Square Smoothing Regularization Matrices with Accurate Boundary Conditions

M. Donatelli[a,*], L. Reichel[b]

[a] *Dipartimento di Scienza e Alta Tecnologia, Università dell'Insubria, Via Valleggio 11, 22100 Como, Italy.*
[b] *Department of Mathematical Sciences, Kent State University, Kent, OH 44242, USA.*

Abstract

This paper is concerned with the solution of large-scale linear discrete ill-posed problems. The determination of a meaningful approximate solution of these problems requires regularization. We discuss regularization by the Tikhonov method and by truncated iteration. The choice of regularization matrix in Tikhonov regularization may significantly affect the quality of the computed approximate solution. The present paper describes the construction of square regularization matrices from finite difference equations with a focus on the boundary conditions. The regularization matrices considered have a structure that makes them easy to apply in iterative methods, including methods based on the Arnoldi process. Numerical examples illustrate the properties and effectiveness of the regularization matrices described.

Keywords: Tikhonov regularization, regularization matrix, boundary conditions

1. Introduction

We are concerned with the computation of an approximate solution of linear systems of equations of the form

$$A\mathbf{x} = \mathbf{b}, \qquad A \in \mathbb{R}^{n\times n}, \qquad \mathbf{x}, \mathbf{b} \in \mathbb{R}^n, \tag{1}$$

with a large matrix A with many singular values of different orders of magnitude close to the origin. In particular, A is severely ill-conditioned and may be singular. Linear systems of equations with a matrix of this kind often are referred to as linear discrete ill-posed problems. They arise, for instance, from the discretization of linear ill-posed problems, such as Fredholm integral equations of the first kind with a smooth kernel. The right-hand side vector $\mathbf{b}$ of

*Corresponding author
Email addresses: marco.donatelli@uninsubria.it (M. Donatelli), reichel@math.kent.edu (L. Reichel)

linear discrete ill-posed problems that occur in applications typically represents measured data and often is contaminated by an unknown error $\mathbf{e} \in \mathbb{R}^n$. We will assume the matrix A to be square; however, our discussion easily can be modified to allow $A \in \mathbb{R}^{m\times n}$ with $m \neq n$. Inconsistent systems (1) are treated as least-squares problems.

Let $\hat{\mathbf{b}} \in \mathbb{R}^n$ denote the unknown noise-free vector associated with $\mathbf{b}$, i.e.,

$$\mathbf{b} = \hat{\mathbf{b}} + \mathbf{e}, \tag{2}$$

and let $\hat{\mathbf{x}}$ be the solution of the unavailable linear system of equations

$$A\mathbf{x} = \hat{\mathbf{b}}, \tag{3}$$

which we assume to be consistent. If A is singular, then $\hat{\mathbf{x}}$ denotes the solution of (3) of minimal Euclidean norm.

Let $A^\dagger$ denote the Moore-Penrose pseudoinverse of A. The solution of (1), given by

$$A^\dagger\mathbf{b} = A^\dagger\hat{\mathbf{b}} + A^\dagger\mathbf{e} = \hat{\mathbf{x}} + A^\dagger\mathbf{e},$$

typically is a useless approximation of $\hat{\mathbf{x}}$ due to severe propagation of the error $\mathbf{e}$. This depends on the large norm of $A^\dagger$. Therefore, one generally replaces the linear system (1) by a nearby problem, whose solution is less sensitive to the error $\mathbf{e}$. This replacement is known as regularization. One of the most popular regularization methods is due to Tikhonov. This method replaces (1) by the minimization problem

$$\min_{\mathbf{x}\in\mathbb{R}^n} \left\{ \|A\mathbf{x} - \mathbf{b}\|^2 + \mu\|L\mathbf{x}\|^2 \right\}, \tag{4}$$

where $\|\cdot\|$ denotes the Euclidean vector norm, the matrix $L \in \mathbb{R}^{k\times n}$, $1 \leq k \leq n$, is referred to as the regularization matrix, and the scalar $\mu > 0$ as the regularization parameter; see, e.g., [3, 14, 15, 18]. We assume that the matrices A and L satisfy

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

where $\mathcal{N}(M)$ denote the null space of the matrix M. Then, for any $\mu > 0$, the minimization problem (4) has the unique solution

$$\mathbf{x}_\mu = (A^TA + \mu L^TL)^{-1}A^T\mathbf{b}, \tag{5}$$

where the superscript ^{T} denotes transposition. The value of μ determines how sensitive $\mathbf{x}_\mu$ is to the error in $\mathbf{b}$ and how small the residual error $\mathbf{b} - A\mathbf{x}_\mu$ is. A suitable value of μ generally is not explicitly known and has to be determined during the solution process.

Minimization problems (4) of small to moderate size can be conveniently solved with the aid of the Generalized Singular Value Decomposition (GSVD) or a related factorization of the matrix pair $\{A, L\}$; see [13, 18]. The regularization matrices discussed in this paper can be applied in this context; however, our primary aim is to develop new square regularization matrices that are convenient to use in an iterative method for the solution of large-scale minimization

problems (4) as well as in iterative methods that regularize by truncated iteration.

Common choices of L for linear systems of equations (1) that are obtained by the discretization of Fredholm integral equations of the first kind in one space-dimension are the identity matrix I, as well as scaled finite difference approximations of the first derivative

$$L = \frac{1}{2}\begin{bmatrix} 1 & -1 & & & & 0 \\ & 1 & -1 & & & \\ & & 1 & -1 & & \\ & & & \ddots & \ddots & \\ 0 & & & & 1 & -1 \end{bmatrix} \in \mathbb{R}^{(n-1)\times n} \tag{6}$$

and of the second derivative

$$L = \frac{1}{4}\begin{bmatrix} -1 & 2 & -1 & & & 0 \\ & -1 & 2 & -1 & & \\ & & \ddots & \ddots & \ddots & \\ 0 & & & -1 & 2 & -1 \end{bmatrix} \in \mathbb{R}^{(n-2)\times n}. \tag{7}$$

These matrices damp fast oscillatory components of a vector $\mathbf{x}$ in (4) more than slowly oscillatory components. The regularization matrices (6) and (7) therefore are referred to as smoothing regularization matrices. Here we think of vectors $\mathbf{x}$ as discretizations of continuous real-valued functions. A vector is said to have a fast oscillatory component if it is the discretization of a function with a fast oscillatory component. The application of a smoothing regularization matrix is beneficial when the desired solution is smooth.

The regularization matrix L should be chosen so that important known features of the desired solution $\hat{\mathbf{x}}$ of (3) can be represented by vectors in $\mathcal{N}(L)$, because then these features are not damped by L in (4). For instance, when the solution is known to be the discretization at equidistant points of a smooth monotonically increasing function, whose graph has small curvature, it may be appropriate to use the regularization matrix (7), because its null space contains the discretization of linear functions.

A Tikhonov regularization problem (4) is said to be in *general form* when $L \neq I$ and in *standard form* when $L = I$. Large-scale regularization problems in standard form can be conveniently solved by Krylov subspace methods, because the Krylov subspace generated is independent of the regularization parameter $\mu > 0$. This is important since the minimization problem (4) typically has to be solved for several values of μ in order to determine an appropriate value. It therefore is attractive to transform large-scale Tikhonov minimization problems in general form (4) to standard form before applying an iterative Krylov subspace method, provided that this transformation can be carried out fairly inexpensively.

We outline the transformation to standard form described by Eldén [15]. Let $L^\dagger$ denote the Moore-Penrose pseudoinverse of L, introduce the A-weighted

pseudoinverse of L,

$$L_A^\dagger = \left(I - (A(I - L^\dagger L))^\dagger A\right) L^\dagger \in \mathbb{R}^{n\times k}, \tag{8}$$

and define the vectors

$$\begin{aligned} \bar{\mathbf{x}} &= L\mathbf{x}, \\ \mathbf{x}^{(0)} &= \left(A(I - L^\dagger L)\right)^\dagger \mathbf{b}, \\ \bar{\mathbf{b}} &= \mathbf{b} - A\mathbf{x}^{(0)}. \end{aligned}$$

Eldén [15] showed that the Tikhonov regularization problem in general form (4) is equivalent to the minimization problem in standard form,

$$\min_{\bar{\mathbf{x}}\in\mathbb{R}^k} \{\|AL_A^\dagger \bar{\mathbf{x}} - \bar{\mathbf{b}}\|^2 + \mu\|\bar{\mathbf{x}}\|^2\}. \tag{9}$$

The solution (5) of (4) can be recovered from the solution $\bar{\mathbf{x}}_\mu$ of (9) according to

$$\mathbf{x}_\mu = L_A^\dagger \bar{\mathbf{x}}_\mu + \mathbf{x}^{(0)}. \tag{10}$$

Moreover,

$$\|AL_A^\dagger \bar{\mathbf{x}}_\mu - \bar{\mathbf{b}}\| = \|A\mathbf{x}_\mu - \mathbf{b}\|.$$

We are concerned with the iterative solution of large-scale problems (9). Despite the complexity (8) of $L_A^\dagger$, the evaluation of matrix-vector products with the matrices $L_A^\dagger$ and $AL_A^\dagger$ is feasible when L is a banded matrix with small bandwidth and has an explicitly known null space, or when L can be diagonalized efficiently by a fast trigonometric transform; see, e.g., [25] for many examples. Under these circumstances, the evaluation of a matrix-vector product with one of the matrices $L_A^\dagger$ or $AL_A^\dagger$ requires only one matrix-vector product evaluation with A.

An alternative to the solution of (9) by an iterative method is to apply the iterative method to the approximate solution of

$$AL_A^\dagger \bar{\mathbf{x}} = \bar{\mathbf{b}}. \tag{11}$$

Regularization is achieved by truncated iteration; iteration is terminated before significant propagation of the error in $\bar{\mathbf{b}}$ occurs.

When the regularization matrix is rectangular, such as (6) and (7), Krylov subspace methods based on partial Lanczos bidiagonalization of $AL_A^\dagger$ can be applied. Each step of Lanczos bidiagonalization requires the evaluation of one matrix-vector product with A and one with A^T. On the other hand, square regularization matrices allow the use of Krylov subspace methods based on the Arnoldi process. Each step of the Arnoldi process only demands the evaluation of one matrix-vector product with A. For many linear discrete ill-posed problems (1), solution methods for (4) based on the Arnoldi process require fewer matrix-vector product evaluations than solution methods based on partial Lanczos bidiagonalization; see [4, 5, 20, 25] for illustrations. We therefore are interested in the derivation of square regularization matrices. Several approaches to

construct square regularization matrices are discussed in [6, 7, 11, 25]. An iterative scheme by Hansen and Jensen [19] allows the application of a rectangular regularization matrix with the Arnoldi process, but does not allow straightforward evaluation of the norm of the residual error associated with each iterate; computation of the residual norm requires the evaluation of an additional matrix-vector product. Knowledge of the residual norm is important when the regularization parameter μ is determined by the discrepancy principle or by the L-curve criterion. A generalized Arnoldi-Tikhonov method that allows a rectangular regularization matrix is also discussed in [16].

The present paper discusses extensions of the finite difference matrices (6) and (7) to square matrices L by imposing boundary conditions. Section 2 first reviews properties of the square regularization matrices obtained by incorporating Dirichlet or Neumann boundary conditions with the finite difference matrix (7), and then turns to the application of antireflective and high-order boundary conditions. All of these boundary conditions have previously been used to modify the blurring matrix in image deblurring problems; see [10, 23, 26]. Discussion of antireflective and high-order boundary conditions in the context of regularization matrices is believed to be new. Section 3 is concerned with some generalizations, such as square regularization matrices obtained by incorporation of boundary conditions to the regularization matrix (6) and invertible regularization matrices. Computed examples are presented in Section 4 and concluding remarks can be found in Section 5.

2. Boundary conditions for second order difference matrices

This section discusses how boundary conditions can be used to determine square regularization matrices from the rectangular matrix (7). First the classical Dirichlet and Neumann boundary conditions are considered. This is followed by discussions on antireflective and high-order boundary conditions, which were introduced in [26] and [10], respectively. We apply the boundary conditions to the standard finite difference discretization of the second derivative and investigate the spectral factorization of the associated finite difference matrix. Our work is inspired by the investigation by Strang [27].

2.1. Classical boundary conditions

Consider the regularization matrix (7). The interior rows (rows 2 to $n-1$) of the new regularization matrices to be determined agree with the corresponding rows of (7). The boundary conditions determine the first and last rows.

We first briefly digress to boundary value problems for the Poisson equation on the interval $[0, \pi]$. The solution u is required to satisfy suitable boundary conditions, such as the Dirichlet conditions $u(0) = u(\pi) = 0$ or the Neumann conditions $u'(0) = u'(\pi) = 0$, where $'$ denotes differentiation. We may, of course, impose the Dirichlet condition at one end point and the Neumann condition at the other one. One also may require u to satisfy linear combinations of the Dirichlet and Neumann conditions, such as $u(0) + u'(0) = 0$.

These boundary conditions have discrete analogues that can be prescribed to solutions $\mathbf{u} = [u_0, \ldots, u_{n-1}]^T \in \mathbb{R}^n$ of the discretized Poisson equation. Here $u_j \approx u(t_j)$, where $t_j = \pi j/n$, $0 \le j \le n-1$. There is an even larger variety of meaningful boundary conditions for the discrete problem than for the continuous one; Strang [27] remarks that "The discrete case has a new level of variety and complexity, often appearing in the boundary conditions."

Dirichlet boundary conditions at both end points correspond to

$$u_{-1} = 0 \qquad \text{and} \qquad u_n = 0.$$

Discretization of the negative second derivative by the standard three-point finite difference stencil and the above boundary conditions yield the second order difference matrix

$$L_D = \begin{bmatrix} 2 & -1 & 0 & \ldots & 0 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ 0 & \ldots & 0 & -1 & 2 \end{bmatrix}. \tag{12}$$

The spectral factorization of L_D is well known. Let $S^{(n)} = [s_{ij}^{(n)}]$ be the discrete sine transform (DST) matrix of order n. Its entries are

$$s_{ij}^{(n)} = \sqrt{\frac{2}{n+1}} \sin\left(\frac{ij\pi}{n+1}\right), \qquad i,j = 1,2,\ldots,n.$$

Then

$$L_D = (S^{(n)})^T \mathrm{diag}(z(\mathbf{t}))S^{(n)}, \tag{13}$$

where

$$\mathbf{t} = [t_1, t_2, \ldots, t_n]^T, \qquad t_j = \frac{j\pi}{n+1}, \qquad j = 1,2,\ldots,n,$$

and

$$z(t) = 2 - 2\cos(t) = -\mathrm{e}^{-\mathrm{i}t} + 2 - \mathrm{e}^{\mathrm{i}t}, \quad \mathrm{i} = \sqrt{-1}. \tag{14}$$

Thus, the eigenvalues of the matrix (12) are a uniform sampling of the function (14) in $(0, \pi)$. In the classical context of Toeplitz matrix theory, this function is the symbol of the matrix (12); it is defined by the entries of an inner row of the matrix.

The discrete analogue of the Neumann boundary condition can be implemented in two ways: we may require symmetry about the mesh point or symmetry about the midpoint. The latter approach corresponds to

$$u_{-1} = u_0 \qquad \text{and} \qquad u_n = u_{n-1}, \tag{15}$$

and leads to the difference matrix

$$L_R = \begin{bmatrix} 1 & -1 & 0 & \ldots & 0 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ 0 & \ldots & 0 & -1 & 1 \end{bmatrix}. \tag{16}$$

This matrix can be diagonalized with the aid of the discrete cosine transform DCT-2; see [27]. We will refer to this transform simply as DCT. It is defined by the matrix $C^{(n)} = [c_{ij}^{(n)}] \in \mathbb{R}^{n\times n}$ with entries

$$c_{ij}^{(n)} = \sqrt{\frac{2-\delta_{i,1}}{n}} \cos\left(\frac{(i-1)(2j-1)\pi}{2n}\right), \qquad i,j = 1,2,\ldots,n,$$

where $\delta_{i,1} = 1$ if $i = 1$ and zero otherwise. The spectral factorization of L_R is given by

$$L_R = (C^{(n)})^T \mathrm{diag}(z(\mathbf{t}))C^{(n)}, \tag{17}$$

where

$$\mathbf{t} = [t_1, t_2, \ldots, t_n]^T, \qquad t_i = \frac{(i-1)\pi}{n}, \qquad i = 1,2,\ldots,n.$$

The matrix L_R has the eigenvalues

$$\lambda_i = [C^{(n)}(L_R\mathbf{e}_1)]_i / [C^{(n)}\mathbf{e}_1]_i = \sqrt{\frac{n}{2-\delta_{i,1}}}\,[C^{(n)}(L_R\mathbf{e}_1)]_i.$$

In particular, $\mathbf{1} = [1,1,\ldots,1]^T$ is an eigenvector associated with the eigenvalue zero. Thus, L_R is singular.

We may apply different boundary conditions at the left and right end points. The matrix

$$L_{DR} = \begin{bmatrix} 2 & -1 & 0 & \ldots & 0 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ 0 & \ldots & 0 & -1 & 1 \end{bmatrix} \tag{18}$$

corresponds to applying the Dirichlet boundary condition at the left end point and the Neumann boundary condition at the right one. Similarly, the matrix

$$L_{RD} = \begin{bmatrix} 1 & -1 & 0 & \ldots & 0 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ 0 & \ldots & 0 & -1 & 2 \end{bmatrix} \tag{19}$$

is obtained when using the Neumann condition at the left end point and the Dirichlet condition at the right one. Both the matrices L_{DR} and L_{RD} are invertible and can be diagonalized by trigonometric transforms; for instance, L_{RD} can be diagonalized by the DCT-8 transform.

The A-weighted pseudoinverse (8) simplifies to L^{-1} when L is invertible. This makes it attractive to use invertible regularization matrices, such as (18) and (19), which were considered in [7]. Clearly, L_{DR} is effective only if the desired solution $\hat{\mathbf{x}}$ vanishes at the left end point, while L_{RD} is effective only if $\hat{\mathbf{x}}$ is zero at the right end point. Similarly, L_D performs well if $\hat{\mathbf{x}}$ vanishes

at both end points. These observations are illustrated by computed results in [7]. Indeed, it is evident from the shape of $\hat{\mathbf{x}}$ that L_D is the best regularization matrix for Examples 1, 3, and 4 of [7], while L_{DR} is the best regularization matrix for Examples 2 and 5. We remark that L_D is very effective for Example 4 in [7], not only because the desired solution $\hat{\mathbf{x}}$ vanishes at both end points, but also because it is a uniform sampling of a scalar multiple of the function $\sin(t)$, $t \in [0, \pi]$. Indeed, the spectral factorization (13) of L_D shows that the eigenvector associated with the smallest eigenvalue, $2 - 2\cos(\frac{\pi}{n+1})$, is given by

$$\left[\sin\left(\frac{\pi}{n+1}\right), \sin\left(\frac{2\pi}{n+1}\right), \ldots, \sin\left(\frac{n\pi}{n+1}\right)\right]^T,$$

which is a uniform sampling of $\sin(t)$, $t \in [0, \pi]$. Therefore, the desired solution is recovered quite accurately.

2.2. Antireflective boundary conditions

The symmetry imposed by Neumann boundary conditions (15) preserves the continuity of the solution at the boundary. Serra Capizzano [26] proposed an antisymmetric extension,

$$u_{-1} - u_0 = u_0 - u_1 \qquad \text{and} \qquad u_n - u_{n-1} = u_{n-1} - u_{n-2},$$

which preserves both the continuity of the solution and of its first derivative at the boundary. For a recent survey on antireflective boundary conditions; see [12]. These boundary conditions give rise to the second difference matrix

$$L_{AR} = \begin{bmatrix} 0 & 0 & 0 & \ldots & 0 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ 0 & \ldots & 0 & 0 & 0 \end{bmatrix}. \tag{20}$$

Zero-padding of the finite difference matrices (6) and (7) has previously been proposed in [25] without using the connection to antireflective boundary conditions.

The matrix L_{AR} is not normal and therefore cannot be diagonalized by a unitary similarity transformation. Nevertheless, it is shown in [1] that the eigenvector matrix of L_{AR} is a modification of a unitary matrix with a structure that makes fast diagonalization possible. Specifically, the eigenvector matrix of L_{AR} can be chosen as

$$T_{AR} = \left[\begin{array}{c|c|c} \frac{1}{\sqrt{n}}\mathbf{1} & \begin{matrix}\mathbf{0}^T \\ S^{(n-2)} \\ \mathbf{0}^T\end{matrix} & \mathbf{p} \end{array}\right], \tag{21}$$

where $S^{(n)}$ denotes the DST matrix of order n as in (13) and

$$\mathbf{p} = [p_1, p_2, \ldots, p_n]^T, \qquad p_i = \sqrt{\frac{n(2n-1)}{6(n-1)}}\left(\frac{i-1}{n-1}\right).$$

The vector $\mathbf{p}$ is a uniform sampling of a linear function scaled so that $\|\mathbf{p}\| = 1$. It is shown in [1] that the inverse of (21) has a structure similar to that of T_{AR}. For the computations of the present paper, we only need to be able to evaluate matrix-vector products with the inverse rapidly. This can be done by applying the Sherman-Morrison-Woodbury formula and the DST. The latter can be applied because columns 2 through $n-1$ of T_{AR} are uniform samplings of sine functions. The computations are analogous to those described in [10]; see also the following subsection. Note that we can safely apply the Sherman-Morrison-Woodbury formula, because T_{AR} is a low-rank correction of a unitary matrix. Therefore, the computations are stable.

The spectral decomposition of L_{AR} is given by

$$L_{AR} = T_{AR} D_{AR} T_{AR}^{-1}, \qquad D_{AR} = \mathrm{diag}(z(\mathbf{t})), \tag{22}$$

where

$$\mathbf{t} = [t_1, t_2, \ldots, t_n]^T, \quad t_i = \frac{i-1}{n-1}\pi, \quad i = 1, 2, \ldots, n-1; \quad t_n = 0.$$

We note that

$$\mathcal{N}(L_{AR}) = \mathrm{span}\{\mathbf{1}, \mathbf{p}\}. \tag{23}$$

Therefore, L_{AR} should be a more effective smoothing matrix than L_D, L_{DR}, and L_{RD} when the desired solution $\hat{\mathbf{x}}$ has a significant linear component.

The regularization matrix L_{AR} can be applied conveniently as follows. We compute the QR factorization of the transpose of the matrix (7). This can be done efficiently in only $\mathcal{O}(n)$ arithmetic floating-point operations. The QR factorization and knowledge of $\mathcal{N}(L_{AR})$ can be used to evaluate matrix-vector products with the Moore-Penrose pseudoinverse $L_{AR}^\dagger$ inexpensively; see [25, Proposition 2.1]. Moreover, matrix-vector products with the matrices $L_A^\dagger$ and $AL_A^\dagger$ can be computed quite cheaply. Hence, we can solve (9) or (11) by an Arnoldi-type method and the main computational expense typically is the evaluation of matrix-vector products with the matrix A.

2.3. High-order boundary conditions

The regularization matrix (20) is well suited for Tikhonov regularization (4) when the desired solution $\hat{\mathbf{x}}$ of (3) has a significant linear component. This section discusses a generalization that is well suited for the situation when $\hat{\mathbf{x}}$ is known to be close to other kinds of functions, such as a quadratic or an exponential function. The basic idea is to replace the samplings of the constant and linear functions in (22) by samplings of functions that are close to $\hat{\mathbf{x}}$. This approach has been applied in [10] in the context of image deblurring to modify the blurring matrix. It illustrates an approach suggested by Strang, who at the very end of [27] writes "*We hope that the eigenvector approach will suggest more new transforms, and that one of them will be fast and visually attractive.*" This section defines a regularization matrix L_H with desirable properties by prescribing its eigenvectors and eigenvalues.

Since we are considering a symmetric regularization matrix (the second difference matrix), we use the cosine basis as our starting point. However, a similar approach can be implemented by applying the Fourier basis when the regularization matrix is not symmetric. Let $C^{(n)}$ denote the DCT matrix of order n as in (17). Up to normalization, the $(j+1)$st column of $(C^{(n-2)})^T$ is $[\cos(jt_1), \cos(jt_2), \ldots, \cos(jt_{n-2})]^T$, where

$$t_i = \frac{2i-1}{2n-4}\pi, \qquad i = 1, 2, \ldots, n-2. \tag{24}$$

It is convenient to also define the grid points

$$t_0 = -\frac{\pi}{2n-4}, \qquad t_{n-1} = \frac{2n-3}{2n-4}\pi. \tag{25}$$

Then the t_i are equidistant nodes in the interval $[t_0, t_{n-1}]$.

Introduce the vectors

$$\mathbf{c}_i = [c_{i,1}, c_{i,2}, \ldots, c_{i,n-2}]^T, \qquad i \in \{0, n-1\},$$

with entries

$$c_{i,j} = \sqrt{\frac{2-\delta_{j,1}}{n-2}}\cos((j-1)t_i).$$

It follows from $t_{n-1} = \pi - t_0$ that $c_{n-1,j} = (-1)^{j-1}c_{0,j}$ for $1 \le j \le n-2$.

Define the vectors

$$\mathbf{h}_1 = [h_{1,0}, h_{1,1}, \ldots, h_{1,n-1}]^T, \qquad \mathbf{h}_2 = [h_{2,0}, h_{2,1}, \ldots, h_{2,n-1}]^T.$$

In our application, these vectors will be samplings of real-valued functions h_1 and h_2 at the n equidistant nodes (24)–(25), i.e.,

$$\mathbf{h}_1 = [h_1(t_0), h_1(t_1), \ldots, h_1(t_{n-1})]^T, \qquad \mathbf{h}_2 = [h_2(t_0), h_2(t_1), \ldots, h_2(t_{n-1})]^T.$$

Introduce the eigenvector matrix for L_H,

$$T_H = \left[\begin{array}{c|c|c} \mathbf{h}_1 & \begin{array}{c}\mathbf{c}_0^T \\ (C^{(n-2)})^T \\ \mathbf{c}_{n-1}^T\end{array} & \mathbf{h}_2 \end{array}\right], \tag{26}$$

where the eigenvectors $\mathbf{h}_1$ and $\mathbf{h}_2$ should be chosen so that the matrix T_H is invertible. The regularization matrix L_H is given by

$$L_H = T_H D_H T_H^{-1}, \tag{27}$$

where $D_H = \mathrm{diag}(z(\mathbf{s}))$ and $z(\mathbf{s})$ is the symbol (14) evaluated at

$$\mathbf{s} = [s_1, s_2, \ldots, s_n]^T, \quad s_i = \frac{i-2}{n-2}\pi, \quad i = 2, 3, \ldots, n-1; \quad s_1 = s_n = 0.$$

The eigenvalues $z(s_1)$ and $z(s_n)$ of L_H vanish. Therefore, the vectors $\mathbf{h}_1$ and $\mathbf{h}_2$ are in $\mathcal{N}(L_H)$. The remaining eigenvalues $z(s_i)$, $i = 2, 3, \ldots, n-1$, of T_H

are the eigenvalues of the matrix L_R of order $n-2$; cf. (16). In particular, $z(s_2) = 0$ and it follows that $\mathbf{1}$ is an eigenvector. Thus,

$$\mathcal{N}(L_H) = \text{span}\{\mathbf{1}, \mathbf{h}_1, \mathbf{h}_2\}. \tag{28}$$

If the functions h_1 and h_2 are smooth, then the spectral decomposition (27) defines a smoothing regularization matrix. We expect the regularization matrix L_H to yield a better approximation of $\hat{\mathbf{x}}$ than the matrix L_{AR} for properly chosen vectors $\mathbf{h}_1$ and $\mathbf{h}_2$. This is illustrated in Section 4.

2.4. Computation with $L_H^\dagger$

This subsection describes how the Tikhonov minimization problem

$$\min_{\mathbf{x}\in\mathbb{R}^n} \left\{ \|A\mathbf{x} - \mathbf{b}\|^2 + \mu \|L_H \mathbf{x}\|^2 \right\} \tag{29}$$

can be solved. The solution method can be used for the regularization matrix L_{AR} as well, and we include this matrix in our discussion.

Let $X \in \{AR, H\}$. The eigenvector matrix T_X for L_X is not orthogonal; however, it is almost orthogonal in the following sense.

Proposition 1. *Let $X \in \{AR, H\}$. Then T_X is a rank-2 modification of an orthogonal matrix. Let $P = [\mathbf{0} \,|\, I_{n-2} \,|\, \mathbf{0}]_{(n-2)\times n}$, where I_{n-2} is the identity matrix of order $n-2$. Then*

$$PT_X^{-1}P^T = PT_X^T P^T + R_X \tag{30}$$

for a matrix R_X with

$$\text{rank}(R_X) = \begin{cases} 0 & \text{if } X = AR, \\ 2 & \text{if } X = H. \end{cases}$$

Proof. The fact that T_X is a rank-2 modification of an orthogonal matrix follows from (21) and (26). Property (30) is a consequence of [1, eq. (14)] and [10, Theorem 3.1]. □

The regularization matrices L_H and L_{AR} have the spectral factorizations

$$L_X = T_X D_X T_X^{-1}, \qquad X \in \{AR, H\} \tag{31}$$

with nonsingular eigenvector matrices T_X; see (22) and (27). The null space of L_X is given by (23) or (28), and is of small dimension.

Proposition 2. *Let $X \in \{AR, H\}$. The Moore-Penrose pseudoinverse of the matrix $L_X \in \mathbb{R}^{n\times n}$ is given by*

$$L_X^\dagger = P_{\mathcal{N}(L_X)}^\perp T_X D_X^\dagger T_X^{-1} P_{\mathcal{R}(T_X D_X)}, \tag{32}$$

where $D_X^\dagger$ denotes the Moore-Penrose pseudoinverse of the matrix D_X, $P_{\mathcal{Y}}$ is the orthogonal projector onto the subspace $\mathcal{Y}$, $P_{\mathcal{Y}}^\perp$ the orthogonal projector onto the orthogonal complement of $\mathcal{Y}$, and $\mathcal{R}(M)$ stands for the range of the matrix M.

Proof. The fact that the matrix $L_X^\dagger$ defined by (32) is a pseudoinverse can be shown in a variety of ways, e.g., by showing that $L_X^\dagger$ satisfies the requirements of the definition in [9]. We present a proof that shows how $L_X^\dagger$ can be applied in computations.

Let $\mathbf{d} \in \mathbb{R}^n$. The solution of minimal Euclidean norm of the least-squares problem

$$\min_{\mathbf{y}\in\mathbb{R}^n} \|L_X\mathbf{y} - \mathbf{d}\| \tag{33}$$

is given by $L_X^\dagger\mathbf{d}$. The least-squares problem (33) is equivalent to the linear system of equations

$$L_X\mathbf{y} = P_{\mathcal{R}(T_X D_X)}\mathbf{d}.$$

We will construct $P_{\mathcal{R}(T_X D_X)}$ by using that $\mathcal{R}(T_X D_X)$ is orthogonal to $\mathcal{N}(D_X T_X^T)$. Let $D_X = \mathrm{diag}[d_1, d_2, \ldots, d_n]$ and define the index set

$$\mathcal{I}_{D_X} = \{j \,:\, d_j = 0\}.$$

This set contains only few elements; $\mathcal{I}_{D_{AR}}$ contains two indices and $\mathcal{I}_{D_H}$ contains three; cf. (22) and (28).

Let $\mathbf{e}_j = [0, \ldots, 0, 1, 0, \ldots, 0]^T$ denote the jth axis vector in $\mathbb{R}^n$. Then

$$\mathcal{N}(D_X T_X^T) = \{T_X^{-T}\mathbf{e}_j \,:\, j \in \mathcal{I}_{D_X}\}.$$

The Gram–Schmidt process can be applied to determine an orthonormal basis for $\mathcal{N}(D_X T_X^T)$. The structure of the matrix T_X makes it possible to determine the required vectors $T_X^{-T}\mathbf{e}_j$ inexpensively. Each vector $T_H^{-T}\mathbf{e}_j$ can be computed in only $\mathcal{O}(n \log n)$ arithmetic floating-point operations; see Subsections 2.2 and 2.3. Moreover, the matrix T_{AR}^{-1} is explicitly known (see [1]) and, therefore, the determination of $T_{AR}^{-T}\mathbf{e}_j$ does not require any computations.

Let the columns of the matrix $Q_{\mathcal{N}(D_X T_X^T)}$ form an orthonormal basis for $\mathcal{N}(D_X T_X^T)$. Then

$$P_{\mathcal{R}(T_X D_X)} = I - Q_{\mathcal{N}(D_X T_X^T)} Q_{\mathcal{N}(D_X T_X^T)}^T. \tag{34}$$

Introduce the Moore-Penrose pseudoinverse

$$D_X^\dagger = \mathrm{diag}[d_1^\dagger, d_2^\dagger, \ldots, d_n^\dagger], \qquad d_j^\dagger = \begin{cases} 1/d_j & j \notin \mathcal{I}_{D_X}, \\ 0 & j \in \mathcal{I}_{D_X}. \end{cases}$$

Then $T_X D_X^\dagger T_X^{-1} P_{\mathcal{R}(T_X D_X)}\mathbf{d}$ is a solution of (33). The minimal-norm solution is orthogonal to

$$\mathcal{N}(L_X) = \mathcal{N}(D_X T_X^{-1}) = \{T_X\mathbf{e}_j \,:\, j \in \mathcal{I}_{D_X}\}. \tag{35}$$

A basis for $\mathcal{N}(L_X)$ can be determined without any computations and be orthonormalized by the modified Gram–Schmidt process. Let the columns of the matrix $Q_{\mathcal{N}(L_X)}$ be made up of the orthonormal basis so obtained. Then

$$P_{\mathcal{N}(L_X)}^\perp = I - Q_{\mathcal{N}(L_X)} Q_{\mathcal{N}(L_X)}^T. \tag{36}$$

It follows that the minimal-norm solution of (33) is given by

$$\mathbf{y} = P_{\mathcal{N}(L_X)}^{\perp} T_X D_X^{\dagger} T_X^{-1} P_{\mathcal{R}(T_X D_X)} \mathbf{d}.$$

This shows (32). □

A matrix-vector product with the matrix $L_X^{\dagger}$ can be evaluated in only $\mathcal{O}(n)$ arithmetic floating-point operations by using the representations (32), (34), and (36), provided that the matrices $Q_{\mathcal{N}(D_X T_X^T)}$ and $Q_{\mathcal{N}(L_X)}$ in (34) and (36) are available. These matrices can be determined in at most $\mathcal{O}(n \log n)$ arithmetic floating-point operations as indicated in the proof of Proposition 2. Moreover, since $\mathcal{N}(L_X)$ is explicitly known, we can evaluate matrix-vector products with the matrices $L_A^{\dagger}$ and $AL_A^{\dagger}$, where $L_A^{\dagger}$ is given by (8), quite efficiently. This makes it feasible to solve (9) and (11) by an iterative method based on the Arnoldi process when the regularization matrix is given by (27).

We note that we may proceed similarly when using the regularization matrix L_{AR} as with the matrix L_H. However, it is generally attractive to instead exploit the structure of the matrix L_{AR} and organize the computations as outlined in the paragraph following equation (23).

3. First order differences and invertible regularization matrices

In the above section, second order difference matrices were used to describe how boundary conditions can be applied to determine square regularization matrices. This approach to define banded square regularization matrices can be applied to difference matrices of any order. However, when regularization matrices associated with odd order differences are modified by imposing Neumann or antireflective boundary conditions, the resulting banded matrices will not be diagonalizable by discrete cosine or sine transforms. When high-order boundary conditions are used for regularization matrices defined by odd order differences, the inner part of the eigenvector matrix T_H in (26) has to be defined with the aid of the discrete Fourier transform instead of the cosine transform [10].

This section briefly describes first order difference matrices, including the construction of invertible square regularization matrices. A related approach to determine invertible second order difference matrices also is commented on.

The following invertible forward and backward finite difference matrices have been proposed in [7]:

$$F_{FR} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ & 1 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & & 1 & -1 \\ 0 & \dots & 0 & 0 & \alpha \end{bmatrix}, \quad F_{BR} = \begin{bmatrix} \alpha & 0 & 0 & \dots & 0 \\ -1 & 1 & 0 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 1 & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}.$$

When $\alpha = 0$, the above matrices are equivalent and correspond to the use of Neumann boundary conditions at both end points. The application of these

matrices with $\alpha = 0$ in iterative methods is discussed in [25, Section 2.1]. Invertible regularization matrices can be obtained by letting $\alpha > 0$. The vanishing eigenvalue is replaced by an eigenvalue α. Regularization matrices F_{FR} and F_{BR} with $\alpha > 0$ are considered in [7]. Following [7], we will use $\alpha = 10^{-8}$ in the computed examples of Section 4. The matrices F_{FR} and F_{BR} cannot be diagonalized by the discrete cosine transform since they are not symmetric.

Invertible second difference matrices can be constructed similarly as F_{FR} and F_{BR}. For instance, we may append and prepend rows to the regularization matrix (7) to obtain

$$L_{AR}(\alpha) = \begin{bmatrix} \alpha & 0 & 0 & \dots & 0 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ 0 & \dots & 0 & 0 & \alpha \end{bmatrix}, \qquad \alpha > 0.$$

Thus, $L_{AR}(0) = L_{AR}$. The two vanishing eigenvalues of $L_{AR}(0)$ are replaced by α.

Turning to high-order boundary conditions, we define, analogously to (27),

$$L_H(\alpha) = T_H D_H(\alpha) T_H^{-1}.$$

Here the matrix T_H is the same as in (27) and the diagonal matrix $D_H(\alpha)$ is obtained by replacing the three vanishing eigenvalues, $z(s_1), z(s_2)$ and $z(s_n)$, of the matrix D_H in (27) by α. In particular, $L_H(0) = L_H$.

4. Numerical examples

We present a few computed experiments where we solve linear discrete ill-posed problems. Most of the examples are from the Regularization Tools package by Hansen [17]. The desired solution $\hat{\mathbf{x}}$ and the error-free right-hand side $\hat{\mathbf{b}}$ are available in all examples. The noise vector $\mathbf{e}$ has normally distributed entries with zero mean. The variance of the entries is chosen so that $\mathbf{e}$ achieves a specified noise level

$$\nu = \frac{\|\mathbf{e}\|}{\|\mathbf{b}\|}.$$

We assume that a bound

$$\|\mathbf{e}\| \leq \varepsilon$$

is available. Then we can apply the discrepancy principle to determine a suitable value of the regularization parameter μ in (4) or a suitable number of steps with the Arnoldi process applied to the solution of (11). All computations were carried out in MATLAB with machine epsilon about $2.22 \cdot 10^{-16}$.

When computing an approximate solution of (1) by truncated iteration, we first determine the matrices necessary to evaluate matrix-vector products for $L_A^\dagger$ and $AL_A^\dagger$ and then carry out suitably many, say p, steps of the range restricted

Arnoldi method described in [21, 22]. The computed approximate solution $\bar{\mathbf{x}}_p$ of (11) lives in the Krylov subspace

$$\mathbb{K}_p(AL_A^\dagger, AL_A^\dagger \bar{\mathbf{b}}) = \text{span}\{AL_A^\dagger \bar{\mathbf{b}}, AL_A^\dagger AL_A^\dagger \bar{\mathbf{b}}, \ldots, (AL_A^\dagger)^p \bar{\mathbf{b}}\}.$$

The discrepancy principle prescribes that the number of steps, p, with the range restricted Arnoldi method should be chosen to be the smallest integer such that

$$\|A\mathbf{x}_p - \mathbf{b}\| = \|AL_A^\dagger \bar{\mathbf{x}}_p - \bar{\mathbf{b}}\| \leq \eta\varepsilon,$$

where, analogously to (10),

$$\mathbf{x}_p = L_A^\dagger \bar{\mathbf{x}}_p + \mathbf{x}^{(0)},$$

and $\eta > 1$ is a user-specified constant. We refer to [21] for further details on the range restricted Arnoldi method and its application to the solution of linear discrete ill-posed problems.

In Tikhonov regularization, we first determine p as above and then choose the regularization parameter $\mu \geq 0$ such that the Tikhonov approximate solution $\mathbf{x}_{p,\mu} \in \mathbb{K}_p(AL_A^\dagger, AL_A^\dagger \bar{\mathbf{b}})$ of (9) satisfies

$$\|AL_A^\dagger \bar{\mathbf{x}}_{p,\mu} - \bar{\mathbf{b}}\| = \eta\varepsilon.$$

Then $\bar{\mathbf{x}}_{p,\mu}$ is transformed to $\mathbf{x}_{p,\mu}$ using (10) to obtain an approximate solution of (4). The value of the regularization parameter is computed with Newton's method. We remark that the computational effort for determining μ is negligible, because all computations are carried out with small matrices and vectors (of order about $p \times p$ and p, respectively); see, e.g., [20] for a description of similar computations.

We compare the performance of several square regularization matrices. The application of high-order boundary conditions requires information about the main components of the solution for defining the vectors $\mathbf{h}_1$ and $\mathbf{h}_2$. We will choose

$$\mathbf{h}_1 = \mathbf{n}_1/\|\mathbf{n}_1\|, \tag{37}$$

$$\mathbf{h}_2 = \mathbf{x}_{AR}/\|\mathbf{x}_{AR}\|, \tag{38}$$

where $\mathbf{n}_i = [1^i, 2^i, \ldots, n^i]^T$ and $\mathbf{x}_{AR}$ denotes the computed approximate solution of (1) obtained when imposing antireflective boundary conditions. Then

$$\mathcal{N}(L_H) = \text{span}\{\mathbf{n}_0, \mathbf{n}_1, \mathbf{x}_{AR}\}.$$

We compute both $\mathbf{x}_{AR}$ and the approximate solution $\mathbf{x}_H$ determined with high-order boundary conditions by either using truncated iteration or Tikhonov regularization based on the range restricted Arnoldi method.

In general, we may know some properties of $\hat{\mathbf{x}}$ that are helpful when determining a suitable regularization matrix L_H from the origin of the problem. Moreover, solving a problem with the regularization matrices $L = I$ or $L = L_{AR}$

Table 1: Example 1: Relative errors in approximate solutions computed by Tikhonov regularization based on the range restricted Arnoldi process for the problem phillips with a linear component added to the solution; $\nu = 1 \cdot 10^{-2}$.

Regularization matrix	p	$\|\mathbf{x}_p - \hat{\mathbf{x}}\|/\|\hat{\mathbf{x}}\|$
$(\hat{C}_2^\dagger D_\delta^{-1})^\dagger$	4	$1.4436 \cdot 10^{-2}$
$L_D(W)$	4	$1.5084 \cdot 10^{-2}$
L_{AR}	3	$8.4420 \cdot 10^{-3}$
L_H	1	$7.9935 \cdot 10^{-3}$
L_{DR}	9	$1.2121 \cdot 10^{-1}$
$L_{AR}(\alpha)$	7	$1.2412 \cdot 10^{-2}$
$L_H(\alpha)$	3	$1.2336 \cdot 10^{-2}$

may reveal further properties that determine the design of L_H. Unless explicitly stated otherwise, we use the latter approach in the computed examples below.

Our comparison also includes the $n \times n$ regularization matrices $(\hat{C}_2^\dagger D_\delta^{-1})^\dagger$ and $L_D(W)$ introduced in [25] and [11], respectively. Here $\hat{C}_2^\dagger$ is the Moore-Penrose pseudoinverse of the circulant matrix associated with finite difference approximation of the second derivative with the three smallest eigenvalues set to zero. Further,

$$D_\delta = \mathrm{diag}[\delta, 1, 1, \ldots, 1, \delta] \in \mathbb{R}^{n\times n}, \qquad \delta = 1 \cdot 10^{-8}.$$

The regularization matrix $(\hat{C}_2^\dagger D_\delta^{-1})^\dagger$ is designed not to damp slow oscillations and linear growth in the computed approximate solution very much; see [25] for details. The matrix $L_D(W)$ improves the properties of L_D defined in (12) by prescribing that the null space contains the range of W. Specifically,

$$L_D(W) = L_D(I - WW^T), \tag{39}$$

where $W = \mathrm{orth}[\mathbf{n}_0, \mathbf{n}_1]$ denotes an $n \times 2$ matrix whose columns are obtained by orthonormalizing the vectors $\{\mathbf{n}_0, \mathbf{n}_1\}$. Thus, $I - WW^T$ is an orthogonal projector.

This section compares different regularization matrices obtained from modifications of the second derivative matrix (7). Analogous results are obtained when modifying matrices associated with other derivatives such as (6).

Example 1. Consider the Fredholm integral equation of the first kind

$$\int_{-6}^{6} \kappa(\tau,\sigma)x(\sigma)d\sigma = g(\tau), \qquad -6 \le \tau \le 6, \tag{40}$$

with kernel and solution given by

$$\begin{aligned} \kappa(\tau,\sigma) &= x(\tau-\sigma), \\ x(\sigma) &= \begin{cases} 1+\cos(\frac{\pi}{3}\sigma), & \text{if } |\sigma| < 3, \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

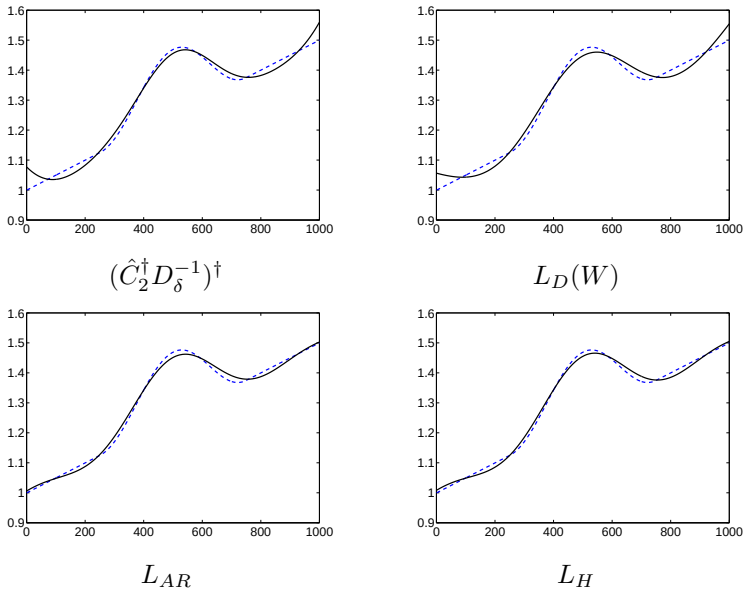

Figure 1: Example 1: Approximate solutions $\mathbf{x}_{p,\mu}$ (continuous graphs) computed by Tikhonov regularization based on the range restricted Arnoldi method for the problem phillips with a linear component added to the solution; $\nu = 1 \cdot 10^{-2}$. The dashed graphs depict the desired solution $\hat{\mathbf{x}}$.

The right-hand side $g(\tau)$ is defined by (40). This integral equation is discussed by Phillips [24]. The MATLAB code phillips in [17] determines a discretization by a Galerkin method with orthonormal box functions. This code also yields a discretization of a scaled solution $\mathbf{x}_0$. The matrix $A \in \mathbb{R}^{1000\times 1000}$ determined by phillips represents the discretized integral operator. We add a discretization of the function $1 + t/2$, $t \in [0, 1]$, to the vector $\mathbf{x}_0$ to obtain a slowly oscillatory and increasing solution $\hat{\mathbf{x}}$. The noise-free right-hand side is given by $\hat{\mathbf{b}} = A\hat{\mathbf{x}}$ and the noise level is $\nu = 1 \cdot 10^{-2}$.

Table 1 shows L_H to give the smallest approximation error. However, all square regularization matrices proposed in this work yield quite accurate approximations of $\hat{\mathbf{x}}$; see Figure 1. □

Example 2. The Fredholm integral equation of the first kind

$$\int_0^\pi \kappa(\sigma,\tau)x(\tau)d\tau = g(\sigma), \qquad 0 \le \sigma \le \frac{\pi}{2}, \tag{41}$$

with $\kappa(\sigma,\tau) = \exp(\sigma\cos(\tau))$, $g(\sigma) = 2\sinh(\sigma)/\sigma$, and solution $x(\tau) = \sin(\tau)$ is discussed by Baart [2]. We use the MATLAB code baart in [17] to determine a discretization by a Galerkin method with 400 orthonormal box functions. The code produces the matrix $A \in \mathbb{R}^{400\times 400}$ and a scaled discrete approximation $\mathbf{x}_0$ of $x(\tau)$. We add a discretization of the function $\frac{1}{37}((t-\frac{\pi}{2})^2 - \frac{t}{9})$, $t \in [0, \pi/2]$, to

Table 2: Example 2: Relative errors in approximate solutions computed by Tikhonov regularization based on the range restricted Arnoldi process for the problem baart with a quadratic component added to the solution; $\nu = 1 \cdot 10^{-6}$.

Regularization matrix	p	$\|\mathbf{x}_p - \hat{\mathbf{x}}\|/\|\hat{\mathbf{x}}\|$
$(\hat{C}_2^\dagger D_\delta^{-1})^\dagger$	2	$3.1862 \cdot 10^{-3}$
$L_D(W)$	3	$3.0572 \cdot 10^{-2}$
L_{AR}	2	$8.1448 \cdot 10^{-4}$
L_H	1	$8.1700 \cdot 10^{-4}$
L_{DR}	6	$1.2456 \cdot 10^{-1}$
$L_{AR}(\alpha)$	5	$3.8880 \cdot 10^{-3}$
$L_H(\alpha)$	5	$3.8889 \cdot 10^{-3}$

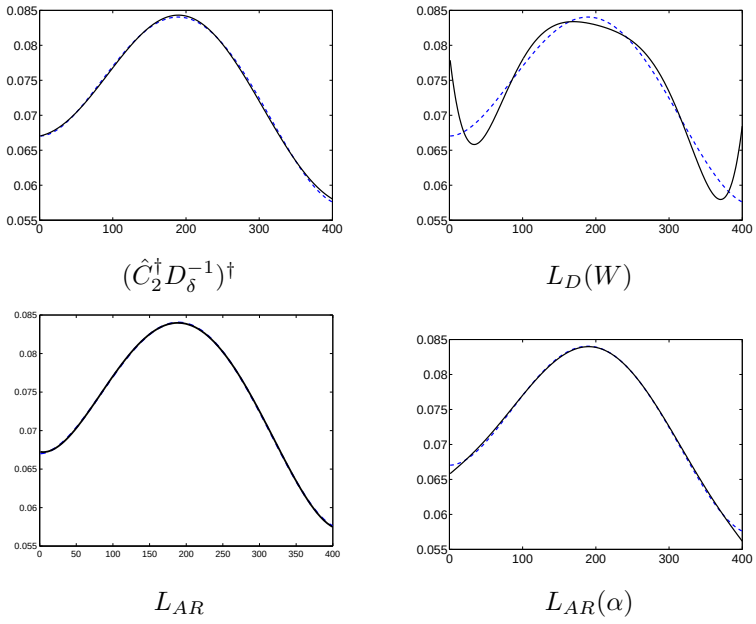

Figure 2: Example 2: Approximate solutions $\mathbf{x}_p$ (continuous graphs) computed by Tikhonov regularization based on the range restricted Arnoldi process for the problem baart with a quadratic component added to the solution; $\nu = 1 \cdot 10^{-6}$. The dashed graphs depict the desired solution $\hat{\mathbf{x}}$.

Table 3: Example 3: Relative errors in approximate solutions computed by Tikhonov regularization based on the range restricted Arnoldi process for the modified problem deriv2 with $\beta = 7$ and $\nu = 1 \cdot 10^{-3}$.

Regularization matrix	p	$\|\mathbf{x}_p - \hat{\mathbf{x}}\|/\|\hat{\mathbf{x}}\|$
$(\hat{C}_2^\dagger D_\delta^{-1})^\dagger$	2	$3.4819 \cdot 10^{-1}$
$L_D(W)$	3	$4.0230 \cdot 10^{-1}$
L_{AR}	3	$2.7193 \cdot 10^{-1}$
L_H	2	$2.6044 \cdot 10^{-1}$
L_{DR}	4	$2.3939 \cdot 10^{-1}$
$L_{AR}(\alpha)$	4	$2.0997 \cdot 10^{-1}$
$L_H(\alpha)$	3	$2.2359 \cdot 10^{-1}$
L_{DAR}	3	$2.1844 \cdot 10^{-1}$
$\hat{L}_H$	1	$5.9357 \cdot 10^{-2}$
$L_D(\hat{W})$	1	$9.2834 \cdot 10^{-2}$

the vector $\mathbf{x}_0$ to obtain a slowly oscillatory solution $\hat{\mathbf{x}}$. The noise-free right-hand side is given by $\hat{\mathbf{b}} = A\hat{\mathbf{x}}$ and the noise level is $\nu = 1 \cdot 10^{-6}$.

Table 2 and Figure 2 show L_{AR} to yield the best approximation of the desired solution $\hat{\mathbf{x}}$. The computed solution is indistinguishable from $\hat{\mathbf{x}}$. Also the regularization matrix L_H gives a very accurate approximation of $\hat{\mathbf{x}}$, and, furthermore, $L = (\hat{C}_2^\dagger D_\delta^{-1})^\dagger$ performs well. We remark that L_{AR} performs well for larger noise levels as well. When the right-hand side vector $\mathbf{b}$ is contaminated by more noise, fewer steps, p, are required with the range restricted Arnoldi method. □

Example 3. We consider the Fredholm integral equation of the first kind

$$\int_0^\beta k(s,t)x(t)dt = \exp(s) + (1-\mathrm{e})s - 1, \quad 0 \le s \le \beta, \tag{42}$$

where

$$k(s,t) = \begin{cases} s(t-1), & s < t, \\ t(s-1), & s \ge t. \end{cases} \tag{43}$$

For $\beta = 1$ this equation is discussed, e.g., by Delves and Mohamed [8]. We discretize the integral equation by a Galerkin method with orthonormal box functions as test and trial functions using a modification of the MATLAB program deriv2 from Regularization Tools [17]. Our program yields a symmetric indefinite matrix $A \in \mathbb{R}^{700\times 700}$ and a scaled discrete approximation $\hat{\mathbf{x}} \in \mathbb{R}^{700}$ of the solution $x(t) = \exp(t)$ of (42). The error-free right-hand side vector is given by $\hat{\mathbf{b}} = A\hat{\mathbf{x}}$, and the noise-level is $\nu = 4 \cdot 10^{-3}$.

We would like to illustrate the performance of the regularization matrices for a rapidly increasing solution and therefore let $\beta = 7$. Several of the regularization matrices discussed in this paper do not perform well; see Table 3.

This example provides an opportunity to illustrate the flexibility of our approach to construct regularization matrices by imposing boundary conditions.

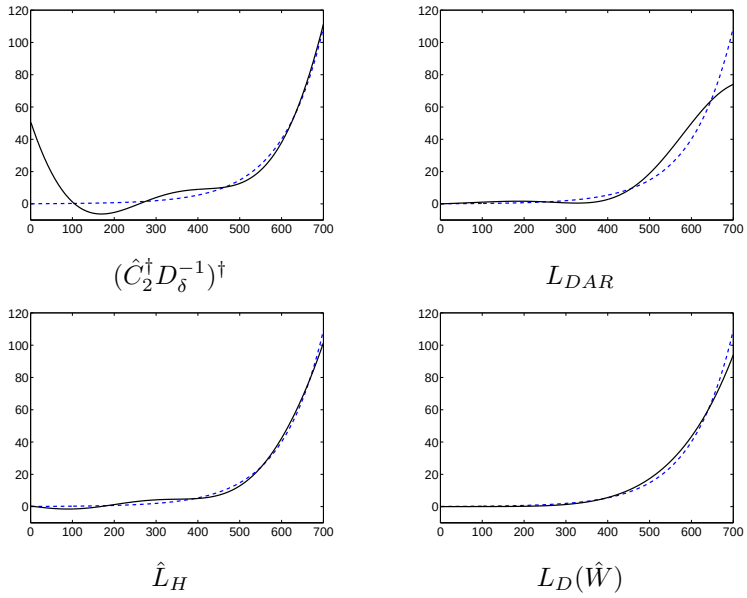

Figure 3: Example 3: Approximate solutions $\mathbf{x}_p$ (continuous graphs) computed by Tikhonov regularization based on the range restricted Arnoldi process for the modified problem deriv2 with $\beta = 7$ and $\nu = 1 \cdot 10^{-3}$. The dashed graphs depict the desired solution $\hat{\mathbf{x}}$.

For instance, requiring Dirichlet boundary conditions on the left-hand side and antireflective boundary conditions on the right-hand side for the second order difference matrix, we obtain

$$L_{DAR} = \begin{bmatrix} 2 & -1 & 0 & \dots & 0 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ 0 & \dots & 0 & 0 & 0 \end{bmatrix}$$

with $\mathcal{N}(L_{DAR}) = \text{span}\{\mathbf{n}_1\}$. Table 3 shows L_{DAR} to give an approximate solution of higher quality than both L_{AR} and L_{DR}.

Figure 3 indicates that the main difficulty when approximating the desired solution $\hat{\mathbf{x}}$ is to capture its rapid growth on the right-hand side. This suggests that it may be possible to determine an accurate approximation of $\hat{\mathbf{x}}$ by using a regularization matrix L_H, defined by (27), that has a vector corresponding to the sampling of a rapidly growing function in its null space. We therefore let $\mathbf{z}$ be the discretization of the rapidly increasing function t^5, $t \in [0, 1]$, and define the regularization matrix L_H with $\mathbf{h}_1$ given by (37) and $\mathbf{h}_2 = \mathbf{z}/\|\mathbf{z}\|$. We denote the regularization matrix so defined by $\hat{L}_H$ in Table 3 and Figure 3. Analogously, we let $\hat{W} = \text{orth}[\mathbf{n}_0, \mathbf{n}_1, \mathbf{z}]$ and define the regularization matrix $L_D(\hat{W})$; cf. (39).

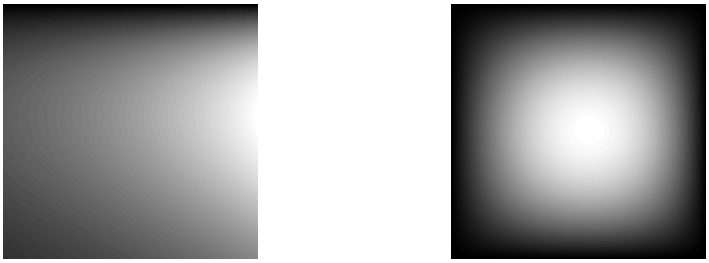

$\widehat{X}$ B

Figure 4: Example 4: The desired solution $\widehat{X}$ and the right-hand side B of the matrix equation (44) obtained with the MATLAB function deriv2 with an exponential solution in the x-direction and a linear plus $\sin(t)/(5t+5)$, $t \in [0,\, \pi]$, solution in the y-direction; $\nu = 1 \cdot 10^{-3}$.

Table 4: Example 4: Relative errors in approximate solutions computed by Tikhonov regularization based on the range restricted Arnoldi process for the two space-dimensional problem of Figure 4; $\nu = 1 \cdot 10^{-3}$.

Regularization matrix	p	$\|\mathbf{x}_p - \hat{\mathbf{x}}\|/\|\hat{\mathbf{x}}\|$
$(\hat{C}_2^\dagger D_\delta^{-1})^\dagger$	67	$1.7865 \cdot 10^{-1}$
$L_D(W)$	27	$1.5311 \cdot 10^{-1}$
L_{AR}	16	$1.0402 \cdot 10^{-1}$
L_H	6	$5.5555 \cdot 10^{-2}$
$L_D(\tilde{W})$	10	$7.1774 \cdot 10^{-2}$

Table 3 and Figure 3 show $\hat{L}_H$ to give the best approximation of $\hat{\mathbf{x}}$. Also the regularization matrix $L_D(\hat{W})$ yields a quite accurate approximation of $\hat{\mathbf{x}}$. □

Example 4. We consider a separable Fredholm integral equation of the first kind in two space-dimensions. In one space-dimension, we have equation (42) with $\beta = 1$, while in the orthogonal direction we consider the equation

$$\int_0^1 k(s,t)x(t)dt = (s^3 - s)/6, \quad 0 \le s \le 1,$$

where $k(s,t)$ is defined in (43) and the solution is $x(t) = t$. Let $A^{(1)} \in \mathbb{R}^{300\times 300}$, $\hat{\mathbf{x}}^{(1)}, \hat{\mathbf{x}}^{(2)} \in \mathbb{R}^{300}$ be computed with the MATLAB program deriv2 from Regularization Tools [17], where $\hat{\mathbf{x}}^{(1)} \in \mathbb{R}^{300}$ is a scaled discretized solution in the first direction and $\hat{\mathbf{x}}^{(2)} \in \mathbb{R}^{300}$ is the scaled discretized solution in the orthogonal direction. We add a discretization of the function $\sin(t)/(5t+5)$, $t \in [0,\, \pi]$, to the latter, and refer to the new solution so obtained also as $\hat{\mathbf{x}}^{(2)}$. Define the matrix $A = A^{(1)} \otimes A^{(1)}$ and the error-free solution in two space-dimensions by $\hat{\mathbf{x}} = \hat{\mathbf{x}}^{(1)} \otimes \hat{\mathbf{x}}^{(2)}$, where $\otimes$ denotes tensor product. The error-free right-hand side vector is given by $\hat{\mathbf{b}} = A\hat{\mathbf{x}}$ and the right-hand side vector $\mathbf{b}$ in (1) is obtained via (2). The noise-level is $\nu = 1 \cdot 10^{-3}$.

$L^{(1)} = L_{AR}$ $L^{(1)} = L_H$ $L^{(1)} = L_D(\tilde{W})$

Figure 5: Example 4: Approximate solutions $X_{p,\mu}$ computed by Tikhonov regularization based on the range restricted Arnoldi method for the problem in two space-dimensions of Figure 4; $\nu = 1 \cdot 10^{-3}$.

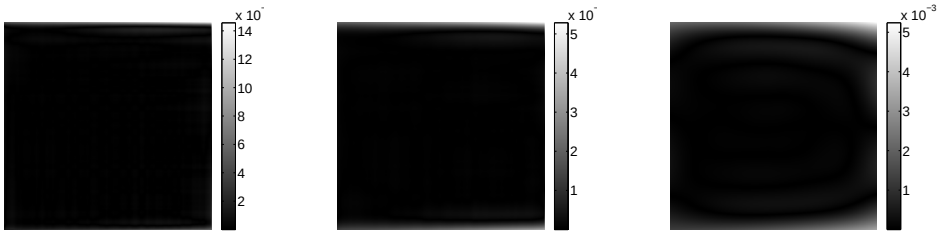

$L^{(1)} = L_{AR}$ $\text{AM} = 8.5022 \cdot 10^{-7}$ $L^{(1)} = L_H$ $\text{AM} = 2.4252 \cdot 10^{-7}$ $L^{(1)} = L_D(\tilde{W})$ $\text{AM} = 4.0479 \cdot 10^{-7}$

Figure 6: Example 4: Component-wise absolute error, $|\widehat{X} - X_{p,\mu}|$, with $X_{p,\mu}$ computed by Tikhonov regularization based on the range restricted Arnoldi method. The arithmetic mean of the absolute error is denoted by AM. The problem is in two space-dimensions; see Figure 4; $\nu = 1 \cdot 10^{-3}$.

A vector $\mathbf{z} \in \mathbb{R}^{300^2}$ can be stored columnwise as a 300×300 matrix; in MATLAB notation $Z = \mathsf{reshape}(\mathbf{z}, [300, 300])$. The structure of the present problem allows us to use this approach to express the linear system of equations (1) as

$$A^{(1)} X A^{(1)} = B, \tag{44}$$

where the matrices $X, \widehat{X}, B, \widehat{B}$ represent discretizations of surfaces. Figure 4 shows the desired solution $\widehat{X}$ and the observed right-hand side B of (44). The graphs are scaled so that black corresponds to the minimum value and white to the maximum value.

Let $L^{(1)}$ be a regularization matrix for a linear discrete ill-posed problem in one space-dimension. Regularization matrices of the form

$$L = \left((L^{(1)})^\dagger \otimes I + I \otimes (L^{(1)})^\dagger \right)^\dagger \tag{45}$$

were proposed in [11] for problems in two space-dimensions. These kinds of matrices are attractive modifications of "natural" regularization matrices

$$L_N = L^{(1)} \otimes I + I \otimes L^{(1)}$$

for problems in two space-dimensions, because $L^\dagger \hat{\mathbf{x}}$ can be evaluated efficiently as $(L^{(1)})^\dagger \widehat{X} + \widehat{X} (L^{(1)})^\dagger$.

The spectral factorization (31) is useful for defining regularization matrices for linear discrete ill-posed problems in two space-dimensions. Consider the matrix

$$L_N = L_X \otimes I + I \otimes L_X = (T_X \otimes T_X)(D_X \otimes I + I \otimes D_X)(T_X \otimes T_X)^{-1}$$

and define

$$L = \left((T_X \otimes T_X)(D_X \otimes I + I \otimes D_X)^\dagger (T_X \otimes T_X)^{-1} \right)^\dagger . \tag{46}$$

Then, if L_X is invertible, $L = L_N$; otherwise L is close to L_N due to Proposition 1. Furthermore, $L^\dagger \hat{\mathbf{x}}$ can be evaluated efficiently in $O(n^2 \log(n))$ arithmetic floating-point operations by using fast transforms.

In this example, we use the regularization matrix L given by (45) for $L^{(1)} \in \{(\hat{C}_2^\dagger D_\delta^{-1})^\dagger, L_D(\cdot)\}$ and L defined by (46) otherwise. We also can use (45) when antireflective boundary conditions are imposed by exploiting the band structure of $L^{(1)}$. However, the formula (46) usually provides more accurate computed approximate solutions for only slightly more computational effort; the number of arithmetic floating-point operations for evaluating $L^\dagger \hat{\mathbf{x}}$ is a factor $O(\log(n))$ larger.

For the matrices L in both (45) and (46), let the columns of $W^{(1)}$ form an orthonormal basis for $\mathcal{N}(L^{(1)})$. Then the columns of $W = W^{(1)} \otimes W^{(1)}$ form an orthonormal basis for $\mathcal{N}(L)$ and the computations can be carried out as described in [11].

For high-order boundary conditions, the basis function $\mathbf{h}_2 \in \mathbb{R}^{300}$ can not, like above, be defined using the solution computed by another method (e.g., by antireflective boundary conditions). We therefore let

$$\mathbf{h}_2 = \mathbf{n}_2/\|\mathbf{n}_2\|.$$

When comparing with the approach in [11], we define $\tilde{W} = \text{orth}[\mathbf{n}_0, \mathbf{n}_1, \mathbf{n}_2]$ and obtain the regularization matrix $L_D(\tilde{W})$.

Table 4 shows $L^{(1)} = L_H$ to give the best approximation of the solution to the error-free problem. The choice $L^{(1)} = (\hat{C}_2^\dagger D_\delta^{-1})^\dagger$ does not give a meaningful approximation, while antireflective boundary conditions are able to provide a fairly accurate restoration. Figure 5 shows some of the best restorations and Figure 6 displays the absolute error of each solution component. The differences in quality of the computed solutions is clearly visible. □

5. Conclusions

We have shown that the use of specific boundary conditions for finite difference approximation of derivatives can be a very effective way to define square smoothing regularization matrices. These matrices allow efficient solution of the associated Tikhonov regularization problem in standard form by an iterative method. The regularization matrices can be designed taking special features of the desired solution into account.

Acknowledgments

LR would like to thank Marco Donatelli for an enjoyable visit to Como, during which work on the present paper was carried out. The authors would like to thank a referee for valuable comments. Research by LR was supported in part by NSF grant DMS-1115385. Research by MD was partially supported by MIUR 2008, grant number 20083KLJEZ, and GNCS 2012 project "Precondizionamento e metodi Multigrid per il calcolo veloce di soluzioni accurate".

References

[1] A. Aricò, M. Donatelli, J. Nagy, S. Serra Capizzano, The Anti-Reflective Transform and Regularization by Filtering, in Numerical Linear Algebra in Signals, Systems, and Control, Lecture Notes in Electrical Engineering, vol. 80, Springer, 2011, pp. 1–21.

[2] M. L. Baart, The use of auto-correlation for pseudo-rank determination in noisy ill-conditioned least-squares problems, IMA Journal of Numerical Analysis 2 (1982) 241–247.

[3] C. Brezinski, M. Redivo-Zaglia, G. Rodriguez, S. Seatzu, Extrapolation techniques for ill-conditioned linear systems, Numer. Math. 81 (1998) 1–29.

[4] D. Calvetti, B. Lewis, L. Reichel, Restoration of images with spatially variant blur by the GMRES method, in Advanced Signal Processing Algorithms, Architectures, and Implementations X, ed. F. T. Luk, Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE), vol. 4116, The International Society for Optical Engineering, Bellingham, WA, 2000, 364–374.

[5] D. Calvetti, B. Lewis, L. Reichel, On the choice of subspace for iterative methods for linear discrete ill-posed problems, Int. J. Appl. Math. Comput. Sci. 11 (2001) 1069–1092.

[6] D. Calvetti, L. Reichel, A. Shuibi, L-curve and curvature bounds for Tikhonov regularization, Numer. Algorithms 35 (2004) 301–314.

[7] D. Calvetti, L. Reichel, A. Shuibi, Invertible smoothing preconditioners for linear discrete ill-posed problems, Appl. Numer. Math. 54 (2005) 135–149.

[8] L. M. Delves, J. L. Mohamed, Computational Methods for Integral Equations, Cambridge University Press, Cambridge, 1985.

[9] C. A. Desoer and B. H. Whalen, A note on pseudoinverses, J. Soc. Indust. Appl. Math. 11 (1963) 442–447.

[10] M. Donatelli, Fast transforms for high order boundary conditions in deconvolution problems, BIT 50 (2010) 559–576.

[11] M. Donatelli, A. Neuman, L. Reichel, Square regularization matrices for large linear discrete ill-posed problems, Numer. Linear Algebra Appl. 19 (2012) 896–913.

[12] M. Donatelli, S. Serra-Capizzano, Antireflective boundary conditions for deblurring problems, Journal of Electrical and Computer Engineering, vol. 2010 (2010) Article ID 241467 (18 pages).

[13] L. Dykes, L. Reichel, On the reduction of Tikhonov minimization problems and the construction of regularization matrices, Numer. Algorithms 60 (2012) 683–696.

[14] H. W. Engl, M. Hanke, A. Neubauer, Regularization of Inverse Problems, Kluwer, Dordrecht, 1996.

[15] L. Eldén, A weighted pseudoinverse, generalized singular values, and constrained least squares problems, BIT 22 (1982) 487–501.

[16] S. Gazzola, P. Novati, Automatic parameter setting for Arnoldi-Tikhonov methods, J. Comput. Appl. Math., 256 (2014), pp. 180–195.

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

[18] P. C. Hansen, Rank-Deficient and Discrete Ill-Posed Problems, SIAM, Philadelphia, 1998.

[19] P. C. Hansen, T. K. Jensen, Smoothing norm preconditioning for regularizing minimum residual methods, SIAM J. Matrix Anal. Appl. 29 (2006) 1–14.

[20] B. Lewis, L. Reichel, Arnoldi-Tikhonov regularization methods, J. Comput. Appl. Math. 226 (2009) 92–102.

[21] A. Neuman, L. Reichel, H. Sadok, Implementations of range restricted iterative methods for linear discrete ill-posed problems, Linear Algebra Appl. 436 (2012) 3974–3990.

[22] A. Neuman, L. Reichel, H. Sadok, Algorithms for range restricted iterative methods for linear discrete ill-posed problems, Numer. Algorithms 59 (2012) 325–331.

[23] M. K. Ng, R. H. Chan, W.-C. Tang, A fast algorithm for deblurring models with Neumann boundary conditions, SIAM J. Sci. Comput. 21 (1999) 851–866.

[24] D. L. Phillips, A technique for the numerical solution of certain integral equations of the first kind, Journal of the ACM 9 (1962) 84–97.

[25] L. Reichel, Q. Ye, Simple square smoothing regularization operators, Electron. Trans. Numer. Anal. 33 (2009) 63–83.

[26] S. Serra Capizzano, A note on anti-reflective boundary conditions and fast deblurring models, SIAM J. Sci. Comput. 25 (2003) 1307–1325.

[27] G. Strang, The discrete cosine transform, SIAM Rev. 41 (1999) 135–147.