

## **Distribution Agreement**

In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

---

Chang Meng

---

Date

Large Scale Inverse Problems: Low-Rank Approximations and Optimization

By

Chang Meng

Doctor of Philosophy

Mathematics

---

James G. Nagy

Advisor

---

Lars Ruthotto

Committee Member

---

Yuanzhe Xi

Committee Member

Accepted:

---

Kimberly Jacob Arriola, Ph.D., MPH

Dean of the James T. Laney School of Graduate Studies

---

Date

Large Scale Inverse Problems: Low-Rank Approximations and Optimization

By

Chang Meng

B.S., Emory University, 2016

Advisor: James G. Nagy, Ph.D.

An abstract of

A dissertation submitted to the Faculty of the  
James T. Laney School of Graduate Studies of Emory University

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in Mathematics

2022

## Abstract

Large Scale Inverse Problems: Low-Rank Approximations and Optimization

By Chang Meng

Inverse problems can be found in a variety of scientific applications, and the development of efficient and reliable methods remain an essential and challenging task. In this thesis, we introduce novel low-rank solvers for linear systems that arise from large scale inverse problems, which are usually ill-posed and require the use of regularization to obtain meaningful solutions. The new methods are developed around the concept of regularization: *i)* the low-rank, Kronecker product based forward model approximation method involves the approximation of a truncated singular value decomposition; and *ii)* the low-rank Krylov subspace methods are based on nuclear norm regularization. We explore the performance of these novel low-rank methods in various imaging applications such as image deblurring, inpainting and computer tomography. Besides applications where the forward model is known and fixed, we also consider an extended application, where the forward model is not exactly known and requires calibration. In this context, we are able to not only apply our new low-rank methods, but also propose a new hybrid machine learning and block coordinate descent algorithm that effectively improves solution accuracy.

Large Scale Inverse Problems: Low-Rank Approximations and Optimization

By

Chang Meng

B.S., Emory University, 2016

Advisor: James G. Nagy, Ph.D.

A dissertation submitted to the Faculty of the  
James T. Laney School of Graduate Studies of Emory University  
in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy  
in Mathematics

2022

## Acknowledgments

Words cannot describe my gratitude for my advisor, Dr. James Nagy, who have supported and inspired me for many years. Without him, I wouldn't have had such a memorable time at Emory. Jim, thank you so much for being my advisor and always being supportive of my choices! I have learned so much from you, and that includes not only knowledge and research, but also how to become a better person. For me, you are both a great professor and a wonderful life mentor.

Also, many thanks to my committee members Lars Ruthotto and Yuanzhe Xi, who have provided so much guidance on my research. I would also like to thank every math teacher and professor I have ever had as a student for showing me the fantastic world of mathematics that shaped me into who I am.

Last but not least, I wouldn't be where I am today without my parents, who have given me unconditional love and support. Mom and dad, thank you for all the sacrifices you have made for me to receive better education, and thank you for being patient with me and letting me take my own road. I am so lucky to be your daughter!

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Contributions of Work . . . . .	5
1.2	Outline of Thesis . . . . .	7
<b>2</b>	<b>Overview of Regularization Methods</b>	<b>9</b>
2.1	Tikhonov Regularization and TSVD . . . . .	10
2.2	Choosing the Regularization Parameter for Tikhonov Regularization . . . . .	14
2.3	Other Regularization Terms . . . . .	16
2.4	Iterative Regularization . . . . .	19
2.4.1	GMRES and LSQR . . . . .	19
2.4.2	Hybrid Krylov Methods . . . . .	22
2.4.3	Flexible Krylov Methods . . . . .	26
<b>3</b>	<b>TSVD approximations</b>	<b>29</b>
3.1	The Kronecker Product . . . . .	30
3.2	Kronecker Product Summation . . . . .	32

3.3	A New TSVD Approximation Method	
	via Kronecker Product Summations . . . . .	34
3.3.1	Approximation Quality . . . . .	37
<b>4</b>	<b>Low Rank Regularization: Derivation of New Methods</b>	<b>42</b>
4.1	Nuclear Norm Regularization . . . . .	44
4.2	Derivation of New Methods . . . . .	46
4.3	Solution of problem (4.8) via Krylov methods . . . . .	49
4.3.1	Methods Based on the GKB Algorithm . . . . .	50
4.3.2	Methods Based on the Arnoldi Algorithm . . . . .	51
4.4	Solution through flexible Krylov subspaces . . . . .	54
4.5	Implementation details . . . . .	58
<b>5</b>	<b>Low Rank Regularization: Numerical Experiments</b>	<b>62</b>
5.1	Low Rank Projection Methods: classical and new approaches . . . . .	64
5.1.1	Low-rank flexible GMRES (LR-FGMRES) and low-rank flexible LSQR (LR-FLSQR) . . . . .	66
5.2	Numerical Examples . . . . .	69
5.2.1	A Note on Regularization Parameters . . . . .	84
<b>6</b>	<b>Extended Application: Model Calibration for Computed Tomogra-</b>	
	<b>phy</b>	<b>88</b>
6.1	Parallel Tomography . . . . .	90
6.1.1	Numerical Examples . . . . .	95
6.2	Fan-Beam Tomography . . . . .	102



6.2.1	Hybrid ML-BCD Method . . . . .	108
6.2.2	Numerical Examples . . . . .	110
<b>7</b>	<b>Conclusions and Future Work</b>	<b>120</b>
	<b>Bibliography</b>	<b>123</b>

# List of Figures

1.1	Observed data. . . . .	2
5.1	<i>Example 1</i> . Relative errors. . . . .	71
5.2	<i>Example 1</i> . Exact, corrupted and reconstructed images. . . . .	72
5.3	<i>Example 1</i> . Zoom-ins of solutions. . . . .	73
5.4	<i>Example 1</i> . Zoom-ins of surface plots. . . . .	74
5.5	<i>Example 1</i> . Singular values. . . . .	75
5.6	<i>Example 2</i> . Relative errors. . . . .	77
5.7	<i>Example 2</i> . Exact and reconstructed images. . . . .	78
5.8	<i>Example 2</i> . Surface plots. . . . .	78
5.9	<i>Example 3</i> . Singular values. . . . .	80
5.10	<i>Example 3</i> ( <b>house</b> ). Relative errors. . . . .	81
5.11	<i>Example 3</i> ( <b>house</b> ). Exact, corrupted and reconstructed images. . . . .	81
5.12	<i>Example 3</i> ( <b>peppers</b> ). Relative errors. . . . .	82
5.13	<i>Example 3</i> ( <b>peppers</b> ). Exact, corrupted and reconstructed images. . . . .	83
5.14	Solutions without regularization. . . . .	87
6.1	Parallel tomography illustration from [36]. . . . .	91

6.2	Test images Brain (left) and Phantom (right). . . . .	95
6.3	Phantom reconstructions (with TwIST). . . . .	98
6.4	Convergence plots for Phantom (with TwIST). . . . .	99
6.5	Convergence plots for Phantom (with Krylov methods). . . . .	100
6.6	Phantom reconstructions (with Krylov methods). . . . .	101
6.7	Fan-beam Tomography illustration. . . . .	103
6.8	Relative errors in geometry parameters. . . . .	111
6.9	Phantom reconstructions using various methods. . . . .	112
6.10	Phantom reconstructions using various methods. . . . .	113
6.11	Examples of different phantoms used. . . . .	114
6.12	Relative errors in geometry parameters. . . . .	115
6.13	Histograms of relative errors. . . . .	117
6.14	Phantom reconstructions using various methods. . . . .	118

# List of Tables

5.1 Comparisons of minimum relative errors. . . . . 86

6.1 Comparisons of PSNRs (with TwIST). . . . . 97

# Chapter 1

## Introduction

Inverse problems are ubiquitous in the areas of engineering, science and medical applications, and accurate solutions often require computationally costly or time consuming methods. This thesis focuses on a class of imaging inverse problems that can be modeled as

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{1.1}$$


where  $\mathbf{A} \in \mathbb{R}^{M \times N}$  models the forward process,  $\mathbf{x} \in \mathbb{R}^N$  is the image of interest,  $\mathbf{b} \in \mathbb{R}^M$  is the observation that is usually contaminated with unknown noise  $\boldsymbol{\eta} \in \mathbb{R}$ , and  $\mathbf{b} = \mathbf{b}^{\text{ex}} + \boldsymbol{\eta}$ . In particular, the desired solution  $\mathbf{x}^{\text{ex}}$  is the exact solution to  $\mathbf{A}\mathbf{x} = \mathbf{b}^{\text{ex}}$ , where  $\mathbf{b}^{\text{ex}}$  is the noise-free observation. The vectors  $\mathbf{x}$  and  $\mathbf{b}$  are obtained by vectorizing the true image  $\mathbf{X} \in \mathbb{R}^{n \times n}$ ,  $n = \sqrt{N}$  and the observed or measured

data  $\mathbf{B}$ , i.e.,

$$\mathbf{x} = \text{vec}(\mathbf{X}) \quad , \quad \mathbf{X} = \text{vec}^{-1}(\mathbf{x}),$$

$$\mathbf{b} = \text{vec}(\mathbf{B}) \quad , \quad \mathbf{B} = \text{vec}^{-1}(\mathbf{b}).$$

In the image deblurring application, the forward model  $\mathbf{A}$  is constructed using a point spread function (PSF), which can be formulated based on knowledge of the physical process, and can be obtained using a precise mathematical expression or through experimentation [33]. For example, a two dimensional Gaussian function [39, 62] can be applied to construct the PSF for blurring caused by atmospheric turbulence. In this case, the observed data  $\mathbf{B} = \text{vec}^{-1}(\mathbf{b})$  is a blurred image; see Figure 1.1(a) for an example. In imaging applications arising from computed tomography (CT), the matrix  $\mathbf{A}$  models the Radon Transform [56] that outputs the projection data obtained from a tomographic scan. For CT, the observed data is a so-called “sinogram”, which stores the projection data; see Figure 1.1(b) for an example.



(a) Blurred image.

(b) CT sinogram.

Figure 1.1: Observed data.

Equivalently, we can rewrite (1.1) as the following least squares problem:

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2. \quad (1.2)$$

Large scale linear systems arising from inverse problems are usually very ill-conditioned, meaning that there might not be a unique solution, and a slight change in the data  $\mathbf{b}$  can lead to a large change in the solution  $\mathbf{x}$ . Hence, regularization is often needed in order to obtain a meaningful solution. Direct regularization methods add a penalty term  $\mathcal{R}(\mathbf{x})$  to the fit-to-data term  $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$ . The regularized optimization problem has the form

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \mathcal{R}(\mathbf{x}), \quad (1.3)$$

where the regularization term  $\mathcal{R}(\mathbf{x})$  is defined as a function of  $\mathbf{x}$  (e.g. a vector norm). By doing so, we would be able to prevent extreme values in  $\mathbf{x}$ . Popular choices for  $\mathcal{R}(\mathbf{x})$  include the  $\ell_2$  norm (Tikhonov regularization) and the  $\ell_1$  norm of  $\mathbf{x}$ . Various types of regularization approaches will be discussed in detail in Chapter 2.

When  $\mathbf{A}$  is not too large, i.e., if its dimensions  $M$  and  $N$  are a few thousands or smaller, direct factorization methods such as the SVD factorization (short for Singular Value Decomposition, introduced in Chapter 2) can be efficiently used to solve the least squares problem (1.2), or (1.3) with Tikhonov regularization. However, as the dimension of  $\mathbf{A}$  increases, the factorizations become increasingly difficult to compute (the complexity for computing the SVD of  $\mathbf{A} \in \mathbb{R}^{N \times N}$  is  $\mathcal{O}(N^3)$ ), but it is still possible to obtain their approximations. For example, while it is difficult

to compute the full SVD of large matrices, we may try to approximate the leading singular vectors and values to form a Truncated SVD (TSVD) factorization. The TSVD can be thought of as a low-rank approximation to matrix  $\mathbf{A}$ , and by discarding the smallest singular values and corresponding singular vectors, we are basically performing regularization by reducing the noise contribution of  $\boldsymbol{\eta}$  in the solution  $\mathbf{x}$ .

An alternative approach for solving the least squares problem (1.2) when  $\mathbf{A}$  is large is through iterative methods. The simplest example of iterative methods would be the Landweber iteration [46], which is essentially gradient descent applied to the function  $f(x) = \|\mathbf{Ax} - \mathbf{b}\|_2^2$ . Another class of iterative methods is called Krylov subspace methods, which look for solutions that belong to a Krylov subspace generated by  $\mathbf{A}$  and  $\mathbf{b}$ . Examples of Krylov methods include Conjugate Gradient, LSQR and GMRES (see [64] and the references therein). What has been observed when solving linear systems with iterative methods is that they exhibit a semi-convergence property (will be illustrated in figures in Section 5.2), that is, the relative errors in the solutions decrease in the first iterations, but will increase in the further iterations when noise starts to corrupt the solutions. Hence, iterations need to be stopped at a suitable iteration index. In literature, hybrid methods (see [8, 12, 21]) have been proposed to combine iterative methods with regularization when needed. These hybrid methods can be particularly useful when combined with a non-traditional regularization term such as TV (total variation) regularization [19], or the low-rank enforcing nuclear norm regularization [18].



## 1.1 Contributions of Work

The main contribution of this thesis is the development of novel low-rank solvers for large scale linear systems. The meaning of “low rank” here is two-fold:

- (1) low-rank approximation of forward model  $\mathbf{A}$ , or
- (2) low-rank regularization for solution  $\mathbf{x}$ .

With respect to low-rank approximation of  $\mathbf{A}$ , we introduce the newly proposed Kronecker product summation based TSVD approximation method [15, 16]. This new method can be used to efficiently approximate a large number of leading singular values and vectors for large, structured square matrices from ill-posed inverse problems in the application of image deblurring. This novel TSVD approximation method is based on a reordering technique that incorporates information from more terms in the Kronecker summation compared to similar methods (see [41, 52]). Furthermore, theoretical bounds for the approximate TSVD operator and the approximate TSVD filtered solution are also developed.

Regarding low-rank regularization for  $\mathbf{x}$ , we describe new solvers for the computation of low-rank approximate solutions to large-scale linear problems. We are mainly concerned with algorithms that solve the so-called nuclear norm regularized problem, where a suitable nuclear norm penalization on the solution is imposed alongside a fit-to-data term expressed in the 2-norm (i.e., setting  $\mathcal{R}(\mathbf{x})$  to be equal to the nuclear norm in (1.3)): this has the effect of implicitly enforcing low-rank solutions. By adopting an iteratively reweighted norm approach, the nuclear norm regularized problem is reformulated as a sequence of quadratic problems, which can then be effi-

ciently solved using Krylov methods, giving rise to an inner-outer iteration scheme. This approach differs from the other solvers available in the literature in that:

- (a) Kronecker product properties are exploited to define the reweighted 2-norm penalization terms;
- (b) efficient preconditioned Krylov methods replace gradient (projection) methods;
- (c) the regularization parameter can be efficiently and adaptively set during the iterations.

Furthermore, we reformulate within the framework of flexible Krylov methods both the new inner-outer methods for nuclear norm regularization and some of the existing Krylov methods incorporating low-rank projections. This results in an even more computationally efficient (but heuristic) strategy, that does not rely on an inner-outer iteration scheme. This is the first time Krylov methods have been used to solve a low-rank, nuclear norm regularized problem.

The new low rank Krylov methods are competitive with other state-of-the-art solvers for low-rank problems including image deblurring, computed tomography and inpainting, and deliver reconstructions of increased quality with respect to other classical Krylov methods. In addition to applying the newly developed methods to (1.3), we also consider another least squares problem

$$\min_{\mathbf{x}, \mathbf{p}} \|\mathbf{A}(\mathbf{p})\mathbf{x} - \mathbf{b}\|_2^2 + \mathcal{R}(\mathbf{x}), \quad (1.4)$$

where  $\mathbf{p}$  is a set of parameters that  $\mathbf{A}$  depends on. An example application in which this problem arises is computed tomography that requires geometry parameter

calibration. To solve (1.4), we could consider the framework of block coordinate descent (BCD), in which (1.4) is alternatively minimized with respect to  $\mathbf{x}$  and  $\mathbf{p}$  in each step. Since the minimization with respect to  $\mathbf{x}$  requires a linear solver, our new low rank solvers can be readily used, and they show better convergence properties than hybrid Krylov solvers applied with traditional regularization terms such as  $\ell_2$ .

We also propose a hybrid BCD method for solving (1.4) by incorporating machine learning. That is, we train a machine learning model  $\hat{\Phi}$  that maps  $\mathbf{b}$  to  $\mathbf{p}$ , and when given an observation  $\mathbf{b}$ , we first use the ML model to make a prediction for  $\mathbf{p}$ , which is then fed into the BCD algorithm as an initial guess. Through this approach, we are able to obtain reconstructions that are of higher quality compared to using BCD or ML alone.

## 1.2 Outline of Thesis

This thesis is organized as follows. In Chapter 2, we provide an overview of regularization methods, including TSVD, Tikhonov and iterative regularization, as they are the foundation to the development of low-rank based methods, which is the main contribution of this thesis. In Chapter 3, we describe the new Kronecker product summation based approach for approximating the TSVD of  $\mathbf{A}$ . In addition to presenting how the approximation is derived, we also provide theoretical results that demonstrate the efficacy of the TSVD, both for the operator  $\mathbf{A}$  and for the approximate solution to (1.1) filtered by the TSVD approximation. In Chapter 4, we introduce how to solve the nuclear norm regularized (NNR) using Krylov methods.

To be specific, we derive the iteratively reweighted norm (IRN) approach that solves a sequence of reformulated quadratic problems using Krylov methods, resulting in an inner-outer iteration algorithm. The IRN method is later reformulated within the flexible Krylov method framework that is considered more efficient (but heuristic). In Chapter 5, results of numerical experiments on several image processing problems are presented. In addition to comparing the newly developed low-rank Krylov methods to standard Krylov subspace methods, we also show comparisons with state-of-the-art solvers for low-rank problems, as well as existing Krylov methods incorporating low-rank projections reformulated as flexible Krylov subspace methods. In Chapter 6, we investigate the imaging application of computed tomography that requires parameter calibration, i.e., the forward model  $\mathbf{A}$  depends on unknown parameters  $\mathbf{p}$ , which need to be calibrated alongside the image reconstruction process. We not only explore the application of the newly proposed low-rank methods in this context, but also investigate a new hybrid machine learning - BCD approach that solves the problem more accurately. We wrap up the thesis in Chapter 7, in which we conclude the findings of this thesis, as well as identify interesting research directions that can be extended from this work.

## Chapter 2

# Overview of Regularization

## Methods

Problems (1.1) and (1.2) are ill-conditioned because for large scale linear systems arising from imaging inverse problems, singular values of  $\mathbf{A}$  usually decay slowly to 0, making the solution  $\mathbf{x}$  very sensitive to noise in the observation  $\mathbf{b}$ . In this chapter, we show how the singular value decomposition (SVD) of  $\mathbf{A}$  is related to the conditioning of the problem, discuss different types of regularization methods, and present techniques of choosing the regularization parameter. We start with a brief introduction of the SVD. Consider the SVD of  $\mathbf{A} \in \mathbb{R}^{M \times N}$ , which takes the form

$$\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T,$$

where  $\mathbf{U} = [\mathbf{u}_1 \ \mathbf{u}_2 \ \cdots \ \mathbf{u}_M] \in \mathbb{R}^{M \times M}$  and  $\mathbf{V} = [\mathbf{v}_1 \ \mathbf{v}_2 \ \cdots \ \mathbf{v}_N] \in \mathbb{R}^{N \times N}$  are orthogonal matrices, and  $\mathbf{\Sigma} \in \mathbb{R}^{M \times N}$  is a diagonal matrix so that  $\mathbf{\Sigma} = \text{diag}(\sigma_1, \sigma_2, \cdots, \sigma_N)$ ,

where its diagonal entries are monotonically non-increasing:  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_N \geq 0$ .

Without loss of generality, we assume that  $M \geq N$ , and the solution to (1.1) can be written as

$$\mathbf{x} = \mathbf{V}\mathbf{\Sigma}^\dagger\mathbf{U}^T\mathbf{b} = \sum_{i=1}^N \frac{\mathbf{u}_i^T\mathbf{b}}{\sigma_i}\mathbf{v}_i = \sum_{i=1}^N \frac{\mathbf{u}_i^T\mathbf{b}^{\text{ex}}}{\sigma_i}\mathbf{v}_i + \sum_{i=1}^N \frac{\mathbf{u}_i^T\boldsymbol{\eta}}{\sigma_i}\mathbf{v}_i,$$

where  $\mathbf{\Sigma}^\dagger \in \mathbb{R}^{N \times M}$  is a diagonal matrix with diagonal entries  $(1/\sigma_1, 1/\sigma_2, \dots, 1/\sigma_N)$ .

We immediately observe that if  $\sigma_i$  is very small (i.e., for large indices  $i$ ) compared to  $\boldsymbol{\eta}$ , then with  $\sigma_i$  on the denominator and  $\boldsymbol{\eta}$  on the numerator, the noise contribution is going to be amplified by small  $\sigma_i$ 's, and may even dominate over the true solution.

Regularization can be thought of as a way to “filter out” the noise contribution in the solution. Let  $\phi_1, \phi_2, \dots, \phi_N$  be the “filter factors” and define the regularized solution to be

$$\mathbf{x}_{\text{reg}} = \sum_{i=1}^N \phi_i \frac{\mathbf{u}_i^T\mathbf{b}}{\sigma_i}\mathbf{v}_i. \quad (2.1)$$

The filter factors  $\phi_i$  should be chosen such that  $0 \leq \phi_i \leq 1$  and  $\phi_i$  tend to 0 as  $i$  increases so that for large  $i$ 's,  $\phi_i$ 's can cancel out the noise contribution caused by small  $\sigma_i$ 's. Two common regularization methods – Tikhonov regularization and Truncated SVD can both be conveniently expressed using filter factors.

## 2.1 Tikhonov Regularization and TSVD

**Tikhonov Regularization** In standard Tikhonov regularization, we have that  $\mathcal{R}(\mathbf{x}) = \lambda\|\mathbf{x}\|_2^2$ ,  $\lambda > 0$  in (1.3). Because of its definition, Tikhonov regularization is

also called  $\ell_2$  regularization. The regularized solution is the unique solution to

$$\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_2^2, \quad (2.2)$$

which is equivalent to solving the following least squares problem

$$\mathbf{x}_\lambda = \arg \min_{\mathbf{x}} \left\| \begin{bmatrix} \mathbf{Ax} \\ \sqrt{\lambda} \mathbf{I} \end{bmatrix} - \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix} \right\|_2^2, \quad (2.3)$$

or its corresponding normal equations

$$(\mathbf{A}^T \mathbf{A} + \lambda \mathbf{I}) \mathbf{x} = \mathbf{A}^T \mathbf{b}.$$

Note that Tikhonov regularization does not require us to compute or approximate the SVD of  $\mathbf{A}$ . However, the regularization parameter  $\lambda$  needs to be chosen beforehand, and ideally, we should take  $\sigma_N \leq \lambda \leq \sigma_1$ . It can be shown that the filter factors for Tikhonov regularization are  $\phi_i = \frac{\sigma_i^2}{\sigma_i^2 + \lambda}$ . By introducing the filter factors, we can express  $\mathbf{x}_\lambda$  as

$$\mathbf{x}_\lambda = \mathbf{V} \boldsymbol{\Sigma}_\lambda^\dagger \mathbf{U}^T \mathbf{b}, \quad (2.4)$$

where  $\boldsymbol{\Sigma}_\lambda^\dagger = \text{diag}(\phi_1/\sigma_1, \dots, \phi_N/\sigma_N) = \text{diag}(\frac{\sigma_1}{\sigma_1^2 + \lambda}, \dots, \frac{\sigma_N}{\sigma_N^2 + \lambda})$ .

Tikhonov regularization may be generalized to define  $\mathcal{R}(\mathbf{x}) = \lambda \|\mathbf{Lx}\|_2^2$ . Sometimes, we introduce the additional regularization matrix  $\mathbf{L}$  because certain properties are desired for the solution  $\mathbf{x}_\lambda$ . For example, we can take  $\mathbf{L}$  to be a discretization of the derivative operator, and in this way, we are implicitly applying a “smoothness”

constraint to the solution. By applying the regularization matrix  $\mathbf{L}$ , the least squares problem is

$$\mathbf{x}_\lambda = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda\|\mathbf{L}\mathbf{x}\|_2^2, \quad (2.5)$$

and equivalently,

$$\mathbf{x}_\lambda = \arg \min_{\mathbf{x}} \left\| \begin{bmatrix} \mathbf{A}\mathbf{x} \\ \sqrt{\lambda}\mathbf{L}\mathbf{x} \end{bmatrix} - \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix} \right\|_2^2. \quad (2.6)$$

The associated normal equations are

$$(\mathbf{A}^T\mathbf{A} + \lambda\mathbf{L}^T\mathbf{L})\mathbf{x} = \mathbf{A}^T\mathbf{b}.$$

Note that (2.5) can be solved by substitution of variable  $\hat{\mathbf{x}} = \mathbf{L}\mathbf{x}$  and  $\mathbf{x} = \mathbf{L}^{-1}\hat{\mathbf{x}}$  under the assumption that  $\mathbf{L}$  is invertible. By doing so, the regularized problem becomes

$$\mathbf{x}_\lambda = \|\mathbf{A}\mathbf{L}^{-1}\hat{\mathbf{x}} - \mathbf{b}\|_2^2 + \lambda\|\hat{\mathbf{x}}\|_2^2, \quad (2.7)$$

and the generalized Tikhonov regularized problem can be reduced to a standard Tikhonov problem.

Both standard and general Tikhonov regularized solutions can be obtained rather easily by using standard linear system solvers. If  $\mathbf{A}$  is small (which is rarely the case), we can use direct solvers – a good example would be MATLAB’s backslash operator applied to (2.3) and (2.6), which uses QR factorization to solve the rectangular system. This is feasible when  $\mathbf{A}$  is small, but in most cases that arise in image processing, the matrix  $\mathbf{A}$  is very large, requiring us to use an iterative solver such as



Krylov subspace methods.

**Truncated SVD.** TSVD is implemented by setting small singular values of  $\mathbf{A}$  to 0 directly. Hence, the TSVD solution is given by

$$\mathbf{x}_{\text{TSVD}} = \sum_{i=1}^k \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i, \quad (2.8)$$

where  $k \leq N$  is the truncation rank that needs to be chosen. This solution is equivalent to applying the following filter factors to (2.1):

$$\phi_i = \begin{cases} 1 & \text{for } i \leq k, \\ 0 & \text{for } i > k. \end{cases}$$

By defining the filter factors, we can rewrite (2.8) as

$$\mathbf{x}_{\text{TSVD}} = \sum_{i=1}^N \phi_i \frac{\mathbf{u}_i^T \mathbf{b}}{\sigma_i} \mathbf{v}_i \quad \text{and} \quad \mathbf{x}_{\text{TSVD}} = \mathbf{V} \boldsymbol{\Sigma}_{\text{TSVD}}^\dagger \mathbf{U}^T \mathbf{b},$$

where  $\boldsymbol{\Sigma}_{\text{TSVD}}^\dagger = \text{diag}(\phi_1/\sigma_1, \dots, \phi_N/\sigma_N) = \text{diag}(1/\sigma_1, \dots, 1/\sigma_k, 0, \dots, 0)$ .

An important question associated with using TSVD is: how should we choose the truncation rank  $k$ ? Unfortunately, there is no perfect answer to this question. The truncation rank  $k$  depends on the individual problem and there is no best way to choose it for all problems. But there are guides that can be used to estimate  $k$ . In Section 2.2, we introduce techniques for selecting  $\lambda$  for Tikhonov regularization, and these techniques can be easily extended to TSVD. In practice, it is extremely

hard to compute the full SVD of large matrices  $\mathbf{A}$ . But we can still apply TSVD regularization by approximating the leading (largest) singular values and vectors, which will be further discussed in Chapter 3.

## 2.2 Choosing the Regularization Parameter for Tikhonov Regularization

Techniques for choosing the regularization parameter for Tikhonov regularization include the L-curve method [34], generalized cross validation and the discrepancy principle [30]. These techniques require evaluating the residual

$$\|\mathbf{A}\mathbf{x}_\lambda - \mathbf{b}\|_2 = \|\mathbf{U}\Sigma\Sigma_\lambda^\dagger\mathbf{U}^T\mathbf{b} - \mathbf{b}\|_2 = \|\Sigma\Sigma_\lambda^\dagger\mathbf{b} - \mathbf{U}^T\mathbf{b}\|_2 = \sum_{i=1}^N \frac{\lambda\hat{b}_i}{\sigma_i^2 + \lambda},$$

where  $\hat{\mathbf{b}} = \mathbf{U}^T\mathbf{b}$ . In addition to going over these three methods, we also explain how to find the optimal regularization parameter, which requires us to know  $\mathbf{x}^{\text{ex}}$  and can be used for analysis purposes, but it is not so realistic in real imaging applications.

**The L-curve method.** This method is first proposed in [34], and requires solving the Tikhonov regularized problem (2.2) repeatedly, each time with a different regularization parameter  $\lambda_k$ . The range of  $\lambda_k$  should depend on the singular spectrum of  $\mathbf{A}$ , and an example would be taking equally spaced points between  $10^{-8}$  and  $10^{-3}$ . For each solution  $\mathbf{x}_\lambda$  obtained by solving the Tikhonov regularized problem with  $\lambda_k$ , we plot the residual  $\mathbf{r}_k = \|\mathbf{A}\mathbf{x}_\lambda - \mathbf{b}\|_2$  against the norm of the solution  $\|\mathbf{x}_\lambda\|_2$ ,

which should result in an L-shaped curve. The  $\lambda_k$  corresponding to the “corner” of the L-curve renders a solution that has a small residual and whose norm is small, corresponding to the fit-to-data term and norm penalty term in the least squares problem (2.2) or (2.7). The `lcorner.m` function from Regularization Tools [29] can be used to locate the corner of the L-curve.

**Generalized Cross Validation (GCV).** As described in [30], the GCV method looks for the regularization parameter  $\lambda$  that minimizes the GCV function

$$G(\lambda) = \frac{N\|\mathbf{Ax}_\lambda - \mathbf{b}\|_2^2}{\left(\text{Tr}(\mathbf{I} - \mathbf{AA}_\lambda^\dagger)\right)^2} = \frac{N\|\Sigma\Sigma_\lambda^\dagger\mathbf{b} - \mathbf{U}^T\mathbf{b}\|_2^2}{\left(\text{Tr}(\mathbf{I} - \Sigma\Sigma_\lambda^\dagger)\right)^2} = \frac{N\left(\sum_{i=1}^N \frac{\hat{b}_i}{\sigma_i^2 + \lambda}\right)^2}{\left(\sum_{i=1}^N \frac{1}{\sigma_i^2 + \lambda}\right)^2},$$

where  $\mathbf{A}_\lambda^\dagger = \mathbf{V}^T\Sigma_\lambda^\dagger\mathbf{U}^T$  as in (2.4), and the trace of  $\mathbf{I} - \Sigma\Sigma_\lambda^\dagger$  can be easily computed since it is a diagonal matrix. The minimizer  $\lambda$  of  $G(\lambda)$  can be found using MATLAB’s `fminbnd` function, by inputting an approximate bound for  $\lambda$ , for example,  $\lambda_{\min} = 0$  and  $\lambda_{\max} =$  an approximation of  $\sigma_1$ .

**The Discrepancy Principle.** To solve the original linear system (1.1) of interest using Tikhonov regularization, we would like our solution  $\mathbf{x}_\lambda$  to satisfy  $\mathbf{Ax}_\lambda = \mathbf{b}^{\text{ex}} + \boldsymbol{\eta}$ , so naturally, we have the relation  $\mathbf{Ax}_\lambda - \mathbf{b} = \boldsymbol{\eta}$ , which is called the discrepancy. The discrepancy principle [30] assigns the regularization parameter  $\lambda$  such that the following equation is satisfied

$$\|\mathbf{Ax}_\lambda - \mathbf{b}\|_2 = \|\boldsymbol{\eta}\|_2 \quad \text{i.e.,} \quad \sum_{i=1}^N \frac{\lambda\hat{b}_i}{\sigma_i^2 + \lambda} - \|\boldsymbol{\eta}\|_2 = 0. \quad (2.9)$$

This approach requires that the norm,  $\|\boldsymbol{\eta}\|_2$ , of the Gaussian white noise  $\boldsymbol{\eta}$  is available to us. Since  $\|\boldsymbol{\eta}\|_2$  cannot be exactly known, in practice the discrepancy principle is implemented as  $\|\mathbf{A}\mathbf{x}_\lambda - \mathbf{b}\|_2 = \mu\|\boldsymbol{\eta}\|_2$ , where  $\mu \gtrsim 1$  is a user chosen parameter, e.g.,  $\mu = 1.01$ . MATLAB's `fzero` function may be used to solve this equation for the unknown  $\lambda$ .

**The Optimal Parameter.** Choosing the optimal regularization parameter requires the true solution  $\mathbf{x}^{\text{ex}}$ , which differentiates this last method from the previous three. We introduce this method in order to provide a basis for comparison with other parameter setting techniques in numerical experiments. To find the best regularization parameter, we try to minimize the difference between  $\mathbf{x}^{\text{ex}}$  and  $\mathbf{x}_\lambda$ :

$$\begin{aligned}
\|\mathbf{x}^{\text{ex}} - \mathbf{x}_\lambda\|_2^2 &= \|\mathbf{x}^{\text{ex}} - \mathbf{V}\boldsymbol{\Sigma}_\lambda^\dagger\mathbf{U}^T\mathbf{b}\|_2^2 \\
&= \|\mathbf{V}^T\mathbf{x}^{\text{ex}} - \boldsymbol{\Sigma}_\lambda^\dagger\mathbf{U}^T\mathbf{b}\|_2^2 \\
&= \|\hat{\mathbf{x}}^{\text{ex}} - \boldsymbol{\Sigma}_\lambda^\dagger\hat{\mathbf{b}}\|_2^2 \\
&= \sum_{i=1}^N \left( \frac{\sigma_i \hat{b}_i}{\sigma_i^2 + \lambda} - \hat{x}_i \right)^2, \tag{2.10}
\end{aligned}$$

where  $\hat{\mathbf{b}} = \mathbf{U}^T\mathbf{b}$  and  $\hat{\mathbf{x}}^{\text{ex}} = \mathbf{V}\mathbf{x}^{\text{ex}}$ . We need to look for the  $\lambda$  that minimizes (2.10), and similar to GCV, we can use MATLAB's `fminbnd` function for this purpose.

## 2.3 Other Regularization Terms

There are many other ways to define the regularization term  $\mathcal{R}(\mathbf{x})$ . Depending on the desired properties in the solution  $\mathbf{x}$ , we can choose  $\mathcal{R}(\mathbf{x})$  accordingly. All regu-

larization methods serve the purpose of reducing the ill-conditioning in the problem.

Popular choices include:

- $\ell_1$  regularization:  $\mathcal{R}(\mathbf{x}) = \|\mathbf{x}\|_1$  is used to enforce sparsity in the solution;
- $\ell_p$  regularization:  $\mathcal{R}(\mathbf{x}) = \|\mathbf{x}\|_p^p$  is a generalization of  $\ell_1$  and  $\ell_2$  regularization;
- Nuclear norm regularization:  $\mathcal{R}(\mathbf{x}) = \|\text{vec}^{-1}(\mathbf{x})\|_* = \|\mathbf{X}\|_*$  is defined as the sum of singular values of  $\mathbf{X} = \text{vec}^{-1}(\mathbf{x})$ , which enforces a low rank constraint implicitly;
- Total variation (TV) regularization:  $\mathcal{R}(\mathbf{x}) = TV(\mathbf{x}) = \sum_{i=1}^N \sqrt{[D_h \mathbf{x}]_i^2 + [D_v \mathbf{x}]_i^2}$ , where  $D_h$  and  $D_v$  denote the finite difference approximations of the horizontal and vertical first derivative operators; it is widely used in image restoration applications, and is very useful in denoising.

While many methods exist for computing the solution to different types of regularized problem, we comment briefly here on one kind of method called “Iteratively Reweighted Norm” (IRN) [25, 61]. The idea is to express different regularization terms  $\mathcal{R}(\mathbf{x})$  in terms of  $\|\mathbf{L}(\mathbf{x})\mathbf{x}\|_2^2$ . For example, for  $\ell_1$  regularization  $\mathbf{L}(\mathbf{x}) = \text{diag}(1/\sqrt{\mathbf{x}})$ , where square root and division are done componentwise.

We immediately notice that the difficulty of this approach lies with the computation of the regularization matrix  $\mathbf{L}(\mathbf{x})$  defined with respect to the solution  $\mathbf{x}$ , which is not available to us. A natural remedy is to use an approximation of  $\mathbf{x}$  to compute  $\mathbf{L}(\mathbf{x})$ , which leads us to Algorithm 1. This algorithm is often referred to as iterative reweighted norm (IRN), which utilizes an inner-outer iteration approach: in the  $i$ th outer iteration, the approximation  $\mathbf{L}_i$  (to be used in place of the true  $\mathbf{L}(\mathbf{x})$ )

is computed using the solution  $\mathbf{x}_{i-1}$  from the previous outer iteration; we then solve the general Tikhonov regularized problem with regularization matrix  $\mathbf{L}_i$  to obtain the updated solution  $\mathbf{x}_i$  by using an iterative method (i.e., inner iterations). Note that some elements of  $\mathbf{x}_i$  may be zero, in which case we may have a division by zero problem depending on the definition of  $\mathbf{L}(\mathbf{x})$  (e.g.,  $\mathbf{L}(\mathbf{x}) = \text{diag}(1/\sqrt{\mathbf{x}})$ ). This can be avoided by adding a small constant to  $\mathbf{x}_i$ .

---

**Algorithm 1** Iteratively Reweighted Norm

---

Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $\mathbf{L}_0 = \mathbf{L}(\mathbf{x}_0)$ ,  $\lambda$   
**for**  $i = 1, 2, \dots$  until a stopping criterion is satisfied **do**  
    Solve for  $\mathbf{x}_i = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{L}_i \mathbf{x}\|_2^2$   
    Update  $\mathbf{L}_i = \mathbf{L}(\mathbf{x}_{i-1})$   
**end for**  
Output:  $\mathbf{x}_i$ 's

---

Many types of regularization terms can be expressed in a similar way, including  $\ell_p$  regularization for a general  $p$  value, TV regularization [19] and nuclear norm regularization [18]. Therefore, the IRN approach can be applied to a wide range of regularized least squares problems in a similar way. In Chapter 4, we elaborate on how to apply IRN to the nuclear norm regularized (NNR) least squares problem which enforces a low rank constraint. In particular, we use Krylov subspace methods (i.e., GMRES and LSQR) as inner solvers for IRN. One disadvantage of the IRN approach is that the inner-outer iteration scheme is not very computationally efficient, since we need to solve many least squares problems in a sequence. So a more efficient alternative – flexible Krylov subspace methods [11, 63] will also be discussed.

## 2.4 Iterative Regularization

Iterative regularization is different from direct regularization that adds a penalization term  $\mathcal{R}(\mathbf{x})$  directly. Instead, it exploits the semi-convergence property of iterative methods, and early stopping is performed when semi-convergence is achieved. That is, we apply an iterative method to the linear system  $\mathbf{A}\mathbf{x} = \mathbf{b}$ , and terminate iterations when a good solution is obtained, before noise starts to contaminate the solution. Early stopping computes a solution that resembles TSVD regularization [31, 69].

A number of iterative methods can be used for this purpose. In this section, we address Krylov subspace methods, namely, GMRES and LSQR (discussed in Section 2.4.1). Implementations of these two methods can be found in MATLAB's library (functions `gmres.m` and `lsqr.m`). Their hybrid counterparts (discussed in Section 2.4.2), i.e., hybrid GMRES and hybrid LSQR [8, 12, 21] can be considered if we want to combine Krylov subspace methods with Tikhonov regularization. Implementations of the hybrid methods can be found in the MATLAB package IR Tools [17]. Note that in this section,  $\mathbf{U}$  and  $\mathbf{V}$  no longer represent the singular vector matrices as in the previous sections of this chapter.

### 2.4.1 GMRES and LSQR

The Generalized Minimum Residual (GMRES) method is proposed by Saad and Schultz in 1986 [65]. It is an iterative method that aims at solving large scale linear systems with square matrices  $\mathbf{A} \in \mathbb{R}^{N \times N}$ . At the  $m$ th iteration of GMRES, the

approximate solution is  $\mathbf{x}_k \in \mathcal{K}_k(\mathbf{A}, \mathbf{b})$  under the assumption that the initial guess  $\mathbf{x}_0 = \mathbf{0}$ , where  $\mathcal{K}_k(\mathbf{A}, \mathbf{b})$  is the Krylov subspace

$$\mathcal{K}_k(\mathbf{A}, \mathbf{b}) = \text{span}\{\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}, \dots, \mathbf{A}^k\mathbf{b}\}.$$

Let  $\mathbf{V}_k = [\mathbf{v}_1, \dots, \mathbf{v}_k] \in \mathbb{R}^{N \times k}$  be the orthonormal basis for the Krylov subspace  $\mathcal{K}_k(\mathbf{A}, \mathbf{b})$ . Specifically,  $\mathbf{v}_1 = \mathbf{b}/\|\mathbf{b}\|$ . Since the solution  $\mathbf{x}_k$  is in  $\mathcal{K}_k(\mathbf{A}, \mathbf{b})$ , it can be written as  $\mathbf{V}_k\mathbf{y}$  for some coefficients  $\mathbf{y}$ . GMRES first uses the Arnoldi algorithm to generate  $\mathbf{V}_k \in \mathbb{R}^{N \times (k+1)}$  whose columns are orthonormal, and an upper Hessenberg matrix  $\mathbf{H}_k \in \mathbb{R}^{(k+1) \times k}$  such that

$$\mathbf{A}\mathbf{V}_k = \mathbf{V}_{k+1}\mathbf{H}_k.$$

Then,

$$\begin{aligned} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 &= \|\mathbf{A}\mathbf{V}_k\mathbf{y} - \mathbf{b}\|_2^2 \\ &= \|\mathbf{V}_{k+1}\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|\mathbf{v}_1\|_2^2 \\ &= \|\mathbf{V}_{k+1}(\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|\mathbf{e}_1)\|_2^2 \\ &= \|\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|\mathbf{e}_1\|_2^2, \end{aligned} \tag{2.11}$$

where  $\mathbf{e}_1 = (1, 0, \dots, 0) \in \mathbb{R}^{k+1}$ . Hence the GMRES solution  $\mathbf{x}_k \in \mathcal{K}_k(\mathbf{A}, \mathbf{b})$  can be easily obtained by solving the projected minimization problem (2.11) for the minimizer  $\mathbf{y}_k$ , which is usually inexpensive to compute. That is, the solution  $\mathbf{x}_k = \mathbf{V}_k\mathbf{y}_k$ , where  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \|\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|\mathbf{e}_1\|_2$ , which leads to the GMRES algorithm



(Algorithm 2).

---

**Algorithm 2** GMRES

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$
  - 2: Take  $\mathbf{v}_1 = \mathbf{b}/\|\mathbf{b}\|_2$
  - 3: **for**  $i = 1, 2, \dots$  until a stopping criterion is satisfied **do**
  - 4: Compute  $\mathbf{w} = \mathbf{A}\mathbf{v}_i$
  - 5: Compute  $h_{ji} = \mathbf{w}^T \mathbf{v}_j$  for  $j = 1, \dots, i$  and set  $\mathbf{w} = \mathbf{w} - \sum_{j=1}^i h_{ji} \mathbf{v}_j$
  - 6: Set  $h_{i+1,i} = \|\mathbf{w}\|_2$ , and if  $h_{j+1,j} \neq 0$ , take  $\mathbf{v}_{i+1} = \mathbf{w}/h_{i+1,i}$
  - 7: **end for**
  - 8: Compute  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \|\mathbf{H}_k \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2^2$  and take  $\mathbf{x}_k = \mathbf{V}_k \mathbf{y}_k$
- 

Another example of Krylov subspace methods is LSQR, which is proposed by Paige and Sanders [54] to solve large and sparse linear systems. Different from GMRES that can only be applied to square matrices  $\mathbf{A} \in \mathbb{R}^{N \times N}$ , LSQR can be used on general rectangular shaped matrices  $\mathbf{A} \in \mathbb{R}^{M \times N}$ , where  $M$  is not necessarily equal to  $N$ . It is based on Golub-Kahan Bidiagonalization (GKB) [24] and the computed solution  $\mathbf{x}_k$  belongs to the Krylov subspace  $\mathcal{K}_k(\mathbf{A}^T \mathbf{A}, \mathbf{A}^T \mathbf{b})$  spanned by columns of  $\mathbf{V}_k$ . The  $k$ th iteration of GKB generates  $\mathbf{V}_k = [\mathbf{v}_1, \dots, \mathbf{v}_k] \in \mathbb{R}^{N \times k}$  and  $\mathbf{U}_{k+1} = [\mathbf{u}_1, \dots, \mathbf{u}_{k+1}] \in \mathbb{R}^{M \times (k+1)}$  with orthonormal columns, and lower bidiagonal matrix  $\mathbf{B}_{k+1} \in \mathbb{R}^{(k+1) \times k}$  such that

$$\mathbf{A}\mathbf{V}_k = \mathbf{U}_{k+1}\mathbf{B}_k.$$

Since the solution  $\mathbf{x} \in \mathcal{R}(\mathbf{V}_k)$ , by letting  $\mathbf{x} = \mathbf{V}_k \mathbf{y}$ , it follows that

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 = \|\mathbf{B}_k \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2^2. \quad (2.12)$$

Steps to achieve (2.12) are omitted because they resemble those for deriving (2.11). Therefore, LSQR computes the solution  $\mathbf{x}_k = \mathbf{V}_k \mathbf{y}_k$  for  $\mathbf{y}_k$  that solves  $\min_{\mathbf{y}} \|\mathbf{B}_k \mathbf{y} - \mathbf{b}\|_2$ . The LSQR algorithm is presented in Algorithm 3.

---

**Algorithm 3** LSQR

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$
  - 2: Take  $\mathbf{u}_1 = \mathbf{b}/\|\mathbf{b}\|_2$
  - 3: **for**  $i = 1, 2, \dots$ , until a stopping criterion is satisfied **do**
  - 4:   Compute  $\mathbf{w} = \mathbf{A}^T \mathbf{u}_i$  and  $\mathbf{w} = \mathbf{w} - \beta_{i-1} \mathbf{v}_{i-1}$  if  $i > 1$
  - 5:   Set  $\alpha_i = \|\mathbf{w}\|_2$  and  $\mathbf{v}_i = \mathbf{w}/\alpha_i$
  - 6:   Compute  $\mathbf{w} = \mathbf{A}^T \mathbf{v}_i$ ,  $\mathbf{w} = \mathbf{w} - \alpha_i \mathbf{u}_i$
  - 7:   Set  $\beta_i = \|\mathbf{w}\|_2$  and  $\mathbf{u}_{i+1} = \mathbf{w}/\beta_i$
  - 8:   Let  $b_{ii} = \alpha_i$  and  $b_{i+1,i} = \beta_i$  and if  $\alpha_i = 0$  or  $\beta_i = 0$ , break
  - 9: **end for**
  - 10:
  - 11: Compute  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \|\mathbf{B}_k \mathbf{y} - \mathbf{b}\|_2$  and take  $\mathbf{x}_k = \mathbf{V}_k \mathbf{y}_k$
- 

## 2.4.2 Hybrid Krylov Methods

Krylov subspace methods such as GMRES are applied to the fit-to-data term  $\|\mathbf{Ax} - \mathbf{b}\|_2^2$  alone, and iterative regularization is done by early stopping. If we want to apply additional regularization such as Tikhonov, we can adopt hybrid Krylov methods. Hybrid GMRES employs the Arnoldi-Tikhonov scheme [19], and rather than adding the regularization term directly to  $\|\mathbf{Ax} - \mathbf{b}\|_2^2$ , it solves a regularized projected problem, i.e., the regularization term is applied to the projected problem  $\|\mathbf{H}_k \mathbf{y} - \mathbf{b}\|_2^2$ . This is legitimate because we have shown that  $\|\mathbf{Ax} - \mathbf{b}\|_2^2 = \|\mathbf{H}_k \mathbf{y} - \mathbf{b}\|_2^2$ , hence

$$\min_{\mathbf{y}} \|\mathbf{AV}_k \mathbf{y} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_2^2$$

is equivalent to

$$\min_{\mathbf{y}} \|\mathbf{H}_k \mathbf{y} - \|\mathbf{b}\| \mathbf{e}_1\|_2^2 + \lambda \|\mathbf{x}\|_2^2, \quad (2.13)$$

where we have substituted  $\mathbf{x} = \mathbf{V}_k \mathbf{y}$ . Therefore, we obtain the hybrid GMRES algorithm, which simply replaces the minimization problem in step 8 of Algorithm 2 by (2.13).

Similar to hybrid GMRES, hybrid LSQR can be considered if additional regularization is needed on top of iterative regularization. In the case of hybrid LSQR, direct regularization is applied to the projected problem (2.12) so that in step 10 of Algorithm 3, the following regularized projected minimization problem is solved instead

$$\mathbf{y}_k = \arg \min_{\mathbf{y}} \|\mathbf{B}_k \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2^2 + \lambda \|\mathbf{y}\|_2^2. \quad (2.14)$$

**Choosing  $\lambda$  for hybrid Krylov methods** To choose a regularization parameter  $\lambda$  for (2.13), one obvious choice would be to use parameter setting techniques described in Section 2.2. However, in the case of discrepancy principle, Gazzola and Novati [20] have proposed a secant approach that can be applied readily to Arnoldi-Tikhonov based methods. The approach exploits the discrepancy principle to look for the  $\lambda$  such that

$$\|\mathbf{A} \mathbf{x}_\lambda - \mathbf{b}\|_2 = \|\mathbf{H}_k \mathbf{y}_\lambda - \|\mathbf{b}\| \mathbf{e}_1\|_2 = \|\boldsymbol{\eta}\|_2$$

holds. While we can still solve for  $\lambda$  by running MATLAB's `fzero` function on (2.9), an alternative approach has been proposed in [20], which can efficiently compute a regularization parameter  $\lambda_i$  for each GMRES iteration, and can be employed as a stopping criteria. This discrepancy principle based approach starts with defining the

discrepancy function

$$\psi_i(\lambda) = \|\mathbf{A}\mathbf{x}_\lambda - \mathbf{b}\|_2 = \|\mathbf{H}_i\mathbf{y}_{i,\lambda} - \|\mathbf{b}\|_2\mathbf{e}_1\|_2, \quad (2.15)$$

where  $\mathbf{y}_{i,\lambda}$  is the solution to (2.13) using regularization parameter  $\lambda$  with  $i = k$ . Assuming that a good estimate  $\epsilon \approx \|\boldsymbol{\eta}\|_2$  is available, we say that the discrepancy principle is satisfied as soon as

$$\psi(\lambda) \leq \mu\epsilon, \quad \text{for } \mu \gtrsim 1.$$

At each iteration  $i$ , consider the linear approximation of  $\psi(\lambda)$

$$\psi_i(\lambda) \approx \psi_i(0) + \beta_i\lambda, \quad \text{where } \beta_i = \frac{\psi_i(\lambda_{i-1}) - \psi_i(0)}{\lambda_{i-1}}.$$

The expression for  $\beta_i$  is obtained from the equation  $\psi_i(\lambda_{i-1}) = \psi_i(0) + \beta_i\lambda_{i-1}$ , where  $\psi_i(\lambda_{i-1})$  is available from the previous iteration, and  $\psi_i(0)$  is the norm of the GMRES residual at iteration  $i$  and can be easily computed with the SVD of  $\mathbf{H}_i$ . To select  $\lambda_i$ , we then force  $\psi_i(\lambda_i) = \psi_i(0) + \beta_i\lambda_i$ , and by substituting  $\beta_i$ , we get

$$\lambda_i = \left| \frac{\mu\epsilon - \psi_i(0)}{\psi_i(\lambda_{i-1}) - \psi_i(0)} \right| \lambda_{i-1}. \quad (2.16)$$

The absolute value is taken to avoid selecting a negative  $\lambda_i$ . Note that the SVD of  $\mathbf{H}_i$ , which was previously used to compute  $\psi_i(0)$ , can again be used here to solve for  $\mathbf{y}_{i,\lambda_i}$ . Using the discrepancy principle, iterations can be terminated as soon as  $\psi_i(\lambda) \leq \mu\epsilon$ . The updated hybrid GMRES using discrepancy principle as a

regularization technique and stopping criteria is summarized in Algorithm 4.

---

**Algorithm 4** Hybrid GMRES with discrepancy principle

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$
  - 2: Take  $\mathbf{v}_1 = \mathbf{b}/\|\mathbf{b}\|_2$
  - 3: **for**  $i = 1, 2, \dots$  until the discrepancy principle is satisfied **do**
  - 4:   Compute  $\mathbf{w} = \mathbf{A}\mathbf{v}_i$
  - 5:   Compute  $h_{ji} = \mathbf{w}^T \mathbf{v}_j$  for  $j = 1, \dots, i$  and set  $\mathbf{w} = \mathbf{w} - \sum_{j=1}^i h_{ji} \mathbf{v}_j$
  - 6:   Set  $h_{i+1,i} = \|\mathbf{w}\|_2$ , and if  $h_{j+1,j} \neq 0$ , take  $\mathbf{v}_{i+1} = \mathbf{w}/h_{i+1,i}$
  - 7:   Select  $\lambda_i$  using (2.16)
  - 8:   Compute  $\mathbf{y}_{i,\lambda_i} = \arg \min_{\mathbf{y}} \|\mathbf{H}_i \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2^2 + \lambda_i \|\mathbf{x}\|_2^2$  and take  $\mathbf{x}_{i,\lambda} = \mathbf{V}_i \mathbf{y}_{i,\lambda_i}$
  - 9: **end for**
  - 10:
- 

The discrepancy based approach for automatically selecting a regularization parameter for each iteration of hybrid-GMRES can also be generalized to hybrid LSQR, for which the discrepancy function is now defined as

$$\psi_i(\lambda) = \|\mathbf{B}_i \mathbf{y}_{i,\lambda} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2,$$

where  $\mathbf{y}_{i,\lambda}$  is the solution to (2.14) with  $i = k$  by using regularization parameter  $\lambda$ .

In each iteration  $i$ ,  $\lambda_i$  is updated using the same formula

$$\lambda_i = \left| \frac{\mu\epsilon - \psi_i(0)}{\psi_i(\lambda_{i-1}) - \psi_i(0)} \right| \lambda_{i-1}. \quad (2.17)$$

The hybrid LSQR method using the discrepancy principle for automatically setting  $\lambda$  and early stopping is shown in Algorithm 5.

---

**Algorithm 5** Hybrid LSQR with discrepancy principle
 

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$
  - 2: Take  $\mathbf{u}_1 = \mathbf{b}/\|\mathbf{b}\|_2$
  - 3: **for**  $i = 1, 2, \dots$ , until a stopping criterion is satisfied **do**
  - 4:   Compute  $\mathbf{w} = \mathbf{A}^T \mathbf{u}_i$  and  $\mathbf{w} = \mathbf{w} - \beta_{i-1} \mathbf{v}_{i-1}$  if  $i > 1$
  - 5:   Set  $\alpha_i = \|\mathbf{w}\|_2$  and  $\mathbf{v}_i = \mathbf{w}/\alpha_i$
  - 6:   Compute  $\mathbf{w} = \mathbf{A}^T \mathbf{v}_i$ ,  $\mathbf{w} = \mathbf{w} - \alpha_i \mathbf{u}_i$
  - 7:   Set  $\beta_i = \|\mathbf{w}\|_2$  and  $\mathbf{u}_{i+1} = \mathbf{w}/\beta_i$
  - 8:   Let  $b_{ii} = \alpha_i$  and  $b_{i+1,i} = \beta_i$  and if  $\alpha_i = 0$  or  $\beta_i = 0$ , break
  - 9:   Select  $\lambda_i$  using (2.17)
  - 10:   Compute  $\mathbf{y}_{i,\lambda_i} = \arg \min_{\mathbf{y}} \|\mathbf{B}_i \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2^2 + \lambda_i \|\mathbf{x}\|_2^2$  and take  $\mathbf{x}_{i,\lambda} = \mathbf{V}_i \mathbf{y}_{i,\lambda_i}$
  - 11: **end for**
- 

### 2.4.3 Flexible Krylov Methods

Flexible Krylov methods are a class of linear solvers that can handle iteration-dependent preconditioners: they were originally introduced in [63] for FGMRES, where a preconditioner for GMRES was allowed to change from one iteration to the next, either because at each iteration the preconditioner is implicitly defined by applying an iterative linear solver, or because the preconditioner can be updated with newly-computed information (see [66] for an overview).

In the framework of regularizing linear solvers, flexible Krylov methods were proposed in [11, 19, 22], where the iteration-dependent “preconditioner” was associated to an iteratively reweighted norm approach to Tikhonov-like regularized problems involving penalization terms expressed in some  $p$ -norm,  $0 < p \leq 1$ . These “preconditioners” have the effect of enforcing specific regularity into the approximation subspace for the solution, rather than accelerating the convergence of the iterative solvers. Leveraging flexible Krylov subspaces in this setting comes with the upside of avoiding restarts of the iterative solver, which is the approach commonly used when

adopting an iteratively reweighted norm method.

Consider first the case of a square  $\mathbf{A} \in \mathbb{R}^{N \times N}$ . In general, starting with  $\mathbf{x}_0 = \mathbf{0}$ , at the  $k$ th iteration, FGMRES updates a partial flexible Arnoldi factorization and computes the  $k$ th approximate solution as follows:

$$\mathbf{AZ}_k = \mathbf{V}_{k+1}\mathbf{H}_k, \quad \mathbf{x}_k = \mathbf{Z}_k\mathbf{y}_k, \quad \text{where } \mathbf{y}_k = \arg \min_{\mathbf{y} \in \mathbb{R}^k} \|\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|_2\mathbf{e}_1\|_2, \quad (2.18)$$

where  $\mathbf{V}_{k+1} = [\mathbf{v}_1, \dots, \mathbf{v}_{k+1}] \in \mathbb{R}^{N \times (k+1)}$  has orthonormal columns,  $\mathbf{H}_k \in \mathbb{R}^{(k+1) \times k}$  is upper Hessenberg, and  $\mathbf{Z}_k = [\mathbf{P}_1\mathbf{v}_1, \dots, \mathbf{P}_k\mathbf{v}_k] \in \mathbb{R}^{N \times k}$  has columns that span the approximation subspace for the solution ( $\mathbf{P}_i$  is an iteration-dependent preconditioner that is applied to  $\mathbf{v}_i$ ). The flexible GMRES method (FGMRES) uses the flexible Arnoldi process [19] to generate iterates of the form  $\mathbf{x}_k = \mathbf{Z}_k\mathbf{y}_k$ , where the vector  $\mathbf{y}_k$  is computed as  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \|\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|_2\mathbf{e}_1\|_2^2$ .

The extension to more general matrices  $\mathbf{A} \in \mathbb{R}^{M \times N}$ , with  $M$  not necessarily equal to  $N$ , can be naturally devised considering the flexible Golub-Kahan (FGK) process [11]. Taking  $\mathbf{x}_0 = \mathbf{0}$  as the initial guess, the  $k$ th FGK iteration updates partial factorizations of the form

$$\mathbf{AZ}_k = \mathbf{U}_{k+1}\mathbf{M}_k \quad \text{and} \quad \mathbf{A}^T\mathbf{U}_{k+1} = \mathbf{V}_{k+1}\mathbf{T}_{k+1}, \quad (2.19)$$

where the columns of  $\mathbf{U}_{k+1} \in \mathbb{R}^{M \times (k+1)}$ ,  $\mathbf{V}_{k+1} \in \mathbb{R}^{N \times (k+1)}$  are orthonormal,  $\mathbf{M}_k \in \mathbb{R}^{(k+1) \times k}$  is upper Hessenberg,  $\mathbf{T}_{k+1} \in \mathbb{R}^{(k+1) \times (k+1)}$  is upper triangular, and  $\mathbf{Z}_k = [\mathbf{P}_1\mathbf{v}_1, \dots, \mathbf{P}_k\mathbf{v}_k] \in \mathbb{R}^{N \times k}$  has columns that span the approximation subspace for the solution ( $\mathbf{P}_i$  is an iteration-dependent preconditioner that is applied to  $\mathbf{v}_i$ ). The

flexible LSQR method (FLSQR) uses the FGK process (2.19) to generate iterates of the form  $\mathbf{x}_k = Z_k \mathbf{y}_k$ , where the vector  $\mathbf{y}_k$  is computed as  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \left\| \mathbf{M}_k \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1 \right\|_2^2$ .



## Chapter 3

# TSVD approximations

In this Chapter, we introduce methods based on Kronecker product summations for approximating the truncated SVD (TSVD) of  $\mathbf{A}$ . Such approximations can be thought of as a low-rank approximation to the matrix being decomposed, and can be used to approximate the solution to the linear system 1.1. The new approximation approach discussed in this chapter is based on the paper jointly published by Garvey, Nagy and the author of this thesis [16]; see also Garvey's Ph.D. dissertation [15]. As described in Chapter 2, the solution to (1.1) can be computed directly if the full SVD of  $\mathbf{A}$  is available. However, computing the full SVD is very expensive both time-wise and storage-wise, and is in general infeasible for large matrices. For example, on a personal computer, MATLAB's `svd` function can compute the full SVD of matrices of sizes up to 4096 by 4096 without much difficulty, but as matrices grow larger, a typical personal computer will run out of memory. That is, it is generally infeasible for image restoration problems larger than 64 by 64 pixels. For larger matrices,

if only a few (e.g., 10) leading singular values and vectors are needed, MATLAB's `svds` function, which uses iterative Golub-Kahan Bidiagonalization [24, 47] for the estimation, can be used. However, `svds` is not attractive in our applications since we may need to compute a few thousand singular values and vectors. Therefore, other efficient approximation methods need to be exploited. The focus of this Chapter is computing TSVD approximations for structured square matrices  $\mathbf{A} \in \mathbb{R}^{N \times N}$  that can be found in the image deblurring application.

### 3.1 The Kronecker Product

Let  $\mathbf{G} \in \mathbb{R}^{n \times n}$  and  $\mathbf{H} \in \mathbb{R}^{n \times n}$  be square matrices with entries  $g_{ij}$  and  $h_{ij}$ . Then the Kronecker product of  $\mathbf{G}$  and  $\mathbf{H}$  is defined as

$$\mathbf{G} \otimes \mathbf{H} = \begin{bmatrix} g_{11}\mathbf{H} & \cdots & g_{1n}\mathbf{H} \\ \vdots & & \vdots \\ g_{n1}\mathbf{H} & \cdots & g_{nn}\mathbf{H} \end{bmatrix} \in \mathbb{R}^{N \times N},$$

where  $N = n^2$ . Useful properties of the Kronecker product include:

- $(\mathbf{A} \otimes \mathbf{B})^T = \mathbf{A}^T \otimes \mathbf{B}^T$ ,
- $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$ ,
- $(\mathbf{B}^T \otimes \mathbf{A})\text{vec}(\mathbf{X}) = \text{vec}(\mathbf{AXB})$ ,

given that the matrices are of appropriate sizes such that the matrix products  $\mathbf{AC}$ ,  $\mathbf{BD}$  and  $\mathbf{AXB}$  can be formed.

Let the SVDs of  $\mathbf{G}$  and  $\mathbf{H}$  be  $\mathbf{G} = \mathbf{U}_G \mathbf{\Sigma}_G \mathbf{V}_G^T$  and  $\mathbf{H} = \mathbf{U}_H \mathbf{\Sigma}_H \mathbf{V}_H^T$  respectively. Suppose that  $\mathbf{A}$  is the Kronecker product of  $\mathbf{G}$  and  $\mathbf{H}$ , then we have that

$$\begin{aligned} \mathbf{A} = \mathbf{G} \otimes \mathbf{H} &= (\mathbf{U}_G \mathbf{\Sigma}_G \mathbf{V}_G^T) \otimes (\mathbf{U}_H \mathbf{\Sigma}_H \mathbf{V}_H^T) \\ &= (\mathbf{U}_G \otimes \mathbf{U}_H) (\mathbf{\Sigma}_G \otimes \mathbf{\Sigma}_H) (\mathbf{V}_G \otimes \mathbf{V}_H)^T, \end{aligned}$$

where we have applied properties of the Kronecker product. Note that  $(\mathbf{U}_G \otimes \mathbf{U}_H)$  is an orthogonal matrix since

$$\begin{aligned} (\mathbf{U}_G \otimes \mathbf{U}_H)^T (\mathbf{U}_G \otimes \mathbf{U}_H) &= (\mathbf{U}_G^T \otimes \mathbf{U}_H^T) (\mathbf{U}_G \otimes \mathbf{U}_H) \\ &= (\mathbf{U}_G^T \mathbf{U}_G) \otimes (\mathbf{U}_H^T \mathbf{U}_H) = \mathbf{I} \otimes \mathbf{I} = \mathbf{I}, \end{aligned}$$

and  $(\mathbf{U}_G \otimes \mathbf{U}_H) (\mathbf{U}_G \otimes \mathbf{U}_H)^T = \mathbf{I}$  can be obtained in the same way. We can also show that  $\mathbf{V}_G \otimes \mathbf{V}_H$  is an orthogonal matrix. Note that  $\mathbf{\Sigma}_G$  and  $\mathbf{\Sigma}_H$  are diagonal matrices with monotonically non-increasing diagonal entries. Therefore,  $\mathbf{\Sigma}_G \otimes \mathbf{\Sigma}_H$  is also a diagonal matrix with nonnegative diagonal entries, but they are not necessarily monotonically decreasing. We can apply a permutation matrix  $\mathbf{P} \in \mathbb{R}^{N \times N}$  to  $(\mathbf{\Sigma}_G \otimes \mathbf{\Sigma}_H)$  to reorder the diagonal so that  $\mathbf{P}^T (\mathbf{\Sigma}_G \otimes \mathbf{\Sigma}_H) \mathbf{P}$  has a decreasing diagonal. As a result, we obtain a proper SVD of  $\mathbf{A}$ :

$$\mathbf{A} = \underbrace{((\mathbf{U}_G \otimes \mathbf{U}_H) \mathbf{P})}_{\mathbf{U}} \underbrace{(\mathbf{P}^T (\mathbf{\Sigma}_G \otimes \mathbf{\Sigma}_H) \mathbf{P})}_{\mathbf{\Sigma}} \underbrace{((\mathbf{V}_G \otimes \mathbf{V}_H) \mathbf{P})^T}_{\mathbf{V}^T}. \quad (3.1)$$

This analysis suggests that if  $\mathbf{A} \in \mathbb{R}^{N \times N}$  can be written as a Kronecker product of two small matrices  $\mathbf{G}, \mathbf{H} \in \mathbb{R}^{n \times n}$ , then we can simply compute the SVDs of

$\mathbf{G}$  and  $\mathbf{H}$  to easily obtain the SVD of  $\mathbf{A}$ . In this way, complexity is reduced to  $\mathcal{O}(n^3) = \mathcal{O}(N^{3/2})$  from  $\mathcal{O}(N^3)$  (if we directly compute the SVD of  $\mathbf{A}$ ).

## 3.2 Kronecker Product Summation

For a general matrix  $\mathbf{A}$ , it is hardly the case that it can be decomposed as  $\mathbf{G} \otimes \mathbf{H}$ . However, we may decompose  $\mathbf{A}$  as a sum of Kronecker products  $\mathbf{G}_i \otimes \mathbf{H}_i$  for  $i = 1, 2, \dots, R$ , i.e.,

$$\mathbf{A} = \sum_{i=1}^R \mathbf{G}_i \otimes \mathbf{H}_i,$$

where the number of terms  $R$  is called the Kronecker rank of  $\mathbf{A}$ . The computational approach to obtain the Kronecker product summation is proposed by Van Loan and Pitsianis [68], and requires taking the SVD of a rearrangement  $\tilde{\mathbf{A}} \in \mathbb{R}^{N \times N}$  of matrix  $\mathbf{A}$ . By doing so, the first term of the sum  $\mathbf{G}_1 \otimes \mathbf{H}_1$  is computed such that

$$\mathbf{G}_1 \otimes \mathbf{H}_1 = \arg \min_{\mathbf{G}, \mathbf{H}} \|\mathbf{A} - \mathbf{G} \otimes \mathbf{H}\|_F^2.$$

So intuitively, the first term  $\mathbf{G}_1 \otimes \mathbf{H}_1$  is the best Kronecker product approximation of  $\mathbf{A}$ . It is the most “significant” term in the summation and contains the most information of  $\mathbf{A}$  compared to the remaining terms. Therefore, a natural SVD approximation of  $\mathbf{A}$  is to take the SVD of  $\mathbf{G}_1 \otimes \mathbf{H}_1$ , i.e., if  $\mathbf{G}_1 = \mathbf{U}_{\mathbf{G}_1} \boldsymbol{\Sigma}_{\mathbf{G}_1} \mathbf{V}_{\mathbf{G}_1}^T$  and

$\mathbf{H}_1 = \mathbf{U}_{\mathbf{H}_1} \boldsymbol{\Sigma}_{\mathbf{H}_1} \mathbf{V}_{\mathbf{H}_1}^T$ , and

$$\begin{aligned} \mathbf{G}_1 \otimes \mathbf{H}_1 &= (\mathbf{U}_{\mathbf{G}_1} \boldsymbol{\Sigma}_{\mathbf{G}_1} \mathbf{V}_{\mathbf{G}_1}^T) \otimes (\mathbf{U}_{\mathbf{H}_1} \boldsymbol{\Sigma}_{\mathbf{H}_1} \mathbf{V}_{\mathbf{H}_1}^T) \\ &= (\mathbf{U}_{\mathbf{G}_1} \otimes \mathbf{U}_{\mathbf{H}_1}) (\boldsymbol{\Sigma}_{\mathbf{G}_1} \otimes \boldsymbol{\Sigma}_{\mathbf{H}_1}) (\mathbf{V}_{\mathbf{G}_1} \otimes \mathbf{V}_{\mathbf{H}_1})^T \\ &= \mathbf{U}_{\mathbf{A}_1} \boldsymbol{\Sigma}_{\mathbf{A}_1} \mathbf{V}_{\mathbf{A}_1}^T, \end{aligned}$$

so the SVD approximation to  $\mathbf{A}$  is

$$\mathbf{A} \approx \mathbf{U}_{\mathbf{A}_1} \boldsymbol{\Sigma}_{\mathbf{A}_1} \mathbf{V}_{\mathbf{A}_1}^T. \quad (3.2)$$

Note that the singular values in  $\boldsymbol{\Sigma}_{\mathbf{A}_1}$  are not yet ordered. This approximation uses only the first term in the Kronecker product summation and discards the information in the remaining terms, so naturally, the accuracy of  $\mathbf{U}_{\mathbf{A}_1} \boldsymbol{\Sigma}_{\mathbf{A}_1} \mathbf{V}_{\mathbf{A}_1}^T$  depends greatly on how closely  $\mathbf{G}_1 \otimes \mathbf{H}_1$  approximates  $\mathbf{A}$ , which depends greatly on the structure of  $\mathbf{A}$ .

To use more Kronecker product terms for the SVD approximation, Nagy and Kamm [41, 52] proposed the following

$$\mathbf{A} \approx \mathbf{U}_{\mathbf{A}_1} \mathring{\boldsymbol{\Sigma}} \mathbf{V}_{\mathbf{A}_1}^T, \quad (3.3)$$

where  $\mathbf{U}_{\mathbf{A}_1}$  and  $\mathbf{V}_{\mathbf{A}_1}$  are the same as in (3.2), and  $\mathring{\boldsymbol{\Sigma}} = \text{diag}(\mathbf{U}_{\mathbf{A}_1}^T \mathbf{A} \mathbf{V}_{\mathbf{A}_1})$ . This approach is better than (3.2) in the sense that it uses more information than just  $\mathbf{G}_1$  and  $\mathbf{H}_1$ . However, the disadvantage is also very clear – the diagonal entries of  $\mathring{\boldsymbol{\Sigma}}$  can be negative, which violates the definition of singular values. Therefore, although

both (3.2) and (3.3) are very computationally efficient, we still want to keep looking for better approximation approaches.

### 3.3 A New TSVD Approximation Method via Kronecker Product Summations

The disadvantages of existing approaches lead to the question of how to use more terms in the Kronecker product summations in order to obtain a more accurate SVD approximation, while strictly adhering to the requirements of the SVD. We now propose an alternative approximation method described in [15, 16]. Starting with the Kronecker sum decomposition  $\mathbf{A} = \sum_{i=1}^R \mathbf{G}_i \otimes \mathbf{H}_i$ , and the SVD of the first term  $\mathbf{G}_1 \otimes \mathbf{H}_1 = \mathbf{U}_{A_1} \boldsymbol{\Sigma}_{A_1} \mathbf{V}_{A_1}^T$ , we have that

$$\begin{aligned}
 \mathbf{A} &= \sum_{i=1}^R \mathbf{G}_i \otimes \mathbf{H}_i \\
 &= \mathbf{U}_{A_1} \boldsymbol{\Sigma}_{A_1} \mathbf{V}_{A_1}^T + \sum_{i=2}^R \mathbf{G}_i \otimes \mathbf{H}_i \\
 &= \mathbf{U}_{A_1} \left( \boldsymbol{\Sigma}_{A_1} + \underbrace{\mathbf{U}_{A_1}^T \left( \sum_{i=2}^R \mathbf{G}_i \otimes \mathbf{H}_i \right) \mathbf{V}_{A_1}}_{\mathbf{W}} \right) \mathbf{V}_{A_1}^T \\
 &= \mathbf{U}_{A_1} \left( \boldsymbol{\Sigma}_{A_1} + \mathbf{W} \right) \mathbf{V}_{A_1}^T.
 \end{aligned}$$

In this new method we propose, we focus on computing the  $k$  largest singular values and corresponding left and right singular vectors (i.e., to obtain a TSVD

approximation). Hence, we need to apply a reordering matrix  $\mathbf{P} \in \mathbb{R}^{N \times N}$  to sort the singular values of  $\Sigma_{\mathbf{A}_1}$  in descending order in the same way as for (3.1). We obtain

$$\begin{aligned} \mathbf{A} &= \mathbf{U}_{\mathbf{A}_1} \mathbf{P} \left( \mathbf{P}^T (\Sigma_{\mathbf{A}_1} + \mathbf{W}) \mathbf{P} \right) (\mathbf{V}_{\mathbf{A}_1} \mathbf{P})^T \\ &= \bar{\mathbf{U}}_{\mathbf{A}_1} \left( \bar{\Sigma}_{\mathbf{A}_1} + \bar{\mathbf{W}} \right) \bar{\mathbf{V}}_{\mathbf{A}_1}^T, \end{aligned}$$

where  $\bar{\mathbf{U}}_{\mathbf{A}_1} = \mathbf{U}_{\mathbf{A}_1} \mathbf{P}$ ,  $\bar{\mathbf{V}}_{\mathbf{A}_1} = \mathbf{V}_{\mathbf{A}_1} \mathbf{P}$ ,  $\bar{\Sigma}_{\mathbf{A}_1} = \mathbf{P}^T \Sigma_{\mathbf{A}_1} \mathbf{P}$  and  $\bar{\mathbf{W}} = \mathbf{P}^T \mathbf{W} \mathbf{P}$ . Since we are interested in the  $k$  largest singular values, we further partition

$$\bar{\Sigma}_{\mathbf{A}_1} = \begin{bmatrix} \bar{\Sigma}_{\mathbf{A}_1, 1:k} & \mathbf{0} \\ \mathbf{0} & \bar{\Sigma}_{\mathbf{A}_1, k+1:N} \end{bmatrix}, \quad (3.4)$$

where  $\bar{\Sigma}_{\mathbf{A}_1, 1:k} \in \mathbb{R}^{k \times k}$  and  $\bar{\Sigma}_{\mathbf{A}_1, k+1:N} \in \mathbb{R}^{N-k \times N-k}$ . We also partition  $\bar{\mathbf{W}} = \begin{bmatrix} \bar{\mathbf{W}}_{11} & \bar{\mathbf{W}}_{12} \\ \bar{\mathbf{W}}_{21} & \bar{\mathbf{W}}_{22} \end{bmatrix}$ ,

where the blocks are of the same dimension as the blocks of  $\bar{\Sigma}_{\mathbf{A}_1}$  in (3.4). Using this notation,

$$\begin{aligned}
\mathbf{A} &= \bar{U}_{A_1} \left( \begin{bmatrix} \bar{\Sigma}_{A_1,1:k} & \mathbf{0} \\ \mathbf{0} & \bar{\Sigma}_{A_1,k+1:N} \end{bmatrix} + \begin{bmatrix} \bar{W}_{11} & \bar{W}_{12} \\ \bar{W}_{21} & \bar{W}_{22} \end{bmatrix} \right) \bar{V}_{A_1}^T \\
&= \bar{U}_{A_1} \left( \begin{bmatrix} \bar{\Sigma}_{A_1,1:k} + \bar{W}_{11} & \mathbf{0} \\ \mathbf{0} & \bar{\Sigma}_{A_1,k+1:N} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \bar{W}_{12} \\ \bar{W}_{21} & \bar{W}_{22} \end{bmatrix} \right) \bar{V}_{A_1}^T \\
&= \bar{U}_{A_1} \left( \begin{bmatrix} U_T \Sigma_T V_T^T & \mathbf{0} \\ \mathbf{0} & \bar{\Sigma}_{A_1,k+1:N} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \bar{W}_{12} \\ \bar{W}_{21} & \bar{W}_{22} \end{bmatrix} \right) \bar{V}_{A_1}^T \\
&= \bar{U}_{A_1} \begin{bmatrix} U_T & \mathbf{0} \\ \mathbf{0} & I \end{bmatrix} \left( \begin{bmatrix} \Sigma_T & \mathbf{0} \\ \mathbf{0} & \bar{\Sigma}_{A_1,k+1:N} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & U_T^T \bar{W}_{12} \\ \bar{W}_{21} V_T & \bar{W}_{22} \end{bmatrix} \right) \begin{bmatrix} V_T^T & \mathbf{0} \\ \mathbf{0} & I \end{bmatrix} \bar{V}_{A_1}^T \\
&= \check{U} \left( \begin{bmatrix} \Sigma_T & \mathbf{0} \\ \mathbf{0} & \bar{\Sigma}_{A_1,k+1:N} \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \widehat{W}_{12} \\ \widehat{W}_{21} & \bar{W}_{22} \end{bmatrix} \right) \check{V}^T, \tag{3.5}
\end{aligned}$$

where we have defined  $\widehat{W}_{12} = U_T^T \bar{W}_{12}$ ,  $\widehat{W}_{21} = \bar{W}_{21} V_T$ ,  $\check{U} = \bar{U}_{A_1} \begin{bmatrix} U_T & \mathbf{0} \\ \mathbf{0} & I \end{bmatrix}$  and

$\check{V} = \bar{V}_{A_1} \begin{bmatrix} V_T & \mathbf{0} \\ \mathbf{0} & I \end{bmatrix}$ . Moreover, we have used the full SVD of  $\mathbf{T} = \bar{\Sigma}_{A_1,1:k} + \bar{W}_{11} = U_T \Sigma_T V_T^T$  in the third step of the derivation. Although computing the full SVD of  $\mathbf{A}$  is impractical when  $N$  is large, it is feasible to compute the full SVD of  $\mathbf{T}$  if  $k$  is moderate. From (3.5), we are able to form a TSVD approximation,  $\check{\mathbf{A}}_{\text{TSVD}}$  of  $\mathbf{A}$

$$\check{\mathbf{A}}_{\text{TSVD}} \approx \check{U}_{1:k} \check{\Sigma}_{1:k} \check{V}_{1:k}^T, \tag{3.6}$$



where  $\check{\mathbf{U}}_{1:k} \in \mathbb{R}^{N \times k}$  and  $\check{\mathbf{V}}_{1:k} \in \mathbb{R}^{N \times k}$  are the submatrices of  $\check{\mathbf{U}}$  and  $\check{\mathbf{V}}$  containing their first  $k$  columns respectively, and  $\check{\Sigma}_{1:k} = \Sigma_T$  for consistency of notation.

### 3.3.1 Approximation Quality

In this subsection, we summarize theoretical results that provide bounds for different types of approximation errors, which is the author's main contribution to this joint work. Before we get started, a few more notations need to be introduced. Since we are bounding errors, true quantities need to be referred to. Let the true SVD of  $\mathbf{A}$  be

$$\mathbf{A} = \mathbf{U}\Sigma\mathbf{V}^T = \begin{bmatrix} \mathbf{U}_{1:k} & \mathbf{U}_{k+1:N} \end{bmatrix} \begin{bmatrix} \Sigma_{1:k} & \mathbf{0} \\ \mathbf{0} & \Sigma_{k+1:N} \end{bmatrix} \begin{bmatrix} \mathbf{V}_{1:k} & \mathbf{V}_{k+1:N} \end{bmatrix}^T,$$

and the true rank- $k$  TSVD approximation of  $\mathbf{A}$  (i.e., the best rank- $k$  approximation of  $\mathbf{A}$ ) is therefore defined as

$$\mathbf{A}_{\text{TSVD}} = \mathbf{U}_{1:k}\Sigma_{1:k}\mathbf{V}_{1:k}^T. \quad (3.7)$$

Sequentially, the true TSVD filtered solution  $\mathbf{x}_{\text{TSVD}}$  and the approximate TSVD solution  $\check{\mathbf{x}}_{\text{TSVD}}$  are

$$\mathbf{x}_{\text{TSVD}} = \mathbf{V}_{1:k}\Sigma_{1:k}^{-1}\mathbf{U}_{1:k}^T\mathbf{b} \quad \text{and} \quad \check{\mathbf{x}}_{\text{TSVD}} = \check{\mathbf{V}}_{1:k}\check{\Sigma}_{1:k}^{-1}\check{\mathbf{U}}_{1:k}^T\mathbf{b} \quad (3.8)$$

We start the error analysis with the bound for the singular vector subspaces,

similar to the results presented by Fierro and Bunch [13] for the case of URV and ULV factorizations. The “signal” subspace of  $\mathbf{A}$  is spanned by  $\mathbf{U}_{1:k}$  and the “noise” subspace is spanned by  $\mathbf{V}_{k+1:N}$ . The distance between the approximated and true signal and noise subspaces are measured by

$$\|\mathbf{U}_{1:k}^T \check{\mathbf{U}}_{k+1:N}\|_2 \quad \text{and} \quad \|\mathbf{V}_{1:k}^T \check{\mathbf{V}}_{k+1:N}\|_2,$$

which represent one measure for the quality of the TSVD approximation.

**Theorem 3.3.1.** *Consider the factorizations of  $\mathbf{A}$  given in equations (3.7) and (3.6), where  $\sigma_i$  denotes a true singular value and  $\check{\sigma}_i$  denotes an approximate singular value.*

*Then*

$$\|\mathbf{U}_{1:k}^T \check{\mathbf{U}}_{k+1:N}\| \leq \frac{\sigma_k \|\widehat{\mathbf{W}}_{21}\| + \|\widehat{\mathbf{W}}_{12}\| \|\boldsymbol{\Sigma}_{K_1, k+1:N} + \overline{\mathbf{W}}_{22}\|}{\sigma_k^2 - \|\boldsymbol{\Sigma}_{K_1, k+1:N} + \overline{\mathbf{W}}_{22}\|^2}$$

*and*

$$\|\mathbf{V}_{1:k}^T \check{\mathbf{V}}_{k+1:N}\| \leq \frac{\sigma_k \|\widehat{\mathbf{W}}_{12}\| + \|\widehat{\mathbf{W}}_{21}\| \|\boldsymbol{\Sigma}_{K_1, k+1:N} + \overline{\mathbf{W}}_{22}\|}{\sigma_k^2 - \|\boldsymbol{\Sigma}_{K_1, k+1:N} + \overline{\mathbf{W}}_{22}\|^2}$$

We refer reders to [16] for the full proof of Theorem 3.3.1. This Theorem is the foundation for proving error bounds for the pseudoinverse  $\mathbf{A}_{\text{TSVD}}$  and approximate solution  $\mathbf{x}_{\text{TSVD}}$  shown below.

**Theorem 3.3.2.** *Consider the true TSVD operator  $\mathbf{A}_{\text{TSVD}}$  defined in (3.7) and its approximation  $\check{\mathbf{A}}_{\text{TSVD}}$  given by (3.6). Let  $\check{\sigma}_i$  be the  $i^{\text{th}}$  diagonal entry of  $\check{\boldsymbol{\Sigma}}_{1:k}$ , and define  $\varphi = (1 + \sqrt{5})/2$ . Then*

$$\frac{\|\mathbf{A}_{\text{TSVD}}^\dagger - \check{\mathbf{A}}_{\text{TSVD}}^\dagger\|}{\|\mathbf{A}_{\text{TSVD}}^\dagger\|} \leq \frac{\varphi}{\check{\sigma}_k} \left( \sigma_1 \|\mathbf{V}_{1:k}^T \check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right).$$

*Proof.* The proof starts with a perturbation result for pseudoinverses presented in [70]: If  $\mathbf{C}$  is an acute perturbation [49] of  $\mathbf{D}$ , with  $\mathbf{D} = \mathbf{C} + \delta\mathbf{C}$ , then

$$\|\mathbf{C}^\dagger - \mathbf{D}^\dagger\| \leq \varphi \|\mathbf{C}^\dagger\| \|\mathbf{D}^\dagger\| \|\delta\mathbf{C}\|. \quad (3.9)$$

Since

$$\begin{aligned} \|\mathbf{A}_{\text{TSVD}} - \check{\mathbf{A}}_{\text{TSVD}}\| &= \|\mathbf{A}\mathbf{V}_{1:k}\mathbf{V}_{1:k}^T - (\mathbf{A}\check{\mathbf{V}}_{1:k}\check{\mathbf{V}}_{1:k}^T - \check{\mathbf{U}}_{k+1:N}\widehat{\mathbf{W}}_{21}\check{\mathbf{V}}_{1:k}^T)\| \\ &\leq \|\mathbf{A}\|\|\mathbf{V}_{1:k}\mathbf{V}_{1:k}^T - \check{\mathbf{V}}_{1:k}\check{\mathbf{V}}_{1:k}^T\| + \|\widehat{\mathbf{W}}_{21}\|, \end{aligned}$$

and it is proved in [24] that  $\|\mathbf{V}_{1:k}\mathbf{V}_{1:k}^T - \check{\mathbf{V}}_{1:k}\check{\mathbf{V}}_{1:k}^T\| = \|\mathbf{V}_{1:k}^T\check{\mathbf{V}}_{k+1:N}\|$ , we get

$$\|\mathbf{A}_{\text{TSVD}} - \check{\mathbf{A}}_{\text{TSVD}}\| \leq \|\mathbf{A}\|\|\mathbf{V}_{1:k}^T\check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\|.$$

$\check{\mathbf{A}}_{\text{TSVD}}$  is an acute perturbation of  $\mathbf{A}_{\text{TSVD}}$ , so by (3.9),

$$\|\mathbf{A}_{\text{TSVD}}^\dagger - \check{\mathbf{A}}_{\text{TSVD}}^\dagger\| \leq \varphi \|\mathbf{A}_{\text{TSVD}}^\dagger\| \|\check{\mathbf{A}}_{\text{TSVD}}^\dagger\| \left( \|\mathbf{A}\|\|\mathbf{V}_{1:k}^T\check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right). \quad (3.10)$$

Dividing both sides of the inequality by  $\|\mathbf{A}_{\text{TSVD}}^\dagger\|$ , we obtain

$$\begin{aligned} \frac{\|\mathbf{A}_{\text{TSVD}}^\dagger - \check{\mathbf{A}}_{\text{TSVD}}^\dagger\|}{\|\mathbf{A}_{\text{TSVD}}^\dagger\|} &\leq \varphi \|\check{\mathbf{A}}_{\text{TSVD}}^\dagger\| \left( \|\mathbf{A}\|\|\mathbf{V}_{1:k}^T\check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right) \\ &= \frac{\varphi}{\check{\sigma}_k} \left( \sigma_1 \|\mathbf{V}_{1:k}^T\check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right). \end{aligned}$$

□

**Theorem 3.3.3.** Consider the true and approximate TSVD filtered solutions  $\mathbf{x}_{TSVD}$  and  $\check{\mathbf{x}}_{TSVD}$  as in (3.8). Further, define the residual  $\mathbf{r} = \mathbf{b} - \mathbf{A}\mathbf{x}_{TSVD}$ . Then

$$\frac{\|\mathbf{x}_{TSVD} - \check{\mathbf{x}}_{TSVD}\|}{\|\mathbf{x}_{TSVD}\|} \leq \frac{\varphi\sigma_1}{\sigma_k\hat{\sigma}_k\sqrt{1 - \frac{\|\mathbf{r}\|^2}{\|\mathbf{b}\|^2}}} \left( \sigma_1 \|\mathbf{V}_{1:k}^T \check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right). \quad (3.11)$$

*Proof.* From inequality (3.10), we obtain

$$\begin{aligned} \|\mathbf{x}_{TSVD} - \check{\mathbf{x}}_{TSVD}\| &\leq \|\mathbf{A}_{TSVD}^\dagger - \check{\mathbf{A}}_{TSVD}^\dagger\| \|\mathbf{b}\| \\ &\leq \varphi \|\mathbf{A}_{TSVD}^\dagger\| \|\check{\mathbf{A}}_{TSVD}^\dagger\| \left( \|\mathbf{A}\| \|\mathbf{A}_{1:k}^T \check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right) \|\mathbf{d}\|. \end{aligned}$$

By the triangle inequality and submultiplicativity of induced norms,

$$\|\mathbf{x}_{TSVD}\| \geq \frac{\|\mathbf{b}\|}{\|\mathbf{A}\|} \sqrt{1 - \frac{\|\mathbf{r}\|^2}{\|\mathbf{b}\|^2}},$$

it follows that

$$\begin{aligned} \frac{\|\mathbf{x}_{TSVD} - \check{\mathbf{x}}_{TSVD}\|}{\|\mathbf{x}_{TSVD}\|} &\leq \frac{\varphi}{\sqrt{1 - \frac{\|\mathbf{r}\|^2}{\|\mathbf{b}\|^2}}} \|\mathbf{A}_{TSVD}^\dagger\| \|\check{\mathbf{A}}_{TSVD}^\dagger\| \|\mathbf{A}\| \left( \|\mathbf{A}\| \|\mathbf{V}_{1:k}^T \check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right) \\ &= \frac{\varphi\sigma_1}{\sigma_k\hat{\sigma}_k\sqrt{1 - \frac{\|\mathbf{r}\|^2}{\|\mathbf{b}\|^2}}} \left( \sigma_1 \|\mathbf{V}_{1:k}^T \check{\mathbf{V}}_{k+1:N}\| + \|\widehat{\mathbf{W}}_{21}\| \right). \end{aligned}$$

□

These two theorems tell us that the quality of the approximate pseudoinverse and the approximate TSVD solution is dependent on three conditions:

- The distance between the approximate and true noise subspaces;

- The ratio of the largest singular value  $\sigma_1$  of  $\mathbf{A}$  to the  $k$ th approximate singular value  $\check{\sigma}_k$ ;
- The size of  $\|\widehat{\mathbf{W}}_{21}\|$ , which depends on whether  $\mathbf{G}_1 \otimes \mathbf{H}_1$  is a good approximation to  $\mathbf{A}$  and whether  $k$  is a good truncation rank.

This newly proposed TSVD approximation method is supported by both theoretical and numerical results, in which the new method computes more accurate solutions than (3.2) and (3.3). Our new method can also approximate singular values and vectors more efficiently than other well-known schemes such as randomized methods [27] and Golub-Kahan bidiagonalization [24, 47]. In this thesis, we omit details on method implementation and numerical experiments, and refer readers to [15] and [16] for a thorough introduction of this new approach.

## Chapter 4

# Low Rank Regularization: Derivation of New Methods

In the Chapter 3, we discussed Kronecker product summations based methods for approximating the TSVD of  $\mathbf{A}$ . TSVD is a singular value filtering type of regularization scheme that eliminates the noise contribution caused by small singular values of  $\mathbf{A}$ , but there is not an explicit regularization term  $\mathcal{R}(\mathbf{x})$  associated with it. In this Chapter, we move on to a different type of regularization – nuclear norm regularization (NNR), for which  $\mathcal{R}(\mathbf{x}) = \|\text{vec}^{-1}(\mathbf{x})\|_* = \|\mathbf{X}\|_*$ , and describe how to solve the nuclear norm regularized problem using Krylov subspace methods and novel preconditioners. We introduce methods that target on both square matrices  $\mathbf{A} \in \mathbb{R}^{N \times N}$  and rectangular matrices  $\mathbf{A} \in \mathbb{R}^{M \times N}$ . These new methods are proposed in [18], which is the author’s joint work with Gazzola and Nagy. We apply nuclear norm regularization to imaging inverse problems that can be formulated as linear

systems

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \tag{4.1}$$

where  $\mathbf{A}$ ,  $\mathbf{x}$  and  $\mathbf{b}$  are defined in the exact same way as in (1.1), in order to implicitly enforce a low rank constraint by penalizing the sum of singular values of the solution. Indeed, two-dimensional images are often assumed to have low-rank or to be well-approximated by low-rank two-dimensional arrays (see [57] and the references therein).

Numerical linear algebra solvers for the estimation of low-rank solutions to linear systems have been developed in the literature, mainly targeting well-posed linear discrete problems, such as those arising when considering the numerical solution of stochastic PDEs (see [48] and the references therein). In particular, the authors of [48] devise a restarted GMRES-like method (RS-LR-GMRES) that involves low-rank projections of the basis vectors of the solution subspace, as well as a low-rank projection of the current solution at the end of each cycle. Since, in general, the basic operations involved in standard GMRES (such as matrix-vector products and vector sums) increase the ranks of the computed quantities, low-rank projections are needed to assure that the computed solution is low-rank.

In the framework of compressive sensing, the authors of [6] consider a modified version of the conjugate gradient method that incorporates appropriate rank-truncation operations. All the methods mentioned so far employ, often in a heuristic way, Krylov subspace methods together with rank-reduction operations (e.g., projections onto a chosen set of low-rank matrices). Since many Krylov subspace methods are iterative regularization methods for (4.1), this brings us to the question of how

incorporating rank-reduction operations would affect the solution of the discrete inverse problem (4.1), with a particular focus on imaging applications.

## 4.1 Nuclear Norm Regularization

Low-rank matrix estimation can be naturally formulated as a nonconvex optimization problem having either:

- (i) a least-squares data fitting term as objective function and a rank constraint; or
- (ii) the rank of  $\mathbf{X} = \text{vec}^{-1}(\mathbf{x})$  as objective function and a constraint on the least-squares data fitting term.

The last instance is commonly referred to as *affine rank minimization problem*, and both formulations are in general NP-hard [57]. In this Chapter, we consider the unconstrained and convex optimization problem

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\text{vec}^{-1}(\mathbf{x})\|_* , \quad (4.2)$$

where  $\lambda > 0$  is a regularization parameter and  $\|\cdot\|_*$  denotes the nuclear norm of  $\text{vec}^{-1}(\mathbf{x}) = \mathbf{X}$ , defined as the sum of the singular values of  $\mathbf{X}$ . If the singular value decomposition (SVD) of  $\mathbf{X}$  is given by  $\mathbf{X} = \mathbf{U}_\mathbf{X} \mathbf{\Sigma}_\mathbf{X} \mathbf{V}_\mathbf{X}^T$ , where  $\mathbf{U}_\mathbf{X}, \mathbf{V}_\mathbf{X} \in \mathbb{R}^{n \times n}$  are orthogonal matrices, and  $\mathbf{\Sigma}_\mathbf{X} \in \mathbb{R}^{n \times n}$  is the diagonal matrix whose diagonal entries



are  $\sigma_1(\mathbf{X}) \geq \dots \geq \sigma_n(\mathbf{X}) \geq 0$ , then

$$\|\mathbf{X}\|_* = \sum_{i=1}^n \sigma_i(\mathbf{X}).$$

Problem (4.2) is referred to as a *nuclear norm regularized (NNR) problem*. In particular, the nuclear norm is a convex function that has been proven to be the best convex lower approximation of the rank function over the set of matrices  $\mathbf{X}$  such that  $\|\mathbf{X}\|_2 \leq 1$  (see [57] and the references therein). The nuclear norm has been used in many applications, such as low-rank matrix completion and compressed sensing; see, e.g., [7, 23, 43, 50, 57], where the constrained formulation of problem (4.2) has also been considered (note that, for a proper choice of  $\lambda > 0$ , constrained and unconstrained formulations are equivalent; see, e.g., [60]). In the framework of compressive sensing, under the assumption that the matrix  $\mathbf{A}$  satisfies a certain null-space property, recovery guarantees for the affine rank minimization problem are proven in [14, 51]. We also consider the following formulation

$$\min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\text{vec}^{-1}(\mathbf{x})\|_{*,p}, \quad \text{where} \quad \|\mathbf{X}\|_{*,p} = \sum_{i=1}^n (\sigma_i(\mathbf{X}))^p, \quad 0 < p \leq 1. \quad (4.3)$$

Problem (4.3) is referred to as *NNRp problem*, and it generalizes problem (4.2) (which is obtained taking  $p = 1$  in (4.3)). The constrained version of (4.3) is already considered in [51], where the authors empirically show an improved recovery performance of the constrained formulation of problem (4.3) with  $p < 1$  with respect to  $p = 1$ . Note, however, that the choice  $p < 1$  in (4.3) results in a nonconvex minimization problem.

## 4.2 Derivation of New Methods

In the following, we will quite often interchange  $\mathbf{x}$  and  $\mathbf{X}$  and, with a slight abuse of notations, we will denote the action of a linear operator on  $\mathbf{x}$  or  $\mathbf{X}$  by  $\mathcal{A}(\mathbf{X}) = \mathcal{A}\mathbf{X} = A\mathbf{x}$ , and the action of the adjoint operator by  $\mathcal{A}^*(\mathbf{Y}) = \mathcal{A}^*\mathbf{Y} = A^T \text{vec}(\mathbf{Y})$ . Define the smooth Schatten- $p$  function as

$$\mathcal{S}_p^\gamma(\mathbf{X}) = \text{Tr}((\mathbf{X}^T \mathbf{X} + \gamma \mathbf{I})^{p/2}), \quad \text{with } \gamma > 0.$$

Note that  $\mathcal{S}_p^\gamma(\mathbf{X})$  is differentiable for  $p > 0$  and convex for  $p \geq 1$ . In particular, for  $p = 1$  and  $\gamma = 0$  (i.e., no smoothing),

$$\mathcal{S}_1^0(\mathbf{X}) = \text{Tr}((\mathbf{X}^T \mathbf{X})^{1/2}) = \|\mathbf{X}\|_*.$$

We start by considering the following smooth approximation to (4.3):

$$\min_{\mathbf{X} \in \mathbb{R}^{n \times n}} \|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_F^2 + \lambda \mathcal{S}_p^\gamma(\mathbf{X}). \quad (4.4)$$

The following derivations are valid for  $p > 0$  (and we keep them generic, being aware that  $p = 1$  approximates (4.2)). The optimality conditions associated to (4.4) read

$$\begin{aligned} 0 &= \nabla_{\mathbf{X}} (\|\mathcal{A}(\mathbf{X}) - \mathbf{B}\|_F^2 + \lambda \mathcal{S}_p^\gamma(\mathbf{X})) \\ &= 2\mathcal{A}^*(\mathcal{A}(\mathbf{X}) - \mathbf{B}) + \lambda p(\mathbf{X}\mathbf{X}^T + \gamma \mathbf{I})^{p/2-1} \mathbf{X}, \end{aligned} \quad (4.5)$$

where we have used that

$$\nabla_{\mathbf{X}} \text{Tr}((\mathbf{X}^T \mathbf{X} + \gamma \mathbf{I})^{p/2}) = p \mathbf{X} (\mathbf{X}^T \mathbf{X} + \gamma \mathbf{I})^{p/2-1} = p (\mathbf{X} \mathbf{X}^T + \gamma \mathbf{I})^{p/2-1} \mathbf{X}.$$

Equivalently, the nonlinear system of equations (4.5) with respect to  $\mathbf{X}$  can be expressed as

$$\begin{aligned} \mathbf{X} &= \left( \mathcal{A}^* \mathcal{A} + \widehat{\lambda} (\mathbf{X} \mathbf{X}^T + \gamma \mathbf{I})^{p/2-1} \right)^{-1} \mathcal{A}^* \mathbf{B} \\ &= \left( \mathcal{A}^* \mathcal{A} + \widehat{\lambda} ((\mathbf{X} \mathbf{X}^T + \gamma \mathbf{I})^{p/4-1/2})^T (\mathbf{X} \mathbf{X}^T + \gamma \mathbf{I})^{p/4-1/2} \right)^{-1} \mathcal{A}^* \mathbf{B}, \end{aligned}$$

with  $\widehat{\lambda} = \lambda p/2$ , which is naturally associated to the following fixed-point iteration scheme

$$\mathbf{X}_{k+1} = \left( \mathcal{A}^* \mathcal{A} + \widehat{\lambda} ((\mathbf{X}_k \mathbf{X}_k^T + \gamma \mathbf{I})^{p/4-1/2})^T (\mathbf{X}_k \mathbf{X}_k^T + \gamma \mathbf{I})^{p/4-1/2} \right)^{-1} \mathcal{A}^* \mathbf{B}, \quad (4.6)$$

which leads to the solution of (4.4). Equivalently,

$$\mathbf{X}_{k+1} = \arg \min_{\mathbf{X}} \left\| \begin{bmatrix} \mathcal{A} \\ \sqrt{\widehat{\lambda}} (\mathbf{X}_k \mathbf{X}_k^T + \gamma \mathbf{I})^{p/4-1/2} \end{bmatrix} \mathbf{X} - \begin{bmatrix} \mathbf{B} \\ \mathbf{0} \end{bmatrix} \right\|_F^2,$$

i.e., (4.6) are the normal equations associated to the penalized least squares problem written above or, equivalently,

$$\mathbf{X}_{k+1} = \arg \min_{\mathbf{X}} \|\mathcal{A} \mathbf{X} - \mathbf{B}\|_F^2 + \widehat{\lambda} \|(\mathbf{X}_k \mathbf{X}_k^T + \gamma \mathbf{I})^{p/4-1/2} \mathbf{X}\|_F^2. \quad (4.7)$$

We now reformulate problem (4.7) in vectorial form.

Let  $\mathbf{U}_{\mathbf{X}_k} \boldsymbol{\Sigma}_{\mathbf{X}_k} \mathbf{V}_{\mathbf{X}_k}^T = \mathbf{X}_k$  be the SVD of  $\mathbf{X}_k$ ; thanks to the invariance of the Frobenius norm under orthogonal transformations, the regularization term in the above problem can be rewritten as

$$\begin{aligned} \|(\mathbf{X}_k \mathbf{X}_k^T + \gamma \mathbf{I})^{p/4-1/2} \mathbf{X}\|_F^2 &= \|\mathbf{U}_{\mathbf{X}_k} (\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{p/4-1/2} \mathbf{U}_{\mathbf{X}_k}^T \mathbf{X}\|_F^2 \\ &= \|(\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{p/4-1/2} \mathbf{U}_{\mathbf{X}_k}^T \mathbf{X} \mathbf{V}_{\mathbf{X}_k}\|_F^2. \end{aligned}$$

Using well-known Kronecker product properties

$$\begin{aligned} \|(\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{p/4-1/2} \mathbf{U}_{\mathbf{X}_k}^T \mathbf{X} \mathbf{V}_{\mathbf{X}_k}\|_F^2 &= \|\text{vec}((\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{p/4-1/2} \mathbf{U}_{\mathbf{X}_k}^T \mathbf{X} \mathbf{V}_{\mathbf{X}_k})\|_2^2 \\ &= \|(\mathbf{V}_{\mathbf{X}_k}^T \otimes ((\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{p/4-1/2} \mathbf{U}_{\mathbf{X}_k}^T)) \mathbf{x}\|_2^2 \\ &= \|(\mathbf{I} \otimes (\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{p/4-1/2}) (\mathbf{V}_{\mathbf{X}_k}^T \otimes \mathbf{U}_{\mathbf{X}_k}^T) \mathbf{x}\|_2^2. \end{aligned}$$

Problem (4.7) is therefore equivalent to

$$\mathbf{x}_{k+1} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \hat{\lambda} \underbrace{\|(\mathbf{I} \otimes (\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{p/4-1/2})}_{=:(\mathbf{W}_p^\gamma)_k} \overbrace{(\mathbf{V}_{\mathbf{X}_k}^T \otimes \mathbf{U}_{\mathbf{X}_k}^T)}^{=:\mathbf{S}_k} \mathbf{x}\|_2^2. \quad (4.8)$$

In the above formulation,  $(\mathbf{W}_p^\gamma)_k$  is a diagonal weighting matrix and  $\mathbf{S}_k$  is an orthogonal matrix; both  $(\mathbf{W}_p^\gamma)_k$  and  $\mathbf{S}_k$  depend on the current approximation  $\mathbf{x}_k$  of the solution  $\mathbf{x}$ . Intuitively, the matrix  $\mathbf{S}_k$  maps  $\mathbf{x}$  into the ‘singular value domain’ of  $\mathbf{X}_k$  (and acts as an iteration-dependent sparsity transform), and the matrix  $(\mathbf{W}_p^\gamma)_k$  assigns suitable weights that allow to approximate a  $p$ -norm of the singular values.

Therefore, the penalization term in (4.8) can be interpreted as a reweighted vectorial 2-norm, with respect to a transformation of the solution  $\mathbf{x}$ . For this reason, the proposed approach is called “IRN-NNR $p$ ” and is summarized in Algorithm 6.

---

**Algorithm 6** IRN-NNR $p$

---

- 1:
  - 2: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $(\mathbf{W}_p^\gamma)_0 = \mathbf{I}$ ,  $\mathbf{S}_0 = \mathbf{I}$
  - 3: **for**  $k = 0, 1, \dots$  until a stopping criterion is satisfied **do**
  - 4:   Solve problem (4.8)
  - 5:   “Decrease”  $\gamma$
  - 6:   Update  $(\mathbf{W}_p^\gamma)_{k+1}$  and  $\mathbf{S}_{k+1}$
  - 7: **end for**
- 

### 4.3 Solution of problem (4.8) via Krylov methods

In this section, we derive new strategies for the efficient solution of the sequence of sub-problems (4.8) appearing in Algorithm 6. First, we rewrite problem (4.8) using an appropriate change of variable as

$$\hat{\mathbf{x}}_{k+1} = \arg \min_{\hat{\mathbf{x}}} \|\mathbf{A}\mathbf{S}_k^T (\mathbf{W}_p^\gamma)_k^{-1} \hat{\mathbf{x}} - \mathbf{b}\|_2^2 + \hat{\lambda} \|\hat{\mathbf{x}}\|_2^2, \text{ with } \hat{\mathbf{x}} = (\mathbf{W}_p^\gamma)_k \mathbf{S}_k \mathbf{x}. \quad (4.9)$$

Note that

$$\mathbf{S}_k^T = \mathbf{S}_k^{-1} = \mathbf{V}_{\mathbf{X}_k} \otimes \mathbf{U}_{\mathbf{X}_k} \quad \text{and} \quad (\mathbf{W}_p^\gamma)_k^{-1} = \mathbf{I} \otimes (\boldsymbol{\Sigma}_{\mathbf{X}_k}^2 + \gamma \mathbf{I})^{1/2-p/4}, \quad (4.10)$$

so that the above transformations (inversion of an orthogonal and a diagonal matrix) are numerically affordable by exploiting properties of Kronecker products. The Tikhonov-regularized problem (4.9) in standard form is equivalent to the Tikhonov-

regularized problem (4.8) in general form. Many Krylov subspace methods based on the Golub-Kahan Bidiagonalization (GKB) or Arnoldi algorithms can be employed to approximate the solution of (4.9), e.g., LSQR and GMRES, as explained in Section 2.4. Moreover, if the regularization parameter  $\hat{\lambda}$  is not known a priori, many efficient strategies to set its value adaptively within the sequence of projected problems can be used (i.e., in the framework of hybrid methods; see Section 2.4.2 and [21, 44]). The matrices  $\mathbf{S}_k$  and  $(\mathbf{W}_p^\gamma)_k^{-1}$  can be formally thought of as preconditioners for the original problem (4.1), whose purpose is to enforce additional regularization into the solution subspace, rather than speeding-up the convergence of linear solvers applied to (4.1).

### 4.3.1 Methods Based on the GKB Algorithm

The  $m$ th step of the GKB algorithm applied to the matrix  $\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)_k^{-1}$  with starting vector  $\mathbf{b}$  (i.e., taking  $\mathbf{x}_0 = \mathbf{0}$ ) can be expressed by the following partial matrix factorizations

$$(\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)_k^{-1})\mathbf{V}_m = \mathbf{U}_{m+1}\mathbf{B}_m \quad \text{and} \quad ((\mathbf{W}_p^\gamma)_k^{-1}\mathbf{S}_k\mathbf{A}^T)\mathbf{U}_{m+1} = \mathbf{V}_{m+1}\mathbf{B}_{m+1}^T, \quad (4.11)$$

where  $\mathbf{U}_j \in \mathbb{R}^{M \times j}$  and  $\mathbf{V}_j \in \mathbb{R}^{N \times j}$  (with  $j = m, m+1$  and  $\mathbf{U}_j\mathbf{e}_1 = \mathbf{b}/\|\mathbf{b}\|_2$ ) have orthonormal columns, and  $\mathbf{B}_{m+1} \in \mathbb{R}^{(m+1) \times (m+1)}$  is lower bidiagonal (with  $\mathbf{B}_m$  obtained by removing the last column of  $\mathbf{B}_{m+1}$ ). The orthonormal columns of  $\mathbf{V}_m$  are

such that

$$\mathcal{R}(\mathbf{V}_m) = \mathcal{K}_m \left( ((\mathbf{W}_p^\gamma)^{-1} \mathbf{S}_k \mathbf{A}^T) (\mathbf{A} \mathbf{S}_k^T (\mathbf{W}_p^\gamma)^{-1}), ((\mathbf{W}_p^\gamma)^{-1} \mathbf{S}_k \mathbf{A}^T) \mathbf{b} \right) .$$

We find an approximate solution of (4.9) by imposing  $\hat{\mathbf{x}} \in \mathcal{R}(\mathbf{V}_m)$ , i.e.,  $\hat{\mathbf{x}}_m = \mathbf{V}_m \mathbf{y}_m$ , where, by exploiting the first decomposition in (4.11) and the properties of the matrices appearing therein,  $\mathbf{y}_m \in \mathbb{R}^m$  is such that

$$\mathbf{y}_m = \arg \min_{\mathbf{y} \in \mathbb{R}^m} \|\mathbf{B}_m \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2^2 + \hat{\lambda}_m \|\mathbf{y}\|_2^2 . \quad (4.12)$$

We used the notation  $\hat{\lambda}_m$  for the regularization parameter to highlight that its value can be adaptively set within the iterations. The approximate solution to problem (4.8) is such that

$$\mathbf{x} = \mathbf{S}_k^T (\mathbf{W}_p^\gamma)^{-1} \hat{\mathbf{x}} \in \mathcal{K}_m \left( (\mathbf{S}_k^T (\mathbf{W}_p^\gamma)^{-2} \mathbf{S}_k) \mathbf{A}^T \mathbf{A}, (\mathbf{S}_k^T (\mathbf{W}_p^\gamma)^{-2} \mathbf{S}_k) \mathbf{A}^T \mathbf{b} \right) . \quad (4.13)$$

Looking at the above approximation subspace for the solution  $\mathbf{x}$ , it is evident that the “preconditioner” acts by first mapping into the “singular value domain” (by applying  $\mathbf{S}_k$ ), enforcing sparsity in the singular values (by reweighting with  $(\mathbf{W}_p^\gamma)^{-2}$ ), and eventually transforming back into the “solution domain” (by applying  $\mathbf{S}_k^T$ ).

### 4.3.2 Methods Based on the Arnoldi Algorithm

If  $\mathbf{A}$  is square, the  $m$ th step of the Arnoldi algorithm applied to the matrix  $\mathbf{A} \mathbf{S}_k^T (\mathbf{W}_p^\gamma)^{-1}$  with starting vector  $\mathbf{b}$  (i.e., taking  $\mathbf{x}_0 = \mathbf{0}$ ) can be expressed by the following partial

matrix factorization

$$(\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1})\mathbf{V}_m = \mathbf{V}_{m+1}\mathbf{H}_m, \quad (4.14)$$

where  $\mathbf{V}_j \in \mathbb{R}^{N \times j}$  (with  $j = m, m + 1$  and  $\mathbf{V}_j \mathbf{e}_1 = \mathbf{b}/\|\mathbf{b}\|_2$ ) has orthonormal columns such that

$$\mathcal{R}(\mathbf{V}_m) = \mathcal{K}_m(\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}, \mathbf{b}),$$

and  $\mathbf{H}_m \in \mathbb{R}^{(m+1) \times m}$  is upper Hessenberg. Similarly to the GKB case, we find an approximate solution of (4.9) by imposing  $\hat{\mathbf{x}} \in \mathcal{R}(\mathbf{V}_m)$  and by solving a projected Tikhonov problem of order  $m$ . The approximate solution to problem (4.8) is such that

$$\mathbf{x} = \mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\hat{\mathbf{x}} \in \mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathcal{K}_m(\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}, \mathbf{b}),$$

where

$$\begin{aligned} & \mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathcal{K}_m(\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}, \mathbf{b}) \\ &= \text{span}\{\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathbf{b}, \dots, (\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathbf{A})^{m-1}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathbf{b}\} \\ &= \mathcal{K}_m(\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathbf{A}, \mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathbf{b}). \end{aligned}$$

Contrarily to the GKB case, we immediately notice that, in this context,  $\mathbf{x}$  does not belong to a meaningful approximation subspace. Indeed, just by looking at the first vector:  $\mathbf{b}$  is in the image space and  $(\mathbf{W}_p^\gamma)^{-1}$  is supposed to act on the singular value space of  $\mathbf{X}_k$ , so  $(\mathbf{W}_p^\gamma)^{-1}\mathbf{b}$  is hard to interpret; furthermore,  $\mathbf{S}_k^T$  is supposed to link the singular value space of  $\mathbf{X}_k$  to the image space, so  $\mathbf{S}_k^T(\mathbf{W}_p^\gamma)^{-1}\mathbf{b}$  is also hard for us to interpret. Although the generated solution subspace is not meaningful for our



applications, it may still have the potential to be a good subspace in other contexts. Similarly to what is proposed in [4, 11], where the Arnoldi algorithm is applied to a regularized problem that enforces sparsity in the wavelet domain, we propose to fix this issue by incorporating  $\mathbf{S}_k$  also as an orthogonal left “preconditioner” for the original system (4.1) so that, by exploiting the invariance of the vector 2-norm under orthogonal transformations, problem (4.9) can be equivalently reformulated as

$$\hat{\mathbf{x}}_{k+1} = \arg \min_{\hat{\mathbf{x}}} \|\mathbf{S}_k(\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)_k^{-1}\hat{\mathbf{x}} - \mathbf{b})\|_2^2 + \hat{\lambda}\|\hat{\mathbf{x}}\|_2^2, \text{ with } \hat{\mathbf{x}} = (\mathbf{W}_p^\gamma)_k\mathbf{S}_k\mathbf{x}. \quad (4.15)$$

The (right and left) preconditioned Arnoldi algorithm applied to problem (4.15) can now be expressed by the following partial matrix factorization

$$(\mathbf{S}_k\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)_k^{-1})\mathbf{V}_m = \mathbf{V}_{m+1}\mathbf{H}_m. \quad (4.16)$$

We find an approximate solution of (4.15) by imposing

$$\hat{\mathbf{x}} \in \mathcal{R}(\mathbf{V}_m) = \mathcal{K}_m(\mathbf{S}_k\mathbf{A}\mathbf{S}_k^T\mathbf{W}^{-1}, \mathbf{S}_k\mathbf{b}),$$

i.e.,  $\hat{\mathbf{x}}_m = \mathbf{V}_m\mathbf{y}_m$ , where, by exploiting (4.16) and the properties of the matrices appearing therein,  $\mathbf{y}_m \in \mathbb{R}^m$  is such that

$$\mathbf{y}_m = \arg \min_{\mathbf{y} \in \mathbb{R}^m} \|\mathbf{H}_m\mathbf{y} - \|\mathbf{b}\|_2\mathbf{e}_1\|_2^2 + \hat{\lambda}_m\|\mathbf{y}\|_2^2. \quad (4.17)$$

Hence

$$\begin{aligned} \mathbf{x} &\in \mathbf{S}_k^T (\mathbf{W}_p^\gamma)_k^{-1} \mathcal{K}_m(\mathbf{S}_k \mathbf{A} \mathbf{S}_k^T (\mathbf{W}_p^\gamma)_k^{-1}, \mathbf{S}_k \mathbf{b}) \\ &= \mathcal{K}_m(\mathbf{S}_k^T (\mathbf{W}_p^\gamma)_k^{-1} \mathbf{S}_k \mathbf{A}, \mathbf{S}_k^T (\mathbf{W}_p^\gamma)_k^{-1} \mathbf{S}_k \mathbf{b}), \end{aligned} \quad (4.18)$$

which is suitable for approximating the solution. The new methods based on the GKB algorithm (for generic matrices) and Arnoldi algorithm (only if  $\mathbf{A} \in \mathbb{R}^{N \times N}$ ) are called “IRN-LSQR-NNR $p$ ” and “IRN-GMRES-NNR $p$ ”, respectively, and are summarized in Algorithm 7.

---

**Algorithm 7** IRN-LSQR-NNR $p$  and IRN-GMRES-NNR $p$

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $(\mathbf{W}_p^\gamma)_0 = \mathbf{I}$ ,  $\mathbf{S}_0 = \mathbf{I}$
  - 2: **for**  $k = 0, 1, \dots$  until a stopping criterion is satisfied **do**
  - 3:   **for**  $m = 1, 2, \dots$  until a stopping criterion is satisfied **do**
  - 4:     Update the factorizations (4.11) and (4.16), respectively
  - 5:     Solve the projected problem (4.12) and (4.17), respectively, tuning  $\hat{\lambda}_m$  if necessary
  - 6:   **end for**
  - 7:   “Decrease”  $\gamma$
  - 8:   Update the new  $(\mathbf{W}_p^\gamma)_{k+1}$  and  $\mathbf{S}_{k+1}$
  - 9: **end for**
- 

## 4.4 Solution through flexible Krylov subspaces

Problem (4.3) reformulated as (4.9) allows us to naturally apply the flexible Golub-Kahan (FGK) and flexible Arnoldi algorithms. Indeed, instead of updating the “preconditioners”  $\mathbf{S}_k$  and  $(\mathbf{W}_p^\gamma)_k$  at the  $k$ th outer iteration of the nested iteration schemes of Algorithm 7, we propose to consider new “preconditioners” as soon as

a new approximation of the solution is available, i.e., at each iteration of a Krylov subspace solver. Therefore, at the  $(i + 1)$ th iteration of the new solvers, the “pre-conditioners”  $(\mathbf{W}_p^\gamma)_i$  and  $\mathbf{S}_i$  are computed as in (4.10), but using the SVD of the  $i$ th approximate solution

$$\mathbf{X}_i = \text{vec}^{-1}(\mathbf{x}_i) = \mathbf{U}_{\mathbf{X}_i} \boldsymbol{\Sigma}_{\mathbf{X}_i} \mathbf{V}_{\mathbf{X}_i}^T, \quad \text{for } i = 1, \dots, k - 1,$$

with  $(\mathbf{W}_p^\gamma)_0 = \mathbf{I}$  and  $\mathbf{S}_0 = \mathbf{I}$ . In order to incorporate iteration-dependent preconditioning, the flexible versions of the Golub-Kahan and Arnoldi factorizations have to be used.

Namely, at the  $i$ th iteration, the new instance of the FGK algorithm updates partial factorizations of the form (2.19), i.e.,  $\mathbf{AZ}_i = \mathbf{U}_{i+1}\mathbf{M}_i$  and  $\mathbf{A}^T\mathbf{U}_{i+1} = \mathbf{V}_{i+1}\mathbf{T}_{i+1}$ , where

$$\mathbf{Z}_i = [\mathbf{S}_0^T(\mathbf{W}_p^\gamma)_0^{-2}\mathbf{S}_0\mathbf{v}_1, \dots, \mathbf{S}_{i-1}^T(\mathbf{W}_p^\gamma)_{i-1}^{-2}\mathbf{S}_{i-1}\mathbf{v}_i], \quad \mathbf{v}_1 = \mathbf{A}^T\mathbf{b}/\|\mathbf{A}^T\mathbf{b}\|_2.$$

Taking  $\mathbf{x}_0 = \mathbf{0}$ , the  $i$ th approximate solution is such that  $\mathbf{x}_i = \mathbf{Z}_i\mathbf{y}_i$ , where

$$\mathbf{y}_i = \arg \min_{\mathbf{y} \in \mathbb{R}^i} \|\mathbf{M}_i\mathbf{y} - \|\mathbf{b}\|_2\mathbf{e}_1\|_2^2 + \hat{\lambda}_i\|\mathbf{y}\|_2^2. \quad (4.19)$$

Note that the subspace for the solution  $\mathcal{R}(\mathbf{Z}_i)$  can be regarded as a generalization of the subspace (4.13) computed when considering preconditioned GKB within the IRN-LSQR-NNR $p$  method. The new method is called “FLSQR-NNR $p$ ”, and is summarized in Algorithm 8.

For  $\mathbf{A} \in \mathbb{R}^{N \times N}$  and  $\mathbf{x}_0 = \mathbf{0}$ , at the  $i$ th iteration, the new instance of the flexible Arnoldi algorithm updates a partial factorization of the form (2.18), with  $k = i$ , and generates

$$\mathbf{Z}_i = [\mathbf{S}_0^T (\mathbf{W}_p^\gamma)_0^{-1} \mathbf{S}_0 \mathbf{v}_1, \dots, \mathbf{S}_{i-1}^T (\mathbf{W}_p^\gamma)_{i-1}^{-1} \mathbf{S}_{i-1} \mathbf{v}_i], \quad \mathbf{v}_1 = \mathbf{b} / \|\mathbf{b}\|_2,$$

where both right and left preconditioners are used analogously to IRN-GMRES-NNR $p$ . The  $i$ th approximate solution is such that  $\mathbf{x}_i = \mathbf{Z}_i \mathbf{y}_i$ , where

$$\mathbf{y}_i = \arg \min_{\mathbf{y} \in \mathbb{R}^i} \|\mathbf{H}_i \mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2^2 + \hat{\lambda}_i \|\mathbf{y}\|_2^2. \quad (4.20)$$

Note that the subspace for the solution  $\mathcal{R}(\mathbf{Z}_i)$  can be regarded as a generalization of the subspace (4.18) computed when considering the preconditioned Arnoldi algorithm within the IRN-GMRES-NNR $p$  method. The new method is called ‘‘FGMRES-NNR $p$ ’’, and is summarized in Algorithm 8.

---

**Algorithm 8** FLSQR-NNR $p$  and FGMRES-NNR $p$

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $(\mathbf{W}_p^\gamma)_0 = \mathbf{I}$ ,  $\mathbf{S}_0 = \mathbf{I}$
  - 2: **for**  $i = 1, 2, \dots$  until a stopping criterion is satisfied **do**
  - 3:   Update a factorization of the form (2.19) and (2.18), respectively, to expand the space  $\mathcal{R}(\mathbf{Z}_i)$
  - 4:   Solve the projected problem (4.19) and (4.20), respectively, tuning  $\hat{\lambda}_i$  if necessary
  - 5:   ‘‘Decrease’’  $\gamma$
  - 6:   Update the new  $(\mathbf{W}_p^\gamma)_i$  and  $\mathbf{S}_i$ , using the SVD  $\mathbf{X}_i = \text{vec}^{-1}(\mathbf{x}_i) = \mathbf{U}_{\mathbf{X}_i} \boldsymbol{\Sigma}_{\mathbf{X}_i} \mathbf{V}_{\mathbf{X}_i}^T$ .
  - 7: **end for**
- 

Note that, although the approach of Algorithm 8 is quite heuristic, it avoids nested iteration cycles and computes only one approximation subspace for the so-

lution of (4.3), where low-rank penalization is adaptively incorporated. Because of this, in many situations, Algorithm 8 computes solutions of quality comparable to the ones computed by Algorithm 7, with a significant reduction in the number of iterations. We should also mention that, in the framework of affine rank minimization problems, [51] outlines an algorithm that avoids inner projected gradient iterations for the solution of each quadratic subproblem in the sequence generated within the IRN strategy.

Finally, we underline that, within the framework of flexible Krylov subspaces, the approximation subspaces  $\mathcal{R}(\mathbf{Z}_i)$  for the  $i$ th approximate solution can be further modified, with some insight into the desired properties of the solution. Indeed, since the  $i$ th basis vector for the solution is of the form

$$\mathbf{z}_i = \mathbf{S}_{i-1}^T (\mathbf{W}_p^\gamma)_{i-1}^{-2} \mathbf{S}_{i-1} \mathbf{v}_i$$

for FLSQR-NNR $p$ , and

$$\mathbf{z}_i = \mathbf{S}_{i-1}^T (\mathbf{W}_p^\gamma)_{i-1}^{-1} \mathbf{S}_{i-1} \mathbf{v}_i$$

for FGMRES-NNR $p$ , one can consider alternative “preconditioners”  $\mathbf{S}_{i-1}$  and  $(\mathbf{W}_p^\gamma)_{i-1}$  that are still effective in delivering low-rank solutions. For instance, focusing on FGMRES, and given  $\mathbf{v}_i = \mathbf{V}_i \mathbf{e}_i$ , where  $\mathbf{V}_i$  is the matrix appearing on the right-hand side of the factorization (2.19), and given the SVD of  $\text{vec}^{-1}(\mathbf{v}_i) = \mathbf{U}_{\mathbf{V}_i} \boldsymbol{\Sigma}_{\mathbf{V}_i} \mathbf{V}_{\mathbf{V}_i}^T$ , one can take

$$\mathbf{S}_{i-1} = \mathbf{V}_{\mathbf{V}_i}^T \otimes \mathbf{U}_{\mathbf{V}_i}^T \quad \text{and} \quad (\mathbf{W}_p^\gamma)_{i-1}^{-1} = \mathbf{I} \otimes (\boldsymbol{\Sigma}_{\mathbf{V}_i})^{1-p/2}, \quad (4.21)$$

and as a result,

$$\begin{aligned} \mathbf{S}_{i-1}\mathbf{v}_i &= \text{vec}(\mathbf{U}_{\mathbf{V}_i}^T \text{vec}^{-1}(\mathbf{v}_i)\mathbf{V}_{\mathbf{V}_i}) = \text{vec}(\boldsymbol{\Sigma}_{\mathbf{V}_i}), \\ (\mathbf{W}_p^\gamma)_{i-1}^{-1}\mathbf{S}_{i-1}\mathbf{v}_i &= \text{vec}((\boldsymbol{\Sigma}_{\mathbf{V}_i})^{1-p/2}\boldsymbol{\Sigma}_{\mathbf{V}_i}) = \text{vec}((\boldsymbol{\Sigma}_{\mathbf{V}_i})^{2-p/2}), \\ \mathbf{S}_{i-1}^T(\mathbf{W}_p^\gamma)_{i-1}^{-1}\mathbf{S}_{i-1}\mathbf{v}_i &= \text{vec}(\mathbf{U}_{\mathbf{V}_i}((\boldsymbol{\Sigma}_{\mathbf{V}_i})^{2-p/2})\mathbf{V}_{\mathbf{V}_i}^T) = \mathbf{z}_i, \end{aligned}$$

so that the singular values of  $\text{vec}^{-1}(\mathbf{v}_i)$  are rescaled: taking  $0 < p \leq 1$ , the power of  $\boldsymbol{\Sigma}_{\mathbf{V}_i}$ ,  $2 - p/2$ , is always larger than 1, which means that large singular values get magnified and small singular values become even smaller. In this way, the gaps between singular values are emphasized and to some extent contribute to the low rank properties of the basis vectors. Similar derivations hold for FLSQR. Hence, methods analogous to LR-FLSQR and LR-FGMRES are obtained, and are called FGMRES-NNR $p(v)$  and FLSQR-NNR $p(v)$ , respectively.

## 4.5 Implementation details

All the methods considered in Section 4.3 and 4.4 are iterative, and therefore at least one suitable stopping criterion should be set for the iterations. When considering hybrid formulations (like the ones in Algorithms 7 and 8), one could simultaneously set a good value for the regularization parameter  $\hat{\lambda}_j$  at the  $j$ th iteration, as well as properly stop the iterations. Strategies for achieving this are already available in the literature (see [17, 21]).

Recall from Section 2.4.2 that, assuming that a good estimate for the norm of the noise  $\boldsymbol{\eta}$  affecting the right-hand-side of (1.1) is available, i.e.,  $\varepsilon \simeq \|\boldsymbol{\eta}\|_2$ , one can

consider the discrepancy principle and stop the iterative scheme at the first iteration  $j$  such that

$$\|\mathbf{b} - \mathbf{A}\mathbf{x}_j\|_2 \leq \mu\varepsilon, \quad \text{where } \mu \gtrsim 1 \text{ is a safety threshold.} \quad (4.22)$$

When running hybrid methods (see Algorithms 7 and 8), we employ the so-called “secant method”, which updates the regularization parameter for the projected problem in such a way that stopping by the discrepancy principle is ensured. We highlight again that the quantities needed to implement the “secant method” (namely, the norm of the residual and the discrepancy associated to (4.9) at each iteration) can be conveniently monitored using projected quantities: this is obvious for IRN-LSQR-NNR $_p$  and FLSQR-NNR $_p$ , as only right-“preconditioning” is employed; it is less obvious for IRN-GMRES-NNR $_p$  and FGMRES-NNR $_p$ , but since the left-“preconditioner” is orthogonal, one can still write

$$\|\mathbf{b} - \mathbf{A}\mathbf{x}_j\|_2 = \|\mathbf{S}_k\mathbf{b} - \mathbf{S}_k\mathbf{A}\mathbf{S}_k^T(\mathbf{W}_p^\gamma)_k^{-1}\hat{\mathbf{x}}_j\|_2 = \|\|\mathbf{b}\|_2\mathbf{e}_1 - \mathbf{H}_j\mathbf{y}_j\|_2.$$

Note that all the methods in Algorithm 7 and 8 can also run with  $\hat{\lambda} = 0$ , and still achieve low-rank approximate solutions: this is because the approximation subspace for the solution incorporates regularizing “preconditioning” (see [28, 32] for details on this approach in the case of smoothing “preconditioning” with finite-difference approximations of derivatives operators). Finally, when dealing with the inner-outer iteration scheme of Algorithm 7, in addition to a parameter choice strategy and stopping criterion for the hybrid projected problems (4.12) and (4.17), one should

also consider a stopping criterion for the outer iterations. We propose to do this by monitoring the norm of the difference of the singular values (normalized by the largest singular value so that  $\sigma_1(\mathbf{\Sigma}_{\mathbf{X}_{k+1}}) = \sigma_1(\mathbf{\Sigma}_{\mathbf{X}_k}) = 1$ ) of two approximations of the solution of (4.3) obtained at two consecutive outer iterations of Algorithm 7, i.e., we stop as soon as

$$\|\text{diag}(\mathbf{\Sigma}_{\mathbf{X}_{k+1}}) - \text{diag}(\mathbf{\Sigma}_{\mathbf{X}_k})\|_2 < \tau_\sigma, \quad k = 1, 2, \dots, \quad (4.23)$$

where  $\text{vec}^{-1}(\mathbf{x}_i) = \mathbf{X}_i = \mathbf{U}_{\mathbf{X}_i} \mathbf{\Sigma}_{\mathbf{X}_i} \mathbf{V}_{\mathbf{X}_i}^T$  ( $i = k, k + 1$ ), and  $\tau_\sigma > 0$  is a user-specified threshold. If no significant changes happen in the rank and singular values of two consecutive approximations of the solution, then (4.23) is satisfied.

We conclude this section with a few remarks about the computational cost of the proposed methods. Note that, if  $\mathbf{A} \in \mathbb{R}^{N \times N}$ , IRN-GMRES-NNR $p$  is intrinsically cheaper than IRN-LSQR-NNR $p$  (since, at each iteration, the former requires only one matrix-vector product with  $\mathbf{A}$ , while the latter requires one matrix-vector product with  $\mathbf{A}$  and one with  $\mathbf{A}^T$ ). However, methods based on the Arnoldi algorithm are typically less successful than methods based on the GKB algorithm for regularization; see [40].

Other key operations for implementing our proposed methods are the computation of the SVDs of relevant quantities, and/or the application of the “preconditioners” in (4.21). Namely, each iteration of FLSQR-NNR $p$ , and FGMRES-NNR $p$  requires the computation of the SVD of an  $n \times n$  matrix, which amounts to  $\mathcal{O}(n^3) = \mathcal{O}(N^{3/2})$  floating point operations. When considering IRN-LSQR-NNR $p$  and IRN-GMRES-NNR $p$ , only the SVD of the approximate solution should be computed once



at each outer iteration. However, each inner iteration of IRN-LSQR-NNR $p$  and IRN-GMRES-NNR $p$ , as well as each iteration of FLSQR-NNR $p$  and FGMRES-NNR $p$ , requires the computation of matrix-vector products of the form  $\mathbf{S}_k^T (\mathbf{W}_p^\gamma)_k^{-1} \mathbf{v}_i$ : this can be achieved within a two-step process, where first the rescaling  $\tilde{\mathbf{v}}_i = (\mathbf{W}_p^\gamma)_k^{-1} \mathbf{v}_i$  is applied with  $\mathcal{O}(N) = \mathcal{O}(n^2)$  floating-point operations, and then  $\mathbf{S}_k^T \tilde{\mathbf{v}}_i = (\mathbf{V}_{\mathbf{X}_k} \otimes \mathbf{U}_{\mathbf{X}_k}) \tilde{\mathbf{v}}_i$  is computed. While a straightforward implementation of the latter would require  $\mathcal{O}(N^2) = \mathcal{O}(n^4)$  floating-point operations, exploiting Kronecker product properties can bring down the cost of this operation to  $\mathcal{O}(n^3) = \mathcal{O}(N^{3/2})$ , by computing  $\mathbf{S}_k^T \tilde{\mathbf{v}}_i = \text{vec}(\mathbf{U}_{\mathbf{X}_k}^T \text{vec}^{-1}(\mathbf{v}_i) \mathbf{V}_{\mathbf{X}_k})$ .

We emphasize that the incorporation of the flexible “preconditioners” does not increase the order of computational complexity and is very practical, since operations are done on matrices of size  $n \times n$  ( $n$  is the dimension of the image). In particular, the full SVD’s of  $n \times n$  matrices can be computed easily with MATLAB’s built-in `svd` function (this is what we used in our numerical experiments); one can also use Lanczos bidiagonalization [47] or randomized SVD [27] to compute the approximate leading singular values and vectors.

## Chapter 5

# Low Rank Regularization: Numerical Experiments

In this Chapter, we present results of numerical experiments on several image processing problems to demonstrate the performance of the new IRN-GMRES-NNR $p$ , IRN-LSQR-NNR $p$ , FGMRES-NNR $p$ , and FLSQR-NNR $p$  methods. Variants of FGMRES-NNR $p$  and FLSQR-NNR $p$  (marked with “(v)”) are also tested. To shorten the acronyms, we omit  $p$  when  $p = 1$ , which means IRN-GMRES-NNR denotes IRN-GMRES-NNR $p$  when  $p = 1$ , etc. Examples are generated using *IR Tools* [17].

In general, we compare the performances of the proposed methods to standard Krylov subspace methods GMRES and LSQR, also used in a hybrid fashion. We also test against the low-rank-projected GMRES (RS-LR-GMRES) method proposed in [48], which leads us to the derivation of alternative low-rank projection solvers LR-FGMRES and LR-FLSQR described in the upcoming section (see Section 5.1). These

solvers can be (re)casted into the framework of flexible Krylov methods and can work with both square and rectangular systems (4.1). The singular value thresholding (SVT) algorithm [7] will also be addressed, which was originally proposed to minimize matrix nuclear norm for low-rank matrix completion problems, and can be extended to problems with linear constraints of the form

$$\min_{\mathbf{x}} \tau \|\text{vec}^{-1}(\mathbf{x})\|_* + \frac{1}{2} \|\text{vec}^{-1}(\mathbf{x})\|_F^2 \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b}, \quad \text{where } \text{vec}^{-1}(\mathbf{x}) = \mathbf{X}. \quad (5.1)$$

The  $k$ th iteration of the SVT algorithm for (5.1) reads

$$\begin{cases} \mathbf{X}_k = \mathcal{D}_\tau(\mathbf{A}^T \mathbf{y}_{k-1}) \\ \mathbf{y}_k = \mathbf{y}_{k-1} + \delta_k (\mathbf{b} - \mathbf{A}\mathbf{x}_k) \end{cases}, \quad (5.2)$$

where  $\delta_k$  is a step size and  $\mathcal{D}_\tau$  is the *singular value shrinkage operator*, defined as

$$\mathcal{D}_\tau(\mathbf{X}) = \mathbf{U}_\mathbf{X} \mathcal{D}_\tau(\boldsymbol{\Sigma}_\mathbf{X}) \mathbf{V}_\mathbf{X}^T, \quad \mathcal{D}_\tau(\boldsymbol{\Sigma}_\mathbf{X}) = \max\{\boldsymbol{\Sigma}_\mathbf{X} - \tau \mathbf{I}, \mathbf{0}\},$$

where  $\mathbf{X} = \mathbf{U}_\mathbf{X} \boldsymbol{\Sigma}_\mathbf{X} \mathbf{V}_\mathbf{X}^T$  is the SVD of  $\mathbf{X}$ ,  $\mathbf{0}$  is a matrix of zeros, and the maximum is taken component-wise. Although (5.1) is not the same problem as (4.2), they are similar in that both penalize the nuclear norm of  $\text{vec}^{-1}(\mathbf{x})$  and they respect the constraint  $\mathbf{A}\mathbf{x} = \mathbf{b}$ .

## 5.1 Low Rank Projection Methods: classical and new approaches

When solving square well-posed linear systems coming from the discretization of some instances of stochastic or time-dependent PDEs, a suitable rearrangement of the solution is expected to be low-rank: for this reason, schemes that incorporate low-rank projections within the basis vectors and the approximate solution obtained by a Krylov method have been proposed in the literature. In the following we summarize the working ideas underlying the so-called restarted low-rank-projected GMRES (RS-LR-GMRES) method proposed in [48].

Recall from Section 2.4.1 that the  $k$ th iteration of GMRES updates the partial Arnoldi factorization and computes the approximate solution as follows:

$$\mathbf{A}\mathbf{V}_k = \mathbf{V}_{k+1}\mathbf{H}_k, \quad \mathbf{x}_k = \mathbf{V}_k\mathbf{y}_k, \quad \text{where } \mathbf{y}_k = \arg \min_{\mathbf{y} \in \mathbb{R}^k} \|\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|_2\mathbf{e}_1\|_2. \quad (5.3)$$

Since matrix-vector products and vector sums of low-rank vectorized matrices increase the rank of the latter, there is no guarantee that the new basis vectors  $\mathbf{v}_k$  for the solution nor the new solution  $\mathbf{x}_k$  are low-rank. To force the basis vector for the solution and the approximate solution to be low-rank, a truncation operator should be incorporated into the GMRES algorithm. Given a vectorized matrix  $\mathbf{c} = \text{vec}(\mathbf{C})$ , and given a desired low-rank  $\kappa$  for  $\mathbf{C}$ , one can define a truncation operator  $\tau_\kappa(\mathbf{c})$  by

the following standard operations:

$$\left[ \begin{array}{l} 1. \text{ Take } \mathbf{C} = \text{vec}^{-1}(\mathbf{c}); \\ 2. \text{ Compute the SVD of } \mathbf{C}, \mathbf{C} = \mathbf{U}_C \mathbf{\Sigma}_C \mathbf{V}_C^T; \\ 3. \text{ Compute } \mathbf{C}_\kappa = \mathbf{U}_C(:, 1 : \kappa) \mathbf{\Sigma}_C(1 : \kappa, 1 : \kappa) \mathbf{V}_C(:, 1 : \kappa)^T; \\ 4. \text{ Take } \tau_\kappa(\mathbf{c}) = \text{vec}(\mathbf{C}_\kappa). \end{array} \right. \quad (5.4)$$

RS-LR-GMRES is a restarted version of the standard GMRES method where the basis vectors for the solution are truncated at each inner iteration, and the solution itself is truncated at the beginning of each outer iteration. Note that truncating the basis vectors does not guarantee that the solution has low rank (which is the reason we still need to truncate the approximate solution). The reason for truncating the basis vectors is to keep the original solution rank from increasing drastically, since it is computed as a linear combination of basis vectors. More precisely, at the  $\ell$ th outer iteration of RS-LR-GMRES, one takes  $\mathbf{v}_1 = \mathbf{r}_{\ell-1} / \|\mathbf{r}_{\ell-1}\|_2$ , where  $\mathbf{r}_{\ell-1} = \mathbf{b} - \mathbf{A}\mathbf{x}_{\ell-1}$ , and, at the  $k$ th inner iteration, one computes

$$\mathbf{v}_k = \tau_\kappa(\mathbf{v}_k), \quad (5.5)$$

and once  $m$  inner iterations are performed, the approximate solution at the  $\ell$ th outer iteration is computed as

$$\mathbf{x}_\ell = \tau_\kappa(\mathbf{x}_{\ell-1} + \mathbf{V}_m \mathbf{y}_m). \quad (5.6)$$

The operations in (5.5) and (5.6) heavily depend on the value  $\kappa$  of the truncated rank, which eventually coincides with the rank of the approximate solution. In the

framework of stochastic PDEs, a suitable estimate for  $\kappa$  can be obtained by first performing coarse-grid computations (see [48] for details, and [45, 67] for similar approaches). As in standard GMRES, RS-LR-GMRES computes a new basis vector for the solution by applying the linear operator  $\mathbf{A}$  to the previous basis vector  $\mathbf{v}_{k-1}$  and orthogonalizing it against the previous basis vectors  $\mathbf{v}_i, i = 1, \dots, k-1$ . However, since the basis vectors are truncated to low rank, the matrix  $\mathbf{V}_k$  does not have orthonormal columns anymore, and  $\mathcal{R}(\mathbf{V}_m)$  is not a Krylov subspace anymore. This remark leads us to the derivation of alternative low-rank projection solvers, which can be (re)casted into the framework of flexible Krylov methods and can work with both square and rectangular systems (4.1).

### 5.1.1 Low-rank flexible GMRES (LR-FGMRES) and low-rank flexible LSQR (LR-FLSQR)

Recall from Section 2.4.3 that starting with  $\mathbf{x}_0 = \mathbf{0}$ , at the  $k$ th iteration, FGMRES updates a partial flexible Arnoldi factorization and computes the  $k$ th approximate solution as follows:

$$\mathbf{AZ}_k = \mathbf{V}_{k+1}\mathbf{H}_k, \quad \mathbf{x}_k = \mathbf{Z}_k\mathbf{y}_k, \quad \text{where } \mathbf{y}_k = \arg \min_{\mathbf{y} \in \mathbb{R}^k} \|\mathbf{H}_k\mathbf{y} - \|\mathbf{b}\|_2 \mathbf{e}_1\|_2, \quad (5.7)$$

where  $\mathbf{V}_{k+1} = [\mathbf{v}_1, \dots, \mathbf{v}_{k+1}] \in \mathbb{R}^{N \times (k+1)}$  has orthonormal columns,  $\mathbf{H}_k \in \mathbb{R}^{(k+1) \times k}$  is upper Hessenberg,  $\mathbf{Z}_k = [\mathbf{P}_1\mathbf{v}_1, \dots, \mathbf{P}_k\mathbf{v}_k] \in \mathbb{R}^{N \times k}$  has columns that span the approximation subspace for the solution and  $\mathbf{P}_i$  is an iteration-dependent preconditioner. In the particular case of low-rank truncation, we set  $\mathbf{P}_i\mathbf{v}_i = \tau_{\kappa_B}(\mathbf{v}_i)$ , i.e., the

truncation operator defined in (5.4), so that  $\text{rank}(\text{vec}^{-1}(\mathbf{Z}_k \mathbf{e}_i)) = \kappa_B$ ,  $i = 1, \dots, k$ . The subscript  $B$  for the truncation rank  $\kappa_B$  suggests that the truncation is done on the original basis vectors  $\mathbf{v}_i$ 's. The resulting algorithm is called ‘‘LR-FGMRES’’, and it is summarized in Algorithm 9. Note that the approximate solution computed as in (2.18) is also truncated to guarantee rank  $\kappa$  (in general, we assume  $\kappa_B \neq \kappa$ ). Also LR-FGMRES is started with  $\mathbf{x}_0 = \mathbf{0}$ , to guarantee that the basis vectors for the solution (rather than a correction thereof) are low-rank.

---

**Algorithm 9** LR-FGMRES

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $\tau_{\kappa_B}$ ,  $\tau_{\kappa}$
  - 2: Take  $\mathbf{v}_1 = \mathbf{b}/\|\mathbf{b}\|_2$
  - 3: **for**  $i = 1, 2, \dots$  until a stopping criterion is satisfied **do**
  - 4:   Compute  $\mathbf{z}_i = \tau_{\kappa_B}(\mathbf{v}_i)$  and  $\mathbf{w} = \mathbf{A}\mathbf{z}_i$
  - 5:   Compute  $h_{ji} = \mathbf{w}^T \mathbf{v}_j$  for  $j = 1, \dots, i$  and set  $\mathbf{w} = \mathbf{w} - \sum_{j=1}^i h_{ji} \mathbf{v}_j$
  - 6:   Compute  $h_{i+1,i} = \|\mathbf{w}\|_2$ , and if  $h_{j+1,j} \neq 0$ , take  $\mathbf{v}_{i+1} = \mathbf{w}/h_{i+1,i}$
  - 7: **end for**
  - 8: Compute  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \|\mathbf{H}_k \mathbf{y} - \mathbf{b}\|_2 \|\mathbf{e}_1\|_2^2$  and take  $\mathbf{x}_k = \tau_{\kappa}(\mathbf{Z}_k \mathbf{y}_k)$
- 

A few remarks are in order. Differently from the  $k$ th iteration in the inner cycle of the RS-LR-GMRES method (5.5), the  $k$ th iteration of LR-FGMRES expands the approximation subspace by modifying (i.e., truncating) the previous orthonormal basis vector for the space  $\mathcal{R}([\mathbf{b}, \mathbf{A}\mathbf{Z}_k])$ . Analogously to RS-LR-GMRES, the basis vectors for the approximate LR-FGMRES solution are all of rank  $\kappa_B$ , are not orthogonal, and do not span a Krylov subspace. Differently from RS-LR-GMRES, the basis vector for the space  $\mathcal{R}([\mathbf{b}, \mathbf{A}\mathbf{Z}_k])$  are orthogonal. Also, the  $k$ th LR-FGMRES approximate solution is obtained by solving an order- $k$  projected least squares problem that is formally analogous to the GMRES one (see (5.3) and (5.7)).

With LR-FGMRES in place, the extension to more general matrices  $\mathbf{A} \in \mathbb{R}^{M \times N}$ , with  $M$  not necessarily equal to  $N$ , can be naturally devised considering the flexible Golub-Kahan (FGK) process [11]. Recall that taking  $\mathbf{x}_0 = \mathbf{0}$  as initial guess, the  $k$ th FGK iteration updates partial factorizations of the form

$$\mathbf{AZ}_k = \mathbf{U}_{k+1}\mathbf{M}_k \quad \text{and} \quad \mathbf{A}^T\mathbf{U}_{k+1} = \mathbf{V}_{k+1}\mathbf{T}_{k+1}, \quad (5.8)$$

where the columns of  $\mathbf{U}_{k+1} \in \mathbb{R}^{M \times (k+1)}$ ,  $\mathbf{V}_{k+1} \in \mathbb{R}^{N \times (k+1)}$  are orthonormal,  $\mathbf{M}_k \in \mathbb{R}^{(k+1) \times k}$  is upper Hessenberg,  $\mathbf{T}_{k+1} \in \mathbb{R}^{(k+1) \times (k+1)}$  is upper triangular,  $\mathbf{Z}_k = [\mathbf{P}_1\mathbf{v}_1, \dots, \mathbf{P}_k\mathbf{v}_k] \in \mathbb{R}^{N \times k}$  has columns that span the approximation subspace for the solution, and  $\mathbf{P}_i$  is an iteration-dependent preconditioner that is applied to  $\mathbf{v}_i$ . In the particular case of low-rank truncation,  $\mathbf{P}_i\mathbf{v}_i = \tau_{\kappa_B}(\mathbf{v}_i)$ , as defined in (5.4), so that  $\text{rank}(\text{vec}^{-1}(\mathbf{Z}_k\mathbf{e}_i)) = \kappa_B$ ,  $i = 1, \dots, k$ . The flexible LSQR method (FLSQR) uses the FGK process (5.8) to generate iterates of the form  $\mathbf{x}_k = \mathbf{Z}_k\mathbf{y}_k$ , where the vector  $\mathbf{y}_k$  is computed as  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \left\| \mathbf{M}_k\mathbf{y} - \|\mathbf{b}\|_2\mathbf{e}_1 \right\|_2^2$ . When rank-truncation of the basis vectors takes place at each iteration, and the final approximate solution is rank-truncated as well, the resulting algorithm is called ‘‘LR-FLSQR’’, and it is summarized in Algorithm 10.

Note that, similarly to RS-LR-GMRES, both LR-FGMRES and LR-FLSQR are quite heuristic. Although the low-rank projection idea can be formulated in the flexible framework, we lack a formal formulation of the problem that is being solved, and also a justification of why they work. Strategies for selecting  $\kappa_B$  and  $\kappa$  are not so clear either. To stabilize the behavior of LR-FGMRES as the iterations proceed, one may consider imposing additional Tikhonov regularization on the projected least-



---

**Algorithm 10** LR-FLSQR
 

---

- 1: Inputs:  $\mathbf{A}$ ,  $\mathbf{b}$ ,  $\tau_{\kappa_B}$ ,  $\tau_{\kappa}$
  - 2: Take  $\mathbf{u}_1 = \mathbf{b}/\|\mathbf{b}\|_2$
  - 3: **for**  $i = 1, 2, \dots$ , until a stopping criterion is satisfied **do**
  - 4:   Compute  $\mathbf{w} = \mathbf{A}^T \mathbf{u}_i$ ,  $t_{ji} = \mathbf{w}^T \mathbf{v}_j$  for  $j = 1, \dots, i-1$
  - 5:   Set  $\mathbf{w} = \mathbf{w} - \sum_{j=1}^{i-1} t_{ji} \mathbf{v}_j$ , compute  $t_{ii} = \|\mathbf{w}\|$  and take  $\mathbf{v}_i = \mathbf{w}/t_{ii}$
  - 6:   Compute  $\mathbf{z}_i = \tau_{\kappa_B}(\mathbf{v}_i)$  and  $\mathbf{w} = \mathbf{A} \mathbf{z}_i$
  - 7:   Compute  $m_{ji} = \mathbf{w}^T \mathbf{u}_j$  for  $j = 1, \dots, i$  and set  $\mathbf{w} = \mathbf{w} - \sum_{j=1}^i m_{ji} \mathbf{u}_j$
  - 8:   Compute  $m_{i+1,i} = \|\mathbf{w}\|$  and take  $\mathbf{u}_{i+1} = \mathbf{w}/m_{i+1,i}$
  - 9: **end for**
  - 10: Compute  $\mathbf{y}_k = \arg \min_{\mathbf{y}} \|\mathbf{M}_k \mathbf{y} - \mathbf{b}\|_2 \|\mathbf{e}_1\|_2^2$  and take  $\mathbf{x}_k = \tau_{\kappa}(\mathbf{Z}_k \mathbf{y}_k)$
- 

squares problem in (2.18), in a hybrid fashion; the same holds for LR-FLSQR (see Sections 4.4 and 5.2 for more details).

## 5.2 Numerical Examples

In this section, three types of imaging inverse problems are presented, and a variety of solvers, including the newly developed low rank methods, are compared. Before getting started, we briefly remark on some details for some of the parameters in these methods.

The Schatten- $p$  function is introduced in Section 4.2 as a smooth approximation for  $\|\cdot\|_{*,p}$ . The smooth approximation allows for further derivations including computation of optimality conditions, where the “smoothing coefficient”  $\gamma$  is crucial. However,  $\gamma$  is not so crucial numerically, and we can set it to 0 without affecting the results (compared to using a very small  $\gamma$ ). However, to be consistent with Algorithms 7 and 8, in our experiments, we have set the initial value of  $\gamma$  to  $10^{-10}$ , and every time we need to decrease  $\gamma$ , we divide the current  $\gamma$  value by 2.

Regarding the comparisons with the low-rank projection methods presented in Section 5.1.1, there are no universal and theoretically informed ways of choosing the truncation ranks for the solutions and for the basis vectors of the solution subspace. Hence, for all test problems, we experiment on a reasonable number of trials, each with different truncation rank choices, and select the best performing rank out of all ranks tested. For simplicity, we consider the same truncation rank for basis vectors and solutions ( $\tau_{\kappa_B} = \tau_{\kappa}$ ). We follow the same process to choose the number of restarts and the number of iterations for each restart for RS-LR-GMRES, as well as the shrinkage threshold  $\tau$  in SVT; strategies to select the step size for SVT are described in [7].

**Example 1: Binary Star.** We consider an image deblurring problem involving a simulated binary star test image of size  $256 \times 256$ : this test image has rank 2. The true image is displayed in the leftmost frame of Figure 5.2. A standard Gaussian blur is applied to the test image, and Gaussian white noise of level  $\|\boldsymbol{\eta}\|_2/\|\mathbf{b}^{\text{ex}}\|_2 = 10^{-3}$  is added. The blurred and noisy images are shown in Figure 5.2, second frame from the left. Due to the presence of noise, the blurred image has full rank. For this example, the blurring operator  $\mathbf{A}$  is square of size  $65536 \times 65536$ , hence GMRES-related methods are used for comparison, namely: GMRES, IRN-GMRES-NNR, FGMRES-NNR, LR-FGMRES and RS-LR-GMRES (i.e., we only consider the case  $p = 1$  here). SVT is also taken into consideration. The truncation rank for LR-FGMRES and RS-LR-GMRES is set to 30 for both basis vectors and approximate solutions (i.e.,  $\tau_{\kappa_B} = \tau_{\kappa} = 30$ ). RS-LR-GMRES is restarted every 40 iterations. The step size for SVT is set to be  $\delta_k = \delta = 2$  and the singular value shrinkage threshold  $\tau$

is 1. Note that, although the true solution has only rank 2, setting truncation rank to 2 for low rank methods produces solutions of worse quality (compared to setting the rank to 30). This might be because of the inherent ill-posedness of the problem, which makes it harder to obtain solutions with desired properties (e.g., with rank 2): indeed, if we do truncate to rank 2, a lot of information about the solution might be lost.

Figure 5.1 displays the histories of relative errors  $\|\mathbf{x}^{\text{ex}} - \mathbf{x}_m\|_2 / \|\mathbf{x}^{\text{ex}}\|_2$  for the first 200 iterations (i.e.,  $m = 1, \dots, 200$ ) of these methods. For IRN-GMRES-NNR, 4 outer cycles were run, each with a maximum of 50 iterations: a new outer cycle is initiated as soon as the discrepancy principle is satisfied in the inner cycle. No additional regularization is used (i.e.,  $\hat{\lambda} = 0$  for all methods).




Figure 5.1: *Example 1*. Relative errors vs. number of iterations for GMRES-based methods and SVT.

We can observe from Figure 5.1 that when the truncation ranks are chosen reasonably, LR-FGMRES and RS-LR-GMRES both produce a less pronounced semi-convergence behavior than GMRES, with LR-FGMRES attaining a smaller relative

error than RS-LR-GMRES. FGMRES-NNR, on the other hand, shows slower semi-convergence than GMRES, but it also converges to a slightly better relative error. IRN-GMRES-NNR behaves especially well in this case, with significantly reduced relative errors even at the end of the second outer cycle. The “jumps” at the beginning of each outer IRN-GMRES-NNR iteration are due to the strategy used for restarts (the older basis vectors are cleared at each restart).

Figure 5.2 displays the exact and the corrupted images, as well as the best reconstructions computed by LR-FGMRES and IRN-GMRES-NNR: these are obtained at the 47th and the 189th (total) iteration of LR-FGMRES and IRN-GMRES-NNR, respectively. By looking at relative errors in Figure 5.1, we see that LR-FGMRES is the second best out of all methods, and yet the quality of the solution is inferior compared to IRN-GMRES-NNR. Compared to the LR-FGMRES solution, the IRN-GMRES-NNR one is a more truthful reconstruction of the exact image: it not only has less artifacts immediately around the stars, but also has less background noise, in the sense that the pixel intensities in the background are closer to the true ones (as it can be seen by looking at the background color).




Figure 5.2: *Example 1*. Exact and corrupted test images, together with the best reconstructions obtained by the LR-FGMRES and the IRN-GMRES-NNR methods.

More details can be spotted if we zoom into the central part ( $51 \times 51$  pixels) of the computed images, as shown in Figure 5.3: here the best LR-FGMRES reconstruction, as well as the IRN-GMRES-NNR reconstructions at the end of the 2nd, 3rd and 4th inner cycles are displayed. It is clear that the IRN reconstructions are improving over each outer cycle, and that even the solution at the end of the 2nd cycle is significantly better than the LR-FGMRES solution, which means that not all four outer iterations need to be run to achieve solutions of superior qualities (even if more outer iterations allow further improvement in the solution).




Figure 5.3: *Example 1.* Zoom-ins of the LR-FGMRES best solution, and the IRN-GMRES-NNR solutions at the end of each inner cycle.

Figure 5.4 displays surface plots of the central part ( $51 \times 51$  pixels) of the test problem data, as well as the best reconstructed images (for RS-LR-GMRES and FGMRES-NNR these are obtained at the 165th and the 63th (total) iterations, respectively). It can be seen that for all the solutions shown here, the reconstructed central two stars approximately have the same intensity, although they are somewhat less intense than in the exact image. These surface plots also confirm our earlier observation that IRN-GMRES-NNR does an exceptional job removing background

noise. In addition, FGMRES-NNR also gives a good background reconstruction.

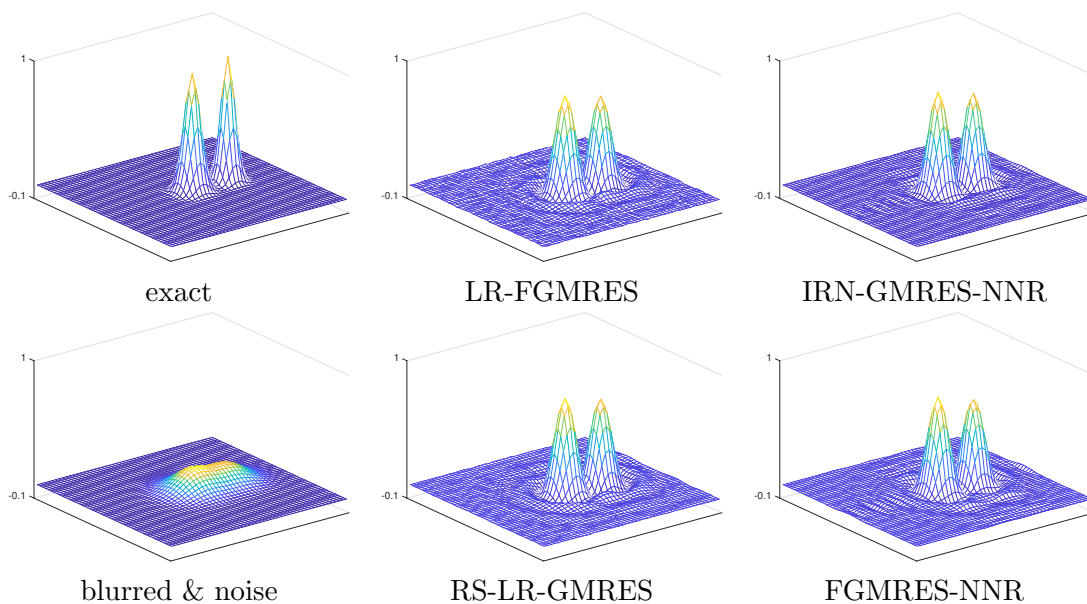


Figure 5.4: *Example 1*. Zoomed-in surfaces of the exact solution and the available data, as well as the best reconstructions obtained by the new GMRES-based methods.

Finally, Figure 5.5 displays the singular values of the best solutions obtained adopting different GMRES-based solvers, as well as the evolution of the singular values of the solution at the end of each inner IRN-GMRES-NNR cycle (matching the reconstructions displayed in Figure 5.3). The singular values are “normalized” (i.e., divided by the largest one), and the graphs are cropped to focus on the relevant values. Looking at the displayed values, we can conclude that the solutions computed by all the low-rank solvers have indeed some low-rank properties, with very quickly-decaying large singular values followed by slowly-decaying smaller singular values. Compared to GMRES, the new FGMRES-NNR and IRN-GMRES-NNR methods give solutions that have a more pronounced low rank, as shown by the large gaps

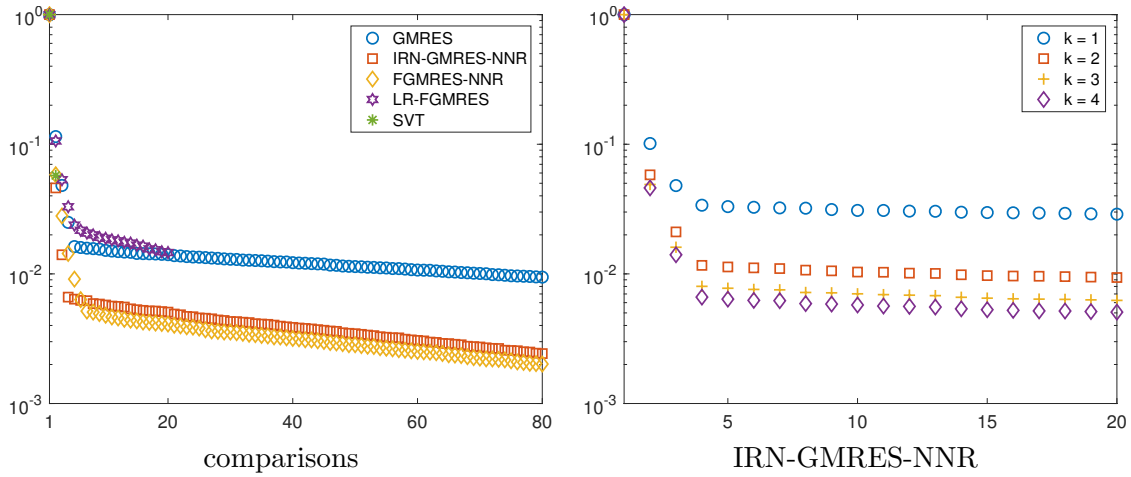


Figure 5.5: *Example 1*. Left frame: normalized singular values of the best solutions computed by each GMRES-based method. Right frame: evolution of the singular values of the solutions computed by IRN-GMRES-NNR at each outer iteration. Singular values less than  $10^{-3}$  are omitted.

between the smaller singular values of the solutions computed by these methods. Regarding IRN-GMRES-NNR, the evolution of the singular values stabilizes as we move toward later outer iterations, which validates the stopping criterion proposed in Section 4.5.

**Example 2: Limited Angle Parallel-Ray Tomography.** We consider a computed tomography (CT) test problem, modeling an undersampled X-ray scan with parallel beam geometry. This is a so called “limited angle” CT reconstruction problem, where the viewing angles for the object span less than 180 degrees. A smooth and rank-4 phantom is considered, as shown in the leftmost frame of Figure 5.7 (note that the yellow straight lines in the northwestern corner do not belong to the phantom; they are shown for later purposes). Gaussian white noise of level  $10^{-2}$  is

added to the data. The coefficient matrix  $\mathbf{A}$  has size  $32942 \times 65536$ . Because of this, among the new solvers, only LR-FLSQR, FLSQR-NNR $p$ , FLSQR-NNR $p(v)$ , and IRN-LSQR-NNR $p$  will be tested, against their standard counterpart LSQR. Recall that FLSQR-NNR $p(v)$  is the FLSQR-NNR $p$  variant that defines the preconditioners using the basis vectors of the solution subspace. The hybrid strategy is not used here, meaning that we set  $\hat{\lambda} = 0$  for all methods. For this test problem, we consider both the values  $p = 1$  and  $p = 0.75$  (recall that, when  $p = 1$ , we omit  $p$  from the notation). The results obtained running the available low-rank solvers SVT and RS-LR-GMRES are shown, too. Note that RS-LR-GMRES only works for square matrices  $\mathbf{A}$ , hence this solver is tested on the normal equations  $\mathbf{A}^T \mathbf{A} \mathbf{x} = \mathbf{A}^T \mathbf{b}$ , which is not the problem solved by the other methods (therefore this comparison may not be completely fair). Parameters for SVT are chosen to be: step size  $\delta_k = \delta = 8 \times 10^{-5}$  and threshold  $\tau = 100$ . RS-LR-GMRES is set to restart every 20 iterations. The truncation rank is 10 for both basis vectors and solutions, and for both the LR-FLSQR and the RS-LR-GMRES methods. The maximum number of iterations is 100 for all methods.

Figure 5.6 displays the history of the relative errors for LSQR, LR-LSQR, FLSQR-NNR $p$ , FLSQR-NNR $p(v)$ , and IRN-LSQR-NNR $p$ , for  $p = 1$  and  $p = 0.75$ . Figure 5.7 displays the exact phantom together with the best reconstructions obtained by LSQR, FLSQR-NNR $p(v)$ , and IRN-LSQR-NNR. Figure 5.8 displays surface plots of the northwestern corner of the exact and reconstructed phantoms ( $64 \times 64$  pixels, as highlighted in the leftmost frame of Figure 5.7).

Looking at relative errors in Figure 5.6, it is obvious that the winners are the



FLSQR-NNR $p(v)$  methods, with both  $p = 1$  and  $p = 0.75$ : they give the lowest relative errors, and the fastest semi-convergences. For this test problem, using a value of  $p < 1$  lowers the relative error of FLSQR-NNR $p(v)$ ; however, the same does not hold for IRN-LSQR-NNR $p$ . Therefore we can conclude that the choice of  $p$  is problem and solver dependent, and using  $p < 1$  does not necessarily improve the quality of the solution. We regard  $p = 1$  as a safe choice for this parameter. Although both the FLSQR-NNR $p(v)$  methods with  $p = 1$  and  $p = 0.75$  perform well, the latter is able to further reduce the noise in the reconstructed solution, especially on the boundary.




Figure 5.6: *Example 2*. Relative errors vs. number of iterations for different solvers. Upper frame: some of the new solvers are compared to the already available solvers. Lower frame: comparisons of different instances of the new solvers (here  $p = 0.75$ ).




Figure 5.7: *Example 2.* Exact phantom and best reconstructions obtained by different solvers.

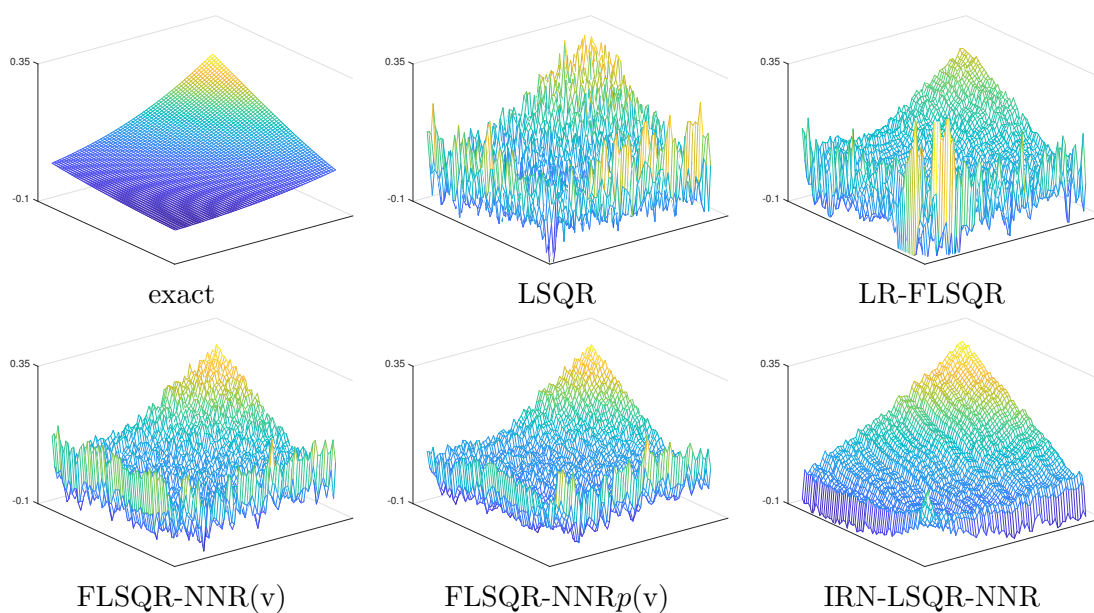


Figure 5.8: *Example 2.* Surface plots of the northwestern corner of the exact phantom (highlighted in Figure 5.7) and the best reconstructed phantoms computed by different solvers.

Looking at all the displayed results, the advantages of our new FLSQR-NNR $p(v)$  and IRN-LSQR-NNR methods are evident. Namely, they produce smooth solutions

that preserve the original concave shape of the exact phantom, and they retain similar intensities of pixels at the same locations of the exact phantom (although the LR-FLSQR solution is smooth within the boundary, it fails to reconstruct intensity at the high point). Differences between FLSQR-NNR $p(v)$  and IRN-LSQR-NNR reconstructions are clear, too: while both are smooth, the IRN-LSQR-NNR reconstruction has a less concave shape compared to that of FLSQR-NNR $p(v)$ , but a smoother boundary.

**Example 3: Inpainting.** We consider two different inpainting test problems. Inpainting is the process of restoring images that have missing or deteriorated parts. These images are likely to have quite a few lost pixels, either in the form of salt and pepper noise, or missing patches with regular or irregular shapes. The two examples considered here are of different nature: the first one has less structured and more randomly distributed missing patches, while the second one has more structured and regularly shaped missing parts. The corrupted images (shown in top-middle frames of Figures 5.11 and 5.13) are constructed by first applying a blur operator, and then superimposing the undersampling pattern to the ideally exact images (shown in the top-left frames of Figures 5.11 and 5.13). We follow this particular order of first blurring and then taking out pixels to simulate the real process of photo-taking. For both these test problems, white noise of level  $10^{-2}$  is added to the data, and we consider purely iterative methods (i.e.,  $\hat{\lambda} = 0$ ). We always take  $p = 1$ , and we run 100 iterations of all the methods.

Firstly, we consider a test problem where 58.2% of the pixels are missing (following some random and not very regular patterns). The exact image is commonly known as

the **house** test image, whose rank is 243 and has a total number of 65536 ( $256 \times 256$ ) pixels; the corrupted image has the same size and number of pixels, but out of which only 27395 are non-zero. A plot of the singular values of the exact image is shown in Figure 5.9(a). Correspondingly, the forward operator  $\mathbf{A}$  is of size  $27395 \times 65536$ , so we have an underdetermined linear system:  $\mathbf{A}$  is obtained by first applying a shaking blur, and by then undersampling the blurred image. This can be easily coded within the *IR Tools* framework.




Figure 5.9: *Example 3*. Singular values of exact test images **house** and **peppers** scaled by the largest singular values respectively.

Figure 5.10 displays the history of the relative errors for LSQR, LR-FLSQR (with truncation of the basis vectors for the solution, as well as the solution, to rank 20), FLSQR-NNR, FLSQR-NNR(v) and IRN-LSQR-NNR. Figure 5.11 displays the exact and corrupted images, together with the best reconstructions obtained by the methods listed above: these correspond to the 16th, 32nd, 67th, 30th and 62nd iterations of LSQR, LR-FLSQR, FLSQR-NNR, FLSQR-NNR(v) and IRN-LSQR-

NNR, respectively (i.e., these are the iterations where the minimum relative error is attained over the total 100 iterations).




Figure 5.10: *Example 3 (house)*. Relative errors vs. number of iterations for different solvers.




Figure 5.11: *Example 3 (house)*. Exact and corrupted images; best reconstructions obtained by standard and new solvers.

Secondly, we consider a test problem similar to the previous one, i.e., we take an exact image commonly known as the **peppers** test image, which has full rank (its singular values are shown in Figure 5.9(b)), and we obtain the forward operator  $\mathbf{A}$  by first applying a shaking blur, and by then undersampling the blurred image. Here the exact image has a total number of 65536 ( $256 \times 256$ ) pixels, and only around 1.3% of pixels are missing and should be inpainted: differently from the previous problem, the missing pixels follow particular patterns (e.g., circles, squares, and rectangles), and this makes the inpainting task somewhat more challenging. Figure 5.12 displays the history of the relative errors for LSQR, LR-FLSQR (with truncation of the basis vectors for the solution, to rank 50), FLSQR-NNR, FLSQR-NNR(v) and IRN-LSQR-NNR. Figure 5.13 displays the exact and corrupted images, together with the best reconstructions obtained by the methods listed above: these correspond to the 11th, 18th, 60th, 33rd and 34th iterations of LSQR, LR-FLSQR, FLSQR-NNR, FLSQR-NNR(v) and IRN-LSQR-NNR, respectively (i.e., these are the iterations where the minimum relative error is attained over the total 100 iterations).




Figure 5.12: *Example 3* (**peppers**). Relative errors vs. number of iterations for different solvers.




Figure 5.13: *Example 3 (peppers)*. Exact and corrupted images; best reconstructions obtained by standard and new solvers.

It is evident that FLSQR-NNR(v) achieves reconstructions of superior quality, including clarity, brightness, and smoothness. Its ability to fill-in missing spots with pixels that are of similar intensity to their surroundings is the best among all methods. The best reconstructions are computed by IRN-LSQR-NNR for the **house** test image, and by FLSQR-NNR for the **pepper** test image: in both cases, these methods are also good at removing noise and restoring missing pixels. However, for both test images, the reconstructions obtained by IRN-LSQR-NNR lack clarity compared to ones obtained by both FLSQR-NNR and FLSQR-NNR(v) methods; compared to the reconstructions obtained by LSQR and LR-FLSQR, they are anyway more desirable in terms of recovered brightness and fill-in of the missing pixels. Moreover, we have seen in these two examples that our newly proposed methods perform very well not

only for low rank, but also for full or nearly full rank image reconstruction, thanks to the regularizing properties of our newly derived “preconditioners”  $(\mathbf{W}_p^\gamma)_k$  and  $\mathbf{S}_k$ . Our methods can also be extensively tested for higher noise levels (for example,  $10^{-1}$ ) and yield similar results. However, for space considerations we are not able to show all of them here.

### 5.2.1 A Note on Regularization Parameters

In the previous examples we have seen that the IRN-NNR methods and the flexible Krylov NNR methods perform exceptionally well on image deblurring, tomography, and inpainting problems, producing superior reconstructions compared to existing methods including SVT, RS-LR-GMRES and the low-rank flexible Krylov methods inspired by RS-LR-GMRES, even without the use of additional regularization. In this section, we explore the effect of additional regularization (i.e., we set  $\hat{\lambda} \neq 0$ ) on the reconstructed images and the corresponding relative errors. In particular, additional regularization allows the new methods to be used in a hybrid fashion. We are going to observe that there is only little to negligible room for the methods to improve when they are used in a hybrid fashion (as their performance is already very good with  $\hat{\lambda} = 0$ ).

We consider three different ways of choosing the regularization parameter  $\hat{\lambda}$ . (i) We take the “secant method” mentioned in Section 4.5, which updates the regularization parameter at each iteration using the discrepancy. (ii) We select the optimal regularization parameter which minimizes the 2-norm of the difference between the exact solution and the regularized solution at each iteration. Namely, when using



standard GMRES and LSQR, at the  $m$ th iteration we seek to minimize with respect to  $\hat{\lambda}$

$$\|\mathbf{x}^{\text{ex}} - \mathbf{x}_{m,\hat{\lambda}}\| = \|\mathbf{V}_m^T \mathbf{x}^{\text{ex}} - \mathbf{V}_m^T \mathbf{x}_{m,\hat{\lambda}}\| = \|\mathbf{V}_m^T \mathbf{x}^{\text{ex}} - \mathbf{y}_{m,\hat{\lambda}}\|;$$

when using the IRN methods we should incorporate the appropriate preconditioners  $(\mathbf{W}_p^\gamma)_k$  and  $\mathbf{S}_k$  and, for all the iterations in the inner iteration cycle corresponding to the  $k$ th outer iteration, we seek to minimize with respect to  $\hat{\lambda}$

$$\|\hat{\mathbf{x}}^{\text{ex}} - \mathbf{V}_m \mathbf{y}_{m,\hat{\lambda}}\| = \|\mathbf{V}_m^T \hat{\mathbf{x}}^{\text{ex}} - \mathbf{V}_m^T \mathbf{V}_m \mathbf{y}_{m,\hat{\lambda}}\| = \|\mathbf{V}_m^T \hat{\mathbf{x}}^{\text{ex}} - \mathbf{y}_{m,\hat{\lambda}}\|, \text{ where } \hat{\mathbf{x}}^{\text{ex}} = (\mathbf{W}_p^\gamma)_k \mathbf{S}_k \mathbf{x}^{\text{ex}}.$$

It is intrinsically difficult to implement this strategy for flexible Krylov subspace methods, because of the complexity of changing preconditioners at each iteration.

(iii) We perform a manual exhaustive search. Namely, we first run the solvers multiple times using various regularization parameters  $\hat{\lambda}$ , starting with a larger range and narrowing down to a smaller range containing the best parameter; we then record the minimum relative errors among all iterations for all values of  $\hat{\lambda}$ , and select the corresponding  $\hat{\lambda}$ . This approach is the most expensive one, and differs from the previous one in that the (optimal) regularization parameter  $\hat{\lambda}$  is fixed for all iterations. Of course, both the second and third approaches require the knowledge of the exact solution and we test them only to investigate the best possible performance of the hybrid approach.

Table 5.1 compares the performances (in terms of minimum relative error achieved by each method) of standard Krylov methods (GMRES and LSQR) and their IRN-NNR and flexible NNR (F-NNR) counterparts, with and without using a hybrid

approach. In this way we can understand how the use of additional regularization affects each solver differently. The three parameter choice methods described above are called “Secant (i)”, “Optimal (ii)” and “Fixed (iii)”, respectively. All the previous examples are considered here. GMRES and its counterparts IRN-GMRES-NNR, FGMRES-NNR are used for Example 1, while LSQR and its counterparts IRN-LSQR-NNR, FLSQR-NNR(v) are used for Examples 2 and 3.

		$\hat{\lambda} = 0$	$\hat{\lambda} \neq 0$	$\hat{\lambda} = 0$	$\hat{\lambda} \neq 0$	$\hat{\lambda} = 0$	$\hat{\lambda} \neq 0$	$\hat{\lambda} = 0$	$\hat{\lambda} \neq 0$
		<i>Example 1</i>		<i>Example 2</i>		<i>Ex. 3 (house)</i>		<i>Ex. 3 (peppers)</i>	
Standard	(i)	0.2995	0.2528	0.1201	0.1389	0.2712	0.2715	0.1141	0.1138
	(ii)	0.2995	0.2268	0.1201	0.1201	0.2712	0.2710	0.1141	0.1138
	(iii)	0.2995	0.2268	0.1201	0.1183	0.2712	0.2710	0.1141	0.1138
IRN-NNR	(i)	0.2081	0.2096	0.0685	0.0696	0.1249	0.1250	0.0964	0.0967
	(ii)	0.2081	0.2292	0.0685	0.0685	0.1249	0.1249	0.0964	0.0964
	(iii)	0.2081	X	0.0685	0.0660	0.1249	X	0.0964	0.0960
F-NNR	(i)	0.2829	0.2658	0.0577	0.0684	0.1035	0.1046	0.0625	0.0618
	(iii)	0.2829	0.2640	0.0577	0.0568	0.1035	X	0.0625	0.0618

Table 5.1: Minimum relative errors without ( $\hat{\lambda} = 0$ ) and with ( $\hat{\lambda} \neq 0$ ) a hybrid approach using various regularization parameter setting techniques: (i) secant; (ii) optimal; (iii) fixed. The mark “X” means that the optimal regularization parameter found by the “Fixed (iii)” method is  $10^{-16}$ , hence there is no need for additional regularization.

It is easy to observe that the use of additional regularization is most effective for the standard GMRES solver, where the minimum relative error is reduced significantly. However, for the other solvers, the hybrid approach does not have a notable advantage over not using regularization. At times the “Fixed (iii)” parameter choice strategy delivers a regularization parameter of the order of  $10^{-16}$ , which is numerically equivalent to not having regularization. This indicates that our new IRN-NNR and F-NNR methods are successful in computing good reconstructions and, even

without additional regularization, they perform much better than standard Krylov methods used in a hybrid fashion (comparing IRN-GMRES-NNR to GMRES in Example 1, and FLSQR-NNR(v) to LSQR in the other examples). Figure 5.14 shows a couple of such comparisons.




Figure 5.14: Reconstructions obtained by standard hybrid Krylov methods and by the new methods without using additional regularization. Left side: zoomed in surface plots of the reconstructions of *Example 1*; right side: reconstructions of *Example 3* (peppers).

## Chapter 6

# Extended Application: Model Calibration for Computed Tomography

In Chapter 5, we explored the effect of Krylov subspace methods, including the newly developed low-rank, nuclear norm based methods (as described in Chapter 4) on a variety of image reconstruction problems: image deblurring, computed tomography, as well as inpainting. All these applications share the property that the forward operator  $\mathbf{A}$  is fixed, and that we are performing minimization with respect to one unknown:  $\mathbf{x}$ .

However, this is not the case for all imaging applications. For example, the advancement of x-ray science and optics technologies has enabled improved spatial resolution and increased sample clarity, but meanwhile, reconstructions of the im-

aged object have become more susceptible to random errors and systematic errors such as drifts in scanning probes, leading to distorted images with decreased resolution, or even misrepresentation of structure. As another example, in the diagnosis and management of infectious diseases such as Covid-19, where portable tomosynthesis machines are brought to the location of patients to help minimize the risk of cross-infection [9, 38], more experimental errors such as inaccurate calibration of geometry parameters (e.g., source to object distance, source orientation) are introduced, making the image acquisition and reconstruction processes more challenging. Due to these uncertainties, forward models that describe the forward process only approximate the reality to some extent, and without proper model calibration, quality of the solutions will be degraded. Using mathematical language – when working on the corresponding inverse problem, we need to take into account an additional set of unknown parameters  $\mathbf{p}$  that the forward operator  $\mathbf{A}$  is dependent on. So instead of (1.3), we need to solve the following problem:

$$\min_{\mathbf{x}, \mathbf{p}} \|\mathbf{A}(\mathbf{p})\mathbf{x} - \mathbf{b}\|_2^2 + \lambda\mathcal{R}(\mathbf{x}). \quad (6.1)$$

A lot of work has been done in various applications, such as scanning position error correction for image reconstruction in the fields of electron tomography [10, 26] and ptychography [3, 37], addressing the importance of accurate system modeling. In the field of X-ray tomography, scientists at Argonne National Laboratory have developed algorithms to calibrate the center of rotation errors [2] using numerical optimization, and drifts in scanning positions [36] using targeted calibration models. In the application where uncertain view angles need to be estimated for computed

tomography in addition to image reconstruction, approaches such as uncertainty quantification have been proposed, and they seek to quantify scanning angles via a model-discrepancy formulation [58, 59]. Though very different from each other in nature, many of these methods share a similar framework called the block coordinate descent (BCD), which alternatively optimizes for the unknown parameters and the image of interest. In this Chapter, we investigate two applications related to computed tomography (CT) reconstruction: parallel and fan-beamed, where the formulation (6.1) is appropriate since model calibration is required, and BCD type methods will be investigated.

## 6.1 Parallel Tomography

We consider the identical parallel-beam x-ray tomography problem as put forward in [36] by Huang, Di and Wild. Scanning positions that shoot parallel x-ray beams through the object are assumed to be equally-spaced with a fixed step size  $\Delta$ . These theoretical scanning positions and theoretical beams are marked by green dots and green dashed lines in Figure 6.1. Mechanical errors in the equipment can result in errors (called drifts) in the scanning positions. As illustrated in Figure 6.1, the true, drifted scanning positions are marked by purple dots, true beams by purple dashed lines, and the amount of drift for the beam  $\tau$  is marked  $\Delta_\tau$ . A range  $[-\Delta_{\max}, \Delta_{\max}]$  exists for possible values of  $\Delta_\tau$ . The existence of these unknown drifts in the scanning positions can cause the image reconstruction quality to be deteriorated, and the larger the amount of drift, the worse the quality. This is because these unknown

drifts cause the true forward operator to be different from the theoretical forward model that is constructed without considering drifts. Aside from drift in scanning positions, additive white Gaussian noise is also present in the sinogram.

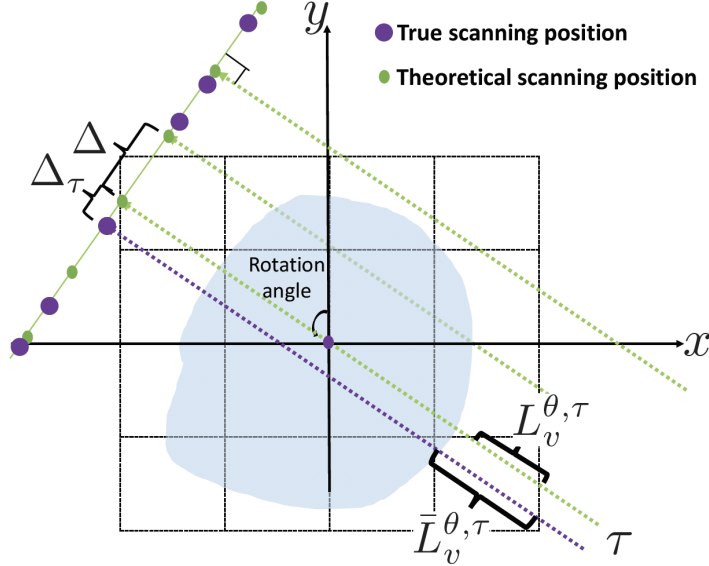


Figure 6.1: Parallel tomography illustration from [36].

The object is an image of size  $n \times n$ , which is vectorized and denoted by  $\mathbf{x} \in \mathbb{R}^N$ . The true forward operator  $\mathbf{A}$  maps the object to the true sinogram  $\mathbf{b}$  (vectorized from  $\mathbf{B}$ ). Note that  $\mathbf{A}$  and  $\mathbf{b}$  are “true” quantities meaning that they are obtained from the experiments and are contaminated by both drift and noise. We denote the theoretical forward operator by  $\mathbf{A}_0$  (the one that we know). The object is rotated  $N_\theta$  times for scanning, each with  $N_\tau$  parallel beams. Hence the sinogram is an image of size  $N_\tau \times N_\theta$ , and it is then vectorized into  $\mathbf{b}$ . Additive Gaussian noise in  $\mathbf{b}$  is denoted by  $\boldsymbol{\eta}$  so that  $\mathbf{b} = \mathbf{b}^{\text{ex}} + \boldsymbol{\eta}$ . For simplicity, it is assumed in [36] that after each rotation of the object, the drifts in the scanning positions remain the same, so the

drifts in beams don't depend on the rotation angle, and there are only  $N_\tau$  unknown drifts. We denote the vector of drifts  $\Delta_\tau$  as  $\boldsymbol{\delta}^\tau$ .

To obtain the reconstruction  $\boldsymbol{x}$ , the following minimization problem is formulated:

$$\min_{\boldsymbol{x}, \boldsymbol{\delta}^\tau} \frac{1}{2} \|\mathbf{A}(\boldsymbol{\delta}^\tau) \boldsymbol{x} - \mathbf{b}\|^2 + \lambda \mathcal{R}(\boldsymbol{x}), \quad (6.2)$$

where  $\mathcal{R}(\boldsymbol{x})$  is the regularization term and  $\lambda > 0$  is the regularization parameter. Due to the presence of unknown drifts, the following algorithm (Algorithm 11) is proposed in [36], which is essentially a block coordinate descent (BCD) algorithm that alternatively optimizes with respect to  $\boldsymbol{x}$  and  $\boldsymbol{\delta}^\tau$  using drift information obtained in each iteration.

---

**Algorithm 11** Optimization Algorithm for Parallel CT Problem

---

- 1: Input:  $\mathbf{b}$ ,  $\mathbf{A}_0$ ,  $\boldsymbol{\delta}_0^\tau = \mathbf{0}$ ,  $\lambda_0$ ,  $k_{\max}$
  - 2: Solve for  $\boldsymbol{x}_0 = \frac{1}{2} \arg \min_{\boldsymbol{x}} \|\mathbf{A}_0 \boldsymbol{x} - \mathbf{b}\|^2 + \lambda_0 \mathcal{R}(\boldsymbol{x})$
  - 3: **for**  $k = 1, \dots, k_{\max}$  **do**
  - 4:   Estimate  $\boldsymbol{\delta}_k^\tau$ , construct  $\mathbf{P}_k = \mathbf{P}(\boldsymbol{\delta}_k^\tau)$ , and update  $\mathbf{A}_k = \mathbf{P}_k \mathbf{A}_{k-1}$
  - 5:   Choose  $\lambda_k$
  - 6:   Solve for  $\boldsymbol{x}_k = \frac{1}{2} \arg \min_{\boldsymbol{x}} \|\mathbf{A}_k \boldsymbol{x} - \mathbf{b}\|^2 + \lambda_k \mathcal{R}(\boldsymbol{x})$
  - 7: **end for**
  - 8: Output:  $\boldsymbol{x}_k$
- 

To construct  $\mathbf{P}_k$  in step 4 of Algorithm 11, firstly, drifts are estimated by matching the true sinogram  $\mathbf{B}$  (i.e., observed sinogram) to the “updated” sinogram  $\mathbf{B}^*$  obtained by reshaping  $\mathbf{A}_{k-1} \boldsymbol{x}_{k-1}$  to matrix form. That is, for the  $\tau$ th beam, the estimated drift is computed by minimizing (among all possible drift values) the difference between the  $\tau$ th column of  $\mathbf{B}$  and an interpolation of corresponding columns of  $\mathbf{B}^*$ . After the drift is estimated for every beam, the interpolation matrix  $\mathbf{P}_k$  is



then computed by a closed-form formula. For more details on constructing  $\mathbf{P}_k$  from  $\delta_k^T$ , we refer the readers to [36].

When it comes to choosing a regularization parameter  $\lambda_k$  in each iteration in Step 5, the paper [36] uses a heuristic approach based on the idea that as  $k$  (the number of iterations) increases,  $\mathbf{x}_k$  becomes a better approximation, and the problem becomes less ill-posed. Therefore, it is natural to decrease  $\lambda_k$  as we move forward. For this approach, a first  $\lambda_0$  and last  $\lambda_{k_{\max}}$  ( $\lambda_0 > \lambda_{k_{\max}}$ ) need to be picked beforehand, and intermediate  $\lambda_k$ 's are chosen so that they are equally spaced. Using MATLAB notation, the  $\lambda_k$ 's for  $k = 1, \dots, k_{\max}$  take values in the vector

$$\mathbf{linspace}(\lambda_0, \lambda_{k_{\max}}, 1 + k_{\max}). \quad (6.3)$$

As we attempt to come up with a more systematic way of choosing  $\lambda_k$ 's, we are immediately reminded of the discrepancy principle based adaptive approach that is used for hybrid Krylov subspace methods [21] described in Section 2.4. A similar idea can be applied here. Firstly, let  $\mathbf{x}_{k,\lambda}$  be the solution to

$$\min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}_k \mathbf{x} - \mathbf{b}\|^2 + \lambda \Phi(\mathbf{x}), \quad (6.4)$$

and define  $\psi_k(\lambda)$  in the same way as (2.15):

$$\psi_k(\lambda) = \|\mathbf{A}_k \mathbf{x}_{k,\lambda} - \mathbf{b}\|. \quad (6.5)$$

Since derivation steps are pretty much the same as in Section 2.4, we omit the

redundant steps here and present the final formula directly. Given noise level  $\epsilon = \|\boldsymbol{\eta}\|_2$  and  $\mu \gtrsim 1$ ,  $\lambda_k$  can be updated as follows:

$$\lambda_k = \left| \frac{\mu\epsilon - \psi_k(0)}{\psi_k(\lambda_{k-1}) - \psi_k(0)} \right| \lambda_{k-1}. \quad (6.6)$$

For both Krylov subspace methods and Algorithm 11,  $\psi_k(\lambda_{k-1})$  can be obtained relatively easily. An advantage of Krylov subspace methods is that the residual  $\psi_k(0) = \|\mathbf{A}_k \mathbf{x}_{k,0} - \mathbf{b}\|$  can be obtained directly without solving for  $\mathbf{x}_{k,0}$ . However, for Algorithm 11, we need to solve for  $\mathbf{x}_{k,0}$  in order to compute  $\psi_k(0)$ . The same solver for step 6 can be used here. And to speed it up, there are a few things we can do. For example, we could

- increase the tolerance for the solver's stopping criterion (since we only need the residual, the solution doesn't need to be very accurate), and
- take  $\mathbf{x}_{k-1,0}$  to be the starting guess for the solver (as observed with this discrepancy-base parameter selecting approach, the value of  $\psi_k(0)$  would gradually stabilize as  $k$  increases).

For step 6 in Algorithm 11, the Two-Step Iterative Shrinkage/Thresholding (TwIST) solver [5] used with TV regularization is considered in [36]. However, other linear solvers, such as Krylov subspace methods (e.g. LSQR) can also be applied in this step. In the upcoming subsection, we will investigate the performance of Algorithm 11 on the parallel tomography reconstruction problem that requires parameter calibration. We will first follow [36] to use TwIST as the linear solver, and explore the effect of the adaptive  $\lambda$  choosing approach (6.6) against the old approach (6.3) when

applied to the BCD algorithm. In addition, we will also observe how solutions differ if we adopt Krylov methods in the linear step.

### 6.1.1 Numerical Examples

**Using TwIST with TV Regularization.** In this part, we demonstrate the performance of Algorithm 11, in which TwIST [5] with TV regularization is used in the linear step (step 6). We perform tests for different drift levels, namely,  $\Delta_{\max} = \max \Delta_{\tau} \in \{\Delta, 3\Delta, 5\Delta\}$ . Recall that  $\Delta_{\tau}$  is drift in beam  $\tau$ , and  $\Delta$  is equal to the theoretical step size between adjacent x-ray beams. For each case, to further test the robustness of the proposed method, we superimpose two levels of Gaussian noise to the sinogram  $\mathbf{b}$  (simulated with the assigned drifts) with standard deviation  $\sigma$  equal to 0.01 and 0.02 respectively. Note that  $\mathbf{b}$  is first “normalized” by its largest element before noise is added, after which  $\mathbf{b}$  is scaled back to its original scale. Two test images of size 256 by 256 are used: Brain and Phantom, and they are displayed in Figure 6.2.




Figure 6.2: Test images Brain (left) and Phantom (right).

We adopt the test metric PSNR (peak signal to noise ratio) to be consistent with

[36]. The larger PSNR is, the better the image reconstruction. The stopping criterion for TwIST is when the number of iterations reaches 500 or when the relative change in the objective function 6.2 is less than or equal to  $10^{-4}$ . And the stopping criterion for the outer iterations (i.e., BCD iterations) is if the change in relative residual is small, i.e.,

$$\left| \frac{1}{2} \|\mathbf{A}_k \mathbf{x}_k - \mathbf{b}\|_2^2 - \frac{1}{2} \|\mathbf{A}_{k-1} \mathbf{x}_{k-1} - \mathbf{b}\|_2^2 \right| \leq 10^{-3}. \quad (6.7)$$

Note that in this experimental setting, in order to use the adaptive  $\lambda$  setting approach (6.6), we need the noise level  $\epsilon = \|\boldsymbol{\eta}\|_2$ , but it is not readily available. We are only given standard deviation of the noise,  $\sigma$ , and an estimate of the noise level  $\epsilon$  can be obtained using the following formula,

$$\epsilon = \|\boldsymbol{\eta}\|_2 \approx b_{\max} \sqrt{N_\tau N_\theta} \sigma, \quad (6.8)$$

where  $b_{\max}$  is the maximum element in  $\mathbf{b}$ . Since this is very likely an overestimate (because it is estimated with  $b_{\max}$ ), in practice, we can actually take  $\mu\epsilon \approx b_{\max} \sqrt{N_\tau N_\theta} \sigma$  without explicitly defining  $\mu \gtrsim 1$  when updating  $\lambda_k$  using (6.6). It is also worth mentioning how we set the initial regularization parameter  $\lambda_0$ . In order to have the initial  $\lambda_0$  reflect well the level of noise from two main sources, the white noise and experimental drifts, in our problem, we propose the following formula for constructing  $\lambda_0$ :

$$\lambda_0 = 0.001(1 + \sigma)(1 + \Delta_{\max}). \quad (6.9)$$

Table 6.1 displays the PSNR values (in dB) for a variety of test cases with different objects, different maximum drift levels and different noise levels. We compare results

of the baseline approach (i.e., solving the linear system directly without calibrating drift) and the heuristic approach (i.e., Algorithm 11 using the heuristic  $\lambda$  choosing method (6.3) against the adaptive regularization parameter setting approach (6.6). The presented PSNR values for the adaptive approach are given by the convergent iterations for each test case. For the heuristic approach, we take the 10th iteration, at which  $\lambda_0$  is decreased to  $0.01\lambda_0$ .

We can observe in Table 6.1 that in general, the adaptive approach outperforms both heuristic approach and baseline approach. However, for a few cases with noise  $\sigma = 0.02$ , the PSNR given by “adaptive” is slightly less than the “heuristic” approach’s PSNR value. This is probably because the estimate for  $\epsilon$  is less accurate for larger values of  $\sigma$ .

Max Drift	$\Delta$			$3\Delta$			$5\Delta$		
Noise $\sigma$	0	0.01	0.02	0	0.01	0.02	0	0.01	0.02
Brain (adap.)	29.43	26.78	24.61	27.01	25.54	23.53	25.46	24.90	23.30
Brain (heur.)	26.99	26.34	25.50	24.44	24.38	23.87	23.57	23.51	23.34
Brain (base.)	26.27	24.34	19.90	20.73	20.22	18.53	18.06	17.80	17.08
Phan (adap.)	26.77	26.05	24.44	22.80	22.59	21.96	21.75	21.54	21.67
Phan (heur.)	24.94	24.98	24.82	21.19	20.90	20.89	19.75	19.86	19.66
Phan (base.)	21.90	21.51	20.53	16.58	16.53	16.23	14.26	14.22	14.04

Table 6.1: Performance comparison of new adaptive (adap.) method with the baseline (base.) and heuristic (heur.) methods for three drift levels and three noise levels on Brain and Phantom (Phan) test images. Test Metric: PSNR (dB).

We visualize the reconstructions of phantom in Figure 6.3. It is noticeable that with drift calibration, we are able to recover the fine feature of the object, and in particular, with the adaptive  $\lambda_k$  approach, the fine details can be recovered to an

enhanced contrast comparing to the old approach in average.




Figure 6.3: Phantom reconstructions given by Algorithm 11 using TwIST. Top row: baseline; middle row: heuristic  $\lambda_k$ 's; bottom row: adaptive  $\lambda_k$ 's.

We further investigate the iterative performance of the two approaches choosing  $\lambda_k$ 's in Figure 6.4. Remarkably, our proposed adaptive way is able to gradually decrease  $\lambda$  until stabilized. Comparing to the heuristic approach, the proposed formula promotes the convergence of the performance in a more automatic way. To perform fair visual comparison, we allow the algorithm to continue until reaching the maximum allowed number of iterations. We indicate the convergence of our proposed adaptive approach on  $\lambda_k$  by a red diamond according to stopping criteria 6.7.

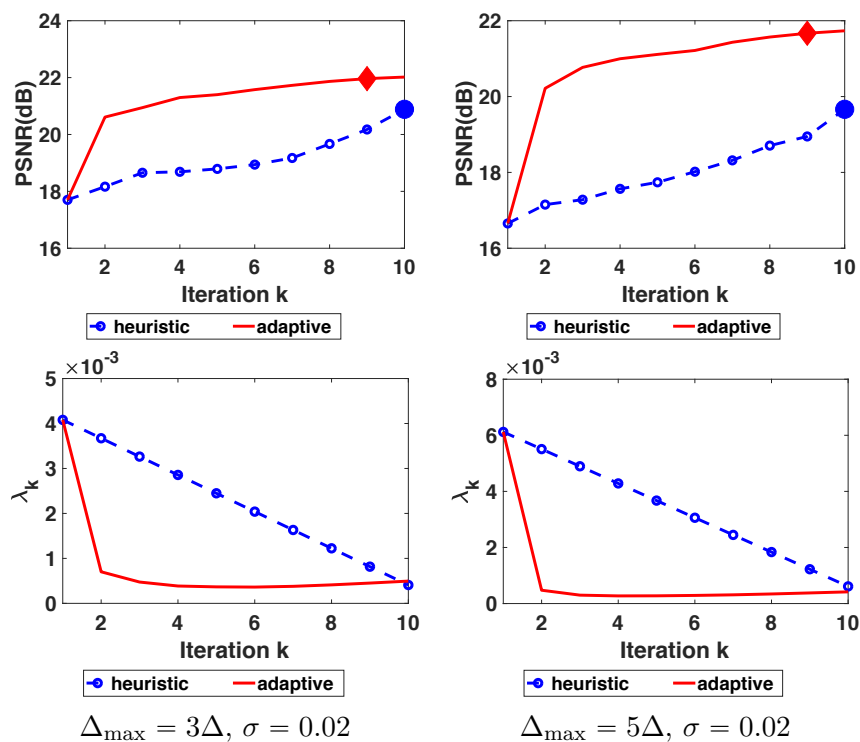


Figure 6.4: Convergence results for Algorithm 11 on Phantom test image. Top row: PSNR; bottom row:  $\lambda_k$ 's

**Using Krylov Subspace Methods.** Since the main contribution of this thesis is the newly developed low-rank Krylov methods, we also put the new methods to test in this parallel tomography calibration problem. Note that we are applying these Krylov methods in step 6 of Algorithm 11, replacing TwIST. Krylov subspace methods are intrinsically different from the previous TwIST solver, so in this part, we compare the performance of the new low-rank methods to other Krylov methods in its family. Namely, we compare the following methods: LSQR (without any form of regularization), LSQR\_tik (LSQR with Tikhonov regularization), LSQR\_htv (LSQR with heuristic TV regularization), as well as the two low rank methods: FLSQR-

NNR(v) (Section 4.4) and LR-FLSQR (Section 5.1.1). Implementations of LSQR\_tik and LSQR\_htv can be found in IR Tools [17].

We test the Krylov solvers using the Phantom test image as seen in Figure 6.2. We fix the noise standard deviation  $\sigma$  to be 0.02, and vary the maximum drift level so that  $\Delta_{\max} = \max \Delta_\tau \in \{\Delta, 3\Delta, 5\Delta\}$ . We set the number of outer iterations to be 20. In terms of inner iterations (i.e., Krylov iterations), we also use 20, but rather than imposing a stopping criterion, we pick the iteration whose solution has the lowest relative error compared against the true solution. We do so because we are interested in comparing the performance of Krylov methods in the best case scenario (i.e., stopping at lowest relative error). Also, the methods LSQR\_tik and LSQR\_htv require regularization parameters, and we follow (6.6) to adaptively pick a fixed  $\lambda_k$  for each outer iteration. This is different from the automatic update of regularization parameters in the inner Krylov iterations.




Figure 6.5: PSNR convergence plots (Phantom) using Krylov methods. For LR-FLSQR, truncation ranks  $\kappa = \kappa_B = 50$ , picked by trial and error.

Figure 6.5 displays the convergence of PSNR values of the various Krylov methods. We can easily observe that out of the 5 Krylov methods, LSQR, LSQR\_tik and LSQR\_htv behave quite similarly for different drift levels, and all three of these






Figure 6.6: Phantom reconstructions using Krylov methods. Upper row: LSQR solutions (no regularization); middle row: LSQR\_htv solutions; bottom row: FLSQR-NNR(v) solutions.

methods are outperformed by the new low-rank Krylov methods FLSQR-NNR(v) and LR-FLSQR. It is also interesting that for the low drift level  $\Delta_{\max} = \Delta$ , the best-performing method is LR-FLSQR, while for the higher drift levels  $\Delta_{\max} = 3\Delta$  and  $5\Delta$ , FLSQR-NNR(v) is the best of all. Furthermore, the improvement in the new low rank methods over traditional Krylov methods is more apparent for the larger drift levels. In spite of the improvements, we should also recognize that even when we are exploiting the Krylov methods to their fullest potential, all Krylov solutions

are still of lower quality compared to TwIST. The best-performing Krylov methods give PSNR values that are 1.5-3.5 lower than TwIST.

In Figure 6.6, we present solutions at the 20th outer iterations of Algorithm 11 using LSQR, LSQR\_hlv, and FLSQR-NNR(v). It is easy to see that FLSQR-NNR(v) renders solutions of the highest clarity compared to the other two – there is less noise in the background, and the phantom shape is better preserved. However, when compared to TwIST reconstructions in Figure 6.3, FLSQR-NNR(v) is still lacking in details (e.g., the three small circles inside the phantom), and are more noisy.

## 6.2 Fan-Beam Tomography

In this section, we describe a different tomography setup – fan-beam tomography. That is, instead of a panel that emits parallel x-ray beams, we now have a point source that emits fan beams, as can be seen in Figure 6.7(a). Ideally, for each scan, the source should be of equal distance to the object and rotates a fixed angle. However, this may not be the case for portable CT machines that are subject to more experimental errors caused by machine transportation. In Figure 6.7(b), the theoretical location of the center scanning position is shown in light green, but during the scanning process it may be shifted to the dark green location. Hence, when reconstructing the imaged object, we also need to take into consideration geometry parameter calibration.

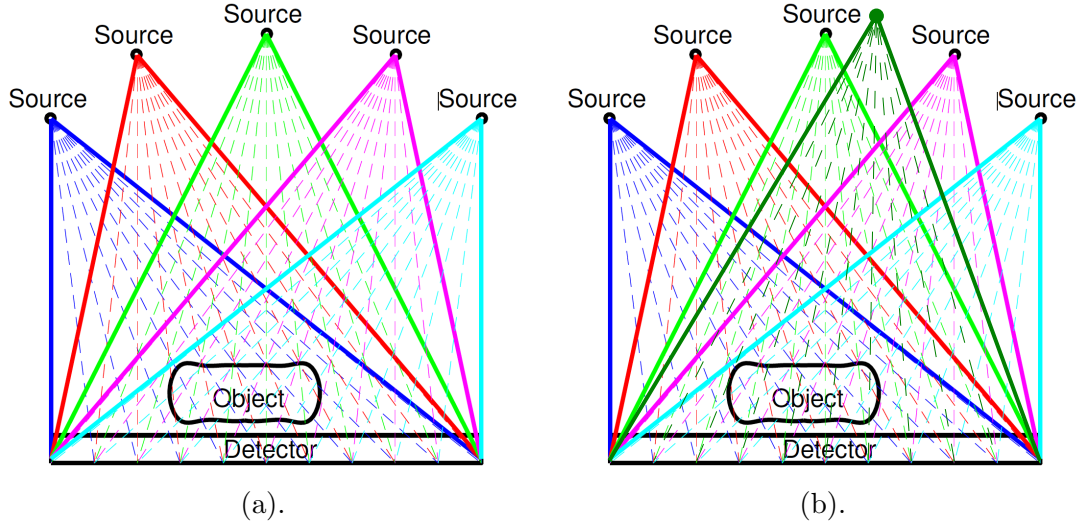


Figure 6.7: Fan-beam Tomography illustration.

Consider the fan-beam CT image reconstruction problem

$$\min_{\mathbf{x}, \mathbf{p}} \|\mathbf{A}(\mathbf{p})\mathbf{x} - \mathbf{b}\|_2^2 + \lambda\mathcal{R}(\mathbf{x}), \quad (6.10)$$

where  $\mathbf{A} \in \mathbb{R}^{M \times N}$  is the forward model that depends on geometry parameters  $\mathbf{p}$ ,  $\mathbf{b} \in \mathbb{R}^M$  is the observed data (sinogram) perturbed by random noise  $\boldsymbol{\eta}$  (i.e.,  $\mathbf{b} = \mathbf{b}^{\text{ex}} + \boldsymbol{\eta}$ ),  $\mathbf{x} \in \mathbb{R}^N$  is the quantity of interest (imaged object), and  $\lambda\mathcal{R}(\mathbf{x})$  is a regularization term. In experimental settings, the geometry parameters  $\mathbf{p}$  may not be exactly known, but we have an estimate, i.e., their theoretical values. As a result, we need to both solve for the solution image  $\mathbf{x}$  and calibrate parameters  $\mathbf{p}$ . Note also that  $\mathbf{b}$  and  $\mathbf{x}$  are vectorized quantities of the sinogram  $\mathbf{B} \in \mathbb{R}^{N_\tau \times N_\theta}$  (where  $N_\tau$  is the number of beams and  $N_\theta$  is the number of scanning angles) and the solution image  $\mathbf{X} \in \mathbb{R}^{n \times n}$ , where  $N = n^2$ .

For this problem, we consider two types of uncertainties in the geometry parameters: source to object distances, and scanning angles. Hence the unknown geometry parameters  $\mathbf{p} \in \mathbb{R}^{2N_p}$  consist of two components  $\mathbf{r}$  and  $\mathbf{d}$ , i.e.,  $\mathbf{p} = [\mathbf{r}; \mathbf{d}]$ , where  $\mathbf{r} \in \mathbb{R}^{N_p} = [r_1; \dots; r_{N_p}]$  are source to object distances with theoretical value  $2n$ , and  $\mathbf{d} \in \mathbb{R}^{N_p} = [\delta_1^\theta; \dots; \delta_{N_p}^\theta]$  are perturbation in scanning angles with theoretical value 0. We assume that  $r_i \in [1.5n, 2.5n]$  ( $n$  is dimension of image) and  $\delta_i^\theta \in [-0.5, 0.5]$ .

Note that technically,  $N_p$  could be equal to the total number of scanning angles and can be as large as 180 or 360. For simplification purposes, we start by choosing  $N_p$  much smaller than total number of scanning angles (e.g.,  $N_p = 3$  or  $N_p = 6$ ). For instance, if we use  $N_\theta = 180$  scanning angles 0:2:359 and  $N_p = 3$ , then  $\mathbf{p} \in \mathbb{R}^6$ ,  $\mathbf{r}, \mathbf{d} \in \mathbb{R}^3$ , and  $(r_i, \delta_i^\theta)$  would be constant for a set of  $180/N_p = 60$  angles.

Similar to what's described in Section 6.1, a block coordinate descent (BCD) optimization algorithm may be used to alternatively solve for  $\mathbf{x}$  and  $\mathbf{p}$  in Problem (6.10). The idea is simple: for the unknowns  $\mathbf{x}$  and  $\mathbf{p}$ , we can fix one and solve for the other, and repeat this process multiple times. Since we are given the theoretical value of  $\mathbf{p}$ , we can use it as an initial guess  $\mathbf{p}_0$ , and start from here to solve for  $\mathbf{x}_0$ , then  $\mathbf{p}_1, \mathbf{x}_1, \mathbf{p}_2, \mathbf{x}_2, \dots$ , etc. The BCD method is summarized in Algorithm 12. Note that our version of BCD (Algorithm 12) is sometimes also called alternating minimization.

The linear steps 2 and 6 of Algorithm 12 can be solved with a linear inverse problem solver such as LSQR. But for the model update step (step 4), since  $\mathbf{A}$  depends nonlinearly on  $\mathbf{p}$ , we need to use a nonlinear solver such as Gauss-Newton based `imfil` [42], and this nonlinear step can be expensive for large values of  $N_p$ .

---

**Algorithm 12** BCD for fan-beam CT problem
 

---

- 1: **Input:**  $\mathbf{b}$ ,  $\mathbf{p}_0$ ,  $\mathbf{A}_0 = \mathbf{A}(\mathbf{p}_0)$
  - 2: Solve for  $\mathbf{x}_0 = \arg \min_{\mathbf{x}} \|\mathbf{A}_0 \mathbf{x} - \mathbf{b}\|^2 + \lambda_0 \|\mathbf{x}\|_2^2$
  - 3: **for**  $k = 1, \dots$  **do**
  - 4:   Solve for  $\mathbf{p}_k = \arg \min_{\mathbf{p}} \|\mathbf{A}(\mathbf{p}) \mathbf{x}_k - \mathbf{b}\|^2$
  - 5:   Update  $\mathbf{A}_k = \mathbf{A}(\mathbf{p}_k)$
  - 6:   Solve for  $\mathbf{x}_k = \arg \min_{\mathbf{x}} \|\mathbf{A}_k \mathbf{x} - \mathbf{b}\|^2 + \lambda_k \|\mathbf{x}\|_2^2$
  - 7: **end for**
  - 8: **return**  $\mathbf{x}_k$
- 

Note that for this step, we can actually reduce problem size by exploiting separability of  $\mathbf{A}(\mathbf{p})$ . That is, the minimization problem in step 4 of Algorithm 12 can be equivalently written as

$$\min_{\mathbf{p}} \left\| \begin{bmatrix} \mathbf{A}(p_1) \\ \mathbf{A}(p_2) \\ \vdots \\ \mathbf{A}(p_{N_p}) \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \vdots \\ \mathbf{b}_{N_p} \end{bmatrix} \right\|^2 = \sum_{i=1}^{N_p} \|\mathbf{A}(p_i) \mathbf{x} - \mathbf{b}_i\|^2. \quad (6.11)$$

By doing so, we divide the big nonlinear problem into  $N_p$  small problems, which can be solved in parallel. The separability property works very well with our proposed BCD framework, but other approaches such as variable projection [53, 55] would not be able to utilize this nice property. On the other hand, the nonlinear solver can depend heavily on the initial guess, and may converge to a local minimum rather than the global minimum, which affects all subsequent  $\mathbf{x}_k$ 's and  $\mathbf{p}_k$ 's. Hence, although we can conveniently define  $\mathbf{p}_0$  as the theoretical value for  $\mathbf{p}$ , it might be worthwhile to look for other ways of defining  $\mathbf{p}_0$  that allows for better convergence of the overall

BCD solver.

Since the nonlinear optimization step is very costly and takes a long time to run, it would be even better if we can find a way to avoid running this step. Inspired by [1], which trains neural networks to learn regularization parameters for inverse problems, we are interested in building machine learning models to learn the unknown CT geometry parameters. In other words, by training a machine learning model  $\hat{\Phi} : \mathbf{b} \rightarrow \mathbf{p}$  that maps sinogram  $\mathbf{b}$  to geometry parameters  $\mathbf{p}$ , we may be able to avoid running the alternating minimization scheme. And as a result, we only need to run one linear step to solve for  $\mathbf{x}$ . This framework is summarized in Algorithm 13.

---

**Algorithm 13** Learning Geometry Parameters

---

- 1: *offline phase*
  - 2: For  $j = 1, \dots, J$ , randomly generate geometry parameters  $\mathbf{p}_j$ , noise  $\boldsymbol{\eta}_j$ , and generate training data  $\mathbf{b}_j = \mathbf{A}(\mathbf{p}_j)\mathbf{x}_j^{\text{ex}} + \boldsymbol{\eta}_j$
  - 3: Train model  $\hat{\Phi}$  which maps observation  $\mathbf{b}$  onto geometry parameters  $\mathbf{p}$
  - 4: *online phase*
  - 5: For new data  $\mathbf{b}_{j'}$ , predict parameters  $\mathbf{p}_{j'} = \hat{\Phi}(\mathbf{b}_{j'})$
  - 6: Solve inverse problem  $\min_{\mathbf{x}} \|\mathbf{A}(\mathbf{p}_{j'})\mathbf{x} - \mathbf{b}\|^2 + \lambda\mathcal{R}(\mathbf{x})$
- 

To train the model  $\hat{\Phi} : \mathbf{b} \rightarrow \mathbf{p}$ , we need the following training data:

- $\hat{\mathbf{B}} \in \mathbb{R}^{J \times M}$  is (feature) matrix with  $J$  rows. Each row is a vectorized sinogram  $\mathbf{b}_j$  of length  $M$ , constructed using random  $\mathbf{r}_j$  and  $\boldsymbol{\delta}_j^\theta$ , with superimposed Gaussian noise, i.e.,  $\mathbf{b}_j = \mathbf{A}(\mathbf{r}_j, \boldsymbol{\delta}_j^\theta)\mathbf{x}_j^{\text{ex}} + \boldsymbol{\eta}_j$ , with  $\mathbf{x}_j^{\text{ex}}$  the true phantom and  $\boldsymbol{\eta}_j$  the random noise vector.
- $\mathbf{R}, \mathbf{D} \in \mathbb{R}^{J \times N_p}$  are matrices of true geometry parameters, with rows  $\mathbf{r}_j$  and  $\boldsymbol{\delta}_j^\theta$  corresponding to sinogram  $\mathbf{b}_j$ .

One important consideration during the training phase is choosing the phantom images  $\mathbf{x}_j^{\text{ex}}$ . If we fix  $\mathbf{x}_j^{\text{ex}}$  to be a constant phantom image for all  $j$ , then we would expect that the trained model only works well for the phantom that the model is trained on. In other words, the model would be best at predicting CT geometry parameters when the phantom being reconstructed is the same phantom that the training data is built upon, and may not generalize well to other phantoms. Since in a real-world CT reconstruction problem, we do not know what the phantom is, using a model that is trained upon a fixed phantom may not render good solutions.

However, if we use different phantoms  $\mathbf{x}_j^{\text{ex}}$ 's when generating the training data, the model may become confused due to the problem of non-uniqueness. It is possible that there exist  $(\mathbf{p}_1, \mathbf{x}_1), (\mathbf{p}_2, \mathbf{x}_2)$  with  $\mathbf{p}_1 \neq \mathbf{p}_2$  and  $\mathbf{x}_1 \neq \mathbf{x}_2$ , such that  $\mathbf{A}(\mathbf{p}_1)\mathbf{x}_1 = \mathbf{A}(\mathbf{p}_2)\mathbf{x}_2 = \mathbf{b}$ . This means that the mapping  $\hat{\Phi} : \mathbf{b} \rightarrow \mathbf{p}$  may not be one-to-one. Therefore, the model may predict parameters  $\mathbf{p}$  that are very far from the true value, thus rendering far-from-true solution  $\mathbf{x}$  in Algorithm 13.

Disturbance-based strategies have been used in applications like Power Systems [35, 71] to help with non-uniqueness. If we translate this strategy into our context, we may apply “disturbances”, i.e., different phantoms and different noise levels to generate many training samples that correspond to the same set of geometry parameters  $\mathbf{p}$ . That is, for each  $\mathbf{p}$ , we generate  $l$  training samples  $\mathbf{b}_i = \mathbf{A}(\mathbf{p})\mathbf{x}_i + \boldsymbol{\eta}_i$  for  $i = 1, \dots, l$  using  $l$  different phantoms and noise levels. We will see how the disturbance-based strategy performs in Section 6.2.2.

### 6.2.1 Hybrid ML-BCD Method

In this Section, we describe a hybrid method that combines machine learning with block coordinate descent for solving (6.10). We will also explain technical details for training the machine learning model.

We have mentioned previously that the machine learning problem of learning geometry parameters when the phantom  $\mathbf{x}_j^{\text{ex}}$  is not fixed in the training data may suffer from non-uniqueness. Although in our experience, using Algorithm 13 with an ML model trained using non-fixed  $\mathbf{x}_j^{\text{ex}}$ 's does not deliver solutions that are too far from the true images, we do notice that relative errors in the predicted parameters on the testing set have increased compared to the case with a fixed  $\mathbf{x}_j^{\text{ex}}$ , but the relative errors are still very low compared to the theoretical values (more on this in Section 6.2.2).

One natural approach that comes to mind is using a hybrid method that combines a machine learning model and BCD. That is, we may first use the trained ML model to predict  $\mathbf{p}$ , then feed the predicted  $\mathbf{p}$  as a starting guess to the BCD algorithm to run a few iterations to refine the solution. This hybrid framework is described in Algorithm 14.

---

**Algorithm 14** ML-BCD hybrid framework

---

- 1: *offline phase*
  - 2: Use offline phase in Algorithm 13 to train ML model  $\hat{\Phi} : \mathbf{b} \rightarrow \mathbf{p}$
  - 3: *online phase*
  - 4: (predict) For new data  $\mathbf{b}_{j'}$ , predict parameters  $\mathbf{p}_{j'} = \hat{\Phi}(\mathbf{b}_{j'})$
  - 5: (refine) Run Algorithm 12 with  $\mathbf{p}_0 = \mathbf{p}_{j'}$  and  $\mathbf{A}_0 = \mathbf{A}(\mathbf{p}_{j'})$
- 

While there are many machine learning models to choose from, we consider one



of the simplest models – a multi-output linear regression model, which can also be thought of as a one-layer neural network. This ML model assumes a linear relationship between input and output, and for simplicity, we take a zero bias term. To build this model, we learn weights  $\mathbf{W} \in \mathbb{R}^{M \times 2N_p}$  by solving the least squares problem

$$\min_{\mathbf{W}} \|\hat{\mathbf{B}}\mathbf{W} - [\mathbf{R} \ \mathbf{D}]\|_F^2. \quad (6.12)$$

Then, given new data  $\mathbf{b}_{j'}$ , we simply need to compute  $\mathbf{p}_{j'} = [\mathbf{r}_{j'}; \delta_{j'}^\theta] = \mathbf{b}_{j'}^T \mathbf{W}$  to obtain the predicted geometry parameters, which can be fed as initial guess  $\mathbf{p}_0$  into the BCD algorithm. We have chosen the loss function to be  $\|\cdot\|_F^2$ , which makes the minimization problem separable (similar to the nonlinear step of BCD). Since  $\mathbf{R} = [\mathbf{R}_1 \cdots \mathbf{R}_{N_p}]$  with column vectors  $\mathbf{R}_i$ , and  $\mathbf{D} = [\mathbf{D}_1 \cdots \mathbf{D}_{N_p}]$  with column vectors  $\mathbf{D}_i$ , (6.12) is equivalent to

$$\sum_{i=1}^{N_p} \min_{\mathbf{w}_{r,i} \in \mathbb{R}^M} \|\hat{\mathbf{B}}\mathbf{w}_{r,i} - \mathbf{R}_i\|_2^2 + \sum_{i=1}^{N_p} \min_{\mathbf{w}_{d,i} \in \mathbb{R}^M} \|\hat{\mathbf{B}}\mathbf{w}_{d,i} - \mathbf{D}_i\|_2^2, \quad (6.13)$$

where  $\mathbf{w}_{r,i}$  and  $\mathbf{w}_{d,i}$  are column vectors of  $\mathbf{W}_R$  and  $\mathbf{W}_D$  such that  $\mathbf{W}_R = [\mathbf{w}_{r,1} \cdots \mathbf{w}_{r,N_p}]$ ,  $\mathbf{W}_D = [\mathbf{w}_{d,1} \cdots \mathbf{w}_{d,N_p}]$ , and  $\mathbf{W} = [\mathbf{W}_R \ \mathbf{W}_D]$ . The sub-problems  $\min_{\mathbf{w}_{r,i}} \|\hat{\mathbf{B}}\mathbf{w}_{r,i} - \mathbf{R}_i\|_2^2$  and  $\min_{\mathbf{w}_{d,i}} \|\hat{\mathbf{B}}\mathbf{w}_{d,i} - \mathbf{D}_i\|_2^2$  are large and non-square linear systems, and we need to use an iterative solver such as LSQR to solve each of them. Algorithm 15 contains more detailed description of the hybrid method.

---

**Algorithm 15** ML-BCD hybrid algorithm
 

---

- 1: *offline phase - training ML model*
  - 2: Generate training data  $\widehat{\mathbf{B}}$ ,  $\mathbf{R}$  and  $\mathbf{D}$
  - 3: Compute ML model weights  $\mathbf{W}$  that is the solution to  $\min_{\mathbf{W}} \|\widehat{\mathbf{B}}\mathbf{W} - [\mathbf{R} \mathbf{D}]\|_F^2$
  - 4: *online phase - prediction and BCD*
  - 5: For new data  $\mathbf{b}_{j'}$ , predict geometry parameters  $\mathbf{p}_{j'} = [r_{j'}; \delta_{j'}^\theta] = \mathbf{b}_{j'}^T \mathbf{W}$
  - 6: Set  $\mathbf{p}_0 = \mathbf{p}_{j'}$  and  $\mathbf{A}_0 = \mathbf{A}(\mathbf{p}_{j'})$ , and solve for  $\mathbf{x}_0 = \arg \min_{\mathbf{x}} \|\mathbf{A}_0 \mathbf{x} - \mathbf{b}\|^2 + \lambda_0 \|\mathbf{x}\|_2^2$
  - 7: Alternatively optimize for  $\mathbf{p}_k$  and  $\mathbf{x}_k$  for  $k = 1, 2, \dots$  as in Algorithm 12.
  - 8: Return  $\mathbf{x}_k$
- 

## 6.2.2 Numerical Examples

In this Section, we present several numerical experiments comparing performances of Algorithms 12, 13 and 15. We have used the `PRtomo` function of `IRTools` [17] to generate all training data in this section. The linear solver used is `IRhybrid_lsqr` from `IRTools` and the nonlinear solver is `imfil` [42]. While there are many linear solvers to choose from, the purpose of the numerical experiments is to compare different algorithms that simultaneously reconstructs the image and calibrates geometry parameters, not comparing performances of different linear solvers, so we will have a fair comparison as long as we fix the linear solver.

**Test 1: 64 by 64 fixed phantom problem with  $N_p = 3$ .** We first consider a  $64 \times 64$  fixed Shepp-Logan phantom as  $\mathbf{x}_j^{\text{ex}} \forall j$  when generating the training data. We use 180 scanning angles, i.e., 0:2:359 and  $N_p = 3$ , so there are 6 unknown geometry parameters in total: 3  $r_i$ 's and 3  $\delta_i^\theta$ 's, with each pair of  $(r_i, \delta_i^\theta)$  corresponding to 60 adjacent scanning angles. As a result, each  $\mathbf{b}_j \in \mathbb{R}^{16380}$ , so  $\widehat{\mathbf{B}} \in \mathbb{R}^{J \times 16380}$ ,  $\mathbf{R}$  and

$\mathbf{D} \in \mathbb{R}^{16380 \times 3}$ . Random white noise with fixed noise level 0.01 is superimposed onto each  $\mathbf{b}_j$ .

We have generated training data  $\hat{\mathbf{B}}$ ,  $\mathbf{R}$  and  $\mathbf{D}$  with  $J = 20000$  training samples, and an additional testing set  $\hat{\mathbf{B}}_{\text{test}}$ ,  $\mathbf{R}_{\text{test}}$ , and  $\mathbf{D}_{\text{test}}$  of 2000 samples. We train ML model weights using the method described in Section 6.2.1, i.e., computing weight matrix  $\mathbf{W} = [\mathbf{W}_R \ \mathbf{W}_D]$  by solving every subproblem of (6.13). We train weights  $\mathbf{W}$  using training sets of different sizes: 2000, 4000, 6000,  $\dots$ , 20000, compute the predicted geometry parameters

$$\mathbf{R}_{\text{pred}} = \hat{\mathbf{B}}_{\text{test}} \mathbf{W}_R \quad \text{and} \quad \mathbf{D}_{\text{pred}} = \hat{\mathbf{B}}_{\text{test}} \mathbf{W}_D,$$

and evaluate testing errors

$$\frac{\|\mathbf{R}_{\text{pred}} - \mathbf{R}_{\text{test}}\|_F}{\|\mathbf{R}_{\text{test}}\|_F} \quad \text{and} \quad \frac{\|\mathbf{D}_{\text{pred}} - \mathbf{D}_{\text{test}}\|_F}{\|\mathbf{D}_{\text{test}}\|_F}$$

on the testing set of 2000 samples. We plot testing errors against training set size in Figure 6.8.




Figure 6.8: Relative errors in  $\mathbf{R}_{\text{pred}}$  (left) and  $\mathbf{D}_{\text{pred}}$  (right) on testing set.

We can observe that as the training set size increases, relative error in the predicted geometry parameters decreases. This is because as the training set becomes larger, it covers more variations in the different combinations of geometry parameters, which makes predictions more accurate. Also, the relative errors in  $\mathbf{R}_{\text{pred}}$  and  $\mathbf{D}_{\text{pred}}$  are very different in magnitude, meaning that it is much easier to predict  $\mathbf{r}$  parameters more accurately.




Figure 6.9: Phantom reconstructions of 3 testing samples that share the same underlying phantom with training set (1 testing sample per row). 1st column: solutions using theoretical parameters (i.e., uncalibrated). 2nd column: solutions using true parameters (i.e. perfect calibration). 3rd column: solutions using ML-predicted parameters (Algorithm 13). 4th column: solutions using 10 BCD iterations (Algorithm 12).

Figure 6.9 shows solutions using theoretical parameters, true parameters, ML-predicted parameters (Algorithm 13), and BCD (Algorithm 12). The ML model weights are trained on a training set of size 20000. We can observe in Figure 6.9 that if we leave the geometry parameters uncalibrated, the solution images will have very poor quality. If we use the BCD scheme (Algorithm 12), the reconstruction quality seems to have improved, but the solutions are still fuzzy and unclear. Using Algorithm 13 to predict geometry parameters using learned ML model weights, we obtain highly accurate solutions, resembling solutions computed using the true parameters (which represents to the case of perfect calibration).

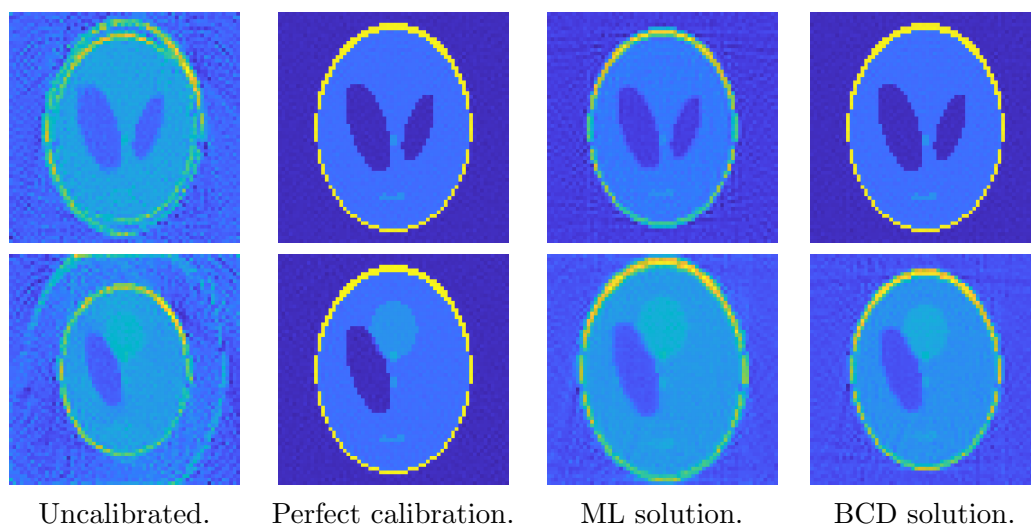


Figure 6.10: Phantom reconstructions of 2 testing samples that have different underlying phantom than training set (1 testing sample per row). Column order same as Figure 6.9.

The model weights trained in this test is based on training data generated with the same phantom. Therefore, if the solution image we are looking for is different from the phantom that the weights are trained on, then the ML model may not gen-

eralize well. As we observe in Figure 6.10, the test samples have different underlying phantoms (different from what's in the training data). As a result, Algorithm 13 no longer yields accurate solutions, and no longer has an advantage over BCD.

**Test 2: Testing Disturbances.** While we want to train models that generalize well to different phantoms, the problem of non-uniqueness may affect the training of machine learning models. In this test, we explore the effect of disturbances by generating the training data using different phantoms  $\mathbf{x}_j^{\text{ex}}$ 's. Examples of such phantoms used are shown in Figure 6.11. It can be seen that each phantom is created by altering the original phantom by randomly removing ellipses and slightly varying ellipse sizes and angles. Random noise of noise level chosen randomly in the range 0.1% to 2% is superimposed to each sample.




Figure 6.11: Examples of different phantoms used.

Two training sets are created. We apply the disturbance strategy in the first training set by generating 5 different phantoms and noise levels for each  $\mathbf{p}_j$ , with 7000  $\mathbf{p}_j$ 's in total, i.e., there are 35000 samples in the data set. As a frame of reference, we compare results given by another training set of the same size, but without disturbances. That is, 35000 samples each with a different  $\mathbf{p}_j$ , phantom and noise level.

We use these two training sets and subsets of the data sets to train ML model weights. We compare the relative errors of the ML-predicted  $\mathbf{p}$  for the testing set consisting of 1000 samples (generated with different samples in the same way as illustrated in Figure 6.11). In Figure 6.12, we show convergence of relative errors in the geometry parameters  $\mathbf{R}_{\text{pred}}$  and  $\mathbf{D}_{\text{pred}}$  predicted by ML models trained using 5000:5000:35000 samples of the two data sets. Note that, for example, with the training set size 15000, one training set consists of 3000  $\mathbf{p}_j$ 's each having 5 disturbances (i.e., phantoms and noise levels), the other training set has 15000  $\mathbf{p}_j$ 's (each having their own phantom and noise level). We also notice that when the model is trained on training data generated using different phantoms, the relative errors in the geometry parameters are around 10 times larger than the relative errors presented in Test 1 due to the variance of phantoms.

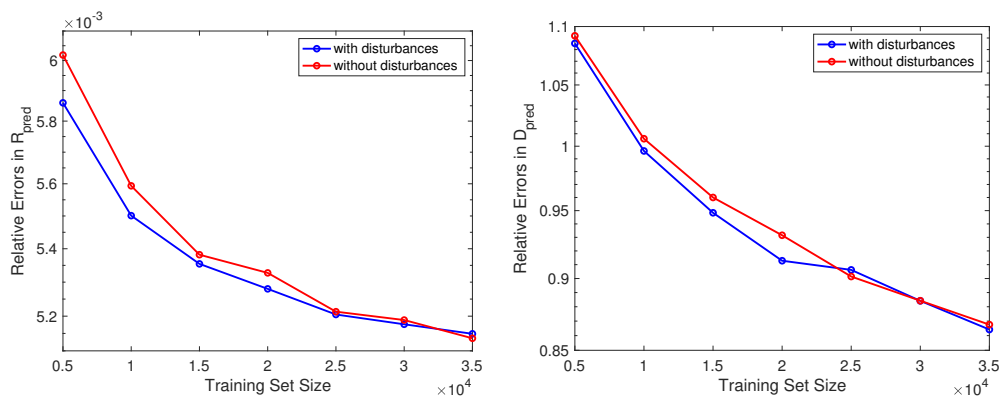


Figure 6.12: Relative errors in  $\mathbf{R}_{\text{pred}}$  and  $\mathbf{D}_{\text{pred}}$  on testing set.

We see in Figure 6.12 that the relative errors decrease for both training sets as training set size increases. If we compare the relative errors vertically, we can see that the errors are not so different for the two training sets. The slight difference may be

caused by randomness in generating the two sets. We believe it's only fair to compare results using training sets of the same size, since, for example, if we compare results given by one training set of 5000 samples with results given by another training set of 5000 samples each with 5 disturbances, the decrease in relative errors may be caused by the increased training set size, not disturbance. Therefore, given the results in this test, it is hard to conclude that disturbances help improve accuracy in the predicted geometry parameters.

**Test 3: 128 by 238 varying phantom problem with  $N_p = 6$ .** We have observed in Tests 1 and 2 that errors in the predicted geometry parameters would increase if the ML model is not trained and tested on the same phantom images. We have also seen that, in this case, although errors in  $\delta^\theta$  are quite large, relative errors in  $\mathbf{r}$  are still relatively low.

In this test, we focus on the case where phantoms vary across training and testing samples. We increase phantom size to  $n = 128$  and number of geometry parameters to  $N_p = 6$ . Hence, with 180 scanning angles 0:2:358, each pair of  $(r_i, \delta_i^\theta)$  will be constant for 30 adjacent angles. and compare solutions of the BCD (Algorithm 12), ML (Algorithm 13) and ML-BCD hybrid (Algorithm 15) methods. The ML weights are trained using the 40000 training samples. We evaluate the ML-BCD hybrid method and BCD method (each with 10 iterations) on a testing set of 100 samples, calculate the relative errors in  $\mathbf{x}_k$ ,  $\mathbf{r}_k$  and  $\delta_i^\theta$  for each BCD iteration  $k$ , and generate the histograms in Figure 6.13.



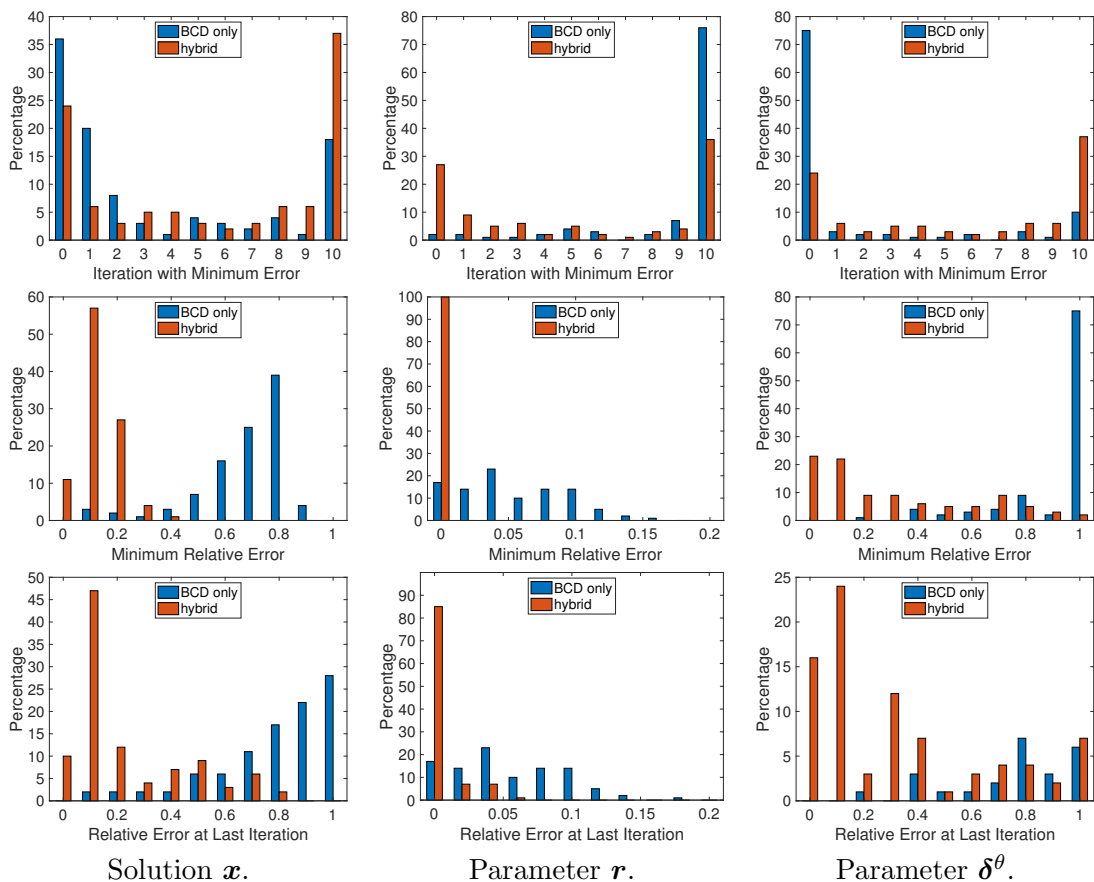


Figure 6.13: 1st row: BCD iteration number with the minimum relative error. 2nd row: minimum relative error across all iterations. 3rd row: relative error at the 10th BCD iteration.

In the first row of Figure 6.13, frequencies of the iteration at which the minimum error occurs are plotted. Note that iteration 0 represents the initial guess, i.e., for BCD, this would be the uncalibrated solution; and for ML-BCD hybrid, this would be obtained directly from using ML-predicted parameters without further refining with BCD. We can draw the conclusion that the hybrid method has better convergence properties for  $\mathbf{x}$ , where for nearly 40% percent of testing samples, BCD reaches the

minimum relative error of  $\boldsymbol{x}$  in the last iteration. While for BCD, in more than 50% percent of testing samples, the minimum relative error in  $\boldsymbol{x}$  is either reached at the initial guess or the first iteration, which means that BCD is not very effective at improving solution accuracy. This is further confirmed in the distribution of relative errors in the second and third rows of Figure 6.13, in which we notice that in general, the relative errors in both the solutions and the calibrated geometry parameters given by the ML-BCD hybrid method are very low compared to BCD.

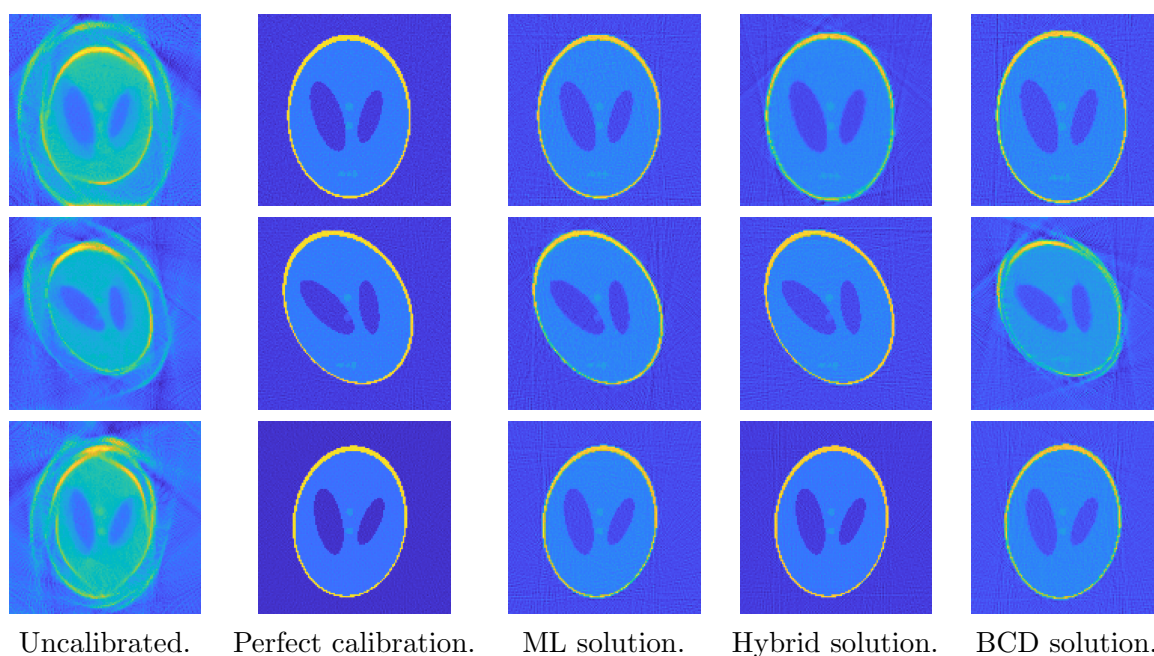


Figure 6.14: Phantom reconstructions of 3 testing samples (1 test sample per row). 1st column: solutions using theoretical parameters (i.e., uncalibrated). 2nd column: solutions using true parameters (i.e. perfect calibration). 3rd column: solutions using ML-predicted parameters (Algorithm 13). 4th column: solutions using ML-BCD hybrid method with 10 BCD iterations (Algorithm 15). 5th column: solutions using 10 BCD iterations (Algorithm 12).

In Figure 6.14, we display solutions of 3 test samples using different methods:

ML-only (Algorithm 13), ML-BCD hybrid (Algorithm 15) and BCD-only (Algorithm 12) (solutions at the last BCD iteration are shown). As a point of reference, we also present the solutions obtained using the true parameters to showcase the best solution possible. The 3 test samples are randomly picked from the 100 testing samples, and we see that for 2 out of 3, the BCD refinement step of the hybrid method has improved solution accuracy over using ML alone, and for these 2 samples, the hybrid solutions have significantly higher accuracy compared to the BCD solutions. For the first test sample, however, the BCD-refinement step of the hybrid method seems to not have improved solution accuracy. This is also expected because we have observed in the histograms of Figure 6.13 that in some cases, the hybrid method may not work so well – but it is evident that ML-BCD hybrid is more advantageous over BCD on average in the tasks of parameter calibration and image reconstruction.

## Chapter 7

# Conclusions and Future Work

In this thesis, we have introduced novel low-rank solvers for large scale linear systems. Firstly, we introduced a new Kronecker product summation based technique for approximating the TSVD for a large matrix  $\mathbf{A}$ . The TSVD can be thought of as a low-rank approximation to the original operator  $\mathbf{A}$ , and by filtering out the small singular values, it has a regularizing effect on the solution  $\mathbf{x}$  to the linear system (1.1). By exploiting a reordering technique, we can effectively utilize more terms in the Kronecker summation for  $\mathbf{A}$ , and produce approximations of higher quality than existing Kronecker product based approximation methods.

Next, we explored the nuclear norm regularized problem (4.2) and derived new solvers, based on Krylov subspace methods, for the computation of approximate low-rank solutions to large-scale linear systems of equations. The starting point of our derivations was an IRN approach to the  $\text{NNR}p$  problem (4.3). In this way, the original problem (4.3) is reduced to the solution of a sequence of quadratic problems,

where an appropriate smoothed linear transformation is introduced to approximate the nondifferentiable nuclear norm regularization term. Our new methods make smart use of Kronecker product properties to reformulate each quadratic problem in the IRN sequence as a Tikhonov-regularized problem in standard form. We developed both Krylov methods with fixed “preconditioners” within an inner-outer iteration scheme, and Krylov methods with flexible iteration-dependent “preconditioners” within a single iteration scheme. Some of these methods can be used in a hybrid framework, so that the Tikhonov regularization parameter can be efficiently, effectively, and adaptively chosen. These new solvers are shown to perform exceptionally well on the test problems such as image deblurring and inpainting, and they give reconstructions of significantly improved quality over existing methods.

In the application of computed tomography, where the forward model  $\mathbf{A}$  is dependent on unknown geometry parameters  $\mathbf{p}$ , model calibration is required at the time of image reconstruction. A block coordinate descent type algorithm can be applied, where the image  $\mathbf{x}$  and unknown parameters  $\mathbf{p}$  are alternatively minimized in each step. For the minimization with respect to  $\mathbf{x}$ , our new low-rank Krylov methods also show improved convergence result and reconstruction quality compared to standard Krylov methods with classical regularization such as  $\ell_2$ . Furthermore, we also proposed an enhanced, hybrid BCD framework that incorporates machine learning for solving the model calibration problem. By training a machine learning model (multi-output linear regression model) that maps the observation  $\mathbf{b}$  to  $\mathbf{p}$  and feeding the predicted parameters as initial guesses to BCD, we are able to improve the accuracy of both the calibrated geometry parameters and the reconstructed image.

Interesting future directions include, for example, extending present low-rank Krylov methods to handle cases where the solution of (4.1) is low-rank but rectangular, i.e.,  $\text{vec}^{-1}(\mathbf{x}) = \mathbf{X} \in \mathbb{R}^{m \times n}$  with  $m \neq n$ . Also, while a solid theoretical justification is provided for IRN-LSQR-NNR $p$  and IRN-GMRES-NNR $p$ , the same is not true for FGMRES-NNR $p$  and FLSQR-NNR $p$ : further analysis will be needed to deeply understand the regularization properties of these flexible solvers. In the context of model calibration, we also see great potential of flexible Krylov subspace methods, since we may be able to design iteration-dependent “preconditioners” so that they contain updated information about both the approximate solution  $\mathbf{x}$  and model parameters  $\mathbf{p}$ . State-of-the-art machine learning techniques could be accommodated to learn the necessary preconditioners that not only enforce regularization properties on the solutions, but also improve model accuracy. Inverse problems can be found in a wide range scientific fields, and is an active, broad and fast-evolving research area. Although this thesis only scratched the surface of what can be done with inverse problems, we hope the new methods proposed here will inspire and promote future research in many other interesting topics.

# Bibliography

- [1] B. M. AFKHAM, J. CHUNG, AND M. CHUNG, *Learning Regularization Parameters of Inverse Problems via Deep Neural Networks*, 2021.
- [2] A. P. AUSTIN, Z. W. DI, S. LEYFFER, AND S. M. WILD, *Simultaneous Sensing Error Recovery and Tomographic Inversion Using an Optimization-Based Approach*, SIAM Journal on Scientific Computing, 41 (2019), pp. B497–B521.
- [3] M. BECKERS, T. SENKBEIL, T. GORNIK, K. GIEWEKEMEYER, T. SALDITT, AND A. ROSENHAHN, *Drift Correction in Ptychographic Diffractive Imaging*, Ultramicroscopy, 126 (2013), pp. 44 – 47.
- [4] M. BELGE, M. E. KILMER, AND E. L. MILLER, *Wavelet Domain Image Restoration with Adaptive Edge-Preserving Regularization*, IEEE Trans. Image Process., 9 (2000), pp. 597–608.
- [5] J. M. BIOUCAS-DIAS AND M. A. T. FIGUEIREDO, *A New TwIST: Two-Step Iterative Shrinkage/Thresholding Algorithms for Image Restoration*, IEEE Transactions on Image Processing, 16 (2007), pp. 2992–3004.

- [6] J. D. BLANCHARD, J. TANNER, AND K. WEI, *CGIHT: Conjugate Gradient Iterative Hard Thresholding for Compressed Sensing and Matrix Completion*, *Information and Inference: A Journal of the IMA*, 4 (2015), pp. 289–327.
- [7] J. CAI, E. CANDÈS, AND Z. SHEN, *A Singular Value Thresholding Algorithm for Matrix Completion*, *SIAM Journal on Optimization*, 20 (2010), pp. 1956–1982.
- [8] D. CALVETTI, S. MORIGI, L. REICHEL, AND F. SGALLARI, *Tikhonov Regularization and the L-Curve for Large Discrete Ill-Posed Problems*, *Journal of Computational and Applied Mathematics*, 123 (2000), pp. 423–446. *Numerical Analysis 2000. Vol. III: Linear Algebra*.
- [9] J. CANT, A. SNOECKX, G. BEHIELS, P. M. PARIZEL, AND J. SIJBERS, *Can Portable Tomosynthesis Improve the Diagnostic Value of Bedside Chest X-ray in the Intensive Care Unit? A Proof of Concept Study*, *European radiology experimental*, (2017), p. 20.
- [10] M. CAO, A. TAKAOKA, H. ZHANG, AND R. NISHI, *An Automatic Method of Detecting and Tracking Fiducial Markers for Alignment in Electron Tomography*, *Microscopy*, 60 (2011), pp. 39–46.
- [11] J. CHUNG AND S. GAZZOLA, *Flexible Krylov Methods for  $\ell_p$  Regularization*, *SIAM Journal on Scientific Computing*, 41 (2019), pp. S149–S171.



- [12] J. CHUNG, J. G. NAGY, AND D. M. O'LEARY, *A weighted-gcv method for lanczos-hybrid regularization.*, *Electronic Transactions on Numerical Analysis*, 28 (2007), pp. 149–167.
- [13] R. D. FIERRO AND J. R. BUNCH, *Bounding the Subspaces from Rank Revealing Two-Sided Orthogonal Decompositions*, *SIAM J. Matrix Anal. Appl.*, 16 (1995), pp. 743–759.
- [14] M. FORNASIER, H. RAUHUT, AND R. WARD, *Low-Rank Matrix Recovery via Iteratively Reweighted Least Squares Minimization*, *SIAM Journal on Optimization*, 21 (2011), pp. 1614–1640.
- [15] C. GARVEY, *Truncated Singular Value Decomposition Approximation for Structured Matrices via Kronecker Product Summation Decomposition*, *Theses and Dissertations*, Emory University, (2018).
- [16] C. GARVEY, C. MENG, AND J. G. NAGY, *Singular Value Decomposition Approximation via Kronecker Summations for Imaging Applications*, *SIAM Journal on Matrix Analysis and Applications*, 39 (2018), pp. 1836–1857.
- [17] S. GAZZOLA, P. C. HANSEN, AND J. G. NAGY, *IR Tools: A MATLAB Package of Iterative Regularization Methods and Large-Scale Test Problems*, *Numerical Algorithms*, 81 (2019), pp. 773–811.
- [18] S. GAZZOLA, C. MENG, AND J. G. NAGY, *Krylov Methods for Low-Rank Regularization*, *SIAM Journal on Matrix Analysis and Applications*, 41 (2020), pp. 1477–1504.

- [19] S. GAZZOLA AND J. G. NAGY, *Generalized Arnoldi-Tikhonov Method for Sparse Reconstruction*, SIAM J. Sci. Comput., 36 (2014), pp. 225–247.
- [20] S. GAZZOLA AND P. NOVATI, *Automatic Parameter Setting for Arnoldi-Tikhonov Methods*, J. Comput. Appl. Math., 256 (2014), pp. 180–195.
- [21] S. GAZZOLA, P. NOVATI, AND M. R. RUSSO, *On Krylov Projection Methods and Tikhonov Regularization*, Electronic Transactions on Numerical Analysis, 44 (2014), pp. 82–123.
- [22] S. GAZZOLA AND M. SABATÉ LANDMAN, *Flexible GMRES for Total Variation Regularization*, BIT Numerical Mathematics, 59 (2019), pp. 721–746.
- [23] D. GOLDFARB AND S. MA, *Convergence of Fixed-Point Continuation Algorithms for Matrix Rank Minimization*, Foundations of Computational Mathematics, 11 (2011), pp. 183–210.
- [24] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, John Hopkins University Press, 4 ed., 2013.
- [25] I. F. GORODNITSKY AND B. D. RAO, *A New Iterative Weighted Norm Minimization Algorithm and its Applications*, in [1992] IEEE Sixth SP Workshop on Statistical Signal and Array Processing, Oct 1992, pp. 412–415.
- [26] D. GÜRSOY, Y. P. HONG, K. HE, K. HUJSAK, S. YOO, S. CHEN, Y. LI, M. GE, L. M. MILLER, Y. S. CHU, V. DE ANDRADE, K. HE, O. COSSAIRT, A. K. KATSAGGELOS, AND C. JACOBSEN, *Rapid Alignment of Nanotomogra-*

- phy Data Using Joint Iterative Reconstruction and Reprojection*, Scientific Reports, 7 (2017), p. 11818.
- [27] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions*, SIAM Review, 53 (2011), pp. 217–288.
- [28] M. HANKE AND P. C. HANSEN, *Regularization Methods for Large-Scale Problems*, Surv. Math. Ind., 3 (1993), pp. 253–315.
- [29] P. C. HANSEN, *REGULARIZATION TOOLS: A Matlab Package for Analysis and Solution of Discrete Ill-Posed Problems*, Numerical Algorithms, 6 (1994), pp. 1–35.
- [30] —, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, Society for Industrial and Applied Mathematics, USA, 1999.
- [31] —, *Discrete Inverse Problems: Insight and Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2010.
- [32] P. C. HANSEN AND T. K. JENSEN, *Smoothing-Norm Preconditioning for Regularizing Minimum-Residual Methods*, SIAM J. Matrix Anal. Appl., 29 (2006), pp. 1–14.
- [33] P. C. HANSEN, J. G. NAGY, AND D. P. O’LEARY, *Deblurring Images: Matrices, Spectra, and Filtering*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2006.

- [34] P. C. HANSEN AND D. P. OLEARY, *The Use of the L-Curve in the Regularization of Discrete Ill-Posed Problems*, SIAM Journal on Scientific Computing, 14 (1993), pp. 1487–1503.
- [35] R. HUANG, R. DIAO, Y. LI, J. SANCHEZ-GASCA, Z. HUANG, B. THOMAS, P. ETINGOV, S. KINCIC, S. WANG, R. FAN, G. MATTHEWS, D. KOSTEREV, S. YANG, AND J. ZHAO, *Calibrating Parameters of Power System Stability Models Using Advanced Ensemble Kalman Filter*, IEEE Transactions on Power Systems, 33 (2018), pp. 2895–2905.
- [36] X. HUANG, S. M. WILD, AND Z. W. DI, *Calibrating Sensing Drift in Tomographic Inversion*, in 2019 IEEE International Conference on Image Processing (ICIP), IEEE, 2019, pp. 1267–1271.
- [37] A. C. HURST, T. B. EDO, T. WALTHER, F. SWEENEY, AND J. M. RODENBURG, *Probe Position Recovery for Ptychographical Imaging*, Journal of Physics: Conference Series, 241 (2010), p. 012004.
- [38] A. JACOBI, M. CHUNG, A. BERNHEIM, AND C. EBER, *Portable Chest X-ray in Coronavirus Disease-19 (Covid-19): A Pictorial Review*, Clinical imaging, 64 (2020), pp. 35–42.
- [39] A. K. JAIN, *Fundamentals of Digital Image Processing*, Prentice-Hall, Inc., USA, 1989.
- [40] T. K. JENSEN AND P. C. HANSEN, *Iterative Regularization with Minimum-Residual Methods*, BIT, 47 (2007), pp. 103–120.

- [41] J. KAMM AND J. G. NAGY, *Kronecker Product and SVD Approximations in Image Restoration*, Linear Algebra and its Applications, 284 (1998), pp. 177–192.
- [42] C. T. KELLEY, *Users' Guide for imfil Version 1.0*, 2011.
- [43] R. H. KESHAVAN, A. MONTANARI, AND S. OH, *Matrix Completion from Noisy Entries*, J. Mach. Learn. Res., 11 (2010), pp. 2057–2078.
- [44] M. E. KILMER AND D. O'LEARY, *Choosing Regularization Parameters in Iterative Methods for Ill-Posed Problems*, SIAM J. Matrix Anal. Appl., 22 (2001), pp. 1204–1221.
- [45] D. KRESSNER AND C. TOBLER, *Low-Rank Tensor Krylov Subspace Methods for Parameterized Linear Systems*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 1288–1316.
- [46] L. LANDWEBER, *An Iteration Formula for Fredholm Integral Equations of the First Kind*, American Journal of Mathematics, 73 (1951), pp. 615–624.
- [47] R. M. LARSEN, *Lanczos Bidiagonalization with Partial Reorthogonalization*, Department of Computer Science, Aarhus University, Technical report, (1998).
- [48] K. LEE AND H. C. ELMAN, *A Preconditioned Low-Rank Projection Method with a Rank-Reduction Scheme for Stochastic Partial Differential Equations*, SIAM J. Sci. Comput., 9 (2017), pp. 828–850.
- [49] Z. LI, Q. XU, AND Y. WEI, *A Note on Stable Perturbations of Moore-Penrose Inverses*, Numer. Linear Algebra Appl., 20 (2011), pp. 18–26.

- [50] S. MA, D. GOLDFARB, AND L. CHEN, *Fixed Point and Bregman Iterative Methods for Matrix Rank Minimization*, Mathematical Programming, 128 (2011), pp. 321–353.
- [51] K. MOHAN AND M. FAZEL, *Iterative Reweighted Algorithms for Matrix Rank Minimization*, J. Mach. Learn. Res., 13 (2012), pp. 3441–3473.
- [52] J. G. NAGY, *Decomposition of Block Toeplitz Matrices into a Sum of Kronecker Products with Applications in Image Restoration*, tech. rep., Southern Methodist University, 1996.
- [53] D. O’LEARY AND B. RUST, *Variable Projection for Nonlinear Least Squares Problems*, Computational Optimization and Applications, 54 (2013), pp. 579–593.
- [54] C. C. PAIGE AND M. A. SAUNDERS, *LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares*, ACM Trans. Math. Softw., 8 (1982), pp. 43–71.
- [55] V. PATHURI-BHUVANA, S. SCHUSTER, AND A. OCH, *Joint calibration and tomography based on separable least squares approach with constraints on linear and non-linear parameters*, in 2020 28th European Signal Processing Conference (EUSIPCO), 2021, pp. 1931–1935.
- [56] J. RADON, *Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten*, Akad. Wiss., 69 (1917), pp. 262–277.

- [57] B. RECHT, M. FAZEL, AND P. PARRILO, *Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization*, SIAM Review, 52 (2010), pp. 471–501.
- [58] N. A. RIIS, Y. DONG, AND P. C. HANSEN, *Computed Tomography Reconstruction with Uncertain View Angles by Iteratively Updated Model Discrepancy*, Journal of Mathematical Imaging and Vision, 63 (2021), pp. 133–143.
- [59] N. A. B. RIIS, Y. DONG, AND P. C. HANSEN, *Computed Tomography with View Angle Estimation Using Uncertainty Quantification*, Inverse Problems, 37 (2021), p. 065007.
- [60] R. T. ROCKAFELLAR, *Convex Analysis*, vol. 28, Princeton university press, 1970.
- [61] P. RODRÍGUEZ AND B. WOHLBERG, *An Efficient Algorithm for Sparse Representations with  $\ell_p$  Data Fidelity Term*, in Proceedings of 4th IEEE Andean Technical Conference (ANDESCON), Cusco, Perú, Oct 2008.
- [62] M. C. ROGGEMANN AND B. WELSH, *Imaging Through Turbulence*, CRC Press, Boca Raton, FL, 1996.
- [63] Y. SAAD, *A Flexible Inner-Outer Preconditioned GMRES Algorithm*, SIAM Journal on Scientific Computing, 14 (1993), pp. 461–469.
- [64] Y. SAAD, *Iterative Methods*, Society for Industrial and Applied Mathematics, 2 ed., 2003.

- [65] Y. SAAD AND M. H. SCHULTZ, *GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems*, Siam Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
- [66] V. SIMONCINI AND D. B. SZYLD, *Recent Computational Developments in Krylov Subspace Methods for Linear Systems*, Numerical Linear Algebra with Applications, 14 (2007), pp. 1–59.
- [67] M. STOLL AND T. BREITEN, *A Low-Rank in Time Approach to PDE-Constrained Optimization*, SIAM J. Sci. Comput., 27 (2015), pp. B1–B29.
- [68] C. F. VAN LOAN AND N. PITSIANIS, *Approximation with Kronecker Products*, in Linear Algebra for Large Scale and Real-Time Applications, Springer, 1993, pp. 293–314.
- [69] C. R. VOGEL, *Computational Methods for Inverse Problems*, Society for Industrial and Applied Mathematics, 2002.
- [70] P.-Å. WEDIN, *Perturbation Theory for Pseudoinverses*, BIT, 13 (1973), pp. 217–232.
- [71] Y. WU, *Model Parameter Calibration in Power Systems*, Graduate College Dissertations and Theses, University of Vermont, (2020).