AIMS Mathematics, 11(11): 37909–37933.
DOI: 10.3934/math.20261465
Received: 12 September 2026
Revised: 05 November 2026
Accepted: 09 November 2026
Published: 13 November 2026

https://www.aimspress.com/journal/Math

Research article

Enlarged GKB for selecting the regularization parameter in an ℓ_1-regularized problem

Abdulmajeed Alsubhi*

Department of Mathematics, Faculty of Science, Islamic University of Madinah, Madinah 42351, Saudi Arabia

* **Correspondence:** Email: ahalsubhi@iu.edu.sa.

Abstract: ℓ_1 regularization is widely used to preserve important features of the desired solution in ill-posed inverse problems. We employed the split Bregman (SB) algorithm, which reformulates the ℓ_1-regularized problem as a sequence of unconstrained optimization problems involving generalized Tikhonov minimization. Selecting the regularization parameter at each SB iteration was crucial for obtaining accurate reconstructions. To address this challenge, we transformed each generalized Tikhonov problem into standard form and determined the regularization parameter using the singular value decomposition (SVD) of a transformed coefficient matrix, thereby improving both stability and efficiency. An enlarged Krylov subspace strategy based on Golub–Kahan bidiagonalization (EGKB) was employed to compute an approximate truncated SVD, with the EGKB implemented in single precision to further reduce computational cost. We extended the unbiased predictive risk estimator (UPRE) to the transformed problem and used it to estimate stabilizing regularization parameters throughout the SB iterations. A theoretical analysis of the computational complexity was presented to demonstrate the efficiency of the proposed framework. Numerical experiments on X-ray tomographic problems demonstrated that EGKB-UPRE achieves reconstruction quality comparable to GSVD-based methods while substantially reducing computational time and storage requirements.

Keywords: ℓ_1 regularization; split Bregman; Golub–Kahan bidiagonalization; unbiased predictive risk estimation; single and double precision
Mathematics Subject Classification: 65F22, 65F10, 68W40

1. Introduction

We are concerned with the solution of linear discrete inverse problems of the form

$$A\mathbf{x}^{\text{true}} \approx \mathbf{b}, \tag{1.1}$$

where $A \in \mathbb{R}^{m\times n}$ is a large ill-conditioned matrix, $\mathbf{x}^{\text{true}} \in \mathbb{R}^{n\times 1}$ is an unknown solution, and $\mathbf{b} \in \mathbb{R}^{m\times 1}$ is measured data typically with Gaussian white noise $\boldsymbol{\eta}$, i.e., $\mathbf{b} = \mathbf{b}^{\text{true}} + \boldsymbol{\eta}$, where $\mathbf{b}^{\text{true}}$ is the exact noise-free data vector that achieves $A\mathbf{x}^{\text{true}} = \mathbf{b}^{\text{true}}$. Here we assume a bound $\|\boldsymbol{\eta}\| \leq \delta$, with $\delta > 0$, is available. Due to the ill-conditioning of A and the presence of noise in $\mathbf{b}$, direct inversion of (1.1) is generally unstable and fails to yield a meaningful approximation of $\mathbf{x}^{\text{true}}$. Therefore, the problem must be replaced by an appropriately stabilized well-posed formulation [1, 2]. For example, a least squares solution of the Tikhonov minimization

$$\mathbf{x}^* = \arg\min_{\mathbf{x}}\{\|A\mathbf{x} - \mathbf{b}\|_2^2 + \mu^2\|L\mathbf{x}\|_2^2\} \tag{1.2}$$

gives an approximation of $\mathbf{x}^{\text{true}}$ in (1.1), where $\mu > 0$ is a regularization parameter that balances the contribution of the fitting of data, $\|A\mathbf{x} - \mathbf{b}\|_2^2$, against the regularization term, $\|L\mathbf{x}\|_2^2$, and $L \in \mathbb{R}^{n\times n}$ is a regularization matrix. If $L = I$, where I is an appropriately sized identity matrix, then (1.2) is in its standard form and μ impacts the norm of the desired solution. When $L \neq I$, as in cases where L is chosen as the discretization of a derivative operator, then (1.2) is in its general form and μ influences the smoothness of the solution.

We assume that

$$\mathcal{N}(A) \cap \mathcal{N}(L) = \{0\},$$

where $\mathcal{N}(T)$ denotes the null space of a matrix T. This condition guarantees that (1.2) has a unique solution for any $\mu > 0$, [3]. For small-scale problems, one could use the generalized SVD (GSVD) of the matrix pair $\{A, L\}$ to obtain a joint decomposition and solve (1.2). The GSVD, however, is computationally unattractive for large problems, as it requires substantial storage and computations. It was shown in [4] that (1.2) is equivalent to the minimization problem in standard form,

$$\arg\min_{\tilde{\mathbf{x}}}\{\|AL_A^\dagger\tilde{\mathbf{x}} - \mathbf{b}\|_2^2 + \mu^2\|\tilde{\mathbf{x}}\|_2^2\}, \tag{1.3}$$

where $L_A^\dagger$ is the A-weighted generalized inverse of L, defined by $L_A^\dagger = (I - (A(I - L^\dagger L))^\dagger A)L^\dagger$. The A-weighted generalized inverse $L_A^\dagger$ permits the standard-form transformation to accommodate more general regularization operators, including rank-deficient ones. Extending the present framework to this more general setting is beyond the scope of this study and is left for future work.

In this work, we restrict the numerical framework to the case in which L is invertible. In this setting, $L_A^\dagger = L^{-1}$, and hence (1.3) reduces to the standard-form problem with $\tilde{A} = AL^{-1}$. The solution of (1.2) is then recovered from the solution of (1.3) as $\mathbf{x} = L^{-1}\tilde{\mathbf{x}}$. Although the Tikhonov problem (1.2) admits a closed-form solution, it generally produces overly smooth reconstructions. Consequently, its effectiveness may be limited when the underlying solution contains sharp edges, as is often the case in practical imaging applications [5, 6].

To accommodate a more general class of problems, we consider the ℓ_1-regularized minimization problem

$$\mathbf{x} = \arg\min_{\mathbf{x}}\{\frac{1}{2}\|A\mathbf{x} - \mathbf{b}\|_2^2 + \beta\|L\mathbf{x}\|_1\}, \tag{1.4}$$

where β is a regularization parameter and L is usually selected as a discrete gradient operator so that $\|L\mathbf{x}\|_1$ represents the total variation of $\mathbf{x}$, see [7]. Unlike the Tikhonov solution in (1.2), the solution of (1.4) generally does not admit a closed-form expression and depends on both L and β. Consequently,

it must be computed iteratively. There are many iterative algorithms to find $\mathbf{x}$, however, all of which require solving an optimization problem at each iteration, see [7–12]. For instance, in [9], the iteratively reweighted norm approximation reformulates the original problem as a sequence of standard-form Tikhonov problems, each solved using a flexible GMRES algorithm. The majorization-minimization framework of [11], which is also employed in [13], generates a sequence of least-squares problems that can be solved efficiently using generalized Krylov subspace methods. When generalized Tikhonov problems arise at each iteration, they can be solved using iterative algorithms, thereby enabling the application of parameter selection techniques for the associated regularization parameters [9, 11, 14]. Comprehensive surveys of hybrid iterative regularization methods and parameter-selection strategies for ill-posed inverse problems can be found in [15]. Here, we use the split Bregman (SB) algorithm [10] with one ℓ_1 regularization term to approximate the solution of (1.4). The inner-generalized Tikhonov problem at each iteration in the SB is solved using the IRcgls algorithm with the original matrix A [16].

It can be difficult to estimate an appropriate regularization parameter in (1.4) when A is large and has no suitable structure [17, 18]. In such cases, the Kronecker product sum approximation to A could be used [19]. We refer to [20–23] for more details on the construction of the Kronecker product sum approximation. However, the complexity introduced by the multiple-term KP structure poses challenges in effectively managing the regularization parameter.

Therefore, we choose to transform the inner–generalized Tikhonov problem arising in the solution of (1.4) into a standard form. This transformation facilitates the use of the SVD of $\tilde{A}$ rather than the GSVD of $\{A, L\}$. If the noise variance in the data is available, the value of the regularization parameter can be determined using the discrepancy principle [24, 25] or the unbiased predictive risk estimation (UPRE) [18]. Otherwise, it can be estimated using methods such as generalized cross-validation (GCV) [25] or the L-curve method [25, 26]. In this paper, we focus on employing the UPRE; however, the methodology can be adapted to fit any parameter selection technique.

However, computing the full SVD of large $\tilde{A}$ is unattractive. Experimental results presented in [27] showed that using the enlarged Krylov space implemented within the Golub–Kahan bidiagonalization algorithm, termed EGKB, significantly improves the quality of the desired TSVD approximation. The benefits of taking extra steps with a Krylov subspace method were noted in [28] for the Arnoldi–Tikhonov regularization. Inspired by this, the EGKB is considered for $\tilde{A}$. The idea of applying projection to the SB is not new; see [29–31] for more discussion. We focus on estimating the regularization parameter for the inner-generalized Tikhonov problem by transforming it to standard form and projecting it onto a low-dimensional subspace, where parameter-choice methods can be applied efficiently. Moreover, we explore the potential of using single precision (SP) for EGKB to reduce both storage requirements and computational costs [32].

1.1. Major contributions

- An implicit transformation of the inner-generalized Tikhonov problem arising in split Bregman iterations into standard form, allowing parameter estimation through the SVD of $\tilde{A} = AL^{-1}$ instead of the GSVD of $\{A, L\}$.
- Development of a projected UPRE criterion for the transformed problem using EGKB-generated TSVD approximations.
- An efficient mixed-precision implementation where EGKB is computed in single precision while the split Bregman iterations remain in double precision.

- A computational complexity analysis demonstrating the reduction in cost compared with GSVD-based parameter selection.
- Numerical validation on tomography problems.

1.2. Organization

The remainder of this paper is structured as follows. Section 2 reviews the essential steps of the SB iterative algorithm. Section 3 presents the EGKB algorithm that is used to form TSVD approximations for any matrix. Section 4 develops the UPRE method to determine the regularization parameter derived from transforming the inner-generalized Tikhonov problem in the SB to standard form. Section 5 provides a theoretical analysis of the computational costs of selecting the regularization parameter in the SB, both using the transformation with the approximation to $\tilde{A}$ and with the GSVD of the pair $\{A, L\}$. The computed results are presented in Section 6, and the conclusions can be found in Section 7.

2. Iterative techniques for ℓ_1-regularized problems

This section is concerned with the solution of the ℓ_1-regularized problem (1.4) obtained using the SB algorithm of Goldstein and Osher [10]. The constrained problem of (1.4) is written as

$$\min_{\mathbf{x},\mathbf{d}}\{\frac{1}{2}\|A\mathbf{x}-\mathbf{b}\|_2^2+\beta\|\mathbf{d}\|_1\} \quad \text{such that} \quad \mathbf{d}=L\mathbf{x}. \tag{2.1}$$

The SB algorithm uses the augmented Lagrangian approach to convert (2.1) to the unconstrained optimization problem

$$(\mathbf{x},\mathbf{d})=\arg\min_{\mathbf{x},\mathbf{d}}\{\frac{1}{2}\|A\mathbf{x}-\mathbf{b}\|_2^2+\frac{\lambda^2}{2}\|L\mathbf{x}-\mathbf{d}\|_2^2+\beta\|\mathbf{d}\|_1\}, \tag{2.2}$$

where $\lambda > 0$ is a penalization parameter. Following [10], the SB algorithm for solving (2.2) computes alternating updates for $\mathbf{x}$ and $\mathbf{d}$ given by

$$\mathbf{x}^{(\ell+1)}=\arg\min_{\mathbf{x}}\{\frac{1}{2}\|A\mathbf{x}-\mathbf{b}\|_2^2+\frac{(\lambda^{(\ell)})^2}{2}\|L\mathbf{x}-\mathbf{h}^{(\ell)}\|_2^2\}, \tag{2.3}$$

$$\mathbf{d}^{(\ell+1)}=\arg\min_{\mathbf{d}}\{\frac{(\lambda^{(\ell)})^2}{2}\|\mathbf{d}-(L\mathbf{x}^{(\ell+1)}+\mathbf{g}^{(\ell)})\|_2^2+\beta^{(\ell)}\|\mathbf{d}\|_1\}, \tag{2.4}$$

where $\mathbf{h}^{(\ell)}=\mathbf{d}^{(\ell)}-\mathbf{g}^{(\ell)}$. The Bregman variable is updated according to

$$\mathbf{g}^{(\ell+1)}=\mathbf{g}^{(\ell)}+L\mathbf{x}^{(\ell+1)}-\mathbf{d}^{(\ell+1)}. \tag{2.5}$$

Throughout, we use the superscript ℓ to represent the value of a variable at the ℓ^{th} SB iteration. Defining $\gamma=\beta^{(\ell)}/(\lambda^{(\ell)})^2$, the update in (2.4) can be written as

$$\mathbf{d}^{(\ell+1)}=\arg\min_{\mathbf{d}}\{\frac{1}{2}\|\mathbf{d}-(L\mathbf{x}^{(\ell+1)}+\mathbf{g}^{(\ell)})\|_2^2+\gamma\|\mathbf{d}\|_1\}. \tag{2.6}$$

The solutions of (2.6) can be written using the shrinkage formulas found in [10, Eqs (4.4)–(4.6)],

$$(\mathbf{d}^{(\ell+1)})_i=\text{shrink}((L\mathbf{x}^{(\ell+1)})_i+\mathbf{g}_i^{(\ell)},\gamma), \tag{2.7}$$

where $\text{shrink}(w,\gamma)=\text{sign}(w)\max(|w|-\gamma,0)$.

In summary, the SB algorithm for the ℓ_1-regularized problem (1.4) is given in Algorithm 1.

Algorithm 1 [10] **Split Bregman method for the ℓ_1-regularized problem** (1.4)

Require: A, $\mathbf{b}$, L, λ, β, $L_{\max}$, τ_{SB}, $I_{\max}$, τ_{IRcgls}
Ensure: $\mathbf{x}^{(\ell)}$
1: Initialize $\ell = 0$, $\mathbf{x}^{(0)} = \mathbf{0}$, $\mathbf{d}^{(0)} = \mathbf{0}$, $\mathbf{g}^{(0)} = \mathbf{0}$
2: **while** not converged and $\ell < L_{\max}$ **do**
3: Update $\mathbf{x}^{(\ell+1)}$ by solving (2.3)
4: Update $\mathbf{d}^{(\ell+1)}$ using (2.7)
5: Update $\mathbf{g}^{(\ell+1)}$ using (2.5)
6: **if** $\|\mathbf{x}^{(\ell+1)} - \mathbf{x}^{(\ell)}\|_2 / \|\mathbf{x}^{(\ell)}\|_2 \leq \tau_{\text{SB}}$ **then**
7: **break**
8: **end if**
9: $\ell = \ell + 1$
10: **end while**

2.1. Computational cost of the SB algorithm

The computational cost of Algorithm 1 is dominated by step 3, namely, the computation of $\mathbf{x}$ as the solution of the generalized Tikhonov problem (2.3) at each SB iteration. As discussed in Section 1, a variety of approaches can be used to compute $\mathbf{x}^{(\ell)}$. However, our focus is on demonstrating the effectiveness of our proposed method for selecting an appropriate λ in (2.3). Therefore, rather than focusing on the choice of solver, we employ the IRcgls algorithm from the IR Tools package [16] to solve the generalized Tikhonov subproblem. The computational cost of IRcgls is primarily determined by the matrix-vector products involving A and $A^\top$. For a matrix $A \in \mathbb{R}^{m\times n}$, this requires approximately $O(4mn\iota)$ flops, where ι denotes the number of iterations to convergence. Since L often corresponds to a sparse finite-difference operator, its contribution to the computational cost in IRcgls is negligible.

To summarize the computational cost of the SB algorithm, let ℓ_{end} denote the number of SB iterations required for convergence and define

$$\iota_{\text{total}} = \sum_{\ell=1}^{\ell_{\text{end}}} \iota_\ell,$$

where ι_ℓ is the number of IRcgls iterations performed at the ℓth SB iteration. Neglecting the costs of the updates (2.5) and (2.7), the overall cost of the SB algorithm is dominated by the matrix-vector products with A and $A^\top$ arising in the solution of the generalized Tikhonov subproblems.

3. Enlarged GKB

The GKB algorithm was introduced in [33] to reduce a general matrix to bidiagonal form. Applying k steps of GKB to the matrix A with starting vector $\mathbf{b}$ generates two related decompositions

$$AQ_k = S_{k+1}C_k,$$
$$A^\top S_{k+1} = Q_{k+1}C_k^\top,$$

where $Q_k \in \mathbb{R}^{n\times k}$ and $S_{k+1} \in \mathbb{R}^{m\times k+1}$ have orthonormal columns, and

$$C_k = \begin{bmatrix} \alpha_1 & & & & \\ t_2 & \alpha_2 & & & \\ & \ddots & \ddots & & \\ & & & t_k & \alpha_k \\ & & & & t_{k+1} \end{bmatrix} \in \mathbb{R}^{(k+1)\times k} \tag{3.1}$$

is a lower bidiagonal matrix.

The SVD of C_k yields a rank-k truncated SVD (TSVD) approximation of A. To improve this approximation, we continue the GKB process beyond the desired rank k and construct an enlarged Krylov subspace of dimension $k_p = k + p$, where p denotes the number of additional GKB steps beyond the desired rank k. This yields the enlarged projected matrix $\overline{A}_{k_p}$ that captures both the dominant and the smallest singular values of A [34], thereby providing useful information about the spectral properties and conditioning of A. Let $\tilde{\sigma}_i$ denote the ith singular value of $\overline{A}_{k_p}$, ordered from largest to smallest. Then, for $i = 1, \ldots, k$, the singular values $\tilde{\sigma}_i$ typically provide accurate approximations to the dominant singular values σ_i of A, whereas the singular values corresponding to indices $i \gg k$ approximate the smallest singular values of A [34].

The strategy of enlarging the projected subspace was proposed in [35, Section 3.3.1], and subsequently employed for regularization parameter estimation in [14] and to improve reconstruction quality in challenging geophysical applications in [36]. More recently, this approach was used in [27] to improve the accuracy of the Kronecker product approximation for general matrices.

The steps of the enlarged GKB (EGKB) approach for finding an approximate rank-k TSVD are given in Algorithm 2.

To preserve the orthogonality of the columns of Q, one-sided reorthogonalization is performed at step 13, following the approach of [37]. While an efficient reorthogonalization approach based on Cholesky QR decomposition was proposed in [38], the projected dimension k_p is assumed to be sufficiently small such that modified Gram–Schmidt (MGS) reorthogonalization is computationally inexpensive and therefore adopted in step 13.

When $k_p = k + p$, the steps of EGKB are applied to B, and a rank-k TSVD approximation is obtained from the enlarged projected matrix, where p is typically small but may increase with k. Selecting an appropriate value of k_p may be accomplished by examining the convergence of the singular values as k_p increases. Such an approach, however, can be inefficient because it requires repeated SVD computations for different values of k_p. Alternatively, [27, 39] suggest determining k based on the monotonic decay of the geometric mean of the products of the elements in (3.1). However, these approaches may lead to underestimating k, which often results in a poor TSVD approximation of B. Therefore, we use the discrepancy principle as the stopping criterion in Algorithm 2. Specifically, the iteration is terminated at the first index k satisfying

$$k = \arg\min_j \{\|C_j \mathbf{y}_j - t_1 \mathbf{e}_1^{(j+1)}\|_2 < \delta\, \tau_{EGKB}\}. \tag{3.2}$$

Here, $\mathbf{e}_1^{(j+1)}$ denotes the first canonical unit vector in $\mathbb{R}^{j+1}$, δ is the noise bound, and $\tau_{\text{EGKB}} > 1$ is a user-defined parameter typically chosen close to 1. In our experiments, we set $\tau_{\text{EGKB}} = 1.01$. Finally, the computed TSVD approximation is truncated to k terms giving the desired $\bar{B}_k$.

Algorithm 2 EGKB: Computation of a rank-k TSVD approximation

Require: $B \in \mathbb{R}^{m\times n}$, $\mathbf{z} \in \mathbb{R}^m$, $k_{\max}$, p, δ, $\tau_{\text{EGKB}} > 1$, prec $\in \{\text{SP}, \text{DP}\}$
Ensure: $\bar{B}_k = \bar{U}_k \bar{\Sigma}_k \bar{V}_k^\top$

$t_1 = \|\mathbf{z}\|_2$, $\mathbf{s}_1 = \mathbf{z}/t_1$, $\mathbf{q} = B^\top \mathbf{s}_1$, $\alpha_1 = \|\mathbf{q}\|_2$, $\mathbf{q}_1 = \mathbf{q}/\alpha_1$, $k = k_{\max}$
for $j = 2, \ldots, k_{\max} + 1$ **do**
    $\mathbf{s}_j = B\mathbf{q}_{j-1} - \alpha_{j-1}\mathbf{s}_{j-1}$
    if prec = SP **then**
        Reorthogonalize $\mathbf{s}_j$ against S_{j-1} using MGS
    end if
    $t_j = \|\mathbf{s}_j\|_2$;
    if $t_j = 0$ **then**
        $k = j - 1$; **break**
    end if
    $\mathbf{s}_j = \mathbf{s}_j/t_j$
    $\mathbf{q}_j = B^\top \mathbf{s}_j - t_j \mathbf{q}_{j-1}$
    Reorthogonalize $\mathbf{q}_j$ against Q_{j-1} using MGS
    $\alpha_j = \|\mathbf{q}_j\|_2$;
    if $\alpha_j = 0$ **then**
        $k = j - 1$; **break**
    end if
    $\mathbf{q}_j = \mathbf{q}_j/\alpha_j$
    Form C_j and compute the projected solution $\mathbf{y}_j$
    if $\|C_j \mathbf{y}_j - t_1 \mathbf{e}_1^{(j+1)}\|_2 \le \tau_{\text{EGKB}}\delta$ **then**
        $k = j$; **break**
    end if
end for
$k_p = \min\{k + p, m, n\}$ and continue GKB to k_p steps to obtain S_{k_p+1}, Q_{k_p}, and C_{k_p}
Compute the SVD $C_{k_p} = U_{k_p} \bar{\Sigma}_{k_p} V_{k_p}^\top$, where $U_{k_p} \in \mathbb{R}^{(k_p+1)\times k_p}$, $\bar{\Sigma}_{k_p} \in \mathbb{R}^{k_p\times k_p}$, and $V_{k_p} \in \mathbb{R}^{k_p\times k_p}$
Form $\bar{U}_{k_p} = S_{k_p+1} U_{k_p} \in \mathbb{R}^{m\times k_p}$, $\bar{V}_{k_p} = Q_{k_p} V_{k_p} \in \mathbb{R}^{n\times k_p}$
Truncate $\bar{U}_{k_p}$, $\bar{\Sigma}_{k_p}$, and $\bar{V}_{k_p}$ to their first k singular components to obtain $\bar{U}_k$, $\bar{\Sigma}_k$, and $\bar{V}_k$
$\bar{B}_k = \bar{U}_k \bar{\Sigma}_k \bar{V}_k^\top$

3.1. Computational cost of $\bar{B}_k$

We present the computational cost of constructing an approximate TSVD of a general matrix $B \in \mathbb{R}^{m\times n}$. We assume that B has no special structure and that the operations with B and $B^\top$ are performed using the standard approach. The cost of computing a rank-k approximate TSVD using EGKB with oversampling parameter p was derived in [36, Eq (11)] and is given by

$$\text{Cost with Algorithm 2: } O\left(4mnk_p + 2mk^2 + 2n(k^2 + k_p^2) + 12k_p^3\right) \approx O(4mnk_p). \tag{3.3}$$

4. Regularization parameter for the SB

In this section, we provide a brief overview of the derivations of UPRE. We then extend the UPRE to estimate the regularization parameter in the SB.

4.1. Unbiased predictive risk estimation

We derive UPRE [18] for the inner-generalized Tikhonov problem (2.3). For notational convenience, we will drop the superscripts ℓ. Let $\mathbf{h} = L\bar{\mathbf{h}}$, and then (2.3) can be written as

$$\mathbf{y}_\lambda = \arg\min_{\mathbf{y}}\{\frac{1}{2}\|A\mathbf{y} - \bar{\mathbf{b}}\|_2^2 + \frac{\lambda^2}{2}\|L\mathbf{y}\|_2^2\}, \quad \bar{\mathbf{b}} = \mathbf{b} - A\bar{\mathbf{h}}, \tag{4.1}$$

with the solution $\mathbf{y}_\lambda = A_{L_\lambda}\bar{\mathbf{b}}$, where A_{L_λ} is the regularized inverse matrix defined by $A_{L_\lambda} = (A^\top A + \lambda^2 L^\top L)^{-1}A^T$. The solution of (2.3) can be obtained from the solution of (4.1) via $\mathbf{x} = \mathbf{y}_\lambda + \mathbf{h}$.

The predictive error for $\mathbf{y}_\lambda$ is defined by

$$\mathbf{p}_\lambda = A\mathbf{y}_\lambda - \bar{\mathbf{b}}^{\text{true}} = AA_{L_\lambda}\bar{\mathbf{b}} - \bar{\mathbf{b}}^{\text{true}} = (A_\lambda^\dagger - I_m)\bar{\mathbf{b}}^{\text{true}} + A_\lambda^\dagger\boldsymbol{\eta}, \tag{4.2}$$

where $\bar{\mathbf{b}}^{\text{true}} = \mathbf{b}^{\text{true}} - A\bar{\mathbf{h}}$ and $A_\lambda^\dagger = AA_{L_\lambda}$ is the influence matrix. The regularized residual of $\mathbf{y}_\lambda$ can be written as

$$\mathbf{r}_\lambda = A\mathbf{y}_\lambda - \bar{\mathbf{b}} = (A_\lambda^\dagger - I_m)\bar{\mathbf{b}} = (A_\lambda^\dagger - I_m)\bar{\mathbf{b}}^{\text{true}} + (A_\lambda^\dagger - I_m)\boldsymbol{\eta}. \tag{4.3}$$

Applying [18, Lemma 7.2] to (4.2) and (4.3), respectively, yields

$$\mathbb{E}\left(\frac{1}{m}\|\mathbf{p}_\lambda\|_2^2\right) = \frac{1}{m}\|(A_\lambda^\dagger - I_m)\bar{\mathbf{b}}^{\text{true}}\|_2^2 + \frac{\rho^2}{m}\text{trace}(A_\lambda^{\dagger\top}A_\lambda^\dagger) \tag{4.4}$$

and

$$\mathbb{E}\left(\frac{1}{m}\|\mathbf{r}_\lambda\|_2^2\right) = \frac{1}{m}\|(A_\lambda^\dagger - I_m)\bar{\mathbf{b}}^{\text{true}}\|_2^2 + \frac{\rho^2}{m}\text{trace}\left((A_\lambda^\dagger - I_m)^\top(A_\lambda^\dagger - I_m)\right), \tag{4.5}$$

where ρ^2 denotes the variance of the additive white Gaussian noise and trace(C) is the trace of a matrix C. In the numerical experiments, where both $\bar{\mathbf{b}}$ and the noise-free data $\mathbf{b}^{\bar{\text{true}}}$ are available, the variance is estimated by

$$\rho^2 = \frac{\|\bar{\mathbf{b}} - \mathbf{b}^{\bar{\text{true}}}\|_2^2}{m}. \tag{4.6}$$

When $\mathbf{b}^{\bar{\text{true}}}$ is unavailable, ρ^2 may be estimated directly from the measured data using standard techniques, such as residual-based estimators or discrepancy-principle-based approaches; see, e.g., [41]. The resulting estimate can then be incorporated into the UPRE functional in place of the true variance, allowing the proposed parameter-selection strategy to be applied even when the noise level is not known a priori.

Using the symmetry of $A_\lambda^\dagger$ and the linearity of the trace, we rewrite (4.5) as

$$\mathbb{E}\left(\frac{1}{m}\|\mathbf{r}_\lambda\|_2^2\right) = \frac{1}{m}\|(A_\lambda^\dagger - I_m)\bar{\mathbf{b}}^{\text{true}}\|_2^2 + \frac{\rho^2}{m}\text{trace}(A_\lambda^{\dagger 2}) - \frac{2\rho^2}{m}\text{trace}(A_\lambda^\dagger) + \rho^2. \tag{4.7}$$

By comparing (4.4) and (4.7), we obtain

$$\mathbb{E}\left(\|\mathbf{p}_\lambda\|_2^2\right) = \mathbb{E}\left(\|\mathbf{r}_\lambda\|_2^2\right) + 2\rho^2\text{trace}(A_\lambda^\dagger) - m\rho^2.$$

By setting $\mathbb{E}(\|\mathbf{r}_\lambda\|_2^2) \approx \|\mathbf{r}_\lambda\|_2^2$, λ_{opt} is selected to minimize the UPRE function as

$$\lambda_{\text{opt}} = \arg\min_\lambda\{U(\lambda) = \|(A_\lambda^\dagger - I_m)\bar{\mathbf{b}}\|_2^2 + 2\rho^2\text{trace}(A_\lambda^\dagger) - m\rho^2\}. \tag{4.8}$$

Similar to the implementation of the GCV approach presented in [3], we determine λ_{opt} as the minimum value from a range of λ values derived from the SVD of $A_\lambda^\dagger$ when $L = I$. If $L \neq I$, one may use the GSVD of $\{A, L\}$ [25, 40] to find an expansion for $A_\lambda^\dagger$. However, for non-sparse large A and L matrices, the computations of the GSVD tend to be unattractive.

4.2. Extension of UPRE to the ℓ_1-regularized problem

In this section, we derive UPRE in terms of the solution of the projected problem based on an efficient TSVD calculated using the EGKB, described in Algorithm 2. First, we transform (4.1) into standard form as

$$\min_{\mathbf{z}}\{\frac{1}{2}\|\tilde{A}\mathbf{z} - \tilde{\mathbf{b}}\|_2^2 + \frac{\lambda^2}{2}\|\mathbf{z}\|_2^2\}, \quad \mathbf{z}_\lambda = L\mathbf{y}_\lambda, \tag{4.9}$$

where

$$\tilde{A} = AL^{-1} \in \mathbb{R}^{m\times n}, \quad \tilde{\mathbf{b}} = \mathbf{b} - \tilde{A}\mathbf{h} \in \mathbb{R}^{m\times 1}. \tag{4.10}$$

Here we extend the derivation of UPRE presented in [14] to encompass the problem in (4.9) in terms of the projected problem based on the regularization parameter ζ. Applying $k + p$ steps of the EGKB in Algorithm 2 on matrix $\tilde{A}$ with initial vector $\tilde{\mathbf{b}}$ and then truncating to the first k terms yields the bidiagonal form $\tilde{A}Q_k = S_{k+1}C_k$. Let $\mathbf{z}_\lambda = Q_k\mathbf{w}_\zeta$, and then by the column orthogonality of Q_k, we have $\|\mathbf{z}_\lambda\|_2 = \|Q_k\mathbf{w}_\zeta\|_2 = \|\mathbf{w}_\zeta\|_2$. This gives the Tikhonov projected problem

$$\min_{\mathbf{w}}\{\frac{1}{2}\|C_k\mathbf{w} - t_1\mathbf{e}_1^{(k+1)}\|_2^2 + \frac{\zeta^2}{2}\|\mathbf{w}\|_2^2\}, \tag{4.11}$$

where $t_1\mathbf{e}_1^{(k+1)} = S_{k+1}^\top\tilde{\mathbf{b}}$, whose solution is given by

$$\mathbf{w}_\zeta = t_1(C_k^\top C_k + \zeta^2 I_k)^{-1}C_k^\top\mathbf{e}_1^{(k+1)}.$$

To find ζ, we need to write the predictive error and regularized residual in terms of the projected problem. Let $\tilde{A}_\zeta^\dagger = (\tilde{A}Q_k)(\tilde{A}Q_k)_\zeta$ be the influence matrix for the projected solution, where $(\tilde{A}Q_k)_\zeta = \left((\tilde{A}Q_k)^\top\tilde{A}Q_k + \zeta^2 I_k\right)^{-1}(\tilde{A}Q_k)^\top$ is the regularized inverse matrix. Then we can rewrite the predictive error and regularized residual as

$$\mathbf{p}_\zeta = \tilde{A}Q_k\mathbf{w}_\zeta - \tilde{\mathbf{b}}^{\text{true}} = \tilde{A}_\zeta^\dagger\tilde{\mathbf{b}} - \tilde{\mathbf{b}}^{\text{true}}, \tag{4.12}$$

$$\mathbf{r}_\zeta = \tilde{A}Q_k\mathbf{w}_\zeta - \tilde{\mathbf{b}} = (\tilde{A}_\zeta^\dagger - I_m)\tilde{\mathbf{b}}, \tag{4.13}$$

where $\tilde{\mathbf{b}}^{\text{true}} = \mathbf{b}^{\text{true}} - A\bar{\mathbf{h}}$. Since the $\tilde{A}_\zeta^\dagger$ is still symmetric, and by comparing (4.12) with (4.2) and (4.13) with (4.3), we have

$$\zeta_{\text{opt}} = \arg\min_\zeta\{U_{\text{full}}(\zeta) = \|(\tilde{A}_\zeta^\dagger - I_m)\tilde{\mathbf{b}}\|_2^2 + 2\rho^2\text{trace}(\tilde{A}_\zeta^\dagger) - m\rho^2\}.$$

The influence matrix $\tilde{A}_\zeta^\dagger$ can be simplified into

$$\begin{aligned}\tilde{A}_\zeta^\dagger &= (\tilde{A}Q_k)\left((\tilde{A}Q_k)^\top(\tilde{A}Q_k) + \zeta^2 I\right)^{-1}(\tilde{A}Q_k)^\top \\ &= S_{k+1}C_k\left((S_{k+1}C_k)^\top S_{k+1}C_k + \zeta^2 I_k\right)^{-1}(S_{k+1}C_k)^\top \\ &= S_{k+1}\left(C_k\left(C_k^\top C_k + \zeta^2 I_k\right)^{-1}C_k^\top\right)(S_{k+1})^\top \\ &= S_{k+1}H_k(\zeta)S_{k+1}^\top,\end{aligned}$$

where $H_k(\zeta)$ is the influence matrix for the projected Tikhonov problem. Using this decomposition and the orthogonality of the columns of S_{k+1}, we obtain

$$\operatorname{trace}(\tilde{A}_\zeta^\dagger) = \operatorname{trace}\left(S_{k+1}H_k(\zeta)S_{k+1}^\top\right) = \operatorname{trace}(H_k(\zeta)).$$

Therefore, we obtain

$$\zeta_{\text{opt}} = \arg\min_\zeta\{U_{\text{full}}(\zeta) = \|t_1(H_k(\zeta) - I_{k+1})\mathbf{e}_1^{(k+1)}\|_2^2 + 2\rho^2\operatorname{trace}(H_k(\zeta)) - m\rho^2\}. \tag{4.14}$$

The parameter ζ_{opt} is calculated without reprojecting the solution $\mathbf{w}_\zeta$ for every ζ back into the original subspace $\mathbf{z}$. Since $\mathbf{w}_\zeta$ solves the problem with $C_k \in \mathbb{R}^{(k+1)\times k}$ and a right-hand-side vector of size $k+1$, we rewrite the UPRE for the projected problem as

$$\zeta_{\text{opt}} = \arg\min_\zeta\{U_{\text{proj}}(\zeta) = \|t_1(H_k(\zeta) - I_{k+1})\mathbf{e}_1^{(k+1)}\|_2^2 + 2\rho^2\operatorname{trace}(H_k(\zeta)) - (k+1)\rho^2\}. \tag{4.15}$$

Note that (4.14) and (4.15) differ only in their final constant terms. In fact,

$$U_{\text{full}}(\zeta) = U_{\text{proj}}(\zeta) - (m-k-1)\rho^2. \tag{4.16}$$

Since this difference is independent of ζ, the two criteria have exactly the same minimizer for a fixed projected subspace

$$\arg\min_\zeta U_{\text{full}}(\zeta) = \arg\min_\zeta U_{\text{proj}}(\zeta). \tag{4.17}$$

Thus, no approximation is introduced by replacing U_{full} with U_{proj} when optimizing over a fixed EGKB subspace. The approximation enters earlier, through the replacement of the original transformed operator $\widetilde{A}$ by its projected representation. Accordingly, the quality of the resulting regularization parameter depends on how accurately the selected Krylov subspace represents the dominant spectral components relevant to regularization.

Let $C_{k_p} = U_{k_p}\Sigma_{k_p}V_{k_p}^\top$ be the economy-size SVD of the enlarged projected matrix, where

$$U_{k_p} \in \mathbb{R}^{(k_p+1)\times k_p}, \quad \Sigma_{k_p} \in \mathbb{R}^{k_p\times k_p}, \quad V_{k_p} \in \mathbb{R}^{k_p\times k_p}.$$

Retaining only the leading k singular triplets yields the rank-k approximation $C_k = U_k\Sigma_k V_k^\top$, where

$$U_k = U_{k_p}(:, 1:k), \quad \Sigma_k = \Sigma_{k_p}(1:k, 1:k) = \operatorname{diag}(\check{\sigma}_1, \ldots, \check{\sigma}_k), \quad V_k = V_{k_p}(:, 1:k). \tag{4.18}$$

Because C_k is rectangular, $p+1$ additional orthogonal residual directions remain outside the retained rank-k singular subspace. Its contribution is independent of ζ and therefore does not affect the minimizer of the projected UPRE functional. Using (4.18), we obtain

$$H_k(\zeta) = U_k \operatorname{diag}\left(\frac{\breve{\sigma}_1^2}{\breve{\sigma}_1^2+\zeta^2}, \ldots, \frac{\breve{\sigma}_k^2}{\breve{\sigma}_k^2+\zeta^2}\right) U_k^\top, \tag{4.19}$$

and hence

$$\operatorname{trace}(H_k(\zeta)) = \sum_{i=1}^{k} \frac{\breve{\sigma}_i^2}{\breve{\sigma}_i^2+\zeta^2}. \tag{4.20}$$

Let $u_{1,i}$ denote the ith component in the first row of the matrix U_k. Then the projected UPRE function can be written as

$$\begin{aligned} &\arg\min_{\zeta}\{\mathrm{UPRE}(\zeta) = \|t_1(\Sigma_k^\dagger - I_k)U_k^\top \mathbf{e}_1^{(k_p+1)}\|_2^2 + 2\rho^2 \sum_{i=1}^{k} \frac{\breve{\sigma}_i^2}{\breve{\sigma}_i^2+\zeta^2} - k\rho^2\} \\ =&\arg\min_{\zeta}\{\mathrm{UPRE}(\zeta) = \sum_{i=1}^{k}\left(\frac{t_1\zeta^2 u_{1,i}}{\breve{\sigma}_i^2+\zeta^2}\right)^2 + 2\rho^2 \sum_{i=1}^{k} \frac{\breve{\sigma}_i^2}{\breve{\sigma}_i^2+\zeta^2} - k\rho^2\}. \end{aligned} \tag{4.21}$$

Since the final term in (4.21) is independent of ζ, it can be omitted without altering the minimizer.

The optimal value of ζ is obtained by evaluating the projected UPRE function over a logarithmically spaced set of candidate values. Since the quantities $\breve{\sigma}_i$ and $u_{1,i}$ are derived from the SVD of the projected bidiagonal matrix C_k, the UPRE criterion can be evaluated entirely in the projected space, thereby avoiding the formation of the dense transformed matrix $\tilde{A}$. The resulting value of ζ is then used as an approximation to the regularization parameter of the original SB problem and, for notational consistency, is denoted by λ hereafter.

Classical UPRE assumes additive white Gaussian measurement noise with variance ρ^2. During the SB iterations, however, the effective right-hand side $\tilde{\mathbf{b}}$ in (4.9) depends on the Bregman variables and, consequently, on the noisy data. Therefore, the UPRE assumptions may not hold exactly at later SB iterations, and their application to the updated projected problems is regarded as an approximation. The effectiveness of this approximation is assessed through the numerical experiments.

5. Computational complexity of the SB

Based on the computational costs provided in Sections 2.1 and 3.1, we show the results for determining the regularization parameter in the inner-generalized Tikhonov problem (2.3) using UPRE for two cases: (i) using the GSVD of $\{A, L\}$ for (4.1), and (ii) using EGKB for the projected problem (4.11). Note that the operation with $\tilde{A} = AL^{-1}$ in (4.11) is performed without explicitly forming $\tilde{A}$. A product with $\tilde{A}$ is evaluated sequentially as

$$\tilde{A}\mathbf{v} = A(L^{-1}\mathbf{v}),$$

whereas a transpose product is evaluated as

$$\tilde{A}^\top \mathbf{u} = L^{-T}(A^\top \mathbf{u}).$$

Thus, the matrices L^{-1} and L^{-T} are neither formed nor stored explicitly; only their actions on given vectors are required. Assuming that these actions can be computed efficiently, their additional cost is small relative to the dominant matrix-vector products with A and $A^\top$. Consequently, the cost of computing a rank-k TSVD approximation via EGKB associated with $\tilde{A}$ is $\mathcal{O}(4mnk_p)$, as given in (3.3).

As discussed in [42], computing the GSVD of $\{A, L\}$ needs approximately $\mathcal{O}(mn^2+2n^3)$. Since A and L remain fixed throughout the SB iterations, the GSVD needs to be computed only once. Empirically, we observe that using either the GSVD components of $\{A, L\}$ or the TSVD components associated with $\tilde{A}$ in (4.21) leads to comparable costs within the UPRE evaluation. We therefore omit the costs common to the UPRE evaluations and focus on the dominant costs associated with the GSVD and EGKB used to determine λ. The dominant cost associated with the SB iterations in all cases is the overall sum of columns 2, 3, and 4 in Table 1.

Table 1. Dominant computational costs for ℓ_{end} iterations of the SB algorithm when the regularization parameter is determined using (i) the GSVD with (4.1) and (ii) EGKB with (4.11), where the rank-k approximation associated with $\tilde{A}$ is obtained using Algorithm 2. The augmented system in (2.3) has size $2n \times n$ when the regularization matrix is of size $n \times n$.

Method	TSVD of $B = \tilde{A}$	GSVD of $\{A, L\}$	Total IRcgls Iterations
(i) GSVD		$\mathcal{O}(mn^2 + 2n^3)$	$\mathcal{O}\left(8n^2\iota_{total}\right)$
(ii) EGKB	$\mathcal{O}(4mnk_p)$		$\mathcal{O}\left(8n^2\iota_{total}\right)$

As shown in Table 1, comparing the dominant GSVD and EGKB costs gives the theoretical cost ratio

$$S_\lambda = \frac{mn^2 + 2n^3}{4mnk_p\ell_\lambda} = \frac{mn + 2n^2}{4mk_p\ell_\lambda}, \tag{5.1}$$

where ℓ_λ denotes the number of EGKB runs required until λ reaches convergence with EGKB. Once λ has converged, it is kept fixed in the subsequent SB iterations; hence, $\ell_\lambda \le \ell_{\text{end}}$. Therefore, the EGKB approach is computationally advantageous whenever

$$4mk_p\ell_\lambda < mn + 2n^2,$$

or equivalently, $S_\lambda > 1$. In particular, when $k_p\ell_\lambda \ll \min(m, n)$, the ratio in (5.1) increases with $m + 2n$, indicating a potentially substantial reduction in computational cost when the rank-k approximation is employed within the SB algorithm.

Regarding storage requirements, the GSVD of $\{A, L\}$ requires approximately $m^2 + 2n^2$ entries. By contrast, the projected approach stores the EGKB bases S_{k_p+1} and Q_{k_p} together with the truncated SVD factors of the projected problem, resulting in a storage cost of $\mathcal{O}(mk_p + nk_p + k^2)$. Since $k_p \ll \min(m, n)$ in practical applications, the projected framework is considerably more memory efficient than the GSVD while simultaneously reducing the computational burden.

Mixing precision for the SB

Using SP has become increasingly popular for large-scale computations due to its significant advantages in storage and floating-point costs. However, this approach often sacrifices accuracy in

the results. While higher precision can be used to maintain accuracy, it is often costly for large-scale computations. By using mixed-precision techniques, one can achieve superior results with high precision while benefiting from the efficient performance of lower precision [32].

Motivated by recent work on mixed precision in [43], we use SP arithmetic to compute the low-rank approximation $\bar{A}_k$ via EGKB and the DP for the SB solution. However, when EGKB is performed in SP, one-sided reorthogonalization may be insufficient to maintain orthogonality. Therefore, full two-sided reorthogonalization using the MGS algorithm is applied at steps 5 and 13 of Algorithm 2 to maintain the accuracy of the SP results. Although this additional reorthogonalization increases the cost of EGKB, particularly as the projected dimension grows, it does not alter the dominant computational cost of Algorithm 2.

It is worth noting that the use of SP reduces the computational cost of DP by approximately a factor of 2 [32], thereby decreasing the computational cost of EGKB relative to cases (i) and (ii) in Table 1. In our implementation, SP computations are performed using the MATLAB function `single`, and the low-rank approximation $\bar{A}_k$ is constructed from `single`$(\tilde{A})$.

Summary of the proposed framework

Before presenting the numerical experiments, we summarize the main steps of the proposed framework as follows.

- Transform the generalized Tikhonov subproblem (2.3) into the standard form (4.9).
- Apply the EGKB method in Algorithm 2 using SP arithmetic to construct an enlarged Krylov subspace for the transformed problem.
- Select the regularization parameter λ from the resulting low-dimensional projected problem using the TSVD-based UPRE criterion (4.21), thereby avoiding explicit computation of the GSVD.
- Solve the generalized Tikhonov subproblem using IRcgls in DP and update the SB variables.
- Repeat these steps until the prescribed SB stopping criterion is satisfied.

Thus, the proposed framework integrates Krylov subspace projection, adaptive regularization-parameter selection, and iterative solution of the generalized Tikhonov subproblem within the SB scheme.

6. Numerical experiments

We present numerical experiments designed to assess both the performance of the EGKB for finding $\bar{A}_k$ for $\tilde{A}$ in the transformation (4.9), as well as the performance of using $\bar{A}_k$ with the UPRE in determining the regularization parameter λ for the SB algorithm. We suppose that both $\mathbf{x}^{\text{true}}$ and $\mathbf{b}^{\text{true}}$ are available for all simulations.

The noisy data $\mathbf{b}$ are generated by adding white Gaussian noise $\boldsymbol{\eta}$ to the exact measurements $\mathbf{b}^{\text{true}}$, where $\boldsymbol{\eta}$ is generated using MATLAB's `randn('seed',1)` function. A fixed random seed is used to ensure the reproducibility of the numerical results. For a prescribed signal-to-noise ratio (SNR), the noise level is defined by

$$\text{SNR} = 20 \log_{10} \frac{\|\mathbf{b}^{\text{true}}\|_2}{\|\boldsymbol{\eta}\|_2}.$$

The variance ρ^2 appearing in the UPRE criterion (4.21) is taken as the variance of the generated Gaussian noise and is estimated in the numerical experiments using (4.6). Thus, the prescribed SNR determines the magnitude of the added noise, while ρ^2 provides the corresponding noise variance required for evaluating the UPRE functional.

To evaluate the computed solution performance qualitatively, we record the relative error (RE) in the solution at iteration ℓ, as given by

$$\mathrm{RE}(\mathbf{x}^{(\ell)}) = \frac{\|\mathbf{x}^{(\ell)} - \mathbf{x}^{\mathrm{true}}\|_2}{\|\mathbf{x}^{\mathrm{true}}\|_2}, \tag{6.1}$$

and the peak SNR (PSNR), which is defined by

$$\mathrm{PSNR}(\mathbf{x}^{(\ell)}) = 10\log_{10}\left(\frac{n\|\mathbf{x}^{\mathrm{true}}\|_\infty^2}{\|\mathbf{x}^{(\ell)} - \mathbf{x}^{\mathrm{true}}\|_2^2}\right), \tag{6.2}$$

where n denotes the total number of pixels in the image. A closer solution $\mathbf{x}^{(\ell)}$ to $\mathbf{x}^{\mathrm{true}}$ has a smaller RE and a larger PSNR.

The SB algorithm requires two convergence tolerances: τ_{SB} for the outer iteration ℓ and τ_{IRcgls} for the inner iteration ι in IRcgls. We stop the IRcgls algorithm when the relative residual for the normal equations (2.3) becomes smaller than the tolerance τ_{IRcgls}:

$$\frac{\left\|\begin{bmatrix} A \\ \lambda L \end{bmatrix}^{\top}\left(\begin{bmatrix} \mathbf{b} \\ \mathbf{d}-\mathbf{g} \end{bmatrix} - \begin{bmatrix} A \\ \lambda L \end{bmatrix}\mathbf{x}^{(\iota)}\right)\right\|_2}{\left\|\begin{bmatrix} A \\ \lambda L \end{bmatrix}^{\top}\begin{bmatrix} \mathbf{b} \\ \mathbf{d}-\mathbf{g} \end{bmatrix}\right\|_2} < \tau_{IRcgls}.$$

Convergence of the SB algorithm is determined using the relative change in $\mathbf{x}$ as:

$$\mathrm{RC}_{SB}(\mathbf{x}^{(\ell)}) = \frac{\|\mathbf{x}^{(\ell)} - \mathbf{x}^{(\ell-1)}\|_2}{\|\mathbf{x}^{(\ell-1)}\|_2}. \tag{6.3}$$

The SB solution is terminated when $\mathrm{RC}_{SB}(\mathbf{x}^{(\ell)}) < \tau_{SB}$. This stopping criterion is particularly useful when the true solution $\mathbf{x}^{\mathrm{true}}$ is unavailable. Although $\mathbf{x}^{\mathrm{true}}$ is known in our numerical experiments, we use $\mathrm{RC}_{SB}(\mathbf{x}^{(\ell)})$ to mimic practical applications and to assess the convergence behavior of the algorithm. Moreover, the stopping tolerance τ_{SB} is chosen according to the noise level in the data. Specifically, higher noise levels permit the use of a larger tolerance, leading to earlier termination of the iteration, whereas lower noise levels require a smaller tolerance to achieve a more accurate solution.

For both the outer SB iterations and the inner IRcgls iterations, maximum iteration limits are imposed, denoted by $L_{\max}$ and $I_{\max}$, respectively. Since the number of SB iterations is expected to remain relatively small compared with the problem size, we set $L_{\max} = 50$. On the other hand, because IRcgls solves the large-scale generalized Tikhonov subproblem (2.3), we choose $I_{\max} = 100$ to ensure that the dominant spectral components of the problem are adequately captured.

The regularization matrix L in (2.3) is defined by

$$L = \begin{bmatrix} 1 & -1 & & & 0 \\ & 1 & -1 & & \\ & & \ddots & \ddots & \\ & & & 1 & -1 \\ 0 & & & 0 & 1 \end{bmatrix} \in \mathbb{R}^{n\times n}. \tag{6.4}$$

In the implementation, L and L^{-1} are not explicitly constructed. Their actions on vectors are computed directly, thereby avoiding the explicit formation and storage of these matrices.

For a small-scale problem, we evaluate the performance of the proposed strategy for selecting λ within the SB algorithm by comparing it with the GSVD-based selection and a fixed optimal value of λ. The optimal λ is determined by testing 100 logarithmically spaced values in the interval $[10^{-2}, 10^{2}]$ and selecting the one that yields the smallest RE.

We record the timing of the SB algorithm without recording it to find the optimal fixed λ. Observe that finding λ by sweeping is usually impractical, as it requires knowledge of the true solution and running the algorithm for multiple choices of λ. When reporting results obtained with the transformation $\tilde{A}$ and GSVD, we include the associated costs of selecting λ and implementing SB. We record the average time over 5 runs.

We consider λ with the EGKB method to be converged and stop updating it if

$$\mathrm{RC}\Big((\lambda^{(\ell)})^2\Big) = \frac{|(\lambda^{(\ell)})^2 - (\lambda^{(\ell-1)})^2|}{|(\lambda^{(\ell-1)})^2|} \tag{6.5}$$

becomes smaller than a fixed tolerance, here $\tau_\lambda = 10^{-3}$, provided that $\lambda^{(\ell-1)} \neq 0$. If $\lambda^{(\ell)}$ converges to a fixed value at early iterations, this can significantly reduce the computational cost of the SB, as it eliminates the need to compute a TSVD of $\tilde{A}$ during the transformation, as indicated in column 2 of Table 1. For a fixed shrinkage parameter γ, once λ is selected, the corresponding value of β is determined by

$$\beta = \gamma\lambda^2. \tag{6.6}$$

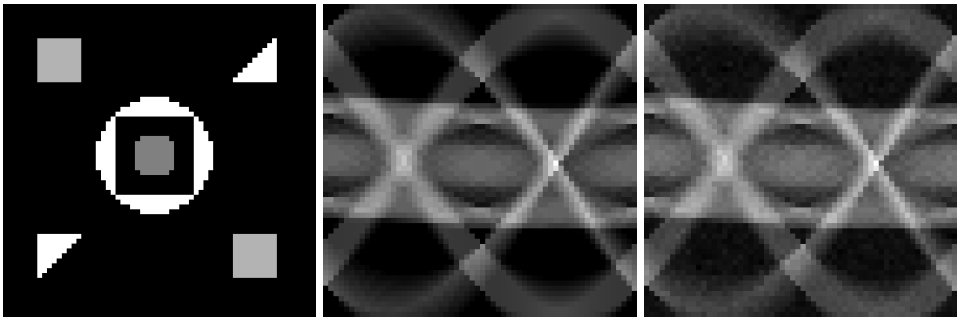
Thus, the proposed procedure should be interpreted as an adaptive SB regularization scheme in which the effective regularization weight $\beta^{(\ell)}$ varies with the UPRE-selected parameter $\lambda^{(\ell)}$, while the shrinkage parameter γ is kept fixed.

All numerical results presented in this work were computed using MATLAB R2025b [44] on a 2023 MacBook Pro equipped with an M3 Pro chip and 18 GB of memory. The software is available from the author upon request.

Example 6.1 (Underdetermined example)**.** *We generate a synthetic image of size* 64×64*, shown in Figure 1(a), and use* `PRtomo` *from [16] to generate a two-dimensional X-ray tomography test problem. This procedure yields a discrete forward operator* $A \in \mathbb{R}^{3600\times4096}$*. The sinogram data, displayed in Figure 1(b), are generated using* 60 *projection angles uniformly distributed between* 0 *and* 179 *together with* 60 *parallel-beam detector elements. The measurements are then contaminated with* 6% *additive white Gaussian noise, corresponding to an SNR of approximately* 24*, producing the noisy sinogram data* $\mathbf{b} \in \mathbb{R}^{3600\times1}$ *shown in Figure 1(c).*

Using the regularization matrix $L \in \mathbb{R}^{4096\times4096}$ defined in (6.4), we apply the EGKB algorithm with $B = \tilde{A} \in \mathbb{R}^{3600\times4096}$ in (4.9) to construct $\bar{A}_k$ in SP. The first k singular values computed by Algorithm 2 with $p = 5$ and $p = 27$ are compared with those obtained using DP, the TSVD, and the standard GKB method [33] without oversampling ($p = 0$). The RE in $\bar{A}_k$, measured in the Frobenius norm and reported in Table 2, indicate that $p = 27$ provides the most accurate approximation. The singular-value plots in Figure 2 further support this choice, showing that the SP implementation of EGKB with $p = 27$ yields approximations of the dominant singular values comparable to those obtained with DP and the TSVD. These results suggest that the use of SP has little effect on the accuracy of the dominant singular

values, provided that sufficient accuracy is maintained in the Krylov basis vectors, consistent with the observations in [45].



(a) True image **(b)** True sinogram **(c)** Noisy sinogram

Figure 1. The true image of 64×64 pixels, true sinogram, and noisy sinogram with 6% additive white Gaussian noise, for the problem in Example 6.1.

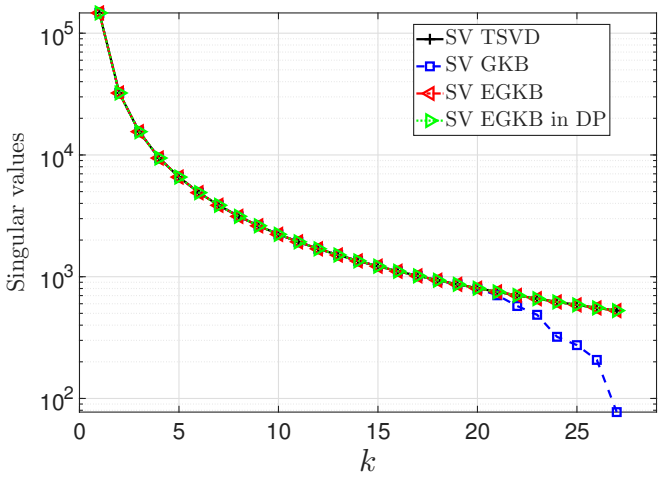


Figure 2. Approximation of the singular values using the true SVD, GKB, and EGKB in SP and DP with $k = 27$ and $p = 27$ for the transformation $\tilde{A}$ in (4.9) in Example 6.1.

Table 2. The SP results of estimating $\bar{A}_k$ for (4.9) in Example 6.1, with $\tilde{A} \in \mathbb{R}^{3600\times4096}$. The results in DP for the corresponding methods are provided in parentheses.

Method	k	p	$\frac{\|\tilde{A}-\bar{A}_k\|_F}{\|\tilde{A}\|_F}$
GKB	27	0	0.0180 (0.0180)
EGKB	27	5	0.0165 (0.0165)
		27	0.0160 (0.0160)

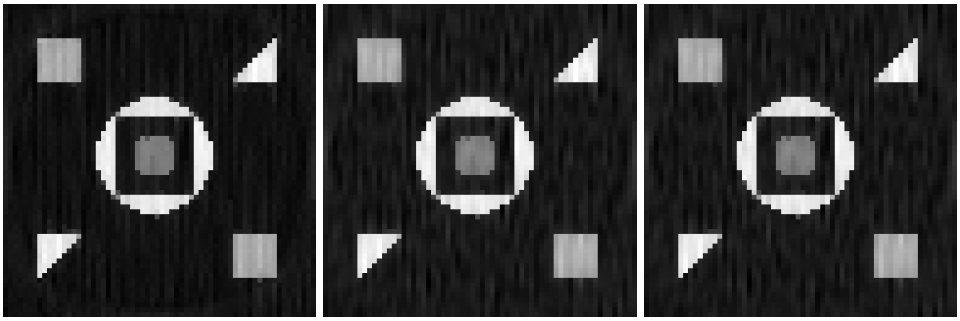
For the proposed method, the regularization parameter λ in (4.9) is determined by replacing $\tilde{A}$ with $\bar{A}_k$. This approximation is computationally efficient and is not expected to significantly affect the parameter selection, since EGKB provides nearly indistinguishable approximations to the first k

singular values of $\tilde{A}$, as illustrated in Figure 2. For the inner-generalized Tikhonov problem (2.3), γ is kept fixed throughout the SB iterations, while λ is selected using UPRE as in (4.21), and β is computed according to (6.6). At each SB iteration, the updated right-hand side yields a new $\bar{A}_k$ and, consequently, an updated λ. Strictly speaking, although the GSVD factors remain unchanged during the SB iterations, the UPRE minimizer depends on the current right-hand side and may therefore vary with ℓ. In this work, the GSVD-based parameter is computed only at the first SB iteration and then kept fixed. Consequently, the GSVD results should be interpreted as a baseline comparison rather than a fully adaptive parameter-selection strategy.

For the SB algorithm, we set $\tau_{SB} = 10^{-2}$ and $\tau_{IRcgls} = 10^{-3}$, and fix $\gamma = 1$, resulting in a monotonically decreasing RC_{SB}, as defined in (6.3). The optimal λ obtained by sweeping is $\lambda = 2.9151$, whereas the GSVD-based approach yields $\lambda = 2.5543$. The reconstructions in Figure 3 show that the EGKB method produces solutions comparable to those computed with the optimal and GSVD methods. The results for SB with these selection methods are displayed in Figure 4. As illustrated in Figure 4(e), EGKB does not select a suitable λ at the first iteration because the solution is still far from convergence and λ is more suited for the noise in the initial steps. The quantitative results, together with the computational times reported in Table 3, show that although EGKB requires more iterations than GSVD, it incurs a lower computational cost while achieving reconstruction quality comparable to that of the optimal approach with a fixed λ. This behavior can be attributed to the rapid stabilization of the regularization parameter selected by (6.5) within the EGKB framework. In particular, the parameter remains unchanged after $\lambda^{(6)}$, as illustrated in Figure 4(d), (e).

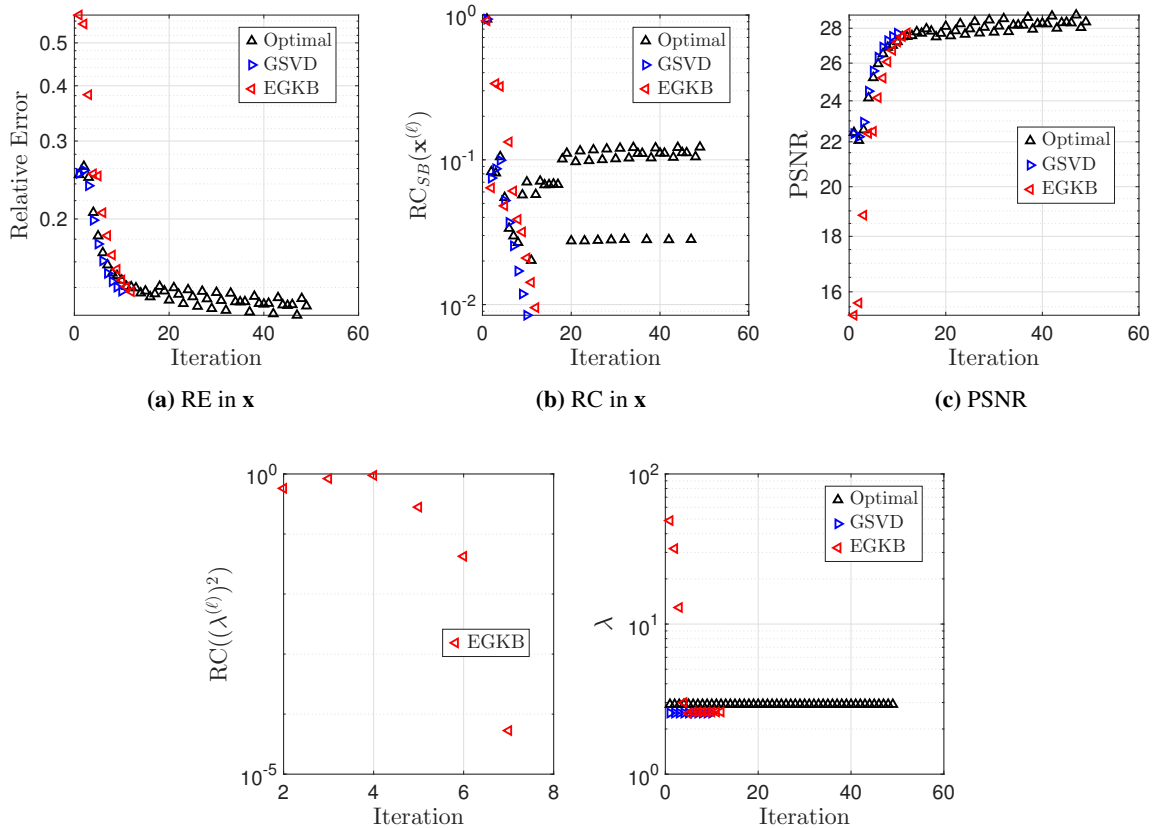
As discussed in Section 5, the dominant cost associated with determining λ for the original problem is the computation of the GSVD of $\{A, L\}$, as reflected in the results reported in column 6 of Table 3. To evaluate the predictive accuracy of the theoretical cost model in (5.1), we compare its estimates with the observed computational cost obtained when λ is determined using either the GSVD of $\{A, L\}$ or the low-rank approximation $\bar{A}_k$. Using the values $m = 3600$, $n = 4096$, $k = 27$, $p = 27$, and $\ell_\lambda = 6$ from Table 3, (5.1) predicts that determining λ via the GSVD of $\{A, L\}$ is approximately ten times more expensive than using the low-rank approximation $\bar{A}_k$. In practice, the timings reported in column 6 of Table 3 show that the GSVD-based approach is nearly eight times more expensive. These results demonstrate that EGKB is considerably more efficient than the GSVD approach, even though the GSVD is computed only once at the first SB iteration, whereas $\bar{A}_k$ is generated through EGKB over 12 outer iterations. Although the SB method requires 12 outer iterations to converge, the EGKB procedure for estimating λ is performed only during the first $\ell_\lambda = 6$ iterations. Thereafter, the converged value of λ is reused, which is consistent with the complexity analysis in (5.1).

The results in column 6 of Table 3 demonstrate that using SP reduces the computational cost of EGKB by approximately a factor of two compared with DP, confirming the computational advantages of using SP arithmetic. Moreover, comparing columns 6 and 7 shows that estimating λ using the GSVD and EGKB requires approximately 20 and 1.5 times, respectively, the cost of solving the SB problem with the original matrix A. Thus, replacing the GSVD with the low-rank approximation $\bar{A}_k$ yields a significant reduction in computational cost, a benefit that becomes increasingly important for large-scale problems.



(a) SB Opt ($\lambda = 2.9151$) **(b)** SB GSVD ($\lambda = 2.5543$) **(c)** SB EGKB

Figure 3. SB solutions at convergence for Figure 1(c). The regularization parameter is selected using the optimal approach, the GSVD, or the EGKB.



(a) RE in **x** **(b)** RC in **x** **(c)** PSNR

(d) RC in λ^2 **(e)** Selection of λ

Figure 4. The corresponding RE, PSNR, RC, and $\lambda^{(\ell)}$ in the outer iterations ℓ in the SB algorithm for the methods applied to Example 6.1. Note that λ is fixed with Optimal and GSVD.

Table 3. The corresponding RE and PSNR, defined in (6.1) and (6.2), respectively, by the SB with λ obtained with A and using the optimal approach, the GSVD, and EGKB, respectively, applied to the problem with SNR ≈ 24 in Example 6.1. The results in DP for the EGKB are provided in parentheses. The best results are shown in boldface.

					Timing		
Method	ℓ_{end}	ι_{total}	RE($\mathbf{x}$)	PSNR($\mathbf{x}$)	λ	SB	Total
Optimal	49	903	**0.13**	**28.37**	–	5.01	5.01
GSVD	**10**	**173**	0.14	27.70	23.26	**1.11**	24.37
EGKB	12 (12)	235 (253)	0.14 (0.14)	27.72 (27.72)	**3.03** (6.01)	1.95 (4.45)	**4.98** (10.46)

While the theoretical cost in (5.1) predicts a larger advantage for estimating λ using $\bar{A}_k$ than is observed in practice, the practical timings are based on SP computations and do not account for factors such as memory access, memory usage, and data movement. These effects, together with the choice of precision, are expected to play an increasingly important role as problem sizes grow, although obtaining reliable comparative measurements is difficult. Nevertheless, the results demonstrate a substantial computational benefit from using $\bar{A}_k$, particularly when the regularized problem is solved repeatedly for multiple right-hand sides. Moreover, Figure 4(a),(b) show comparable reductions in RE and RC with respect to the SB iterations. Therefore, estimating λ from $\bar{A}_k$ provides an efficient and effective alternative for large-scale problems where computing the GSVD of $\{A, L\}$ is impractical.

Example 6.2 (Overdetermined example)**.** *We consider the* `SheppLogan` *image of size* 256×256*, shown in Figure 5(a), and use* `PRtomo` *from [16] to generate a two-dimensional X-ray tomography test problem. This procedure yields a discrete forward operator* $A \in \mathbb{R}^{88200\times 65536}$*. The sinogram data, displayed in Figure 5(b), are generated using* 245 *uniformly distributed projection angles between* 0 *and* 179*, together with* 360 *parallel-beam detector elements. The measurements are then contaminated with* 10% *additive white Gaussian noise, corresponding to an SNR of approximately* 20*, producing the noisy sinogram data* $\mathbf{b} \in \mathbb{R}^{88200\times 1}$ *shown in Figure 5(c).*

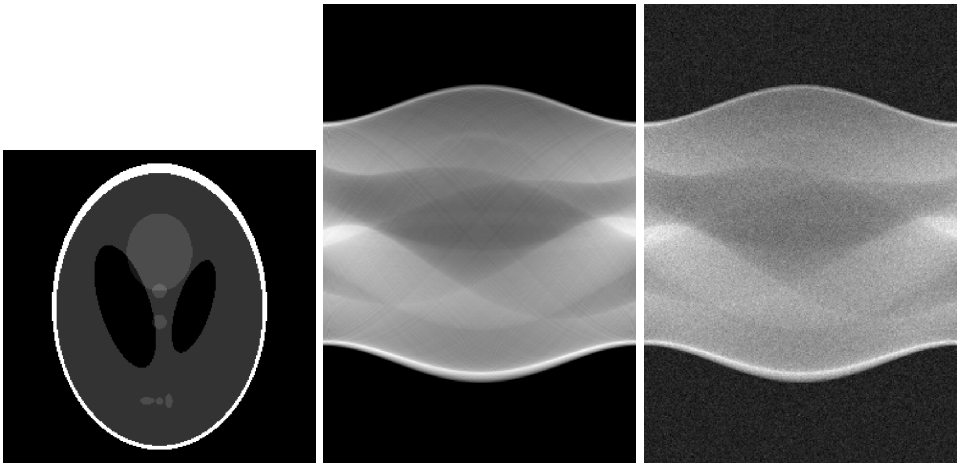
The regularization matrix $L \in \mathbb{R}^{65536\times 65536}$ is defined similarly to that in (6.4). Because the size of this problem is large, we construct $\bar{A}_k$ without explicitly forming the large matrix $\widetilde{A}$. Thus, a comparison with the true TSVD of $\widetilde{A}$ is not provided. The parameter γ was set to 2 based on preliminary numerical experiments. In particular, values of $\gamma < 1$ generally resulted in poorer reconstruction quality, whereas larger values, e.g., $\gamma > 10$, led to faster convergence at the expense of reconstruction quality. Hence, $\gamma = 2$ was selected as a reasonable compromise between reconstruction quality and convergence behavior. The SB stopping tolerance was set to $\tau_{SB} = 10^{-3}$. Notably, using $\tau_{IRcgls} = 10^{-3}$ yields reconstruction quality comparable to that obtained with $\tau_{IRcgls} = 10^{-4}$, while reducing the computational effort. The computational requirements of the GSVD and optimal methods make them less attractive for this simulation. The EGKB reconstruction is shown in Figure 6, while the corresponding results are presented in Figure 7. The quantitative results and computational timings are summarized in Table 4.

From these results, we can see that EGKB in SP and DP yield reconstructions with comparable RE and PSNR. As reported in column 6 of Table 4, EGKB in SP is considerably faster than in DP, saving approximately 37% of the computational time. Although this problem is larger than the one

discussed in Example 6.1 and the construction of $\bar{A}_k$ requires significant computational effort, the gain in computation time for selecting λ and obtaining the solution for the SB is reasonable.

We conduct an additional experiment with 15% white Gaussian noise, corresponding to SNR ≈ 16, to assess the sensitivity of the method to increased noise. As shown in the second row of Table 4, the higher noise level results in a deterioration of the reconstruction quality, as reflected by the larger RE values and lower PSNR values. Nevertheless, the EGKB-based approach remains effective, and the SP and DP implementations continue to produce comparable reconstruction results. Therefore, we conclude only that the proposed method exhibits similar qualitative behavior for the two noise levels considered in this study.

Our results demonstrate the effectiveness of the stopping criterion (3.2) for determining k in EGKB. The parameters p, γ, and τ_{IRcgls} were selected as described above, with values chosen independently for each problem based on its numerical characteristics. Although the experiments focus on X-ray tomographic reconstruction problems, the proposed approach is general and can be readily applied to other inverse problems, including image deblurring, signal restoration, and MRI reconstruction.



(a) True image **(b)** True sinogram **(c)** Noisy sinogram

Figure 5. The true image of 256×256 pixels, true sinogram, and noisy sinogram with 10% Gaussian noise, for the problem in Example 6.2.

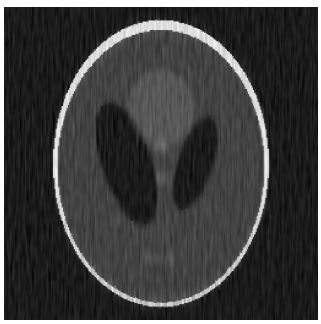
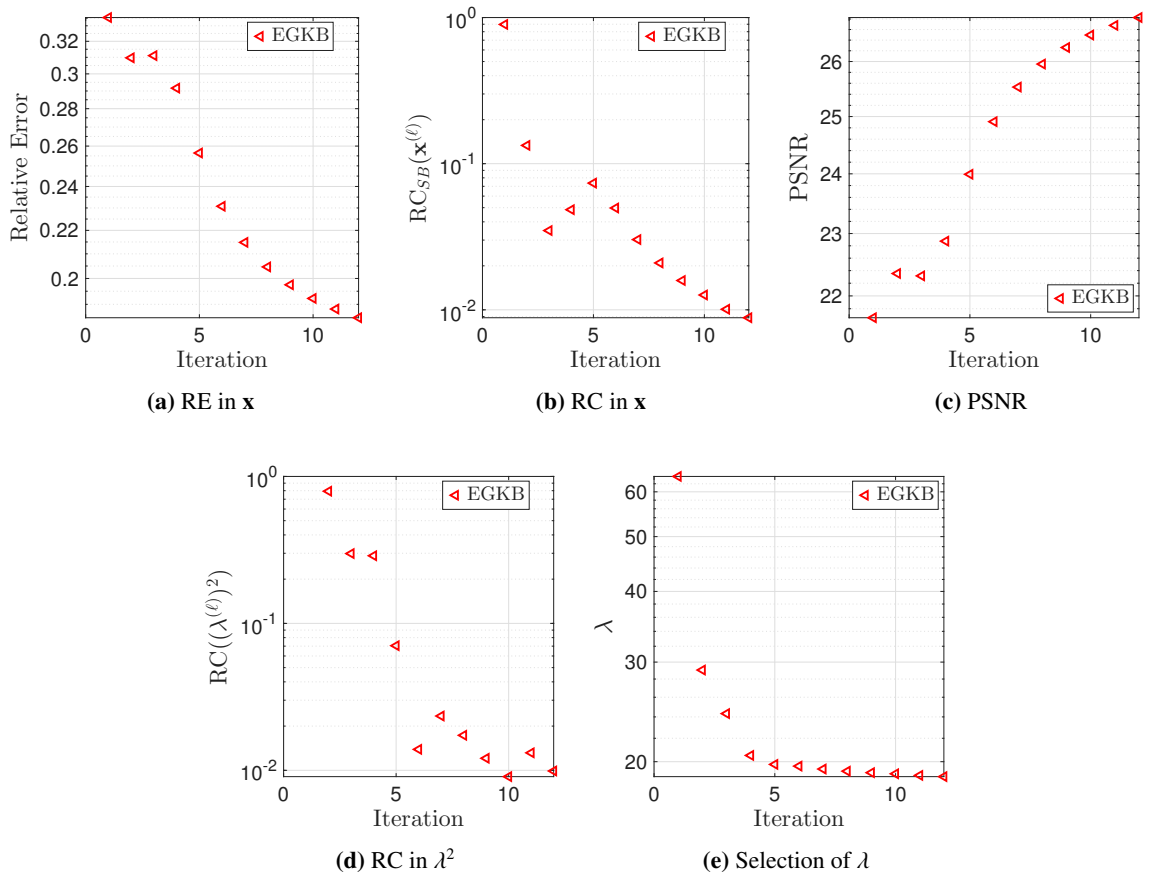


Figure 6. SB solution with EGKB at convergence for Figure 5(c).



(a) RE in **x** **(b)** RC in **x** **(c)** PSNR

(d) RC in λ^2 **(e)** Selection of λ

Figure 7. The corresponding RE, PSNR, RC, and $\lambda^{(\ell)}$ in the outer iterations ℓ in the SB algorithm for the EGKB applied to Example 6.2.

Table 4. The corresponding RE and PSNR values, defined in (6.1) and (6.2), respectively, obtained by the SB method with λ selected using EGKB for the problem in Example 6.2. The first and second rows correspond to SNR $\approx$ 20 and SNR $\approx$ 16, respectively. The results in DP for EGKB are provided in parentheses. The best results are shown in boldface.

						Timing		
Method	Noise	ℓ_{end}	ι_{total}	RE(**x**)	PSNR(**x**)	λ	SB	Total
EGKB	10%	12 (12)	178 (178)	**0.19** (**0.19**)	26.83 (26.83)	**471.43** (741.54)	**427.76** (675.29)	**899.19** (1416.83)
	15%	32 (28)	574 (499)	0.39 (**0.38**)	20.39 (20.54)	**2571** (3011)	**2484** (2896)	**5055** (5907)

7. Conclusions

This paper presents an efficient strategy for solving the inner-generalized Tikhonov regularization problems arising in the split Bregman (SB) framework. The proposed approach first transforms each generalized Tikhonov problem into standard form and then projects it onto a low-dimensional Krylov subspace. To improve parameter estimation accuracy, we use an enlarged Golub–Kahan bidiagonalization (EGKB) method that enriches the Krylov subspace. The regularization parameter

is then selected using a computationally efficient TSVD-based UPRE criterion, avoiding the explicit computation of the GSVD. The resulting SB iterates are computed using the IRcgls algorithm applied to the original problem formulation.

We provide a theoretical analysis of the computational costs of the proposed method and compare them with those of GSVD-based approaches. Numerical experiments demonstrate that the EGKB-UPRE framework achieves reconstruction quality comparable to that obtained with GSVD-based parameter estimation. Furthermore, although single-precision arithmetic does not reduce the number of iterations, it lowers the computational cost and storage requirements per iteration without compromising reconstruction accuracy.

Overall, the proposed approach offers an effective and computationally attractive alternative to GSVD-based methods for large-scale inverse problems, particularly when explicit GSVD computations are impractical.

Future work will explore extending the framework to other regularization parameter selection techniques, such as GCV and discrepancy-principle-based approaches, as well as its application to broader classes of inverse problems and sparsity-promoting regularization models. The present framework is restricted to invertible L; extensions to rank-deficient regularization operators via the A-weighted generalized inverse $L_A^\dagger$, including those arising in total-variation regularization, remain for future investigation.

Use of Generative-AI tools declaration

The author declares he has not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

The author extends his appreciation to the Deanship of Scientific Research, Islamic University of Madinah, Saudi Arabia, for funding this research work.

Conflict of interest

The author declares no conflict of interest.

References

1. Å. Björck, *Numerical methods for least squares problems*, SIAM, Philadelphia, 1996.
2. G. H. Golub, C. F. Van Loan, *Matrix computations*, Johns Hopkins University Press, Baltimore, 2013.
3. P. C. Hansen, Regularization tools: A MATLAB package for analysis and solution of discrete ill-posed problems, *Numer. Algorithms*, **6** (1994), 1–35. https://doi.org/10.1007/BF02149761
4. L. Eldén, A weighted pseudoinverse, generalized singular values, and constrained least squares problems, *BIT Numer. Math.*, **22** (1982), 487–502. https://doi.org/10.1007/BF01934412
5. M. Fuhry, L. Reichel, A new Tikhonov regularization method, *Numer. Algorithms*, **59** (2012), 433–445. https://doi.org/10.1007/s11075-011-9498-x

6. G. H. Golub, P. C. Hansen, D. P. O'Leary, Tikhonov regularization and total least squares, *SIAM J. Matrix Anal. Appl.*, **21** (1999), 185–194. https://doi.org/10.1137/S0895479897326432
7. L. I. Rudin, S. Osher, E. Fatemi, Nonlinear total variation based noise removal algorithms, *Physica D*, **60** (1992), 259–268. https://doi.org/10.1016/0167-2789(92)90242-F
8. A. Chambolle, An algorithm for total variation minimization and applications, *J. Math. Imaging Vis.*, **20** (2004), 89–97. https://doi.org/10.1023/B:JMIV.0000011325.36760.1e
9. S. Gazzola, M. S. Landman, Flexible GMRES for total variation regularization, *BIT Numer. Math.*, **59** (2019), 721–746. https://doi.org/10.1007/s10543-019-00750-x
10. T. Goldstein, S. Osher, The split Bregman method for ℓ_1-regularized problems, *SIAM J. Imaging Sci.*, **2** (2009), 323–343. https://doi.org/10.1137/080725891
11. A. Lanza, S. Morigi, L. Reichel, F. Sgallari, A generalized Krylov subspace method for $\ell_p - \ell_q$ minimization, *SIAM J. Sci. Comput.*, **37** (2015), S30–S50. https://doi.org/10.1137/140967982
12. C. R. Vogel, M. E. Oman, Iterative methods for total variation denoising, *SIAM J. Sci. Comput.*, **17** (1996), 227–238. https://doi.org/10.1137/0917016
13. M. Pasha, A. K. Saibaba, S. Gazzola, M. I. Español, E. de Sturler, A computational framework for edge-preserving regularization in dynamic inverse problems, *Electron. T. Numer. Anal.*, **58** (2023), 486–516. https://doi.org/10.1553/etna_vol58s486
14. R. A. Renaut, S. Vatankhah, V. E. Ardestani, Hybrid and iteratively reweighted regularization by unbiased predictive risk and weighted GCV for projected systems, *SIAM J. Sci. Comput.*, **39** (2017), B221–B243. https://doi.org/10.1137/15M1037925
15. J. Chung, S. Gazzola, Computational methods for large-scale inverse problems: A survey on hybrid projection methods, *SIAM Rev.*, **66** (2024), 205–284. https://doi.org/10.1137/21M1441420
16. S. Gazzola, P. C. Hansen, J. G. Nagy, IR Tools: A MATLAB package of iterative regularization methods and large-scale test problems, *Numer. Algorithms*, **81** (2019), 773–811. https://doi.org/10.1007/s11075-018-0570-7
17. P. C. Hansen, J. G. Nagy, D. P. O'Leary, *Deblurring images: Matrices, spectra, and filtering*, SIAM, Philadelphia, 2006. https://doi.org/10.1137/1.9780898718874
18. C. R. Vogel, *Computational methods for inverse problems*, SIAM, Philadelphia, 2002. https://doi.org/10.1137/1.9780898717570
19. C. Van Loan, N. Pitsianis, *Approximation with Kronecker products*, Technical Report, Cornell University, 1992.
20. J. G. Nagy, *Decomposition of block Toeplitz matrices into a sum of Kronecker products with applications in image restoration*, Technical Report, Citeseer, 1996.
21. J. Kamm, J. G. Nagy, Optimal Kronecker product approximation of block Toeplitz matrices, *SIAM J. Matrix Anal. Appl.*, **22** (2000), 155–172. https://doi.org/10.1137/S0895479898345540
22. J. G. Nagy, M. K. Ng, L. Perrone, Kronecker product approximations for image restoration with reflexive boundary conditions, *SIAM J. Matrix Anal. Appl.*, **25** (2003), 829–841. https://doi.org/10.1137/S0895479802419580

23. B. F. Sweeney, M. I. Español, R. A. Renaut, Efficient decomposition-based algorithms for ℓ_1-regularized inverse problems with column-orthogonal and Kronecker product matrices, *Inverse Probl.*, **41** (2025), 075016. https://doi.org/10.1088/1361-6420/adec12
24. T. Bonesky, Morozov's discrepancy principle and Tikhonov-type functionals, *Inverse Probl.*, **25** (2008), 015015. https://doi.org/10.1088/0266-5611/25/1/015015
25. P. C. Hansen, *Rank-deficient and discrete ill-posed problems: Numerical aspects of linear inversion*, SIAM, Philadelphia, 1998. https://doi.org/10.1137/1.9780898719697
26. P. C. Hansen, D. P. O'Leary, The use of the L-curve in the regularization of discrete ill-posed problems, *SIAM J. Sci. Comput.*, **14** (1993), 1487–1503. https://doi.org/10.1137/0914086
27. A. Alsubhi, R. A. Renaut, Enlarged GKB and RSVD for KP approximations illustrated for mixed precision ℓ_1 image restoration, *Numer. Algorithms*, **101** (2026), 2245–2267. https://doi.org/10.1007/s11075-025-02086-w
28. B. Lewis, L. Reichel, Arnoldi–Tikhonov regularization methods, *J. Comput. Appl. Math.*, **226** (2009), 92–102. https://doi.org/10.1016/j.cam.2008.05.003
29. A. Buccini, M. Pasha, L. Reichel, Linearized Krylov subspace Bregman iteration with nonnegativity constraint, *Numer. Algorithms*, **87** (2021), 1177–1200. https://doi.org/10.1007/s11075-020-01004-6
30. M. Alotaibi, A. Buccini, L. Reichel, Krylov subspace split Bregman methods, *Appl. Numer. Math.*, **184** (2023), 371–390. https://doi.org/10.1016/j.apnum.2022.10.009
31. A. Buccini, F. Chen, M. Pasha, L. Reichel, Krylov subspace based FISTA-type methods for linear discrete ill-posed problems, *Numer. Linear Algebr. Appl.*, **32** (2025), e2610. https://doi.org/10.1002/nla.2610
32. N. J. Higham, T. Mary, Mixed precision algorithms in numerical linear algebra, *Acta Numer.*, **31** (2022), 347–414. https://doi.org/10.1017/S0962492922000022
33. G. Golub, W. Kahan, Calculating the singular values and pseudo-inverse of a matrix, *J. Soc. Ind. Appl. Math. Ser. B: Numer. Anal.*, **2** (1965), 205–224. https://doi.org/10.1137/0702016
34. C. C. Paige, M. A. Saunders, LSQR: An algorithm for sparse linear equations and sparse least squares, *ACM T. Math. Software*, **8** (1982), 43–71. https://doi.org/10.1145/355984.355989
35. R. M. Larsen, Lanczos bidiagonalization with partial reorthogonalization, *DAIMI Rep. Ser.*, 1998. https://doi.org/10.7146/dpb.v27i537.7070
36. R. A. Renaut, J. D. Hogue, S. Vatankhah, S. Liu, A fast methodology for large-scale focusing inversion of gravity and magnetic data using the structured model matrix and the 2-D fast Fourier transform, *Geophys. J. Int.*, **223** (2020), 1378–1397. https://doi.org/10.1093/gji/ggaa372
37. H. D. Simon, H. Zha, Low-rank matrix approximation using the Lanczos bidiagonalization process with applications, *SIAM J. Sci. Comput.*, **21** (2000), 2257–2274. https://doi.org/10.1137/S1064827597327309
38. I. Yamazaki, S. Tomov, J. Dongarra, Mixed-precision Cholesky QR factorization and its case studies on multicore CPU with multiple GPUs, *SIAM J. Sci. Comput.*, **37** (2015), C307–C330. https://doi.org/10.1137/14M0973773

39. S. Gazzola, E. Onunwor, L. Reichel, G. Rodriguez, On the Lanczos and Golub–Kahan reduction methods applied to discrete ill-posed problems, *Numer. Linear Algebr. Appl.*, **23** (2016), 187–204. https://doi.org/10.1002/nla.2020

40. P. C. Hansen, Regularization, GSVD and truncated GSVD, *BIT Numer. Math.*, **29** (1989), 491–504. https://doi.org/10.1007/BF02219234

41. P. C. Hansen, *Discrete inverse problems: Insight and algorithms*, SIAM, Philadelphia, PA, 2010. https://doi.org/10.1137/1.9780898718836

42. Z. Bai, The CSD, GSVD, their applications and computations, *Inst. Math. Appl.*, 1992.

43. M. Baboulin, O. Kaya, T. Mary, M. Robeyns, Mixed precision iterative refinement for low-rank matrix and tensor approximations, *SIAM J. Sci. Comput.*, **47** (2025), A2906–A2935. https://doi.org/10.1137/23M1576992

44. The MathWorks, Inc., *MATLAB Version R2025b*, Natick, MA, USA, 2025.

45. C. Boutsikas, P. Drineas, I. C. F. Ipsen, Small singular values can increase in lower precision, *SIAM J. Matrix Anal. Appl.*, **45** (2024), 1518–1540. https://doi.org/10.1137/23M1557209



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0)