectors to determine an accurate approximate solution when approximating A by block Lanczos vectors than when using singular triplets. On the other hand, since the singular triplets are independent of the right-hand side vector and the block Lanczos vectors are not, some examples require fewer block Lanczos vectors than singular triplets. Our interest in using block Lanczos vectors instead of singular vectors stems from the fact that the former are cheaper to compute.

The bottom part of Table 1 reports results obtained for nonsymmetric linear discrete ill-posed problems (6). Here, the block Golub–Kahan bidiagonalization method is compared to TSVD. This table shows that conclusions similar to those for symmetric matrices are obtained.

Table 1 shows the smallest achievable error. However, in real-world applications the exact solution is not known. We therefore complement these table with Table 2, which shows experiments in which the computed solutions are determined with the aid of the discrepancy principle. The matrices are of order 1000×1000 (1024×1024 for the `tomo` matrix), the block size is $p = 10$, and the truncation parameter $k = k_{\text{method}}$ is determined by applying the discrepancy principle (12).

Regularization by truncated iteration is not reliable, in general, for block methods. Therefore, in Table 2 the reduced problem is solved by Tikhonov regularization as was discussed in Section 2.3. The upper part of Table 2 shows that the solutions determined by using a few steps of the block Lanczos tridiagonalization are as accurate approximations of X_{true} as the solutions $\breve{X}_k$ computed with the aid of the full truncated eigenvalue decomposition method, while being much cheaper to evaluate. Similarly, the bottom part of Table 2 shows that the block Golub–Kahan bidiagonalization method produces solutions that are equivalent in quality to those obtained by TSVD, but are much cheaper to compute. Table 3 test different values of the block size p; the matrix size is 1000 for the `gravity` test matrix, and 1024 for the `tomo` problem.

TABLE 2 Comparison of the quality of computed solutions that are determined by truncated block Lanczos (BLT) and truncated eigenvalue decomposition (TEIG) methods (upper table), and by truncated BGKB and truncated singular value decomposition methods (bottom table). The truncation indexes k_{BLT}, k_{TEIG}, k_{BGKB}, and k_{TSVD}, are determined by the discrepancy principle (12). The test matrix is of size 1000×1000 for `gravity`, and of size 1024×1024 for `tomo`.

Noise level	Matrix	k_{BLT}	μ	E_{BLT}	k_{TEIG}	E_{TEIG}
	`deriv2`	7	1.24×10^{-5}	4.72×10^{-3}	88	5.08×10^{-3}
10^{-6}	`gravity`	3	8.24×10^{-4}	7.75×10^{-4}	14	5.61×10^{-4}
	`phillips`	4	2.05×10^{-3}	2.57×10^{-4}	26	2.87×10^{-4}
	`deriv2`	4	2.26×10^{-4}	1.03×10^{-2}	18	1.08×10^{-2}
10^{-4}	`gravity`	3	1.48×10^{-2}	3.66×10^{-3}	8	4.91×10^{-3}
	`phillips`	3	2.10×10^{-2}	1.61×10^{-3}	10	1.17×10^{-3}
	`deriv2`	2	2.81×10^{-3}	2.25×10^{-2}	4	1.82×10^{-2}
10^{-2}	`gravity`	3	1.28×10^{-2}	1.85×10^{-2}	6	1.41×10^{-2}
	`phillips`	2	2.60×10^{-2}	9.71×10^{-3}	6	9.02×10^{-3}
Noise level	Matrix	k_{BGKB}	μ	E_{BGKB}	k_{TSVD}	E_{TSVD}
	`heat`	7	1.44×10^{-5}	2.10×10^{-2}	76	2.10×10^{-2}
10^{-6}	`lotkin`	2	4.14×10^{-5}	1.74×10^{-1}	10	1.68×10^{-1}
	`tomo`	101	3.52×10^{-6}	3.08×10^{-2}	1018	2.09×10^{-2}
	`heat`	5	4.91×10^{-4}	2.94×10^{-2}	34	3.10×10^{-2}
10^{-4}	`lotkin`	1	3.17×10^{-3}	2.39×10^{-1}	6	2.41×10^{-1}
	`tomo`	89	3.34×10^{-2}	8.68×10^{-2}	1002	9.38×10^{-2}
	`heat`	2	8.11×10^{-3}	5.81×10^{-2}	12	6.29×10^{-2}
10^{-2}	`lotkin`	1	2.99×10^{-1}	3.42×10^{-1}	2	3.47×10^{-1}
	`tomo`	9	$1.80 \times 10^{+00}$	1.71×10^{-1}	656	1.94×10^{-1}

It is well known that block algorithms perform better than vector implementations on modern computers endowed with optimized basic linear algebra software. To illustrate this fact, we applied both the Lanczos and the block-Lanczos methods to the solution of p symmetric random linear systems of size 1000, letting $p = 5, 10, \dots, 50$. We let both the implementations of the Lanczos methods perform all the iterations allowed, that is, $1000/p$ for the block version and 1000 for the standard Lanczos method. The same was done for a 2000×1000 random linear least-squares problem, by applying BGKB and the LSQR methods. The computing times are reported in Figure 6. No breakdown occurred during the tests. The two graphs show that, as expected, while the execution time increases for the vector methods as the number of linear systems grows, the timings for the block algorithms first decreases, as the block size increases, and then stabilizes. Indeed, the time required for a block or a vector operations are roughly equivalent, and the number of iterations performed by the block algorithms decreases as the block size p increases.

TABLE 3 Comparison of the quality of computed solutions that are determined by the BLT and TEIG methods (upper table), and by the truncated BGKB and truncated singular value decomposition methods (bottom table), with different block sizes. The truncation indexes k_{BLT}, k_{TEIG}, k_{BGKB}, and k_{TSVD}, are determined by the discrepancy principle (12). The test matrix is of size 1000×1000 for `gravity`, and of size 1024×1024 for `tomo`.

Matrix	Noise level	Block size	μ	k_{BLT}	E_{BLT}	k_{TEIG}	E_{TEIG}
		10	1.42×10^{-2}	3	3.66×10^{-3}	8	4.91×10^{-3}
`gravity`	10^{-4}	20	1.46×10^{-2}	2	3.72×10^{-3}	10	3.26×10^{-3}
		30	1.45×10^{-5}	2	3.88×10^{-3}	10	3.33×10^{-3}

Matrix	Noise level	Block size	μ	k_{BGKB}	E_{BGKB}	k_{TSVD}	E_{TSVD}
		10	4.79×10^{-2}	75	4.22×10^{-2}	990	4.78×10^{-2}
`tomo`	10^{-4}	20	8.54×10^{-2}	48	4.72×10^{-2}	987	5.32×10^{-2}
		30	6.73×10^{-2}	32	4.27×10^{-2}	983	4.55×10^{-2}

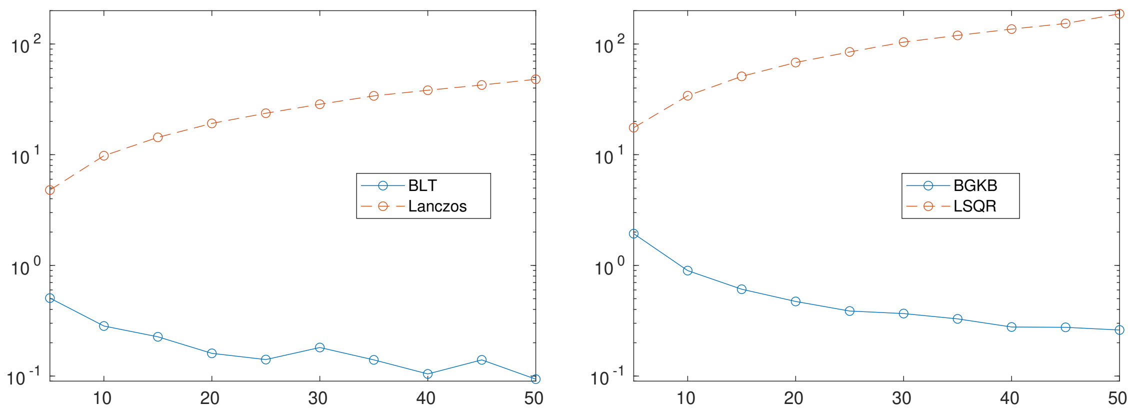

FIGURE 6 Computing times in seconds for solving p square 1000×1000 random symmetric linear systems by the Lanczos and the BLT methods (graph on the left), and for solving p random least squares problems of size 2000×1000 by the BGKB and LSQR methods (graph on the right), for $p = 5, 10, \dots, 50$.

Our last example is concerned with deblurring a color image. This example is generated with IR Tools[6]. We take as the true image, X_{true}, a subimage of the `tissue.png` test image available in MATLAB's Image Processing Toolbox. According to the RGB encoding, a color image can be represented as an array of $N \times N$ pixels in each one of the three channels representing red, green, and blue; see[13]. For this example, $N = 256$. We assume that each color channel of X_{true} has been contaminated by the same shaking blur having a Kronecker product structure. This is the so-called "within-channel" blur; we assume that there is no "cross-channel" blur. Under these assumptions, an approximation of X_{true} can be obtained by regularizing a block linear system of the form (1), where $n = N^2 = 65536$ and $p = 3$. More specifically, $X = [x^{(1)}, x^{(2)}, x^{(3)}]$, $B = [b^{(1)}, b^{(2)}, b^{(3)}] \in \mathbb{R}^{n \times 3}$, where $x^{(i)}$, $b^{(i)} \in \mathbb{R}^n$ are the vectorized images that appear in the ith channel, $i = 1, 2, 3$; see[2] for more details. The blurring matrix $A = K_1 \otimes K_2$ is generated by the following MATLAB instructions from IR Tools:

```
A = PRblurshake(256,opt); Kall = kronApprox(A); K1 = Kall.a{1}; K2 = Kall.b{1};
```

The image encoded in B is contaminated by Gaussian white noise E of level $\|E\|_F / \|B_{\text{true}}\|_F = 10^{-2}$. Exact and corrupted images are displayed in the leftmost and central frames of Figure 7, respectively.

The leftmost frame of Figure 8 displays, in logarithmic scale, the upper bound given in (23) as a function of the number of iterations. Despite this problem being large-scale, the quantities on the right-hand side of (23) can be easily computed by exploiting the Kronecker product structure of A.

The remaining frames of Figure 8 display the values of the relative error and the regularization parameter versus the number of iterations, for both the regularization method based on BGKB used together with Tikhonov regularization (see Section 2.3)

(a) (b) (c)

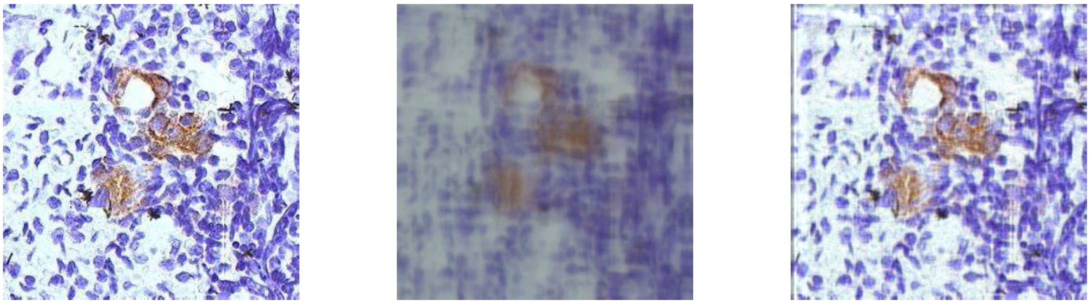

FIGURE 7 Color image deblurring test problem. (a) exact image; (b) blurred & noisy image; (c) restored image computed by BGKB and Tikhonov regularization (relative error 1.22×10^{-1}, regularization parameter 4.94×10^{-2}).

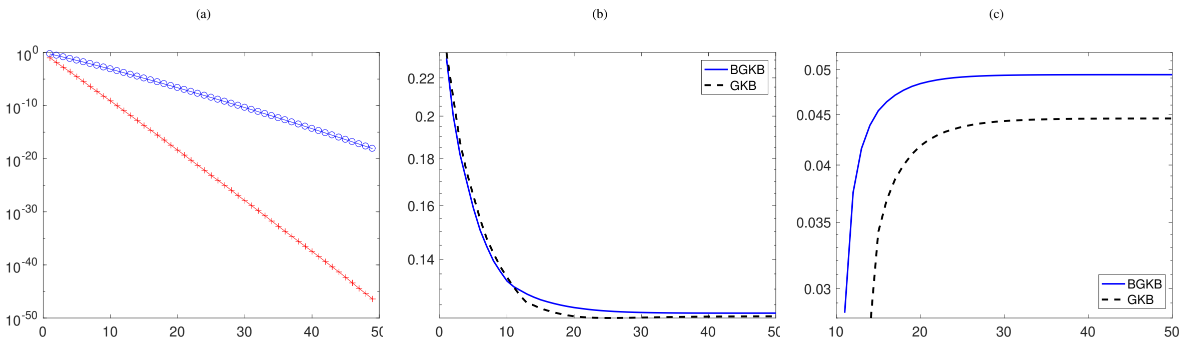

FIGURE 8 Color image deblurring test problem. (a) bound in (23) versus number of iterations (the left-hand and right-hand sides of (23) are represented by crosses and circles, respectively); (b) relative errors versus number of iterations for methods based on BGKB and the classical GKB; (c) regularization parameters versus number of iterations for methods based on BGKB and the classical GKB.

and for a classical regularization method based on GKB, i.e., Golub–Kahan bidiagonalization with block size one, and Tikhonov regularization; see, e.g.,[3, 6, 7] for discussions of this solution method. Running the methods based on GKB and BGKB took 5.4 and 1.7 seconds, respectively (note that, in order to compare approximation subspaces of the same dimension, 150 GKB and 50 BGKB iterations were performed).

5 | CONCLUSION

This paper applies a few steps of the block Lanczos or the block Golub–Kahan bidiagonalization methods to large discrete ill-posed problem to determine the solution by solving a projected problem of fairly small size. The eigenvalues or singular values of the projected matrix are shown to be accurate approximations of the corresponding largest eigenvalues (in absolute value) or singular values of the discrete ill-posed problem, respectively. The same result holds for the corresponding eigenvectors and singular vectors. This suggests that in order to determine a solution of a given large discrete ill-posed problem, it often suffices to use a partial Lanczos block tridiagonalization or a partial Golub–Kahan block bidiagonalization, instead of computing partial spectral or singular value decompositions. This is advantageous because the computation of a partial Lanczos block tridiagonalization or a partial Golub–Kahan block bidiagonalization is much cheaper. Computed examples provide illustrations.

ACKNOWLEDGEMENTS

The authors would like to thank the two anonymous referees for their insightful comments that lead to improvements of the presentation. The work of SG was partially supported by EPSRC, under grant EP/T001593/1. Work by LR was supported in part by NSF grants DMS-1720259 and DMS-1729509. The work of GR was partially supported by the Fondazione di Sardegna 2017 research project "Algorithms for Approximation with Applications [Acube]", the INdAM-GNCS research project "Tecniche numeriche per l'analisi delle reti complesse e lo studio dei problemi inversi", and the Regione Autonoma della Sardegna research project "Algorithms and Models for Imaging Science [AMIS]" (RASSR57257, intervento finanziato con risorse FSC 2014-2020 - Patto per lo Sviluppo della Regione Sardegna). This study does not have any conflicts to disclose.

References

1. A. H. Bentbib, M. El Guide, and K. Jbilou, *The block Lanczos algorithm for linear ill-posed problems*, Calcolo, 54 (2017), pp. 711–732.

2. A. H. Bentbib, M. El Guide, K. Jbilou, E. Onunwor, and L. Reichel, *Solution methods for linear discrete ill-posed problems for color image restoration*, BIT Numer. Math., 58 (2018), pp. 555–578.

3. D. Calvetti and L. Reichel, *Tikhonov regularization of large linear problems*, BIT Numer. Math., 43 (2003), pp. 263–283.

4. C. Fenu, L. Reichel, and G. Rodriguez, *GCV for Tikhonov regularization via global Golub–Kahan decomposition*, Numer. Linear Algebra Appl., 23 (2016), pp. 467–484.

5. C. Fenu, L. Reichel, G. Rodriguez, and H. Sadok, *GCV for Tikhonov regularization by partial SVD*, BIT Numer. Math., 57 (2017), pp. 1019–1039.

6. S. Gazzola, P. C. Hansen, and J. G. Nagy, *IR Tools: A MATLAB package of iterative regularization methods and large-scale test problems*, Numer. Algorithms, 81 (2019), pp. 773–811.

7. S. Gazzola, P. Novati, and M. R. Russo, *On Krylov projection methods and Tikhonov regularization*, Electron. Trans. Numer. Anal., 44 (2015), pp. 83–123.

8. S. Gazzola, E. Onunwor, L. Reichel, and G. Rodriguez, *On the Lanczos and Golub–Kahan reduction methods applied to discrete ill-posed problems*, Numer. Linear Algebra Appl., 23 (2016) pp. 187–204.

9. S. Gazzola and M. Sabatè Landman, *Krylov methods for inverse problems: Surveying classical, and introducing new, algorithmic approaches*, GAMM-Mitteilungen (2020). https://doi.org/10.1002/gamm.202000017.

10. G. H. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed., Johns Hopkins University Press, Baltimore, 2013.

11. P. C. Hansen, *Rank-Deficient and Discrete Ill-Posed Problems*, SIAM, Philadelphia, 1998.

12. P. C. Hansen, *Regularization tools version 4.0 for MATLAB 7.3*, Numer. Algorithms, 46 (2007) 189–194.

13. P. C. Hansen, J. Nagy, and D. P. O'Leary, *Deblurring Images: Matrices, Spectra, and Filtering*, SIAM, Philadelphia, 2006.

14. S. Kindermann, *Convergence analysis of minimization-based noise level-free parameter choice rules for linear ill-posed problems*, Electron. Trans. Numer. Anal., 38 (2011), pp. 233–257.

15. S. Kindermann and K. Raik, *A simplified L-curve method as error estimator*, Electron. Trans. Numer. Anal., 53 (2020), pp. 217–238.

16. M. Kubínová and K. M. Soodhalter, *Admissible and attainable convergence behavior of block Arnoldi and GMRES*, SIAM J. Matrix Anal. Appl., 41 (2020), pp. 464–486.

17. V. A. MOROZOV, *On the solution of functional equations by the method of regularization*, Soviet Math. Dokl., 7 (1966), pp. 414–417.

18. L. REICHEL AND G. RODRIGUEZ, *Old and new parameter choice rules for discrete ill-posed problems*, Numer. Algorithms, 63 (2013), pp. 65–87.

$\square$