

RELAXED REGULARIZATION FOR LINEAR INVERSE PROBLEMS*

NICK LUIKEN[†] AND TRISTAN VAN LEEUWEN[†]

Abstract. We consider regularized least-squares problems of the form $\min_x \frac{1}{2} \|Ax - b\|_2^2 + \mathcal{R}(Lx)$. Recently, Zheng et al. [*IEEE Access*, 7 (2019), pp. 1404–1423], proposed an algorithm called Sparse Relaxed Regularized Regression (SR3) that employs a splitting strategy by introducing an auxiliary variable y and solves $\min_{x,y} \frac{1}{2} \|Ax - b\|_2^2 + \frac{\kappa}{2} \|Lx - y\|_2^2 + \mathcal{R}(x)$. By minimizing out the variable x , we obtain an equivalent optimization problem $\min_y \frac{1}{2} \|F_\kappa y - g_\kappa\|_2^2 + \mathcal{R}(y)$. In our work, we view the SR3 method as a way to approximately solve the regularized problem. We analyze the conditioning of the relaxed problem in general and give an expression for the SVD of F_κ as a function of κ . Furthermore, we relate the Pareto curve of the original problem to the relaxed problem, and we quantify the error incurred by relaxation in terms of κ . Finally, we propose an efficient iterative method for solving the relaxed problem with inexact inner iterations. Numerical examples illustrate the approach.

Key words. inverse problems, optimization, machine learning, regularization, sparsity, total variation

AMS subject classifications. 65F22, 65F10, 15A18, 15A29

DOI. 10.1137/20M1348091

1. Introduction. Inverse problems are problems where a certain quantity of interest has to be determined from indirect measurements. In medicine, well-known examples include magnetic resonance imaging (MRI) [44], CT [28], and ultrasound imaging [6] where the objective is to obtain images of the interior of the human body. In the geosciences, inverse problems arise in seismic exploration and seismology [42], where the interest lies in exploring the elastic properties of the different layers of our planet. Other examples include tomography [3, 5, 33], radar imaging [7], remote sensing [34, 39], astrophysics [38], and more recently, machine learning [19].

Inverse problems are challenging for a number of reasons. There may be limited data available, or the data may be corrupted by noise. The data sets are generally very large, and the underlying model is generally not well defined for retrieving the quantity of interest. Therefore, inverse problems often have to be *regularized*, meaning prior information has to be added. They can be posed in the following way:

$$(1.1) \quad \min_x \frac{1}{2} \|Ax - b\|_2^2 + \mathcal{R}(Lx),$$

where $A \in \mathbb{R}^{m \times n}$ is the linear forward operator, $\mathcal{R}(\cdot)$ is the regularization term, and $L \in \mathbb{R}^{p \times n}$ the regularization operator. The latter two encode the prior information about x . In our work, we focus on $\mathcal{R}(\cdot) = \lambda \|\cdot\|_p^p$, or, equivalently, $\mathcal{R}(\cdot) = \delta_{\|\cdot\|_p \leq \tau}(\cdot)$, which is the indicator function of the set $\|\cdot\|_p \leq \tau$. By equivalent we mean that for every τ there is a λ such that the solutions of the two problems coincide [2]. A direct solution to the problem above is generally not possible, either because a closed-form

*Received by the editors June 26, 2020; accepted for publication (in revised form) March 8, 2021; published electronically May 20, 2021.

<https://doi.org/10.1137/20M1348091>

Funding: The work of the first author was supported by the Delphi consortium. The work of the second author was supported by the Netherlands Organization for Scientific Research (NWO) as part of research programme 613.009.032.

[†]Mathematical Institute, Utrecht University, Utrecht, 3584 CD, The Netherlands (n.a.luiken@uu.nl, t.vanleeuwen@uu.nl).

solution does not exist or because evaluating the direct solution is too computationally expensive. Therefore, we have to resort to iterative methods to solve the problem, with most algorithms being designed for specific choices of p and L .

Traditionally, $p = 2$, called Tikhonov regularization, is a popular choice, because the objective function is differentiable and allows for a closed-form expression of the solution of (1.1) in terms of A , L , and λ . For this class of problems, Krylov based algorithms have been proven very effective [9, 10, 17, 18, 24, 29, 30, 31, 32, 45]. These methods generally exploit the fact that a closed-form solution exists by constructing a low dimensional subspace from which an approximate solution is extracted.

The choice $p = 1$ has gained popularity in recent years because it gives sparse solutions while still yielding a convex objective. Sparsity is important in a number of applications, like compressed sensing [11], seismic imaging [27], image restoration [36], and tomography [25]. However, the objective is no longer differentiable and the aforementioned Krylov methods do not apply. If $L = I$, a proximal gradient method (sometimes referred to as iterative soft thresholding—ISTA) [12] can be applied, iteratively updating the solution via

$$x_{k+1} = \text{prox}_{\alpha\lambda\|\cdot\|_1} (x_k - \alpha A^T(Ax_k - b)),$$

where $\alpha \in (0, \|A\|_2^{-2})$ is the stepsize and the proximal operator is the soft thresholding operator, which can be efficiently evaluated. Generally, ISTA achieves a sublinear rate of convergence of $\mathcal{O}(1/k)$ (unless $m \geq n$ and A has full rank, in which case we have a linear rate of convergence). FISTA (fast iterative soft thresholding algorithm) [4] is a faster version of ISTA that generally achieves a sublinear rate of $\mathcal{O}(1/k^2)$.

If $L = I$ the optimization problem is said to be in standard-form and for any other L the algorithm is in general form. If L is full-rank and has no nullspace, the optimization problem can be put into standard-form via the change of variables $y = Lx$. Instead of the matrix A , we get AL^\dagger . In such cases we can apply the (F)ISTA method directly at the expense of having to evaluate L^\dagger . In some applications, we have $L^\dagger = L^T$ (e.g., when L is a tight frame). If L has a nontrivial nullspace the algorithm can still be put in standard-form by the standard-form transformation [15, 25], but this is nontrivial, because the nullspace has to be accounted for.

If $L \neq I$ and we cannot easily transform the problem to standard form, the proximal operator is no longer easy to evaluate in general and FISTA may no longer be attractive. An example of this class of problems is total variation (TV) regularization, where L is the discretization of the gradient, which gives blocky solutions. A popular algorithm for this class of problems is the alternating direction method of multipliers, ADMM [8]. ADMM solves (1.1) by forming the *augmented Lagrangian*

$$\min_{x,y} \max_z \frac{1}{2} \|Ax - b\|_2^2 + \lambda \|y\|_p^p + z^T (Lx - y) + \frac{\rho}{2} \|Lx - y\|_2^2$$

and alternately minimizing over the variables x and y , and the Lagrange multiplier z . The strength of ADMM is that it can closely approximate the solution of any convex sparse optimization problem. However, convergence can be slow [8].

If $p < 1$, the emphasis on sparsity of the solution is stronger than for the case $p = 1$. However, the objective function is no longer convex which makes it more difficult to solve.

Recently, a unifying algorithm was proposed that allows the efficient approximation of the solution of any problem of the form (1.1), called Sparse Relaxed Regularized Regression (SR3) [43]. This algorithm makes use of a splitting strategy by introducing an auxiliary variable y and yields

$$(1.2) \quad \min_{x,y} \frac{1}{2} \|Ax - b\|_2^2 + \frac{\kappa}{2} \|Lx - y\|_2^2 + \mathcal{R}(y).$$

By minimizing out x , we obtain a new optimization problem of the form:

$$(1.3) \quad \bar{y}_\kappa = \arg \min_y \frac{1}{2} \|F_\kappa y - g_\kappa\|_2^2 + \mathcal{R}(y),$$

where

$$F_\kappa = \begin{pmatrix} \kappa^{1/2} (I - \kappa L H_\kappa^{-1} L^T) \\ \kappa A H_\kappa^{-1} L^T \end{pmatrix} \text{ and} \\ g_\kappa = \begin{pmatrix} \kappa^{1/2} L H_\kappa^{-1} A^T b \\ b - A H_\kappa^{-1} A^T b \end{pmatrix}, H_\kappa = A^T A + \kappa L^T L.$$

The solution to (1.2) is then given by

$$(1.4) \quad \bar{x}_\kappa = H_\kappa^{-1} (\kappa L^T \bar{y}_\kappa + A^T b).$$

This solution is then used as an approximation of the solution of (1.1). In [43] the particular case with $L^T L = I$ is analyzed. Using the SVD of A , the singular values of F_κ were calculated, showing a relation between the condition number of F_κ and A depending on κ . In short, the result shows that a small κ improves the conditioning of F_κ and as $\kappa \rightarrow \infty$ the condition numbers are the same, because the original optimization problem is obtained.

For the implementation of SR3, it is not necessary to form the operator F_κ , as was shown in [43]. The authors propose the following algorithm for solving the relaxed problem

$$(1.5) \quad x_{k+1} \leftarrow (A^T A + \kappa L^T L)^{-1} (A^T b + \kappa L^T y_k),$$

$$(1.6) \quad y_{k+1} \leftarrow \text{prox}_{\alpha \mathcal{R}} (y_k - \alpha \kappa (y_k - L x_{k+1})),$$

which for the particular choice $\alpha = 1/\kappa$ simplifies to

$$(1.7) \quad x_{k+1} \leftarrow (A^T A + \kappa L^T L)^{-1} (A^T b + \kappa L^T y_k),$$

$$(1.8) \quad y_{k+1} \leftarrow \text{prox}_{1/\kappa \mathcal{R}} (L x_{k+1}).$$

This method has several advantages when applied to solving inverse problems that we highlight in the examples below.

1.1. Motivating examples. Below we show some typical examples encountered in various areas of science to which SR3 can be applied. The problems we tackle are of the form

$$(1.9) \quad \min_x \frac{1}{2} \|Ax - b\|_2^2 \quad \text{s.t.} \quad \|Lx\|_1 \leq \tau.$$

The main tasks are to solve this for a given value of τ and to find an appropriate value of τ . The latter is achieved by picking the corner of the Pareto curve (sometimes called the L-curve) $\phi(\tau) = \min_{\|x\|_p \leq \tau} \|Ax - b\|_2$. Comparing a proximal gradient method to SR3, we show the residual as a function of τ , the optimal reconstruction, and the convergence history in terms of the primal-dual gap. These examples show two favorable aspects of SR3 over the conventional proximal gradient method: i) SR3 converges (much) faster for any fixed value of τ and ii) the corners of both Pareto-curves coincide, allowing us to effectively use SR3 to estimate τ .

Spiky deconvolution ($m = n$, $L = I$). Consider a deconvolution problem where A is a Toeplitz matrix that convolves the input with a bandlimited function,

$$a_{ij} = w(t_i - t_j),$$

where $w(t) = (1 - (t/\sigma)^2)e^{-(t/\sigma)^2}$ and $t_i = i \cdot h$. We take $n = 101$, $h = 1/n$, and $\sigma = 0.05$. The results are shown in Figure 1.1.

Compressed sensing ($m < n$, $L = I$). Here, the goal is to recover a sparse signal from compressive samples. The forward operator is a random matrix with independent and identically distributed (i.i.d.) normally distributed entries. We take $n = 101$ and $m = 20$. The results are shown in Figure 1.2.

TV ($m = n$, $L = D$). Consider a deconvolution problem where A is a Toeplitz matrix that convolves the input with a bandlimited function,

$$a_{ij} = w(t_i - t_j),$$

where $w(t) = e^{-(t/\sigma)^2}$ and $t_i = i \cdot h$. L is a finite-difference discretization of the first-order derivative with Neumann boundary conditions. We take $n = 101$, $h = 1/n$, and $\sigma = 0.05$. The results are shown in Figure 1.3.




FIG. 1.1. *Spiky deconvolution example.* The left figure shows the Pareto curve, the middle figure shows the solution, and the right figure shows the primal-dual gap as a function of the number of iterations. The gray line in the middle figure shows the minimum norm solution.




FIG. 1.2. *Compressed sensing.* A signal is reconstructed from very few samples, which requires sparse regularization. The left figure shows the Pareto curve, the middle figure shows the solution, and the right figure shows the primal-dual gap as a function of the number of iterations. The gray line in the middle figure shows the minimum norm solution.




FIG. 1.3. *TV example. Here, the solution has a blocky structure. The left figure shows the Pareto curve, the middle figure shows the solution, and the right figure shows the primal-dual gap as a function of the number of iterations. The gray line in the middle figure shows the minimum norm solution.*

1.2. Contributions. In this paper we set out to further analyze the SR3 method proposed in [43] and analyze in detail the observations made in the above examples. Our contributions are the following.

Conditioning of F_κ for general L . We extend the analysis of [43] and derive the SVD of F_κ for general L . We show how the singular values and the condition number of F_κ are related to the generalized singular values of (A, L) . As a by-product, we show that SR3 implicitly makes a standard-form transformation [15] of (1.1).

Approximation of the Pareto curve. We show that that the Pareto curve corresponding to the relaxed problem (1.2) always underestimates the Pareto curve of the original problem (1.1) and that the error is of order $\mathcal{O}(\kappa^{-2})$. A by-product of this result is a better understanding of the Pareto curve for general p and an intuitive explanation of the observation that the corners of the relaxed original Pareto curves coincide.

Inexact solves. We propose an inexact inner-outer iterative version of the SR3 algorithm where the regularized least-squares problem (1.7) is solved approximately using a Krylov-subspace method. In particular, we propose an automated adaptive stopping criterion for the inner iterations.

1.3. Outline. In section 2 we analyze the operator F_κ . We derive the SVD of F_κ and analyze the limiting cases $\kappa \rightarrow \infty$ and $\kappa \rightarrow 0$. Our main results are a characterization of the singular values of F_κ and showing that SR3 implicitly applies a standard-form transformation. In section 3, we relate the Pareto curve of SR3 to the Pareto curve of the original problem and derive an error bound in terms of κ . Next, section 4 is concerned with the implementation of SR3. We propose two ingredients that make SR3 suitable for large-scale applications. In section 5, we conduct our numerical experiments and verify the theoretical results from section 2. Moreover, we numerically investigate the influence of κ on the convergence rate. Finally, in section 6, we draw our conclusions.

2. Analysis of SR3. In this section we analyze some of the properties of the operator F_κ . We will characterize the singular values of F_κ for general L and analyze the limits $\kappa \rightarrow 0$ and $\kappa \rightarrow \infty$. First, we will treat some preliminaries needed for understanding what happens in the limit $\kappa \rightarrow \infty$.

2.1. The generalized SVD. The central tool in our analysis is the generalized SVD (GSVD) of (A, L) . The definition of the GSVD depends on the size of the

matrices and the dimensions of the matrices relative to each other. We use the definitions for the case $A \in \mathbb{R}^{m \times n}$ and $L \in \mathbb{R}^{p \times n}$ where $m \geq n$, $p < n$ or $m < n$, $p > n$ because this corresponds to the examples we use in our experiments.

DEFINITION 2.1 (GSVD). *Let $A \in \mathbb{R}^{m \times n}$ and $L \in \mathbb{R}^{p \times n}$. The GSVD of (A, L) is given by $A = U\Sigma X$, $L = V\Gamma X$, where*

$$\Sigma = \begin{bmatrix} \Sigma_p & 0 \\ 0 & I_{n-p} \\ 0 & 0 \end{bmatrix}, \quad \Gamma = \begin{bmatrix} \Gamma_p & 0 \end{bmatrix} \quad \text{for } m \geq n, p \leq n,$$

and

$$\Sigma = \begin{bmatrix} 0 & \Sigma_m \end{bmatrix}, \quad \Gamma = \begin{bmatrix} I_{n-m} & 0 \\ 0 & \Gamma_m \\ 0 & 0 \end{bmatrix} \quad \text{for } m < n, p > n.$$

The matrices Σ_r and Γ_r (where $r = p$ or $r = m$) are $r \times r$ diagonal matrices satisfying $\Sigma_r^T \Sigma_r + \Gamma_r^T \Gamma_r = I_r$, X is invertible and U and V are orthonormal. Moreover, we have the following ordering of the diagonal elements σ_i of Σ and γ_i of Γ :

$$\begin{aligned} 0 &\leq \gamma_r \leq \cdots \leq \gamma_1 \leq 1, \\ 0 &\leq \sigma_1 \leq \cdots \leq \sigma_r \leq 1. \end{aligned}$$

The decomposition of A and L in the GSVD share similar properties to the SVD. The number of nonzero entries of Σ and Γ give the rank of A and L , respectively. If r_A is the rank of A and r_L is the rank of L then the last $r - r_A$ columns, corresponding to Σ_r , of U form a basis for the range of A and the first r_L columns, corresponding to Γ_r , of V form a basis for the range of L . The first $r - r_A$ columns, corresponding to Σ_r , of X^{-1} form a basis for the nullspace of A and the last $r - r_L$ columns, corresponding to Γ_r , of X^{-1} form a basis for the nullspace of L .

2.2. Standard-form transformation. The standard-form transformation (see, e.g., [15, 22]) makes a substitution $y = Lx$ such that $x = x_{\mathcal{M}} + x_{\mathcal{N}}$, where

$$(2.1) \quad \bar{x}_{\mathcal{M}} = L_A^\dagger \bar{y}, \quad \bar{y} = \arg \min_y \frac{1}{2} \|AL_A^\dagger y - b\|_2^2 + \mathcal{R}(y), \quad L_A^\dagger = (I - (A(I - L^\dagger L))^\dagger A) L^\dagger$$

and

$$(2.2) \quad \bar{x}_{\mathcal{N}} = (A(I - L^\dagger L))^\dagger b.$$

The operator L_A^\dagger is called the A -weighted pseudo-inverse. The transformation splits the solution into two parts: one part in the range of L , $L_A^\dagger y$, and one part in the nullspace of L , $x_{\mathcal{N}}$. The operator L_A^\dagger makes the two parts A -orthogonal. The parts $L_A^\dagger y$ and $x_{\mathcal{N}}$ are then obtained by two independent optimization problems. If L is invertible $L_A^\dagger = L^{-1}$ and if $p > n$ and L has full rank we have $L_A^\dagger = L^\dagger$. Hence, if $L^T L = I$, the standard-form is achieved by simply applying L^T .

In terms of the GSVD of (A, L) , the standard-form transformation has a much simpler form. The operator L_A^\dagger can be written in terms of the GSVD as

$$L_A^\dagger = X^{-1} \Gamma^\dagger V^T,$$

and hence (2.1) can be written as

$$(2.3) \quad \bar{x}_{\mathcal{M}} = X^{-1}\Gamma^\dagger V^T \bar{y}, \quad \bar{y} = \arg \min_y \frac{1}{2} \|U\Sigma\Gamma^\dagger V^T y - b\|_2^2 + \mathcal{R}(y).$$

Similarly, (2.2) can be written in terms of the GSVD as

$$(2.4) \quad \bar{x}_{\mathcal{N}} = X^{-1} \begin{bmatrix} 0 & 0 \\ 0 & I_{p-r_L} \end{bmatrix} U^T b.$$

2.3. The SVD of F_κ . In this section we derive the SVD of F_κ in terms of the GSVD of (A, L) .

THEOREM 2.2. *Let $F_\kappa = Y\Lambda Z^T$ be the SVD of F_κ . Let the GSVD of $\begin{bmatrix} A \\ L \end{bmatrix} = \begin{bmatrix} U\Sigma \\ V\Gamma \end{bmatrix} X$. Then*

$$Y = \begin{bmatrix} \kappa^{1/2} V \tilde{\Sigma}_{\kappa, I}^{1/2} & \kappa V \tilde{\Sigma}_{\kappa, I}^{-1/2} \Gamma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Sigma^T \\ \kappa U \Sigma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T \tilde{\Sigma}_{\kappa, I}^{-1/2} & -\kappa^{-1/2} U \tilde{\Sigma}_{m, \kappa}^{1/2} \end{bmatrix}$$

$$\Lambda = \begin{bmatrix} \tilde{\Sigma}_\kappa^{1/2} \\ 0 \end{bmatrix}$$

$$Z = V,$$

where $\tilde{\Sigma}_\kappa = \kappa (I_p - \kappa \Gamma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T)$, $\tilde{\Sigma}_{\kappa, m} = \begin{bmatrix} \tilde{\Sigma}_\kappa & 0 \\ 0 & I_{m-p} \end{bmatrix}$ if $m \geq n \geq p$ and $\tilde{\Sigma}_{\kappa, m} = \tilde{\Sigma}_{\kappa, I}$ if $m < n \leq p$, and the square root denotes the entry wise square root. If $p > n$ the diagonal matrix $\tilde{\Sigma}_\kappa$ will have zeros on the diagonal. We denote $\tilde{\Sigma}_{\kappa, I}$ to be the matrix $\tilde{\Sigma}_\kappa$ where the zeros have been replaced by ones.

Proof. Using the GSVD of (A, L) we have $H_\kappa^{-1} = X^{-1}(\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} X^{-T}$, and hence $LH_\kappa^{-1}L^T = V\Gamma(\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1}\Gamma^T V^T$. Given the fact that V is orthonormal and $\Gamma(\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1}\Gamma^T$ is a diagonal matrix the above expression is the SVD of $LH_\kappa^{-1}L^T$, and we obtain the expressions for Λ and Z . To obtain Y , we first partition $Y = \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix}$. We have

$$F_\kappa F_\kappa^T = Y\Lambda\Lambda^T Y^T$$

$$\iff \begin{bmatrix} \kappa (I - \kappa LH_\kappa^{-1}L^T)^2 & \kappa\sqrt{\kappa} (I - \kappa LH_\kappa^{-1}L^T) LH_\kappa^{-1}A^T \\ \kappa\sqrt{\kappa} AH_\kappa^{-1}L^T (I - \kappa LH_\kappa^{-1}L^T) & \kappa^2 AH_\kappa^{-1}LL^T H_\kappa^{-1}A^T \end{bmatrix}$$

$$= \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{21} & Y_{22} \end{bmatrix} \begin{bmatrix} \tilde{\Sigma}_\kappa & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} Y_{11}^T & Y_{21}^T \\ Y_{12}^T & Y_{22}^T \end{bmatrix} = \begin{bmatrix} Y_{11} \tilde{\Sigma}_\kappa Y_{11}^T & Y_{11} \tilde{\Sigma}_\kappa Y_{21}^T \\ Y_{21} \tilde{\Sigma}_\kappa Y_{11}^T & Y_{21} \tilde{\Sigma}_\kappa Y_{21}^T \end{bmatrix}.$$

Plugging in the GSVD gives

$$F_\kappa F_\kappa^T = \begin{bmatrix} \kappa^{-1} V \tilde{\Sigma}_\kappa^2 V^T & \sqrt{\kappa} V \tilde{\Sigma}_\kappa \Gamma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Sigma^T U^T \\ \sqrt{\kappa} U \Sigma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T \tilde{\Sigma}_\kappa V^T & \kappa^2 U \Sigma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T \Gamma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Sigma^T U^T \end{bmatrix}.$$

Solving for Y_{11} gives

$$Y_{11} = \kappa^{-1/2} V \tilde{\Sigma}_{\kappa, I}^{1/2}.$$

Using this in the upper right part gives

$$Y_{21} = \kappa U \Sigma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T \tilde{\Sigma}_{\kappa, I}^{-1/2}.$$

To solve for Y_{12} and Y_{22} , we use

$$YY^T = \begin{bmatrix} Y_{11}Y_{11}^T + Y_{12}Y_{12}^T & Y_{11}Y_{21}^T + Y_{12}Y_{22}^T \\ Y_{21}Y_{11}^T + Y_{22}Y_{12}^T & Y_{21}Y_{21}^T + Y_{22}Y_{22}^T \end{bmatrix} = \begin{bmatrix} I_p & 0 \\ 0 & I_m \end{bmatrix}.$$

The upper left part yields

$$Y_{12} = \kappa V \tilde{\Sigma}_{\kappa, I}^{-1/2} \Gamma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Sigma^T.$$

The upper right part yields

$$Y_{22} = -\kappa^{-1/2} U \tilde{\Sigma}_{\kappa, m}^{1/2}. \quad \square$$

Note that the singular values are ordered in ascending order. We have the following corollary.

COROLLARY 2.3. *If $m \geq n$ and $p < n$ the singular values of F_κ are given by*

$$\psi_i(F_\kappa) = \sqrt{\frac{\sigma_{n-i+1}^2}{\sigma_{n-i+1}^2/\kappa + \gamma_{n-i+1}^2}}.$$

If $m < n$ and $p > n$ the singular values of F_κ are given by

$$\psi_i(F_\kappa) = \begin{cases} \sqrt{\kappa} & \text{if } i \leq p - r_L, \\ \sqrt{\frac{\sigma_{m-i+1}^2}{\sigma_{m-i+1}^2/\kappa + \gamma_{m-i+1}^2}} & \text{if } p - r_L < i \leq p - r_L + r_A, \\ 0 & \text{if } i > p - r_L + r_A. \end{cases}$$

The question arises whether there is a direct relation between the singular values of A and the σ_i . The answer is no, but we do, however, have the following result from [20].

THEOREM 2.4 ([20, Thm. 2.4]). *Let $\psi_i(A)$ and $\psi_i(L)$ denote the singular values of A and L , respectively, and let σ_i and γ_i denote the nonzero entries of the matrices Σ and Γ , respectively. Then for all $\sigma_i, \gamma_i \neq 0$*

$$\begin{aligned} \left\| \begin{bmatrix} A \\ L \end{bmatrix}^\dagger \right\|_2^{-1} &\leq \frac{\psi_{r-i+1}(A)}{\sigma_i} \leq \left\| \begin{bmatrix} A \\ L \end{bmatrix} \right\|_2, \\ \left\| \begin{bmatrix} A \\ L \end{bmatrix}^\dagger \right\|_2^{-1} &\leq \frac{\psi_i(L)}{\gamma_i} \leq \left\| \begin{bmatrix} A \\ L \end{bmatrix} \right\|_2. \end{aligned}$$

Remark 2.5. This result shows that, if the operator A has quickly decaying singular values, the σ_i will have the same behavior; see also [22, p. 24]. This is an important result because it shows how the ill-conditioning of A transfers over to F_κ . Note that if $\sigma_i \approx 0$ we have $\gamma_i \approx 1$ and the singular values of

$$\psi_i(F_\kappa) = \sqrt{\frac{\sigma_{r-i+1}}{\sigma_{r-i+1}/\kappa + \gamma_{r-i+1}}} \approx \sqrt{\frac{\sigma_{r-i+1}}{\sigma_{r-i+1}/\kappa + 1}} \approx 0.$$

Hence, if the operator A is severely ill-posed, this ill-posedness is inherited by the operator F_κ .

2.4. Limiting cases.

2.4.1. The limit $\kappa \rightarrow \infty$ if $p < n$. If $L = I$ the limit $\kappa \rightarrow \infty$ yields the original optimization problem. However, if $L \neq I$, it is not immediately clear what happens in the limit $\kappa \rightarrow \infty$ due to the presence of the operator L . In this section we derive this limit using the GSVD of (A, L) . We will show that in the limit $\kappa \rightarrow \infty$ SR3 applies a standard-form transformation. We will proceed as follows. Recall that the variable x in SR3 is given by

$$(2.5) \quad \bar{x}_\kappa = H_\kappa^{-1} (\kappa L^T \bar{y}_\kappa + A^T b) = \kappa H_\kappa^{-1} L^T \bar{y}_\kappa + H_\kappa^{-1} A^T b := x_1 + x_2,$$

consisting of the two parts x_1 and x_2 . We will now show that, in the limit $\kappa \rightarrow \infty$, SR3 applies a standard-form transformation, by showing that x_1 and x_2 defined in (2.5) satisfy

$$(2.6) \quad x_1 = \bar{x}_\mathcal{M}, \quad x_2 = \bar{x}_\mathcal{N},$$

where $x_\mathcal{M}$ and $x_\mathcal{N}$ are determined by the standard-form transformation, given by (2.3) and (2.4), respectively.

Given the GSVD of (A, L) , the matrix F_κ and the vector g_κ are given by

$$(2.7) \quad F_\kappa = \begin{bmatrix} \sqrt{\kappa} V \left(I_p - \kappa \Gamma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T \right) V^T \\ \kappa U \Sigma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T V^T \end{bmatrix},$$

and

$$(2.8) \quad g_\kappa = \begin{bmatrix} \sqrt{\kappa} V \Gamma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Sigma^T U^T b \\ U \left(I_m - \Sigma (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Sigma^T \right) U^T b \end{bmatrix}.$$

As $\kappa \rightarrow \infty$ we have

$$F_\kappa \rightarrow \begin{bmatrix} 0 \\ U \Sigma \Gamma^\dagger V^T \end{bmatrix} \quad \text{and} \quad g_\kappa \rightarrow \begin{bmatrix} 0 \\ b \end{bmatrix}.$$

Hence, as $\kappa \rightarrow \infty$, we obtain

$$(2.9) \quad \bar{y}_\kappa = \arg \min_y \frac{1}{2} \|U \Sigma \Gamma^\dagger V^T y - b\|_2^2 + \mathcal{R}(y).$$

Using the GSVD, we have

$$H_\kappa^{-1} = X^{-1} (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} X^{-T},$$

and hence as $\kappa \rightarrow \infty$ we have

$$(\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \rightarrow \begin{bmatrix} 0 & 0 \\ 0 & I_{p-r_L} \end{bmatrix}.$$

Hence,

$$(2.10) \quad H_\kappa^{-1} \rightarrow X^{-1} \begin{bmatrix} 0 & 0 \\ 0 & I_{p-r_L} \end{bmatrix} X^{-T}.$$

Recall that the last columns of X are a basis for the nullspace of L and hence H_κ projects onto the nullspace of L . Using the GSVD of (A, L) we see that

$$\lim_{\kappa \rightarrow \infty} x_1 := \lim_{\kappa \rightarrow \infty} H_\kappa^{-1} A^T b = X^{-1} \begin{bmatrix} 0 & 0 \\ 0 & I_{p-r_L} \end{bmatrix} U^T b,$$

which is equivalent to the nullspace component from (2.4).

We now show that x_1 corresponds to the part in the range of L . We have

$$x_1 := \kappa H_\kappa^{-1} L^T \bar{y} = \kappa X^{-1} (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T V^T \bar{y}.$$

The elements of the diagonal matrix $\kappa (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T$ are

$$\begin{cases} \frac{\gamma_i}{\sigma_i^2/\kappa + \gamma_i^2} & \text{if } i \leq r_L, \\ 0 & \text{if } i > r_L, \end{cases}$$

and as $\kappa \rightarrow \infty$

$$\begin{cases} \frac{1}{\gamma_i} & \text{if } i \leq r_L, \\ 0 & \text{if } i > r_L. \end{cases}$$

Hence, as $\kappa \rightarrow \infty$

$$\kappa (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma)^{-1} \Gamma^T \rightarrow \Gamma^\dagger,$$

and thus

$$\kappa H_\kappa^{-1} L^T \rightarrow X^{-1} \Gamma^\dagger V^T = L_A^\dagger.$$

The limit for the component x_1 is now given by

$$\lim_{\kappa \rightarrow \infty} x_1 = X^{-1} \Gamma^\dagger V^T \bar{y}_\kappa = L_A^\dagger \bar{y}_\kappa,$$

where \bar{y}_κ solves

$$\bar{y}_\kappa = \arg \min_y \frac{1}{2} \|U \Sigma \Gamma^\dagger V^T y - b\|_2^2 + \mathcal{R}(y),$$

which is equivalent to (2.3).

2.4.2. The limit $\kappa \rightarrow \infty$ if $p > n$. If $p > n$, the limit $\kappa \rightarrow \infty$ is a bit more subtle. For large κ , we have

$$(2.11) \quad F_\kappa \sim \begin{bmatrix} V \begin{bmatrix} 0_{r_L \times r_L} & 0 \\ 0 & \sqrt{\kappa} I_{p-r_L} \end{bmatrix} V^T \\ U \Sigma \Gamma^\dagger V^T \end{bmatrix}, \quad \text{and } g_\kappa \sim \begin{bmatrix} 0_{p \times 1} \\ b \end{bmatrix}.$$

Hence, for large κ , SR3 solves a system of the form

$$\begin{bmatrix} \sqrt{\kappa} V_{p-r_L} V_{p-r_L}^T \\ U \Sigma \Gamma^\dagger V^T \end{bmatrix} y = \begin{bmatrix} 0_{p \times 1} \\ b \end{bmatrix},$$

where V_{p-r_L} are the *last* $p-r_L$ columns of V , which means that $V_{p-r_L} V_{p-r_L}^T = \mathcal{P}_{\mathcal{N}(L^T)}$. Because $V_{p-r_L} V_{p-r_L}^T y = 0$, the solution has no parts in $\mathcal{N}(L^T)$ and is restricted to the subspace $\mathcal{R}(L)$. The bottom part of F_κ is equal to the case $p < n$ and hence corresponds to matrix AL_A^\dagger . Let \bar{y}_{std} be the solution to the standard-form transformed system. Then, as $\kappa \rightarrow \infty$, the minimizer \bar{y}_κ of SR3 satisfies

$$(2.12) \quad \bar{y}_{\text{std}} = \mathcal{P}_{\mathcal{R}(L)} \bar{y}_\kappa.$$

However, looking at the original formulation in (1.2), we see that as $\kappa \rightarrow \infty$ we have

$$y = Lx,$$

which means that $y \in \mathcal{R}(L)$. Hence, condition (2.12) is immediately satisfied, and the solutions are the same.

2.5. The limit $\kappa \rightarrow 0$. The limit $\kappa \rightarrow 0$ is much easier to derive. Recall that

$$\bar{x}_\kappa = H_\kappa^{-1} (A^T b + \kappa L^T \bar{y}_\kappa).$$

As $\kappa \rightarrow 0$ we have $\kappa H_\kappa^{-1} L^T \bar{y}_\kappa \rightarrow 0$ and $H_\kappa \rightarrow (A^T A)^{-1}$. Hence $\lim_{\kappa \rightarrow 0} x_\kappa = (A^T A)^{-1} A^T b$ which is the unregularized minimum norm solution.

2.6. Relation to the standard-form transformation.

2.6.1. The case $p \leq n$. We have shown that as $\kappa \rightarrow \infty$ SR3 implicitly applies a standard-form transformation and that as $\kappa \rightarrow 0$ the system is unregularized. The question arises what happens for finite $\kappa > 0$. To show what happens, we rewrite the singular values of F_κ as

$$\psi_i(F_\kappa) = \sqrt{\frac{\sigma_{r-i+1}^2}{\sigma_{r-i+1}^2/\kappa + \gamma_{r-i+1}^2}} = \sqrt{\frac{\sigma_{r-i+1}^2/\gamma_{r-i+1}^2}{\frac{\sigma_{r-i+1}^2/\gamma_{r-i+1}^2}{\kappa} + 1}} = \sqrt{\frac{\psi_i^2(AL_A^\dagger)}{\psi_i^2(AL_A^\dagger)/\kappa + 1}}.$$

This is equivalent to equation 9 in [43], where it was shown that if $L^T L = I$,

$$\psi_i(F_\kappa) = \frac{\psi_i^2(A)}{\psi_i^2(A)/\kappa + 1}.$$

This shows that SR3 is applied to the matrix AL_A^\dagger . This leads to the following theorem.

THEOREM 2.6. *Let $p \leq n$. The following diagram commutes.*

$$\begin{array}{ccc} \min_x \frac{1}{2} \|Ax - b\|_2^2 + \lambda \mathcal{R}(Lx) & \xrightarrow{\text{SR3}} & \begin{array}{l} \bar{y}_\kappa = \arg \min_y \frac{1}{2} \|F_\kappa y - g_\kappa\|_2^2 + \lambda \mathcal{R}(y) \\ \bar{x}_\kappa = H_\kappa^{-1} (\kappa L^T \bar{y}_\kappa + A^T b) \end{array} \\ \downarrow & & \downarrow \uparrow \\ \begin{array}{l} \bar{z}_\kappa = \arg \min_z \frac{1}{2} \|AL_A^\dagger z - b\|_2^2 + \lambda \mathcal{R}(z) \\ \bar{x}_\kappa = L_A^\dagger \bar{z}_\kappa + x_N \end{array} & \xrightarrow{\text{SR3}} & \begin{array}{l} \bar{y}_\kappa = \arg \min_y \frac{1}{2} \|F_\kappa y - g_\kappa\|_2^2 + \lambda \mathcal{R}(y) \\ \bar{z}_\kappa = H_\kappa^{-1} (\kappa \bar{y}_\kappa + (AL_A^\dagger)^T b) \\ \bar{x}_\kappa = L_A^\dagger \bar{z}_\kappa + x_N \end{array} \end{array}$$

2.6.2. The case $p > n$. If $p > n$ the situation is different. Recall that the singular values of F_κ are given by

$$\psi_i(F_\kappa) = \begin{cases} \sqrt{\kappa} & \text{if } i \leq p - r_L, \\ \sqrt{\frac{\psi_i^2(AL_A^\dagger)}{\psi_i^2(AL_A^\dagger)/\kappa + 1}} & \text{if } p - r_L < i \leq p - r_L + r_A, \\ 0 & \text{if } i > p - r_L + r_A. \end{cases}$$

The singular values for F_κ when SR3 is applied to AL_A^\dagger are given by

$$\psi_i(F_\kappa) = \begin{cases} \sqrt{\frac{\psi_i^2(AL_A^\dagger)}{\psi_i^2(AL_A^\dagger)/\kappa + 1}} & \text{if } i \leq r_A, \\ 0 & \text{if } i > r_A. \end{cases}$$

Hence, there are extra singular values $\sqrt{\kappa}$ when SR3 is applied to the general-form system as opposed to the standard-form system. The difference may be seen from the expression (2.7). We have

$$\kappa\Gamma(\Sigma^T\Sigma + \kappa\Gamma^T\Gamma)\Gamma^T = \begin{cases} \begin{bmatrix} I_{n-r_A} & 0 & 0 \\ 0 & \kappa\Gamma_m(\Sigma^T\Sigma + \kappa\Gamma_m^T\Gamma_m)\Gamma_m^T & 0 \\ 0 & 0 & 0 \end{bmatrix} & \text{if } p > n, \\ \begin{bmatrix} I_{p-r_A} & 0 \\ 0 & \kappa\Gamma_m(\Sigma^T\Sigma + \kappa\Gamma_m^T\Gamma_m)\Gamma_m^T \end{bmatrix} & \text{if } p \leq n. \end{cases}$$

Hence, the top part of F_κ is different. Before we state our theorem let us introduce some notation. For the general-form problem, let the function φ be defined as the spectral cut-off function that makes the first $p - r_L$ singular values of F_κ zero. Similarly, for the standard-form transformed problem, let ϱ be defined as the function that makes $p - r_L$ singular values that are 0 equal to $\sqrt{\kappa}$ and accordingly permutes the SVD. We then have $\varphi \circ \varrho = \text{Id}$. We have the following theorem.

THEOREM 2.7. *Let $p > n$. The following diagram commutes.*

$$\begin{array}{ccc} \min_x \frac{1}{2} \|Ax - b\|_2^2 + \lambda\mathcal{R}(Lx) & \xrightarrow{\text{SR3}} & \begin{array}{l} \bar{y}_\kappa = \arg \min_y \frac{1}{2} \|F_\kappa y - g_\kappa\|_2^2 + \lambda\mathcal{R}(y) \\ \bar{x}_\kappa = H_\kappa^{-1}(\kappa L^T \bar{y}_\kappa + A^T b) \end{array} \\ \downarrow & & \varphi \uparrow \varrho \\ \begin{array}{l} \bar{z}_\kappa = \arg \min_z \frac{1}{2} \|AL_A^\dagger z - b\|_2^2 + \lambda\mathcal{R}(y) \\ \bar{x}_\kappa = L_A^\dagger \bar{z}_\kappa + x_N \end{array} & \xrightarrow{\text{SR3}} & \begin{array}{l} \tilde{y}_\kappa = \arg \min_y \frac{1}{2} \|\tilde{F}_\kappa y - g_\kappa\|_2^2 + \lambda\mathcal{R}(y) \\ \tilde{z}_\kappa = H_\kappa^{-1}(\kappa \tilde{y}_\kappa + (AL_A^\dagger)^T b) \\ \tilde{x}_\kappa = L_A^\dagger \tilde{z}_\kappa + x_N \end{array} \end{array}$$

3. Approximating the value function. In this section we quantify the distance between the Pareto curve of the original problem and the Pareto curve of the relaxed problem in terms of κ . We first describe the value function of the problem and then present our theorem.

The value function of an optimization problem expresses the value of the objective at the solution as a function of the other parameters. Using the standard-form transformation, we can, without loss of generality, consider the standard-form value function:

$$\phi_\kappa(\tau) = \min_y \|F_\kappa y - g_\kappa\|_2 \quad \text{s.t.} \quad \|y\|_p \leq \tau.$$

3.1. Value function for $\kappa \rightarrow \infty$. We have seen that for $\kappa \rightarrow \infty$, we retrieve the unrelaxed problem with value function

$$\phi_\infty(\tau) = \min_y \|Ay - b\|_2 \quad \text{s.t.} \quad \|y\|_p \leq \tau.$$

Following [40] we obtain the following (computable) upper and lower bounds for the value function

$$b^T \tilde{r} - \tau \|A^T \tilde{r}\|_q \leq \phi_\infty(\tau) \leq \|\tilde{r}\|_2,$$

where \tilde{y} is any feasible point (i.e., $\|\tilde{y}\|_p \leq \tau$), and $\tilde{r} = b - A\tilde{y}$ is the corresponding residual and $p^{-1} + q^{-1} = 1$. Moreover, by [40, Col. 2.2] the derivative of the value function is given by

$$\phi'_\infty(\tau) = -\|A^T \bar{r}\|_q / \|\bar{r}\|_2,$$

with $\bar{r} = b - A\bar{y}$ and $\bar{y} = \arg \min_{\|y\|_p \leq \tau} \|Ay - b\|_2$.

To gain some insight in the behavior of the value function, we consider ϕ_∞ and ϕ'_∞ at $\tau = 0$ and $\tau = \tau_* = \|A^\dagger b\|_p$:

$$\phi_\infty(0) = \|b\|_2, \quad \phi'_\infty(0) = -\|A^T b\|_q / \|b\|_2,$$

$$\phi_\infty(\tau_*) = \|(I - AA^\dagger)b\|_2, \quad \phi'_\infty(\tau_*) = 0.$$

This immediately suggests that ϕ_∞ decreases linearly near $\tau = 0$ (the zero solution) and flattens out near $\tau = \tau_*$ (the unconstrained minimizer). Since ϕ_∞ is known to be convex, its second derivative is always positive and will gradually bend the curve from decreasing to flat. How fast this happens and whether one can expect the typical L-shape depends on how fast the curve decreases initially. We can bound $\phi'_\infty(0)$ as follows. We let $b = Ay$ and find

$$\|A^T b\|_q = \|A^T Ay\|_q \geq C_q \|A^T Ay\|_2 \geq C_q \|A^\dagger\|^2 \|y\|_2,$$

where C_q is a constant that exists due to the equivalence of norms. Furthermore,

$$\|b\|_2 = \|Ay\|_2 \leq \|A\|_2 \|y\|_2.$$

From this we get

$$\phi'_\infty(0) \leq -C_q \kappa_2(A) \|A^\dagger\|_2,$$

with $\kappa_2(A) = \|A\|_2 \|A^\dagger\|_2$ the condition number of A . We thus expect a steep slope for ill-conditioned problems, giving rise for the characteristic *L*-shape of the curve. While this behavior is well-established for $p = 2$ where it can be analyzed using the SVD of A [22], this analysis gives us new insight in the behavior of the Pareto curve for ill-posed problems for general p . An example for $p = 1$, $L = I$ is shown in Figure 3.1.




FIG. 3.1. Pareto curve for an ill-posed problem; the matrix A is diagonal with elements $e^{-(i-1)/2}$ for $i = 1, 2, \dots, 10$; $b = Ax$ with $x = (1, 1, \dots, 1)$. The tangent lines at $\tau = 0$ and $\tau = \tau_*$ are shown in black.

3.2. Relaxed value function. We now present our theorem on the distance between the Pareto curve of the original problem and the Pareto curve of the relaxed problem.

THEOREM 3.1. *The distance between the Pareto curve of the original problem and the Pareto curve of the relaxed problem is given by*

$$(\phi_\kappa(\tau))^2 - (\phi_\infty(\tau))^2 = -\kappa^{-1} \|A^T(b - A\bar{y}_\kappa)\|_2^2 + \mathcal{O}(\kappa^{-2}),$$

where \bar{y}_κ is the solution of the relaxed problem. In particular, we have

$$\phi_\kappa(\tau) \leq \phi_\infty(\tau).$$

Proof. Let $\epsilon = \kappa^{-1}$. The relaxed value function can be expressed as

$$\phi_\epsilon(\tau) = \min_y \|F_\epsilon y - g_\epsilon\|_2 \quad \text{s.t.} \quad \|y\|_p \leq \tau.$$

For $\epsilon < \|A\|_2^2$ we can expand $H_\epsilon^{-1} = \epsilon I - \epsilon^2 A^T A + \mathcal{O}(\epsilon^3)$ and get

$$F_\epsilon = \begin{pmatrix} A - \epsilon A A^T A + \mathcal{O}(\epsilon^2) \\ \epsilon^{1/2} A^T A + \mathcal{O}(\epsilon^{3/2}) \end{pmatrix}, \quad g_\epsilon = \begin{pmatrix} b - \epsilon A^T b + \mathcal{O}(\epsilon^2) \\ \epsilon^{-1/2} A^T b + \mathcal{O}(\epsilon^{3/2}) \end{pmatrix}.$$

Introduce

$$f(\epsilon) = (\phi_\epsilon(\tau))^2 = \min_{x,y} \|Ax - b\|_2^2 + \epsilon^{-1} \|x - y\|_2^2 \quad \text{s.t.} \quad \|y\|_p \leq \tau.$$

We have $f(0) = \min_{\|y\|_p \leq \tau} \|Ay - b\|_2^2 = (\phi_0(\tau))^2$. Furthermore

$$f'(\epsilon) = -\epsilon^{-2} \|\bar{x}_\epsilon - \bar{y}_\epsilon\|_2^2,$$

where $\bar{x}_\epsilon = H_\epsilon^{-1}(A^T b + \epsilon^{-1} \bar{y}_\epsilon)$ and \bar{y}_ϵ is the optimal y . With this we find

$$(3.1) \quad (\phi_\epsilon(\tau))^2 - (\phi_0(\tau))^2 = \epsilon f'(\eta) = -\epsilon \eta^{-2} \|\bar{x}_\eta - \bar{y}_\eta\|_2^2.$$

We conclude that $\phi_\epsilon(\tau) \leq \phi_0(\tau)$. Alternatively, we can express

$$(3.2) \quad (\phi_\epsilon(\tau))^2 - (\phi_0(\tau))^2 = -\epsilon^{-1} \|\bar{x}_\epsilon - \bar{y}_\epsilon\|_2^2 + \mathcal{O}(\epsilon^2).$$

For small ϵ we get

$$f'(\epsilon) = -\|A^T(b - A\bar{y}_\epsilon)\|_2^2 + \mathcal{O}(\epsilon).$$

Plugging this expression into (3.2) gives the desired result. \square

Remark 3.2. Theorem 3.1 can be used to explain the behavior of the Pareto curves observed in the examples in section 1.1:

- The error gets smaller for large τ . For an unconstrained problem we have $\|A^T(b - A\bar{y}_\kappa)\|_2 = 0$ as $\kappa \rightarrow \infty$. An example is shown in Figure 3.2.
- The elbow of the Pareto curves coincide; ϕ_∞ decreases fast initially for ill-posed problems (cf. Figure 3.1) while ϕ_κ decreases less fast due to the implicit regularizing effect of the relaxation. Since $0 \leq \phi_\kappa \leq \phi_\infty$, the relaxed Pareto curve is pushed down and is therefore likely to have the elbow at the same location as ϕ_∞ .




FIG. 3.2. Pareto curve for an ill-posed problem; the matrix A is diagonal with elements $e^{-(i-1)/2}$ for $i = 1, 2, \dots, 10$; $b = Ax$ with $x = (1, 1, \dots, 1)$. The approximations for various values of ϵ are shown as well.

4. Implementation. Recall from the introduction that we implement SR3 as follows:

$$(4.1) \quad x_{k+1} \leftarrow (A^T A + \kappa L^T L)^{-1} (A^T b + \kappa L^T y_k),$$

$$(4.2) \quad y_{k+1} \leftarrow \text{prox}_{1/\kappa \mathcal{R}}(Lx_k).$$

The last equation shows that for the choice $\mathcal{R}(\cdot) = \lambda \|\cdot\|_p^p$ there is a relation between the parameters κ and λ . More specifically, λ depends on κ , and hence we write $\lambda(\kappa)$. Given the optimal λ_* , we have $\lambda(\kappa) = \lambda_* \cdot \kappa$. Note that if we use the constrained formulation (1.9), the dependence on the stepsize is lost because the proximal operator is the indicator function, and there is no relation between τ and κ .

The computational bottleneck is in the first step, which is the solution to the large-scale linear system

$$(4.3) \quad (A^T A + \kappa L^T L) x_k = A^T b + \kappa L^T y_{k-1}.$$

To avoid explicitly forming $A^T A$ and $L^T L$, we instead solve the following minimization problem

$$(4.4) \quad \min_x \left\| \begin{bmatrix} A \\ \sqrt{\kappa} L \end{bmatrix} x - \begin{bmatrix} b \\ \sqrt{\kappa} y_{k-1} \end{bmatrix} \right\|_2^2,$$

with LSQR.

We will numerically investigate how only partially solving (4.4) affects the convergence of SR3. This has been investigated for ADMM in [13, 14, 1]. The convergence of FISTA with an inexact gradient has been analyzed in [37]. The key message is that the error has to go down as the iterations increase.

In our implementation, we propose two extra ingredients to make SR3 suitable for large-scale problems: warm starts and inexact solves of (4.4). Both ingredients are also used in the implementation of ADMM [8]. However, we propose a new stopping criterion for the inexact solves of (4.4).

A *warm start* is a technique used in inner-outer schemes, where the solution of the previous inner iteration serves as an initial guess to the new inner iteration. That is, we solve

$$(4.5) \quad \min_x \left\| \begin{bmatrix} A \\ \sqrt{\kappa}L \end{bmatrix} x - \left(\begin{bmatrix} b \\ \sqrt{\kappa}y_{k-1} \end{bmatrix} - \begin{bmatrix} A \\ \sqrt{\kappa}L \end{bmatrix} x_{k-1} \right) \right\|_2^2.$$

By *inexact solves* we mean finding an approximate solution to (4.5). The level of inexactness is determined by the difference between the true solution and the inexact solution. There are various ways in which one can solve the optimization problem inexactly. One way is to simply determine a maximum number of iterations. However, the number of iterations to solve (4.4) can vary strongly per outer iteration. Moreover, we may not want to solve the inner system with high precision in the first few outer iterations, because this does not result in significant improvement in the next outer iteration. Recently, the authors in [41] proposed a criterion to determine the amount of inexactness for inner-outer schemes. The idea is to stop the inner iteration once the difference in the resulting outer iterate becomes stagnant. Let x_k denote the current inner iterate and $y_k = \text{prox}_{1/\kappa\mathcal{R}}(Lx_k)$ the resulting outer iterate by applying the proximal operator. Then the authors in [41] propose to stop the inner iterations if

$$(4.6) \quad \|x_{k+1} - x_k\| < \rho \|y_{k+1} - y_k\|$$

for some user defined constant ρ . We propose a similar criterion, namely, to stop if

$$(4.7) \quad \frac{\|y_{k+1} - y_k\|}{\|y_k\|} < \epsilon$$

for some user defined threshold ϵ . The index k refers to the iteration of the iterative method applied to the inner iteration. This yields the proposed implementation of SR3, shown in Algorithm 4.1. Note that in line 4 of the algorithm we use the LSQR algorithm, and we build on the Krylov subspace from the previous step.

Algorithm 4.1 Implementation of SR3

Require: Operators A and L , the data b , and the parameters κ , λ and ϵ .

Ensure: Approximate solution x_k .

```

1: while  $\|x_{k+1} - x_k\| > \delta$  do
2:    $l = 0$ .
3:   while  $\frac{\|\tilde{y}_{l+1} - \tilde{y}_l\|}{\|\tilde{y}_l\|} > \epsilon$  do           ▷ Run LSQR. We do not restart LSQR every
      iteration!
4:      $x_l = \arg \min_{x \in \mathcal{K}_l} \left( \begin{bmatrix} A \\ \sqrt{\kappa}L \end{bmatrix}, \begin{bmatrix} b - Ax_k \\ \sqrt{\kappa}(y_0 - Lx_k) \end{bmatrix} \right) \left\| \begin{bmatrix} A \\ \sqrt{\kappa}L \end{bmatrix} x \right.$ 
5:          $\left. - \left( \begin{bmatrix} b \\ \sqrt{\kappa}y_k \end{bmatrix} - \begin{bmatrix} A \\ \sqrt{\kappa}L \end{bmatrix} x_k \right) \right\|_2^2$ .
6:      $\tilde{y}_{l+1} = \text{prox}_{1/\kappa\mathcal{R}}(Lx_l)$ .           ▷ Prospective update
7:      $l = l + 1$ .
8:   end while
9:    $y_k = \tilde{y}_l$ .
10:   $k = k + 1$ .
11: end while

```

It is important to note that the influence of κ on the outer iteration is different from the influence of κ on the inner iteration. The improved conditioning of the matrix F_κ pertains to the convergence of the outer iteration. The convergence of the inner iteration is completely determined by the properties of the matrix H_κ^{-1} . It is important to note that using the GSVD of (A, L) we get

$$H_\kappa = A^T A + \kappa L^T L = X^T (\Sigma^T \Sigma + \kappa \Gamma^T \Gamma) X,$$

but this is not the SVD of H_κ , because X is not orthonormal. Therefore, the matrix $\Sigma^T \Sigma + \kappa \Gamma^T \Gamma$ does not tell us anything about the convergence rate when solving linear systems involving H_κ .

5. Numerical experiments. In this section we verify the results from section 2 numerically. Furthermore, we implement Algorithm 4.1 and test it on two examples. We use two examples that are regularized by TV regularization, which we solve in its constrained form, i.e.,

$$\min_x \|Ax - b\|_2^2 \quad \text{s.t.} \quad \|Lx\|_1 \leq \tau.$$

5.1. Examples. We will use two examples that are very different in nature in terms of their singular values. For both examples, we will show how their spectra are changed as a function of κ by applying SR3 and how this relates to the inner and outer iterations. After that, we will show how our inexact SR3 greatly reduces the total number of iterations. We do not add noise to the data.

Gravity surveying. The first example is the gravity example from the regu toolbox [23, 21]. This example models gravity surveying. An unknown mass distribution that generates a gravity field is located in the subsurface, and the measured data is related to the gravity field via a Fredholm integral of the first kind, i.e.,

$$b(s) = \int_{\Omega} k(s, t)x(t)dt.$$

The variable $x(t)$ is the mass density at the location t in the subsurface, and $b(s)$ is the gravity field at location s at the surface. The kernel is given by

$$k(s, t) = d(d^2 + (s - t)^2)^{-3/2},$$

where d is the depth. The integral is discretized using the midpoint quadrature rule and yields a symmetric Toeplitz matrix A that is square and severely ill-posed. We have chosen an $x(t)$ that is piecewise constant, and hence we regularize the problem with TV regularization. The operator $L = D$, where D is the first-order finite difference discretization, i.e.,

$$D = \begin{bmatrix} -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & -1 & 1 \end{bmatrix} \in \mathbb{R}^{(n-1) \times n}.$$

The operator is underdetermined and its nullspace has dimension 1. We choose $n = 512$. The true gravity profile is shown in Figure 5.1.

Tomography. Our second example is the tomography example PRtomo from the IR Tools toolbox [16] (see also [26]), which models parallel tomography. It models x-ray attenuation tomography, often referred to as computerized tomography (CT). Parallel rays at different angles penetrate an object. The rays are attenuated at a rate proportional to the length of the ray and the density of the object. The i th ray can be modeled as

$$b_i = \sum_{j \in \mathcal{S}_i} a_{ij} x_j.$$


The set \mathcal{S} denotes the set of pixels that are penetrated, a_{ij} denotes the length of the i th ray through the j th pixel and x_j is the attenuation coefficient. This is a two-dimensional (2D) example where the matrix A is underdetermined and the singular values decay mildly. Again, we use TV regularization for the reconstruction. For 2D regularization, the operator $L = \begin{bmatrix} I \otimes D \\ D \otimes I \end{bmatrix}$. Hence, the operator L is overdetermined and has a nullspace of dimension 1. We choose 18 angles between 0 and 180 degrees and discretize the image on a 128×128 pixel grid. This means that $A \in \mathbb{R}^{3258 \times 16384}$. Our experiments are on the Shepp–Logan phantom, shown in Figure 5.1.

Parameters. For our experiments, we have adapted the implementation of the accelerated proximal gradient algorithm from [35] for SR3 and use the same stopping criterion for the proximal gradient algorithm. For the inexact stopping criterion for the inner iteration we choose $\epsilon = 10^{-6}$. For the exact SR3 method, we let LSQR run to convergence with the standard tolerance of 10^{-6} . For τ , we choose the optimal value $\tau = \|Lx_{\text{true}}\|_1$.

5.2. Singular values of F_κ . In this section we show the singular values of F_κ for the gravity and the tomography example. For the tomography example, the generalized singular values are calculated on a 64×64 grid to reduce computational time, instead of the 128×128 grid for our experiments. We show the generalized singular values, i.e., the singular values of AL_A^\dagger , and the singular values of F_κ for different values of κ for the gravity example in Figure 5.2.

Note that irrespective of the value of κ , the matrix F_κ remains severely ill-posed. For the tomography example, A is not severely ill-posed. The singular values decay only mildly and the situation is different. In this case, for small κ ,

$$\psi_i(F_\kappa) = \sqrt{\frac{\sigma_{r-i+1}^2}{\sigma_{r-i+1}^2/\kappa + \gamma_{r-i+1}^2}} \approx \sqrt{\frac{\sigma_{r-i+1}^2}{\sigma_{r-i+1}^2/\kappa}} = \sqrt{\kappa}.$$



Gravity profile



Shepp–Logan phantom

FIG. 5.1.

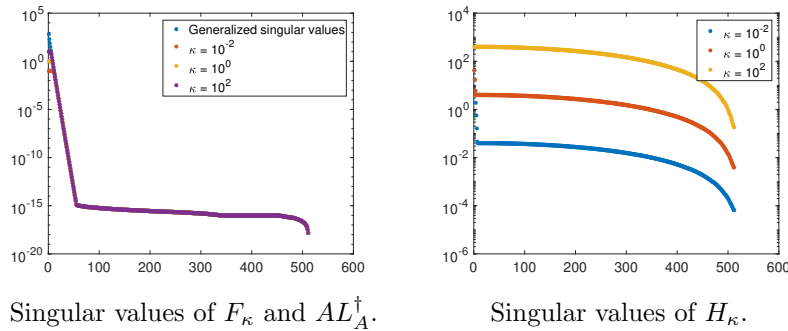



FIG. 5.2. Spectral properties of F_κ and H_κ for the gravity example. Left figure: We show the singular values of AL_A^\dagger and the singular values of F_κ for different values of κ . Note that the singular values of F_κ have a very similar structure to the singular values of AL_A^\dagger . Right figure: The singular values of the matrix H_κ .



^aThe matrix is numerically rank deficient and we have truncated the SVD.

FIG. 5.3. Spectral properties of F_κ and H_κ for the tomography example. The left figure shows the singular values of F_κ . Recall that the first $p - r_L$ singular values of F_κ are $\sqrt{\kappa}$. The right figure shows the singular values of H_κ . There is an inverse relation between the condition number of H_κ and F_κ as a function of κ .


Hence, for small κ the singular values of $F_\kappa \approx \sqrt{\kappa}$ and the condition number is 1. As $\kappa \rightarrow \infty$ we have seen that $\psi_i(F_\kappa) \rightarrow \sigma_{r-i+1}/\gamma_{r-i+1}$. We show the singular values, the generalized singular values, and the singular values of F_κ in Figure 5.3. Note that for this example, the conditioning of the matrix F_κ is improved.

5.3. The Pareto curves. In Figure 5.4 we show the Pareto curves for the original problem and SR3 for both our examples.

As we explained in section section 3, the corner of the Pareto of the original problem and SR3 is likely to be in the same place. This is confirmed by Figure 5.4.


5.4. The influence of κ on the number of iterations. To investigate the influence of κ , we show the amount of inner and outer iterations for varying values of κ and the total number of iterations. The results are shown in Figure 5.5 and Figure 5.6. As we have stated before, the improved convergence rate due to an improved conditioning of κ pertains to the outer iterations. The effect of κ on the convergence of the inner iteration may be completely opposite.

For the gravity example, we see that the amount of inner iterations varies very little as κ increases and even goes up a little bit. This is not unexpected, because



Pareto curves for the gravity example. Pareto curves for the tomography example.


FIG. 5.4. The left figure shows the Pareto curves for the gravity example. The right figure shows the Pareto curves for the tomography example. The x-axis is τ and the y-axis is $\|A\bar{x}_\kappa - b\|_2$.



Inner iterations versus outer iterations.

Total number of iterations.

FIG. 5.5. The left figure shows the inner and outer iterations for varying κ for the gravity example. The right figure shows the total number of iterations.



Inner iterations versus outer iterations.

Total number of iterations.

FIG. 5.6. The left figure shows the inner and outer iterations for varying κ for the tomography example. The right figure shows the total number of iterations.

the decay of the singular values changes very little as κ increases; see Figure 5.2. The amount of outer iterations goes down rapidly as κ decreases, something that is not expected from the distribution of the singular values. This shows that the distribution of the singular values is not the sole property explaining the convergence behavior.

For the tomography example we see a clear trade-off between inner and outer iterations. From Figure 5.3 we clearly see that as the condition number of F_κ

decreases, the condition number of H_κ increases. This explains that, as the amount of inner iterations goes down with increasing κ , the amount of outer iterations goes down.

5.5. Inexact SR3. In this section we compare the error and the total number of iterations for SR3 and inexact SR3 as a function of κ . The results are shown in Figure 5.7.

We see that the total number of iterations needed is greatly reduced by implementing the automated stopping criterion. Another important contribution is that the stopping criterion seems to mitigate the influence of κ on the total amount of iterations. Figure 5.8 and Figure 5.9 show some reconstructions for different values of κ .

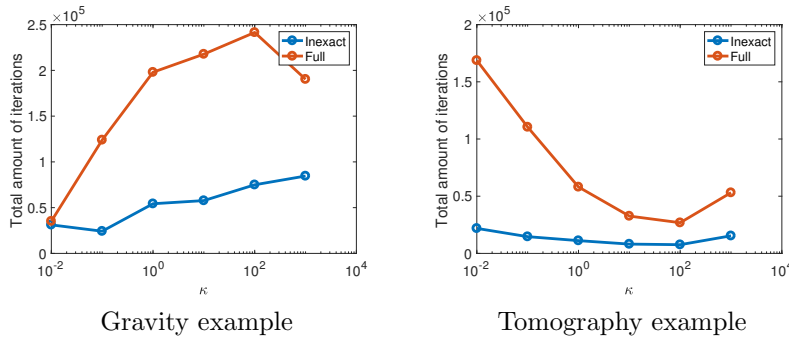


FIG. 5.7. Comparison of the total number of iterations for SR3 and inexact SR3 as a function of κ . Note that the axes are on a log-log scale.

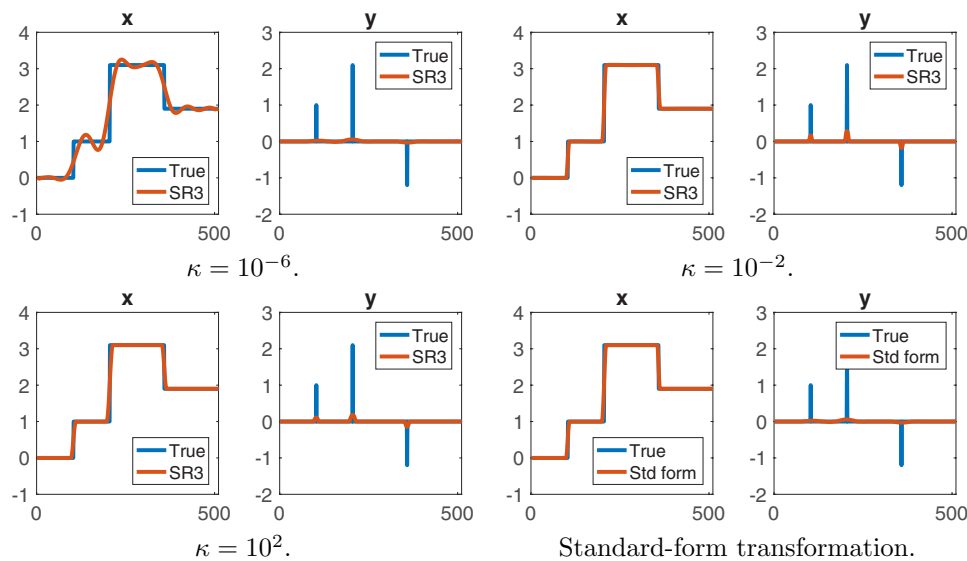


FIG. 5.8. Solution to the gravity example for different κ . We have chosen λ by hand to yield the best reconstruction. We show both \bar{x}_κ and \bar{y}_κ .




FIG. 5.9. Solution to the tomography example for different κ . We have chosen λ by hand to yield the best reconstruction.

6. Conclusion and outlook. In this paper we have analyzed the method SR3 which was introduced in [43]. We have extended Theorem 1 from [43] about the singular values of F_κ to the general form case. We have shown that SR3, as $\kappa \rightarrow \infty$, implicitly applies a standard-form transformation, and that for finite $\kappa > 0$, the singular values of F_κ are related to the standard-form transformed operator.

In section 3 we have shown that the distance between the Pareto curve of the original problem and the Pareto curve of the relaxed problem is of $\mathcal{O}(1/\kappa^2)$ plus the norm of the gradient, which depends on κ .

In section 4 we have presented our implementation of the inexact SR3 algorithm, where we have proposed an automated stopping criterion for the inner iterations.

In our numerical experiments in section 5 we have compared the SR3 algorithm for two example problems with very different spectra. The gravity example is a severely ill-posed problem, and we have shown numerically that the convergence of inner iterations is not affected much by κ , but the convergence of the outer iteration is. For the tomography example we saw a trade-off: As κ decreases the outer iterations converge rapidly, but the number of inner iterations is large. We have shown that our automated stopping criterion greatly reduces the number of iterations needed.

For future research it would be interesting to further investigate the relation between the Pareto curve of the original problem and of the relaxed problem. Specifically, it would be great if we could prove that the corner of the curves are in the same place, something that we have only been able to show qualitatively through Theorem 3.1. This would lead to automatic selection of the regularization parameter λ .

Another interesting topic of research is the selection of κ . As we have seen in our experiments, the choice of κ strongly influences the number of iterations needed

for SR3, although this is largely mitigated by the inexact stopping criterion. The relation between the tolerance for the stopping criterion and κ should also be further investigated.

Acknowledgments. The authors would like to thank Dr. Michiel Hochstenbach and Dr. Ajinkya Kadu for fruitful discussions.

REFERENCES

- [1] M. M. ALVES, J. ECKSTEIN, M. GEREMIA, AND G. M. JEFFERSON, *Relative-error inertial-relaxed inexact versions of Douglas-Rachford and ADMM splitting algorithms*, *Comput. Optim. Appl.*, 75 (2020), pp. 389–422.
- [2] A. ARAVKIN, J. BURKE, AND M. FRIEDLANDER, *Variational properties of value functions*, *SIAM J. Optim.*, 23 (2013), pp. 1689–1717, <https://doi.org/10.1137/120899157>.
- [3] S. R. ARRIDGE AND J. C. SCHOTLAND, *Optical tomography: Forward and inverse problems*, *Inverse Problems*, 25 (2009), 123010, <https://doi.org/10.1088/0266-5611/25/12/123010>.
- [4] A. BECK AND M. TEOULLE, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, *SIAM J. Imaging Sci.*, 2 (2009), pp. 183–202.
- [5] M. BERTERO AND P. BOCCACCI, *Introduction to Inverse Problems in Imaging*, Institute of Physics Publishing, Bristol, 1998.
- [6] A. BESSON, J. THIRAN, AND Y. WIAUX, *Imaging from Echoes: On Inverse Problems in Ultrasound*, Ecole Polytechnique Fédérale de Lausanne, 2019.
- [7] B. BORDEN, *Mathematical problems in radar inverse scattering*, *Inverse Problems*, 18 (2001), p. R1, <https://doi.org/10.1088/0266-5611/18/1/201>.
- [8] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed optimization and statistical learning via the Alternating Direction Method of Multipliers*, *Found. Trends Mach. Learn.*, 3 (2011), pp. 1–122.
- [9] D. CALVETTI AND L. REICHEL, *Tikhonov regularization of large linear problems*, *BIT*, 43 (2003), pp. 261–281.
- [10] D. CALVETTI, G. SPALETTA, L. REICHEL, AND F. SGALLARI, *An L-ribbon for large underdetermined linear discrete ill-posed problems*, *Numer. Algorithms*, 25 (2000), pp. 89–107.
- [11] E. J. CANDÉS, J. ROMBERG, AND T. TAO, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, *IEEE Trans. Inform. Theory*, 52 (2006), pp. 489–509.
- [12] I. DAUBECHIES, M. DEFRISE, AND C. DE MOL, *An iterative thresholding algorithm for linear inverse problems with a sparsity constraint*, *Comm. Pure Appl. Math.*, 57 (2004), pp. 1413–1457.
- [13] J. ECKSTEIN AND W. YAO, *Approximate ADMM algorithms derived from Lagrangian splitting*, *Comput. Optim. Appl.*, 68 (2017), pp. 363–405.
- [14] J. ECKSTEIN AND W. YAO, *Relative-error approximate versions of Douglas–Rachford splitting and special cases of the ADMM*, *Math. Program.*, 170 (2018), pp. 417–444.
- [15] L. ELDÉN, *A weighted pseudoinverse, generalized singular values, and constrained least squares problems*, *BIT*, 22 (1982), pp. 487–502, <https://doi.org/10.1007/BF01934412>.
- [16] S. GAZZOLA, P. C. HANSEN, AND J. G. NAGY, *IR tools: A MATLAB package of iterative regularization methods and large-scale test problems*, *Numer. Algorithms*, 81 (2019), pp. 773–811.
- [17] S. GAZZOLA, P. NOVATI, AND M. R. RUSSO, *On Krylov projection methods and Tikhonov regularization*, *Electron. Trans. Numer. Anal.*, 44 (2015), pp. 83–123.
- [18] G. H. GOLUB AND U. VON MATT, *Tikhonov Regularization for Large Scale Problems*, in *Proceedings of the Workshop on Scientific Computing*, 1997.
- [19] I. GOODFELLOW, Y. BENGIO, A. COURVILLE, AND F. BACH, *Deep Learning*, MIT Press, Cambridge, MA, 2016.
- [20] P. C. HANSEN, *Regularization, GSVD and truncated GSVD*, *BIT*, 29 (1989), pp. 491–504.
- [21] P. C. HANSEN, *Regularization tools: A MATLAB package for analysis and solution of discrete ill-posed problems*, *Numer. Algorithms*, 6 (1994), pp. 1–35.
- [22] P. C. HANSEN, *Rank-deficient and Discrete Ill-posed Problems: Numerical Aspects of Linear Inversion*, SIAM, Philadelphia, 1998.
- [23] P. C. HANSEN, *Deconvolution and regularization with Toeplitz matrices*, *Numer. Algorithms*, (2002), pp. 323–378.
- [24] P. C. HANSEN, *Discrete Inverse Problems: Insight and Algorithms*, SIAM, Philadelphia, 2010.

- [25] P. C. HANSEN, *Oblique projections and standard-form transformations for discrete inverse problems*, Numer. Linear Algebra Appl., 20 (2013), pp. 250–258.
- [26] P. C. HANSEN AND J. S. JØRGENSEN, *Air tools II: Algebraic iterative reconstruction methods, improved implementation*, Numer. Algorithms, 79 (2018), pp. 107–137.
- [27] G. HENNENFENT AND F. J. HERRMANN, *Simply denoise: Wavefield reconstruction via jittered undersampling*, Geophysics, 73 (2008), pp. 19–28.
- [28] G. HERMAN, *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*, Springer, New York, 2009.
- [29] M. HOCHSTENBACH AND L. REICHEL, *An iterative method for Tikhonov regularization with a general linear regularization operator*, J. Integral Equations, 22 (2010), pp. 463–480.
- [30] M. HOCHSTENBACH, L. REICHEL, AND X. YU, *A Golub-Kahan-type reduction method for matrix pairs*, J. Sci. Comput., 65 (2015), pp. 767–789.
- [31] M. E. KILMER, P. C. HANSEN, AND M. I. ESPANOL, *A projection-based approach to general-form Tikhonov regularization*, SIAM J. Sci. Comput., 29 (2007), pp. 315–330.
- [32] M. E. KILMER AND D. P. O’LEARY, *Choosing regularization parameters in iterative methods for ill-posed problems*, SIAM J. Matrix Anal. Appl., 22 (2001), pp. 1204–1221.
- [33] F. NATTERER, *The Mathematics of Computerized Tomography*, SIAM, Philadelphia, 2001.
- [34] C. J. NOLAN AND M. CHENEY, *Synthetic aperture inversion*, Inverse Problems, 18 (2002), pp. 221–235, <https://doi.org/10.1088/0266-5611/18/1/315>.
- [35] B. O’DONOGHUE, *APG*, <https://github.com/bodono/apg>, 2016.
- [36] L. I. RUDIN AND S. OSHER, *Total variation based image restoration with free local constraints*, in Proceedings of 1st International Conference on Image Processing, vol. 1, 1994, pp. 31–35.
- [37] M. SCHMIDT, N. L. ROUX, AND F. R. BACH, *Convergence rates of inexact proximal-gradient methods for convex optimization*, in Advances in Neural Information Processing Systems 24, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, eds., Curran Associates, Red Hook, NY, 2011, pp. 1458–1466.
- [38] J.-L. STARCK, *Sparsity and inverse problems in astrophysics*, J. Phys. Conf. Ser., 699 (2016), 012010, <https://doi.org/10.1088/1742-6596/699/1/012010>.
- [39] J. TAMMINEN, *Inverse problems and uncertainty quantification in remote sensing*, in Proceedings of the ESA Earth Observation Summer School on Earth System Monitoring and Modeling, 2012.
- [40] E. VAN DEN BERG AND M. FRIEDLANDER, *Probing the Pareto frontier for basis pursuit solutions*, SIAM J. Sci. Comput., 31 (2008), pp. 890–912.
- [41] T. VAN LEEUWEN AND A. ARAVKIN, *Non-smooth Variable Projection*, preprint, arXiv:1601.05011, 2020, <https://arxiv.org/abs/1601.05011>.
- [42] J. VIRIEUX AND S. OPERTO, *An overview of Full-Waveform Inversion in exploration geophysics*, Geophysics, 74 (2009), pp. 1–26.
- [43] P. ZHENG, T. ASKHAM, S. L. BRUNTON, J. N. KUTZ, AND A. Y. ARAVKIN, *A unified framework for sparse relaxed regularized regression: SR3*, IEEE Access, 7 (2019), pp. 1404–1423, <https://doi.org/10.1109/ACCESS.2018.2886528>.
- [44] L. ZHI-PEI AND P. C. LAUTERBUR, *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*, SPIE Optical Engineering Press, Bellingham, WA, 2000.
- [45] I. ZWAAN AND M. HOCHSTENBACH, *Multidirectional subspace expansion for one-parameter and multiparameter Tikhonov regularization*, SIAM J. Sci. Comput., 70 (2017), pp. 990–1009.