Computers & Geosciences 137 (2020) 104427

Contents lists available at ScienceDirect

Computers and Geosciences

journal homepage: www.elsevier.com/locate/cageo

Research paper

Comparing RSVD and Krylov methods for linear inverse problems

Nick Luiken [*,1], Tristan van Leeuwen [1]

Mathematical Institute, Utrecht University, Budapestlaan 6, 3584CD, Utrecht, The Netherlands

ARTICLE INFO

Keywords:
Tikhonov regularization
Regularization parameter
Model order reduction
Krylov subspaces
Randomized singular value decomposition

ABSTRACT

In this work we address regularization parameter estimation for ill-posed linear inverse problems with an ℓ_2 penalty. Regularization parameter selection is of utmost importance for all of inverse problems and estimating it generally relies on the experience of the practitioner. For regularization with an ℓ_2 penalty there exist a lot of parameter selection methods that exploit the fact that the solution and the residual can be written in explicit form. Parameter selection methods are functionals that depend on the regularization parameter where the minimizer is the desired regularization parameter that should lead to a good solution. Evaluation of these parameter selection methods still requires solving the inverse problem multiple times. Efficient evaluation of the parameter selection methods can be done through model order reduction. Two popular model order reduction techniques are Lanczos based methods (a Krylov subspace method) and the Randomized Singular Value Decomposition (RSVD). In this work we compare the two approaches. We derive error bounds for the parameter selection methods using the RSVD. We compare the performance of the Lanczos process versus the performance of RSVD for efficient parameter selection. The RSVD algorithm we use is based on the Adaptive Randomized Range Finder algorithm which allows for easy determination of the dimension of the reduced order model. Some parameter selection also require the evaluation of the trace of a large matrix. We compare the use of a randomized trace estimator versus the use of the Ritz values from the Lanczos process. The examples we use for our experiments are two model problems from geosciences.

1. Introduction

Inverse problems are ubiquitous in the earth-sciences with applications ranging from seismology to seismic exploration. Often, these inverse problems are *ill-posed*, meaning that a unique, stable solution does not exist. Regularization is needed to render the problem well-posed. In this paper we discuss finite-dimensional, linear inverse problems which can be posed as

$$\min_{\mathbf{m}} \|G\mathbf{m} - \mathbf{d}\|_2^2 + \lambda \|L\mathbf{m}\|_2^2, \tag{1}$$

where $G \in \mathbb{R}^{m\times n}$ is the *forward operator*; $\mathbf{d} \in \mathbb{R}^m$ denotes the data; $\mathbf{m} \in \mathbb{R}^n$ are the parameters of interest; $L \in \mathbb{R}^{p\times n}$ is the *regularization operator* and $\lambda \in \mathbb{R}_+$ is the *regularization parameter*. The regularization operator incorporates the prior information needed to make the problem uniquely solvable. Without loss of generality, we assume that $L = I$, as every problem of the form (1) can be transformed to this form (Hansen, 2013). The regularization parameter balances the prior information and information from the data. The solution can be written in closed form, given by

$$\widehat{\mathbf{m}}_\lambda = G_\lambda \mathbf{d}, \tag{2}$$

with

$$G_\lambda = (G^T G + \lambda I)^{-1} G^T. \tag{3}$$

Although this expression is convenient for derivations, in practice it is usually not feasible to form the matrix G_λ explicitly. Therefore, the solution $\widehat{\mathbf{m}}_\lambda$ is usually approximated using an iterative solver. A major issue in solving (1) is the selection of the regularization parameter λ. Methods for selecting the regularization parameter are called *parameter selection methods*. A complete overview and comparison of parameter selection methods is given in Bauer and Lukas (2011). Parameter selection methods, generally, rely on repeatedly solving (1) and selecting the value of λ that satisfies some auxiliary criteria. These criteria usually involve minimizing a functional $V(\lambda)$ whose evaluation involves the solution of (1). This yields a $\widehat{\lambda}$,

$$\widehat{\lambda} = \arg\min_{\lambda} V(\lambda). \tag{4}$$

Solving the inverse problem even once is costly and therefore, finding the optimal parameter $\widehat{\lambda}$ is computationally intensive. In order to

* Corresponding author.
E-mail addresses: n.a.luiken@uu.nl (N. Luiken), t.vanleeuwen@uu.nl (T. van Leeuwen).
URLs: https://www.uu.nl/staff/NALuiken/0 (N. Luiken), https://www.uu.nl/staff/TvanLeeuwen (T. van Leeuwen).

[1] Nick Luiken and Tristan van Leeuwen have contributed to all parts of this work.

https://doi.org/10.1016/j.cageo.2020.104427
Received 23 July 2018; Received in revised form 19 January 2020; Accepted 29 January 2020
Available online 1 February 2020

overcome this computational drawback, various methods for approximating $V(\lambda)$ have been proposed.

1.1. Approach

The general approach is to approximate $V(\lambda)$ in such a way that it is cheaper to evaluate and thus allows for efficient estimation of $\hat{\lambda}$. Evaluating V involves two main tasks; evaluating a weighted norm of a given vector

$$\mathbf{u}^T f_\lambda(A)\,\mathbf{u}, \tag{5}$$

and computing the trace of a matrix function

$$\text{trace}\left(f_\lambda(A)\right). \tag{6}$$

Here A is a positive semi-definite matrix which is either G^TG or GG^T and $\mathbf{u} = \mathbf{d}$ or $\mathbf{u} = G^T\mathbf{d}$. $G \in \mathbb{R}^{m\times n}$ and for notational simplicity we write $A \in \mathbb{R}^{d\times d}$, where d is either m or n. An obvious approach is to replace G by a low-rank approximation G_k and use this to approximate (5) and (6). An important aspect is the approximation error, its influence on the approximation of V and ultimately on the estimated $\hat{\lambda}$. For most applications, it is not feasible to explicitly form a (truncated) Singular Value Decomposition (SVD) of G, so we will need to approximate the truncated SVD in order to obtain a reduced order model G_k. Traditionally, Krylov subspace method have been very popular for this purpose (Kilmer and O'Leary, 2001). These methods can be used to find upper and lower bounds for (5) as well (Golub and von Matt, 1995, 1997). Recently, randomized techniques have gained popularity. The Randomized Singular Value Decomposition (RSVD) was used for defining the reduced forward operator G_k (Wei et al., 2016; Xiang and Zou, 2013, 2015). The trace (6) can be estimated using randomized trace estimation (Hutchinson, 1990; Ubaru et al., 2017).

1.2. Contributions and outline

In this paper we compare the use of Krylov based subspace methods versus the use of RSVD for solving discrete inverse problems, specifically with regard to selecting the regularization parameter. We briefly present some parameter selection methods and review the Lanczos process and the RSVD and cite the relevant literature. For the Lanczos process we focus on the fact that we can obtain lower and upper bounds for the parameter selection methods. We provide error bounds for the parameter selection methods when approximated using the TSVD and the RSVD. We also discuss the use of Hutchinson's trace estimator for the parameter selection methods. We present a theorem that provides probabilistic bounds for the trace estimation combined with a low dimensional approximation with the Lanczos process, based on the work by Ubaru et al. (2017). We also show that obtaining guarantees for the accuracy of the trace estimator is too computationally expensive. We compare the randomized trace estimator to estimating the trace using the Ritz values obtained from the Lanczos process or the RSVD. In our numerical examples we present two examples from geosciences. The first example is a severely ill-posed problem and the second is a mildly ill-posed underdetermined problem. We discuss the performance of the Lanczos process and an adaptive RSVD algorithm for parameter selection and discuss the performance of the randomized trace estimator and the Lanczos/RSVD based trace estimator.

The paper is organized as follows. In Section 2 we review the necessary theory on parameter selection methods and the Lanczos process and the RSVD. In Section 3 we show template algorithms to obtain a lower dimensional approximation for two parameter selection methods. In Section 4 we discuss the performance of the algorithms for two model problems from geosciences. Lastly, in Section 5 we draw our conclusions.

Table 1
Parameter selection methods.

Method	$V(\lambda)$	Noise estimate
DP	$\left(\lambda^2\mathbf{d}^T\left(GG^T+\lambda I\right)^{-2}\mathbf{d}-\delta^2 m\right)^2$	Yes
GCV	$\dfrac{\lambda^2\mathbf{d}^T\left(GG^T+\lambda I\right)^{-2}\mathbf{d}}{(m^{-1}\text{trace}(I-GG_\lambda))^2}$	No
Reginska's rule	$\mathbf{d}^TG\left(G^TG+\lambda I\right)^{-2}G^T\mathbf{d}\cdot\lambda^2\mathbf{d}^T\left(GG^T+\lambda I\right)^{-2}\mathbf{d}$	No
QO	$\lambda^2\mathbf{d}^TG\left(G^TG+\lambda I\right)^{-4}G^T\mathbf{d}$	No

2. Theory

2.1. Parameter selection methods

In this section we review the parameter selection methods that we use in this work. Among the parameter selection methods there is an important distinction to be made: methods that require knowledge about the noise level in the data and/or the underlying model and methods that do not. Methods that do not require any knowledge on the model and/or the noise level on the data are sometimes called heuristic methods. We consider a standard deterministic approach, for a Bayesian approach to solving linear inverse problems see, e.g. Zunino and Mosegaard (2019), Neupauer and Borchers (2001) and Mejer Hansen and Mosegaard (2008). The main method we cover that requires knowledge on the noise level is the Discrepancy Principle (DP). We use a stochastic Gaussian noise model of the form

$$\mathbf{d} = \mathbf{d}_{\text{true}} + \xi, \tag{7}$$

where $\xi \sim \mathcal{N}(0, \delta^2 I)$ is uncorrelated Gaussian noise with mean zero and variance δ^2 and $\mathbf{d}_{\text{true}}$ is defined as $\mathbf{d}_{\text{true}} := G\mathbf{m}$, i.e., the true noiseless data. The methods we cover that do not require any knowledge on the noise level or the model are Generalized Cross Validation (GCV), Reginska's rule and the Quasi Optimality Criterion (QO). An overview of the functionals V corresponding to each method is shown in Table 1.

In order the express these in terms of the weighted norm (5) and trace (6) we use the following identities:

$$\|\hat{\mathbf{m}}_\lambda\|^2 = \mathbf{d}^TG\left(G^TG+\lambda I\right)^{-2}G^T\mathbf{d}, \tag{8}$$

and

$$\|G\hat{\mathbf{m}}_\lambda - \mathbf{d}\|^2 = \lambda^2\mathbf{d}^T\left(GG^T+\lambda I\right)^{-2}\mathbf{d}. \tag{9}$$

A derivation of these identities is included in appendix (Appendix A). Below, we briefly discuss each method in detail.

2.1.1. Reginska's rule

Reginska's rule (Reginska, 1996) is a variant of the well-known L-curve (Hansen, 1992). We choose to use Reginska's rule because it allows for a easier evaluation of the optimal λ by minimizing

$$V_{\text{RR}(\alpha)}(\lambda) = \mathbf{d}^TG\left(G^TG+\lambda I\right)^{-2}G^T\mathbf{d}\cdot\lambda^2\mathbf{d}^T\left(GG^T+\lambda I\right)^{-2}\mathbf{d} \tag{10}$$

It has been proven in Reginska (1996) that if the L-curve has maximal curvature at $\hat{\lambda}$ and has a tangent with slope $\hat{\alpha}$, then $V_{RR(\hat{\alpha})}$ has a minimizer at $\hat{\lambda}$. In practice, α is generally chosen to be 1.

2.1.2. Generalized cross validation

Generalized Cross Validation (GCV) was first introduced by Golub et al. (1979) as a method for choosing the regularization parameter and is an alternative to UPRE (Section 2.4) when the noise level is not known. It is important to note that although the noise level need not be known, there is an underlying assumption of a white Gaussian noise model (Wahba, 1990). The GCV estimates the optimal λ by minimizing

$$V_{\text{GCV}}(\lambda) = \frac{\lambda^2\mathbf{d}^T\left(GG^T+\lambda I\right)^{-2}\mathbf{d}}{(m^{-1}\text{tr}(I-GG_\lambda))^2}. \tag{11}$$

The idea behind GCV is that it tries to estimate λ in such a way that the data is explained well, while preventing overfitting. It is known the GCV has desirable statistical properties, but that it deals poorly with correlated noise. It also tends to undersmooth solutions. For further issues we refer the reader to Golub et al. (1979), Hansen (1992), Wahba (1990), Varah (1983) and Hansen (2010). There exist a few variants of the GCV that are in a sense weighted forms of the GCV and overcome some of the drawbacks of the GCV. All variants have been shown to be more stable than the GCV (Chung et al., 2008; Lukas, 2006, 2008) in the sense that they emphasize the generally flat minimum of the GCV by making it more pronounced. The Unbiased Predictive Risk Estimator (UPRE) (Vogel, 2002), also known as Mallow's C_p (Mallows, 1973), is based on the predictive risk. It is in a sense the predecessor of the GCV, as the GCV was developed as a noise-free alternative to the UPRE (Wahba, 1990).

2.1.3. The Discrepancy Principle

The Discrepancy Principle is an easy to use method that was first introduced by Morozov (1984a). The optimal λ found by the Discrepancy Principle is the λ for which the residual equals the noise level, i.e.

$$\mathbf{d}^T \left(GG^T + \lambda I\right)^{-2} \mathbf{d} = \eta\delta^2 m,$$

where $\eta \geq 1$ is a user-defined constant. The parameter η is introduced to prevent oversmoothing of the solution. We can cast this into the desired form by introducing

$$V_{\mathrm{DP}}(\lambda) = \left(\mathbf{d}^T \left(GG^T + \lambda I\right)^{-2} \mathbf{d} - \eta\delta^2 m\right)^2.$$

It is known that the Discrepancy Principle generally tends to oversmooth the solution (Hansen and O'Leary, 1993), i.e. the value for λ is too large. Another drawback is that the estimate of the noise level has to be accurate, and that small errors in the estimate can lead to large deviations in the solution (Hansen, 1998).

2.1.4. Quasi-optimality criterion

The quasi-optimality criterion is one of the first heuristic parameter choice criteria (Bakushinskii, 1981; Leonov, 1978, 1991; Morozov, 1984b). The λ estimated by the quasi-optimality criterion is the minimizer of

$$V_{QO}(\lambda) = \lambda^2 \mathbf{d}^T G \left(G^T G + \lambda I\right)^{-4} G^T \mathbf{d}. \tag{12}$$

For a derivation of this expression we refer the reader to Engl et al. (1996).

2.2. Model order reduction and trace estimation

In this section we review various methods for approximation quantities of the form

$$W(A) = \mathbf{w}^T f_\lambda(A)\mathbf{w},$$

and

$$T(A) = \mathrm{trace}\left(f_\lambda(A)\right),$$

where $f_\lambda(x) = (x+\lambda)^{-p}, p \in \mathbb{N}$ and $A \in \mathbb{R}^{d\times d}$ is a symmetric positive semi-definite (SPSD) matrix. We define a matrix function in the conventional sense. Given the eigenvalue decomposition $A = Q\Lambda Q^T$, the function is defined as $Qf_\lambda(\Lambda)Q^T$, where $f_\lambda(\Lambda)$ is a diagonal matrix with $f_\lambda(\lambda_i)$ as its entries.

2.2.1. Truncated SVD

In this section we provide bounds for the parameter selection rules based on the Truncated SVD (Hansen, 1987). They will be the basis for the error bounds derived for the RSVD which will be presented in Section 2.3.1.

Theorem 1. *Let $W(A) = \mathbf{w}^T f_\lambda(A)\mathbf{w}$ and $T(A) = \mathrm{trace}\left(f_\lambda(A)\right)$. Let $A = \sum_{i=1}^d \sigma_i^2 \mathbf{u}_i\mathbf{u}_i^T$ where $(\sigma_i^2, \mathbf{u}_i)$ denotes an eigenpair of A and let $A_k = \sum_{i=1}^k \sigma_i^2 \mathbf{u}_i\mathbf{u}_i^T$. Then the relative errors are bounded by*

$$\frac{\left|W(A) - W(A_k)\right|}{W(A)} \leq (d-k)\frac{p}{\lambda^{p+1}}\frac{\sigma_{k+1}}{f_\lambda(\sigma_d)}. \tag{13}$$

$$\frac{\left|T(A) - T(A_k)\right|}{T(A)} \leq \frac{(d-k)f_\lambda(\sigma_{k+1})}{\sum_{i=1}^d f_\lambda(\sigma_i)}. \tag{14}$$

Proof. Using a standard error estimate using the Taylor expansion we obtain:

$$\begin{aligned}\left|W(A) - W(A_k)\right| &\leq \sup_x \left|f_\lambda'(x)\right| \left|\mathbf{w}^T \left(A - A_k\right)\mathbf{w}\right| \\ &= \sup_x \left|f_\lambda'(x)\right| \left|\mathbf{w}^T U_{d-k+1}\Sigma_{d-k+1}U_{d-k+1}^T\mathbf{w}\right| \\ &\leq \sup_x \left|f_\lambda'(x)\right| (d-k)\|\mathbf{w}\|^2\sigma_{k+1}\end{aligned}$$

We have an explicit expression for $\sup_x \left|f_\lambda'(x)\right|$, given by:

$$\sup_x \left|f_\lambda'(x)\right| = \sup_x \left|-p(x+\lambda)^{-p-1}\right| = \frac{p}{\lambda^{p+1}}.$$

A different bound can be obtained by making the following observation:

$$\begin{aligned}\mathbf{w}^T f_\lambda(A)\mathbf{w} &= \mathrm{trace}(\mathbf{w}^T f_\lambda(A)\mathbf{w}) \\ &= \mathrm{trace}(f_\lambda(A)\mathbf{w}\mathbf{w}^T) \\ &= \mathrm{trace}(U f_\lambda(\Sigma)U^T\mathbf{w}\mathbf{w}^T).\end{aligned}$$

Now using the fact that $\mathbf{w}\mathbf{w}^T$ is a rank 1 matrix with eigenvalue $\|\mathbf{w}\|^2$, we can use von Neumann's trace inequality to obtain:

$$\|\mathbf{w}\|^2 f_\lambda(\sigma_d) \leq \mathbf{w}^T f_\lambda(A)\mathbf{w} \leq \|\mathbf{w}\|^2 f_\lambda(\sigma_1).$$

Putting both inequalities together we obtain (13). For $T(A)$ we get

$$\frac{\left|T(A) - T(A_k)\right|}{T(A)} = \frac{\sum_{i=k+1}^d f_\lambda(\sigma_i)}{\sum_{i=1}^d f_\lambda(\sigma_i)} \leq \frac{(d-k)f_\lambda(\sigma_{k+1})}{\sum_{i=1}^d f_\lambda(\sigma_i)} \quad \square$$

It is important to note that the above error estimate depends on λ. As $\lambda \to 0$, $f_\lambda' \to \infty$. However, given a certain $\lambda > 0$, there exists a bound for the derivative, but it will become large for small λ. This means that there is an inverse relation between k and λ: for large λ k can be small, whereas for small λ, k has to be large.

2.2.2. Krylov methods and Gauss quadrature

The approach makes use of the fact that the quantity $\mathbf{w}^T f(A)\mathbf{w}$ can be written as an integral with a certain measure, i.e.

$$\mathbf{w}^T f_\lambda(A)\mathbf{w} = \int_a^b f_\lambda(x)\mathrm{d}\omega(x), \tag{15}$$

where ω is a piecewise constant measure with discontinuities at the eigenvalues of A. A short intuitive explanation of this equality is given in (Appendix B). The integral can be approximated by a quadrature rule of the form

$$\int_a^b f_\lambda(x)\mathrm{d}\omega(x) = \sum_{i=1}^k w_i f_\lambda(x_i) + E_k(f_\lambda) := I_k(f_\lambda) + E_k(f_\lambda), \tag{16}$$

where $I_k(f)$ denotes the approximation with k nodes and $E_k(f)$ the associated error. The w_i are the weights and the x_i are the nodes. The weights and nodes for the Gauss quadrature rule are chosen such that the quadrature rule is exact for all polynomials of degree $2k$. It can be shown that the there is no quadrature rule that is exact for all polynomials of order larger than $2k$. A variant, the Gauss–Radau rule, fixes one node, which means that the Gauss–Radau rule is exact for polynomials up to degree $2k-1$. The errors for the k-point Gauss rule (E_k) and the k-point Gauss–Radau rule ($\widetilde{E}_k$) are given by (Stoer and Bulirsch, 1983; Golub and Meurant, 2009):

$$E_k(f) = \frac{f^{(2k)}(\xi_1)}{(2k)!}\sum_{i=1}^m \mathbf{u}_i^T\mathbf{w}\left[\prod_{j=1}^k(\sigma_i^2 - \theta_j^{(k)})\right]^2, \tag{17}$$

$$\widetilde{E}_k(f) = \frac{f^{(2k-1)}(\xi_2)}{(2k-1)!}\sum_{i=1}^{m}\mathbf{u}_i^T\mathbf{w}(\sigma_i^2-a)^2\left[\prod_{j=2}^{k}(\sigma_i^2-\theta_j^{(k)})\right]^2. \tag{18}$$

Recall that the parameter selection methods are functions of the form $f_\lambda(x) = (x+\lambda)^{-p}$ where $p \in \mathbb{N}$, typically, $p = 1, 2$ or 4. The derivatives for this class of functions are:

$$f_\lambda^{(2k)}(x) = (-1)^{(2k)}q(q+1)\cdots(q+2k-1)(x+\lambda)^{-(q+2k)} > 0 \tag{19}$$

$$f_\lambda^{(2k-1)}(x) = (-1)^{(2k-1)}q(q+1)\cdots(q+2k-2)(x+\lambda)^{-(q+2k-1)} < 0 \tag{20}$$

The nodes and weights for the Gauss quadrature are obtained by the eigendecomposition of the tridiagonal matrix T_k, which obtained by Lanczos tridiagonalization with starting vector $\mathbf{w}$. Let $T_k = Q\Lambda Q^T$, then the nodes of the quadrature are given by the eigenvalues and the weights are given by the first entry of the corresponding eigenvector.

2.3. Evaluating the Gauss and Gauss–Radau rule

Let T_k be the tridiagonal matrix obtained by the Lanczos process with starting vector $\mathbf{w}$. For a general form $\mathbf{w}^T f(A)\mathbf{w}$ the k-point Gauss quadrature rule is given by (Calvetti et al., 1999):

$$I_k(f) = \sum_{i=1}^{k} w_i f(x_i) = \|\mathbf{w}\|^2\sum_{i=1}^{k} f(\lambda_i)(\mathbf{e}_1^T Q\mathbf{e}_i)^2 \tag{21}$$

$$= \|\mathbf{w}\|^2\mathbf{e}_1^T Q^T f(\Lambda)Q\mathbf{e}_1 \tag{22}$$

$$= \|\mathbf{w}\|^2\mathbf{e}_1^T f(T_k)\mathbf{e}_1. \tag{23}$$

The functions that have to be evaluated are either functions of the form $\mathbf{w}^T G(G^TG+\lambda I)^{-p}G^T\mathbf{w}$ or $\mathbf{w}^T(GG^T+\lambda I)^{-p}\mathbf{w}$. For functions of the form $\mathbf{w}^T(GG^T+\lambda I)^{-p}\mathbf{w}$ the matrix T_k is obtained by using the Lanczos bidiagonalization process with starting vector $\mathbf{b}$. Let B_k denote the lower bidiagonal matrix obtained by the Lanczos bidiagonalization algorithm and let $\bar{B}_k$ be B_k with its last column removed. Then the Gauss and Gauss–Radau are obtained by (Calvetti et al., 1999):

$$I_k(f) = \|\mathbf{w}\|^2\mathbf{e}_1^T f(B_kB_k^T)\mathbf{e}_1, \tag{24}$$

$$\tilde{I}_k(f) = \|\mathbf{w}\|^2\mathbf{e}_1^T f(\overline{B}_{k-1}\overline{B}_{k-1}^T)\mathbf{e}_1. \tag{25}$$

For functions of the form $\mathbf{w}^T G(G^TG+\lambda I)^{-p}G^T\mathbf{w}$ the matrix T_k can still be obtained by the Lanczos lower bidiagonalization process, however, it has to be slightly modified. Let B_k be the lower bidiagonal matrix obtained by the Lanczos bidiagonalization process. Let $B_k = Q\widetilde{B}_k$ be the QR decomposition of B_k. Then the Gauss and Gauss–Radau rules are obtained by (Calvetti et al., 1999):

$$I_k(f) = \|\mathbf{w}\|^2\mathbf{e}_1^T f(\widetilde{B}_k\widetilde{B}_k^T)\mathbf{e}_1, \tag{26}$$

$$\tilde{I}_k(f) = \|\mathbf{w}\|^2\mathbf{e}_1^T f(\overline{\widetilde{B}}_{k-1}\overline{\widetilde{B}}_{k-1}^T)\mathbf{e}_1. \tag{27}$$

The QR decomposition can be carried out in $\mathcal{O}(k)$ steps. Alternatively, the matrix $T_k = \widetilde{B}_k\widetilde{B}_k^T$ for the Gauss and Gauss–Radau rule for functions of the form $\mathbf{w}^T G(G^TG+\lambda I)^{-p}G^T\mathbf{w}$ may be obtained by the Lanczos upper bidiagonalization algorithm (Golub and von Matt, 1997).

2.3.1. Randomized SVD

In this section we present the algorithms that are used to compute the RSVD. Moreover, we provide error bounds for the parameter selection methods. Although the RSVD algorithm has been used before for the purpose of solving discrete ill-posed problems, see e.g. Xiang and Zou (2013, 2015) and Vatankhah et al. (2018), the algorithms presented there are fixed rank algorithms in the sense that they return an RSVD given an a-priori target rank. Here, we use an two-step algorithm from Halko et al. (2011) which similar to the Lanczos algorithm is iterative in nature. The first step is to extract a good approximation to the range of G, which is done iteratively. The second step is to extract the RSVD. The first step of the algorithm, called the Adaptive Randomized Range Finder (Halko et al., 2011, algorithm 4.2), is presented in 1. The second step is shown in 2. We show the RSVD algorithm, taken from Halko et al. (2011), in Algorithm 2. We now present the error bounds for the parameter selection methods for the RSVD.

Algorithm 1 Randomized Range Finder (Algorithm 4.2 from Halko et al., 2011)

Require: General matrix $G \in \mathbb{R}^{m\times n}$, tolerance ϵ and an integer r.
Ensure: Matrix Q_k s.t. $\|G - Q_kQ_k^TG\| < \epsilon$ holds with probability at least $1-\min\{m,n\}10^{-r}$.

```
Draw a standard normally distributed matrix Ω ∈ R^{n×r}.
Compute Y = GΩ.
Set j = 0. Q_0 is empty.
while max{‖y^(j+1)‖, ..., ‖y^(j+r)‖} > δ/(10√(1/2π)) do
    j = j+1.
    y^(j) = y^(j) − Q_{j−1} Q_{j−1}^T y^(j).
    q^(j) = y^(j)/‖y^(j)‖.
    Q_j = [Q_{j−1} q^(j)].
    Draw a random vector ω^(j+r).
    y^(j+r) = (I − Q_j Q_j^T) G ω^(j+r).
    Orthogonalize y^(j+1), ..., y^(j+r−1) against q^(j).
end while
```

Algorithm 2 RSVD algorithm (Algorithm 5.1 from Halko et al., 2011)

Require: General matrix $G \in \mathbb{R}^{m\times n}$, tolerance ϵ and an integer r.
Ensure: $G \approx U_k\Sigma_l V_k^T$, U and V are orthonormal and Σ_k diagonal.

```
Compute Q_k using algorithm 1.
Compute B = Q_k^T G.
Compute the SVD of B: B = Ũ_k Σ_k V_k^T.
Compute U = Q_k Ũ.
```

Theorem 2 (*Adapted from Halko et al., 2011, Corollary 10.9*)**.** *Let $W(A) = \mathbf{w}^T f_\lambda(A)\mathbf{w}$ and $T(A) = trace\left(f_\lambda(A)\right)$. Let $A = \sum_{i=1}^n \sigma_i^2\mathbf{u}_i\mathbf{u}_i^T$ where $(\sigma_i^2, \mathbf{u}_i)$ denotes an eigenpair of A. Let $\widetilde{A}_k = \widetilde{U}_k\widetilde{\Sigma}_k\widetilde{V}_k^T$ with $\widetilde{S}_k = diag\left(\widetilde{\sigma}_1,\dots,\widetilde{\sigma_k}\right)$ be the RSVD of A given by Algorithm 2. Then the relative errors are bounded by*

$$\frac{|W(A)-W(A_k)|}{W(A)} \le \frac{p}{\lambda^{p+1}}\frac{\left(1+8\sqrt{(k+p)p\log p}\right)\sigma_{k+1}+3\sqrt{k+p}\left(\sum_{j>k}\sigma_j^2\right)^{1/2}}{f_\lambda(\sigma_m)} \tag{28}$$

with failure probability at most $6p^{-p}$,

$$\frac{|T(A)-T(A_k)|}{T(A)} \le \frac{p}{\lambda^{p+1}}\frac{\left(\sum_{i=1}^k(\sigma_i^2-\widetilde{\sigma}_i^2)+\sum_{i=k+1}^m\sigma_i^2\right)}{\sum_{i=1}^m f_\lambda(\sigma_i)}. \tag{29}$$

Proof. The proof is similar to the proof of Theorem 1xcept that the errors between A and A_k are now determined by the RSVD algorithm. The error bound is directly taken from Halko et al. (2011, Corollary 10.9). □

Note that the RSVD in this theorem is the RSVD of A, which is either GG^T or G^TG. In practice we use the RSVD of G.

2.3.2. Randomized trace estimation

In this section we discuss estimating the trace. The trace of a symmetric positive definite matrix A can be estimated by a randomization approach, using Hutchinson's trace estimator (Hutchinson, 1990), given by:

$$T(A) \approx T_N(A) := \frac{1}{N}\sum_{i=1}^{N}\mathbf{v}_i^T f_\lambda(A)\mathbf{v}_i, \tag{30}$$

This estimator is an unbiased estimator for the trace (Golub and von Matt, 1995, theorem 1). The entries of the vector $\mathbf{v}_i$ are chosen according to a uniform distribution on the interval $[0, 1]$. Let t_i denote the ith number drawn from this distribution, then the entries of $\mathbf{v}$ are given by

$$v_i = \begin{cases} +1 \text{ if } t_i \geq 1/2 \\ -1 \text{ if } t_i < 1/2 \end{cases} \tag{31}$$

The vectors $\mathbf{v}$ drawn from the distribution (31) are referred to as Rademacher vectors. For a matrix function, we can estimate its trace by the quantity $\mathbf{v}^T f_\lambda(A)\mathbf{v}$. To increase the accuracy of the estimator, the trace can be estimated by $V_N^T f_\lambda(A) V_N$, where V_N is a matrix with N columns and each column is of the form (31). We cannot bound the trace exactly, but there exist probabilistic bounds for the trace estimator. In Ubaru et al. (2017) a probabilistic bound for combined randomized trace estimation and model order reduction through Krylov subspaces is presented. The authors present an a priori bound for the combined randomized trace estimator and Gauss quadrature. In our case, we have to rely on an a priori bound for the randomized trace estimator. However, the accuracy of the Gauss quadrature to the trace estimator is estimated a posteriori based on how close the lower and upper bound are. We state a theorem similar to theorem 4.1 in Ubaru et al. (2017). Our approximation is either the upper or the lower bound denoted by $\tilde{I}_k$ and I_k. The key to obtaining bounds is to split the error into two parts:

$$\left|T(A) - T_N(A_k)\right| \leq \left|T(A) - T_N(A)\right| + \left|T_N(A) - T_N(A_k)\right|.$$

The first term concerns the accuracy of the trace estimator itself. The second term is approximated using the lower and upper bounds by using the Lanczos process. For the first term there exist standard probabilistic bounds, (Ubaru et al., 2017; Hutchinson, 1990). The second term is bounded in Ubaru et al. (2017) for general functions using a different error bound for the Gauss quadrature. Here, we present an a posteriori bound based on the lower and upper bounds. We have

$$\left|T_N(A) - T_N(A_k)\right| \leq \left|\tilde{I}_k - I_k\right|.$$

To obtain a useful bound we now require

$$\left|\tilde{I}_k - I_k\right| \leq \frac{\epsilon}{2} T(A).$$

Using the fact that $f_\lambda(x) = (x+\lambda)^{-p}$, $p > 0$, we require

$$\left|\tilde{I}_k - I_k\right| \leq d\frac{\epsilon}{2} f_\lambda(\sigma_1) \leq \frac{\epsilon}{2} T(A).$$

This leads to the following theorem.

Theorem 3 (*Adapted from Ubaru et al., 2017*)**.** *Choose* $N \geq (24/\epsilon^2)\log(2/\eta)$ *number of starting Rademacher vectors. Carry out* k *iterations of the Lanczos process such that*

$$\left|\tilde{I}_k - I_k\right| \leq d\frac{\epsilon}{2} f_\lambda(\sigma_1).$$

Then the output $T_N(A_k)$ *is such that:*

$$Pr\left[\left|T(A) - T_N(A_k)\right| \leq \epsilon\left|T(A)\right|\right] \geq 1 - \eta. \tag{32}$$

Of course, σ_1 is not available. However, we can estimate σ_1 using the first Ritz value θ_1. Unfortunately though, there is no estimate available for the quantity $|\sigma_1 - \theta_1|$ that does not depend on the singular values of A. By standard convergence theory for the Lanczos process, we do know that the largest Ritz value converges to the largest singular value first. Therefore, we expect the approximation to be quite accurate and use θ_1 instead of σ_1 to calculate the error bound. It should be noted, however, that a theorem of this form is not particularly useful for this application. In the limit, i.e. $\epsilon, \eta \to 1$, we have $N \gtrsim 16$ already, which can be prohibitively expensive. Preferably, we would like to use very few random vectors. We will investigate the impact of increasing N for small N (roughly 1–10) in our numerical experiments. It has been reported before in Bai et al. (1996) that $N = 1$ has the optimal trade-off between computational complexity and accuracy. In our numerical experiments we investigate the influence of increasing N for small N.

2.3.3. Obtaining the solution

It is important be able to evaluate the parameter selection methods quickly in order to obtain a suitable λ. However, we are ultimately interested in the solution to the problem and the question arises whether we can obtain the solution to the problem quickly using the reduced order model for the parameter selection method. For the RSVD this is trivial: we simply use the RSVD we have also used to evaluate the parameter selection method. For the Lanczos process we do the same thing. Theorem 3.1 from Kilmer and O'Leary (2001) shows that the solution obtained from Lanczos process for the norm of the solution is the same as the solution from Conjugate Gradient applied to $(G^TG + \lambda I)\mathbf{m} = G^T\mathbf{d}$. This can be obtained easily from the B_k obtained from evaluating the norm of the residual. Hence, for every parameter selection method we can easily obtain a solution to the problem with an estimate for λ.

2.4. Computational costs

We compare the computational costs for Lanczos bidiagonalization to the presented RSVD algorithm in terms of FLOPs. We start with Lanczos bidiagonalization. The costs for the standard Lanczos bidiagonalization algorithm for a matrix $G \in \mathbb{R}^{m\times n}$ are

$$k\text{nnz}(A)(m+n) + 5(m+n), \tag{33}$$

where the first term is for the matrix–vector multiplication and the second term is for various subtractions, divisions and taking the norm of vectors. It should be noted that the Lanczos bidiagonalization algorithm is known to be unstable, i.e. the orthogonal bases lose orthogonality, and may require reorthogonalization (Golub and Van Loan, 2013). It has been shown that a loss of orthogonalization has a strong influence on the estimated eigenvalues, but the effect on the solution of a linear system is small (Hansen, 1998, page 158). The costs for Algorithm 1 are reported in Halko et al. (2011), section 6.2, and are

$$kmR + k\text{nnz}(A)n + k^2m, \tag{34}$$

where the first term is the cost for generating Ω, the second term is the cost of matrix–vector multiplication and the third term is the cost for the orthogonalization of Q, in this case done by the Gram–Schmidt algorithm. If, for numerical stability, we use Householder reflectors, this cost would increase to roughly $2k^2m - \frac{2}{3}k^3$ (Golub and Van Loan, 2013, section 5.2.2). The costs for extracting the SVD are $\mathcal{O}(mk^2)$ with the addition of $2mnk$ FLOPs for the multiplications with Q. The big advantage of the RSVD is the possibility to easily parallelize the computation and the fact that only one pass over the data is needed. The ability to parallelize makes that, although the amount of matrix–vector multiplications may be similar, the RSVD algorithm is faster in terms of computational time. When access to the matrix A is prohibitively expensive the RSVD is certainly the desired option. For an in depth discussion on this topic see Halko et al. (2011), section 6.2.

3. Algorithms

In this section we show a blueprint for an algorithm based on either Lanczos quadrature or the RSVD for selecting the regularization parameter. We wish to make some small clarifying notes. We use the sampling for λ in order to be able to detect if there is a minimizer and to easily check how close the upper and lower bounds are. If we find a minimizer where the upper and lower bounds are not close enough, we resample around the minimizer. We always check whether the minimizer is not at the boundary of the sampled λ. In this case we simply resample again. Minimizing only the upper bound halves the amount of evaluations we have to do for the sampled λ. Evaluating for a given λ is cheap, as it involves solving systems involving B_k.

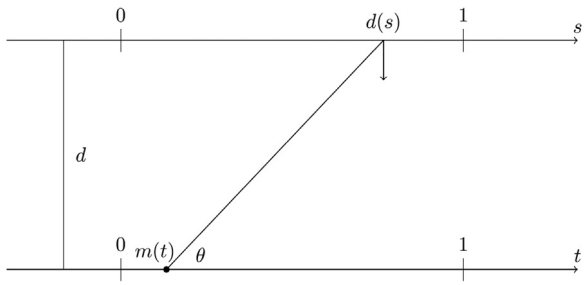

Fig. 1. Setup for the gravity problem. Figure is taken from Hansen (2010).

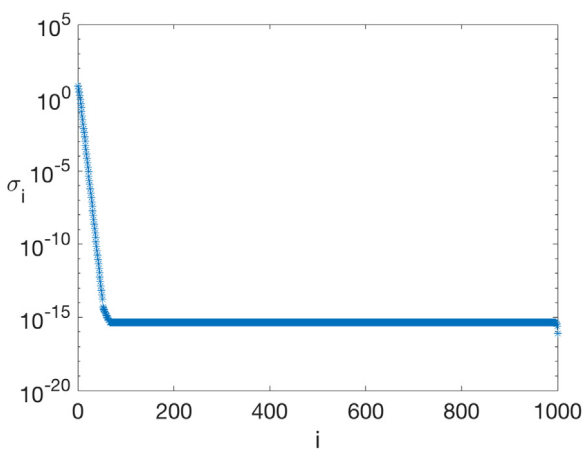

Fig. 2. Singular values of the matrix G for the gravity problem where $m = n = 1000$.

Table 2
Comparison of quadrature bounds versus the RSVD for $\epsilon = 10^{-1}$. $m = n = 1000$. Results are averages plus-minus one standard deviation over 10 different noise realizations.

Method	Lanczos		RSVD	
	$\frac{\|\lambda - \hat{\lambda}\|}{\|\lambda\|}$	k	$\frac{\|\lambda - \hat{\lambda}\|}{\|\lambda\|}$	k
GCV	$1.3 \cdot 10^{-2} \pm 1.2 \cdot 10^{-2}$	11.9 ± 1	$1.1 \cdot 10^{-2} \pm 1.2 \cdot 10^{-2}$	15
Reginska	$2.4 \cdot 10^{-3} \pm 2.5 \cdot 10^{-3}$	9.5 ± 0.7	$4.0 \cdot 10^{-4} \pm 9.0 \cdot 10^{-5}$	15
QO	$8.3 \cdot 10^{-3} \pm 8.8 \cdot 10^{-3}$	9 ± 0	$8.5 \cdot 10^{-3} \pm 5.1 \cdot 10^{-3}$	15
DP	$6.1 \cdot 10^{-4} \pm 1.6 \cdot 10^{-3}$	7.5 ± 1	$1.4 \cdot 10^{-2} \pm 1.3 \cdot 10^{-2}$	15

Table 3
Comparison of quadrature bounds versus the RSVD for $\epsilon = 10^{-2}$. $m = n = 1000$. Results are averages plus-minus one standard deviation over 10 different noise realizations.

Method	Lanczos		RSVD	
	$\frac{\|\lambda - \hat{\lambda}\|}{\|\lambda\|}$	k	$\frac{\|\lambda - \hat{\lambda}\|}{\|\lambda\|}$	k
GCV	$1.2 \cdot 10^{-3} \pm 1.1 \cdot 10^{-3}$	14.2 ± 1.4	$1.3 \cdot 10^{-3} \pm 1.8 \cdot 10^{-3}$	18
Reginska	$9.5 \cdot 10^{-4} \pm 1.4 \cdot 10^{-3}$	9.5 ± 0.5	$1.6 \cdot 10^{-5} \pm 1.1 \cdot 10^{-5}$	18
QO	$1.1 \cdot 10^{-3} \pm 2.2 \cdot 10^{-3}$	9.6 ± 0.5	$3.2 \cdot 10^{-5} \pm 1.7 \cdot 10^{-5}$	18
DP	$7.5 \cdot 10^{-6} \pm 1.4 \cdot 10^{-5}$	8.4 ± 1	$1.6 \cdot 10^{-4} \pm 4.7 \cdot 10^{-5}$	18

4. Numerical experiments

4.1. Gravity surveying

We consider the classical example of gravity surveying, see e.g. Hansen (2010). Let $m(t)$ be the gravity field at location t and $d(s)$ be the measured force at the surface at location s. Let h denote the depth of the gravity field. We then have the following relation

$$d(s) = \int_0^1 \frac{h}{(h^2 + (s-t)^2)^{3/2}} m(t) \mathrm{d}t.$$

Algorithm 3 Quadrature bounds for Reginska's rule.

Require: The data $\mathbf{d}$, matrix G and tolerance ϵ and a range of $\lambda \in [\lambda_{\min}, \lambda_{\max}]$.
Ensure: $\lambda^U_{\text{Reginska}}$ and $\lambda^L_{\text{Reginska}}$ with relative error ϵ.

```
while Not converged do
    Carry out a step of Lanczos bidiagonalization yielding B_{k+1,k} and B_{k,k}.
    Compute the matrices B̃_{k,k} and B̃_{k,k-1}.
    Compute the upper bound for the norm of the solution and the norm of the residual:
        ub_s(λ) = ||b||^2 e_1^T (B̃_{k,k-1} B̃_{k,k-1}^T + λI)^{-2} e_1
        ub_r(λ) = λ^2 ||b||^2 e_1^T (B_{k,k} B_{k,k}^T + λI)^{-2} e_1
    if upper bound yields a minimizer then
        Calculate lower bound at the minimizer.
        if relative error is smaller then ε then
            Compute λ^U_Reginska := min_λ ub^(k)_Reginska(λ) and λ^L_Reginska := min_λ lb^(k)_Reginska(λ).
            if |ub_Reginska(λ^U_Reginska) − lb_Reginska(λ^U_Reginska)| < ε and |λ^U_Reginska − λ^L_Reginska| < ε then
                break
            end if
        else
            Resample λ around the minimizer.
        end if
    end if
    k → k + 1.
end while
```

Table 4
Comparison of quadrature bounds versus the RSVD for $\epsilon = 10^{-3}$. $m = n = 1000$. Results are averages plus-minus one standard deviation over 10 different noise realizations.

Method	Lanczos		RSVD	
	$\frac{\|\lambda - \hat{\lambda}\|}{\|\lambda\|}$	k	$\frac{\|\lambda - \hat{\lambda}\|}{\|\lambda\|}$	k
GCV	$2.5 \cdot 10^{-4} \pm 2.1 \cdot 10^{-4}$	15.5 ± 1.7	$7.8 \cdot 10^{-4} \pm 2.4 \cdot 10^{-3}$	21
Reginska	$2.4 \cdot 10^{-5} \pm 2.5 \cdot 10^{-5}$	10 ± 0	$6.4 \cdot 10^{-7} \pm 7.9 \cdot 10^{-7}$	21
QO	$9.0 \cdot 10^{-5} \pm 1.2 \cdot 10^{-4}$	9.9 ± 0.7	$3.1 \cdot 10^{-5} \pm 1.9 \cdot 10^{-5}$	21
DP	$2.5 \cdot 10^{-8} \pm 1.8 \cdot 10^{-7}$	9.9 ± 0.6	$4.3 \cdot 10^{-5} \pm 5.3 \cdot 10^{-5}$	21

The problem is to retrieve $m(t)$ from measurements $d(s)$. We show the setup in Fig. 1. Because this is a Fredholm integral operator of the first kind, the problem of retrieving $m(t)$ is ill-posed. Specifically, the gravity surveying problem is severely ill-posed due to the severe decay of the singular values, as can be observed from Fig. 2. We regularize the problem using standard form Tikhonov regularization. We show the approximation error and dimension of the lower dimensional space for $\epsilon = 10^{-1}, 10^{-2}$ and 10^{-3} in Tables 2, 3 and 4 respectively. Because we use a probabilistic measure to check for convergence the size of Q_k may vary with different realizations.

In Fig. 3 we show the performance of the randomized trace estimator for varying N, where we have averaged over 10 realizations of the random vectors $\mathbf{u}$. We show the average of the 10 realizations and the dotted lines indicate one standard deviation. It is clear that with increasing N we obtain a better approximation on average. Moreover, the standard deviation drastically decreases. However, twenty random vectors is generally too computationally expensive and with regard to Theorem 3, does not give us any strong guarantees on how close it will be to the true trace. In Fig. 4 we show the trace estimator using the Ritz values. The accuracy of the trace estimator using the Ritz values rapidly increases as k increases. For large values of λ the trace is

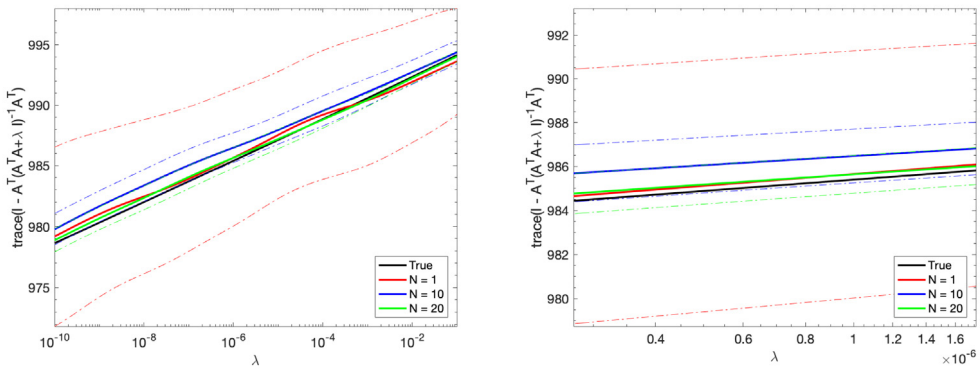

(a) Full view (b) Zoom

Fig. 3. Randomized trace estimator for increasing N for the gravity problem. We show the average for 10 random realizations and the dotted line are plus minus one standard deviation.

Algorithm 4 Quadrature bounds for GCV with randomized trace estimator.

Require: Data $\mathbf{d}$ and $U \in \mathbb{R}^{n\times N}$, matrix G and tolerance ϵ and a range of $\lambda \in [\lambda_{\min}, \lambda_{\max}]$.
Ensure: λ^U_{GCV} and λ^L_{GCV} with relative error ϵ.

while Not converged **do**
  Carry out a step of Lanczos bidiagonalization for starting vectors $\mathbf{b}, \mathbf{u}_1, \dots, \mathbf{u}_k$ yielding $B_{k+1,k}$ and $B_{k,k}$ for every starting vector.
  Sample λ. We choose 10 values on a log scale between $\lambda_{\min}$ and $\lambda_{\max}$.
  Compute the lower and upper bound for the norm of the residual and the trace estimators:

$$\text{ub\_r}(\lambda) = \lambda^2 \|\mathbf{b}\|^2 \mathbf{e}_1^T \left(B_{k,k}B_{k,k}^T + \lambda I\right)^{-2} \mathbf{e}_1$$

$$\text{ub\_u}_i(\lambda) = \|\mathbf{u}_i\|^2 \mathbf{e}_1^T \left(\widetilde{B}_{k,k}\widetilde{B}_{k,k}^T + \lambda I\right)^{-1} \mathbf{e}_1$$

  if upper bound yields a minimizer **then**
    Calculate lower bound at the minimizer.
    if relative error is smaller then ϵ **then**
      Compute $\lambda^U_{\text{GCV}} := \min_\lambda \text{ub}^{(k)}_{\text{GCV}}(\lambda)$ and $\lambda^L_{\text{GCV}} := \min_\lambda \text{lb}^{(k)}_{\text{GCV}}(\lambda)$.
      if $|\text{ub}_{\text{GCV}}(\lambda^U_{\text{GCV}}) - \text{lb}_{\text{GCV}}(\lambda^U_{\text{GCV}})| < \epsilon$ and $|\lambda^U_{\text{GCV}} - \lambda^L_{\text{GCV}}| < \epsilon$ **then**
        break
      end if
    else
      Resample λ around the minimizer.
    end if
  end if
  $k \to k+1$.
end while

Algorithm 5 Quadrature bounds for GCV with Ritz value based trace estimator.

Require: Data $\mathbf{d}$ and $U \in \mathbb{R}^{n\times N}$, matrix G and tolerance ϵ and a range of $\lambda \in [\lambda_{\min}, \lambda_{\max}]$.
Ensure: λ^U_{GCV} and λ^L_{GCV} with relative error ϵ.

while Not converged **do**
  Carry out a step of Lanczos bidiagonalization for starting vector $\mathbf{d}$ yielding $B_{k+1,k}$ and $B_{k,k}$.
  Sample λ. We choose 10 values on a log scale between $\lambda_{\min}$ and $\lambda_{\max}$.
  Compute the upper bound for the norm of the residual: $\text{ub\_r}(\lambda) = \lambda^2 \|\mathbf{b}\|^2 \mathbf{e}_1^T \left(B_{k,k}B_{k,k}^T + \lambda I\right)^{-2} \mathbf{e}_1$
  Estimate the trace using the Ritz values $B_k B_k^T = U\Theta V^T$, $\Theta = \text{diag}(\theta_1, \dots, \theta_k)$. $\text{trace}\left(I - A(A^TA + \lambda I)^{-1}A^T\right) \approx \lambda \sum_{i=1}^k \frac{1}{\theta_i + \lambda} + n - k := T(A_k)$
  if upper bound yields a minimizer **then**
    Calculate error with upper bound from previous iteration.
    if relative error is smaller then ϵ **then**
      Compute $\lambda^U_{\text{GCV}} := \min_\lambda \text{ub}^{(k)}_{\text{GCV}}(\lambda)$ and $\lambda^L_{\text{GCV}} := \min_\lambda \text{lb}^{(k)}_{\text{GCV}}(\lambda)$.
      if $|\text{ub}^{(k)}_{\text{GCV}}(\lambda^U_{\text{GCV}}) - \text{ub}^{(k-1)}_{\text{GCV}}(\lambda^U_{\text{GCV}})| < \epsilon$ and $|\lambda^U_{\text{GCV}} - \lambda^L_{\text{GCV}}| < \epsilon$ **then**
        break
      end if
    else
      Resample λ around the minimizer.
    end if
  end if
  $k \to k+1$.
end while

Algorithm 6 RSVD for any parameter selection method.

Require: Data $\mathbf{d}$, matrix G, a tolerance ϵ and an integer r.
Ensure: $\hat{\lambda}$ such that the inequality (28) holds.

1: Compute Q_k using the Adaptive Randomized Range Finder algorithm.
2: Compute the RSVD of $Q_k^T A = U_k \Sigma_k V_k^T$. Obtain $A \approx \widetilde{A}_k = Q_k U_k \Sigma_k V_k^T$.
3: Use $\widetilde{A}_k$ to evaluate the parameter selection methods.

well approximated early, but for small λ we need more iterations. It is important to note that we are interested in approximating the trace well around the optimal λ, which is unlikely to be very small. For $k = 30$ we have already obtained a near perfect approximation of the trace. The extra work needed in calculating the Ritz values is small, as we are computing the SVD of a $k \times k$ symmetric tridiagonal matrix. For $k = 14$ we have already obtained a satisfactory approximation of the trace, as the trace around the optimal λ is well approximated. This is due to the fact that the spectrum decays very quickly, as can be seen from Fig. 2.

4.2. Cross-well tomography

We consider classical linear cross-well tomography, an example taken from the AIR tools package (Hansen and Saxild-Hansen, 2012). We show the setup of the problem in Fig. 5. On the right are the sources and on the left are the receivers. We show the rays travelling from one source to all receivers. The data are the traveltimes from source i to

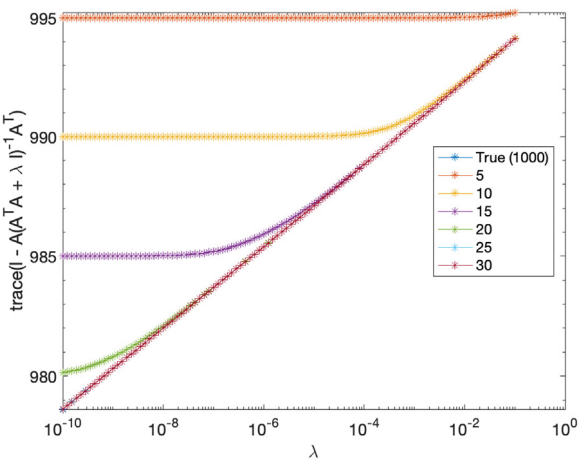

Fig. 4. Approximation of the trace using the Ritz values for increasing k.

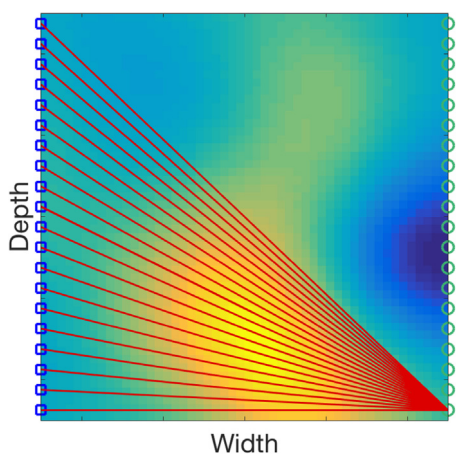

Fig. 5. Linear cross-well tomography.

receiver j and the goal is to reconstruct the well. We show the data and the well in Fig. 6.

Typically, for this setup we have far less sources and receivers than gridpoints. This means that the problem is underdetermined. The well has a smooth structure, and since the problem is underdetermined, we use general form Tikhonov regularization where L is the discrete Laplace operator: this enforces a smooth reconstruction. The problem is mildly ill-posed due to the fact that the singular values decay mildly,

Table 5
Comparison of quadrature bounds versus the RSVD. $m = 400$ and $n = 2500$. Results are averages plus-minus one standard deviation over 10 different noise realizations.

Method	Lanczos	
	$\frac{\|\lambda - \hat{\lambda}\|}{\|\lambda\|}$	k
GCV	$2.4 \cdot 10^{-1} \pm 1.7 \cdot 10^{-1}$	15 ± 2.5
Reginska	$3.1 \cdot 10^{-3} \pm 6.5 \cdot 10^{-3}$	9.4 ± 0.8
DP	$1.4 \cdot 10^{-3} \pm 7.6 \cdot 10^{-4}$	20.7 ± 4

as can be observed from Fig. 8. For the noise level we use $\delta = 10^{-1}$. For the Adaptive Randomized Range Finder we use a modified scheme based on the RSVD for underdetermined problems from Xiang and Zou (2015). Instead of using $A\Omega$ we use ΩA, or equivalently, $A^T\Omega$, to obtain the RSVD. For the Adaptive Randomized Range Finder we use the same parameters as for the gravity problem. Interestingly, the Adaptive Randomized Range Finder does not converge until we have obtained the full QR decomposition. We show the true errors $\|G - Q_k Q_k^T G\|_F$ for all k in Fig. 7. The performance of the Lanczos method is shown in Table 5. The Quasi-Optimality Criterion did not yield a minimizer, hence we have omitted this rule from the results. Notice that although the trace estimator seems rather accurate, there is still a considerable error compared to the optimal λ. This does not mean that the solution will necessarily be bad though.

We compare the randomized trace estimator versus the approximation based on the Ritz values for the GCV in Fig. 9. We show the accuracy of the randomized trace estimator for varying N in Fig. 10 and the approximation of the trace based on the Ritz values for increasing k in Fig. 11.

5. Conclusion

In this paper we have compared the use of the Lanczos process for parameter selection methods versus the use of the RSVD. We have derived bounds for the parameter selection methods when estimated using the RSVD. We have presented a theorem that provides probabilistic bounds for trace estimation combined with a low dimensional approximation obtained by the Lanczos method. This theorem provides us with certain guarantees in terms of accuracy. However, these guarantees require too many computations. We have compared the use of Lanczos quadrature and RSVD for two model problems from geosciences: gravity surveying and linearized cross-well tomography. We have also compared the use of Hutchinson's trace estimator versus the trace estimator based on the estimates for the singular values from the Lanczos process and the RSVD. The gravity surveying problem is severely ill-posed and we have shown that, for this problem, the

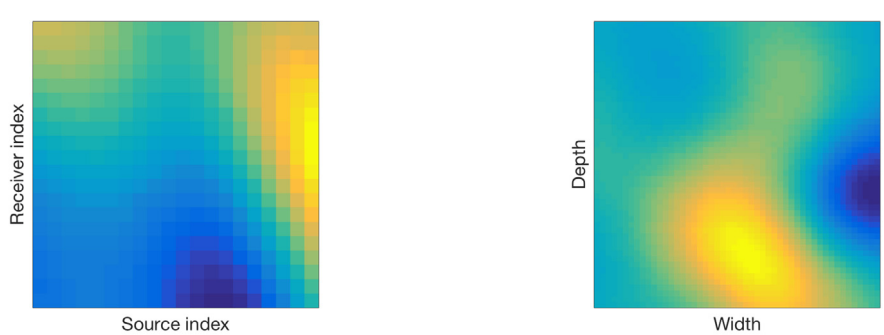

(a) Traveltimes: entry (i,j) indicates the the traveltime from source i to receiver j.

(b) The ground truth.

Fig. 6. Traveltimes and the well for the tomography problem.

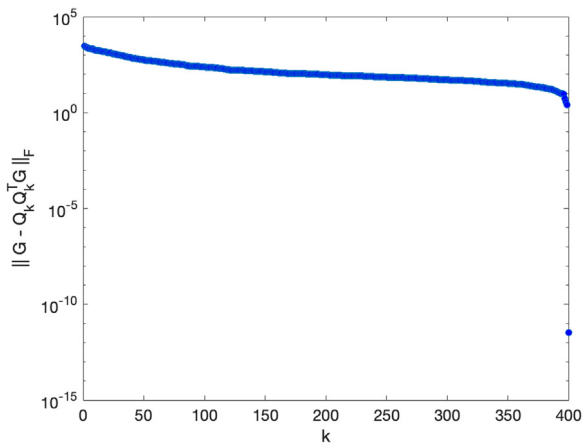

Fig. 7. $\|G - Q_k Q_k^T G\|_F$ for all k.

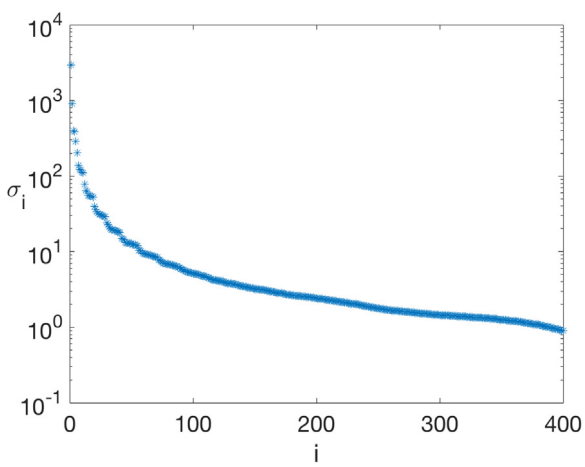

Fig. 8. Singular values of the matrix GL^{-1} for the tomography problem.

Lanczos quadrature method and the RSVD yield comparable results. We have also shown that the trace estimator based on the Ritz values of the Lanczos process (or the estimated singular values of the RSVD) outperforms the randomized trace estimator for the GCV. For the tomography problem, which is a mildly ill-posed underdetermined problem, we have shown that the RSVD failed to provide a satisfactory low dimensional approximation to evaluate the parameter selection methods. The Lanczos quadrature method was able to provide a lower dimensional approximation. The key difference is that, due to the fact that we obtain lower and upper bounds for the parameter selection methods, we obtain a lower dimensional model given the λ estimated by the parameter selection method. For the tomography problem this is a great advantage, because the optimal λ is quite large. A large λ allows for a lower dimensional approximation than a small λ, something which is reflected by the error bounds for the Lanczos quadrature method, and the bounds derived by us for the RSVD. For the tomography problem, we have shown that Hutchinson's trace estimator gives a far better approximation of the trace for small k than using the estimates obtained by the Lanczos procedure.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The first author is financially supported by the DELPHI consortium, The Netherlands and would like to thank the sponsors for their support. The second author is financially supported by the Netherlands Organization for Scientific Research (NWO) as part of research programme 613.009.032. Both authors would like to thank the reviewers for their valuable comments.

Appendix A

We give s short derivation of (9). We have

$$\|G\hat{\mathbf{m}}_\lambda - \mathbf{d}\| = \| \left(G(G^TG + \lambda I)^{-1}G^T - I\right)\mathbf{d}\|. \tag{A.1}$$

We now use the following relation:

$$(G^TG + \lambda I)^{-1}(G^TG + \lambda I)G^T = G^T \tag{A.2}$$

$$\iff (G^TG + \lambda I)^{-1}G^T(GG^T + \lambda I) = G^T \tag{A.3}$$

$$\iff (G^TG + \lambda I)^{-1}G^T = G^T(GG^T + \lambda I)^{-1} \tag{A.4}$$

Plugging this into (A.1) gives

$$\|G\hat{\mathbf{m}}_\lambda - \mathbf{d}\| = \| \left(GG^T(GG^T + \lambda I)^{-1} - I\right)\mathbf{d}\|. \tag{A.5}$$

Using the relation

$$(GG^T + \lambda I)(GG^T + \lambda I)^{-1} = I \tag{A.6}$$

$$\iff GG^T(GG^T + \lambda I)^{-1} = I - \lambda(GG^T + \lambda I)^{-1} \tag{A.7}$$

Plugging this into (A.5) yields the desired result

$$\|G\hat{\mathbf{m}}_\lambda - \mathbf{d}\|^2 = \lambda^2\mathbf{d}^T(GG^T + \lambda I)^{-2}\mathbf{d}. \tag{A.8}$$

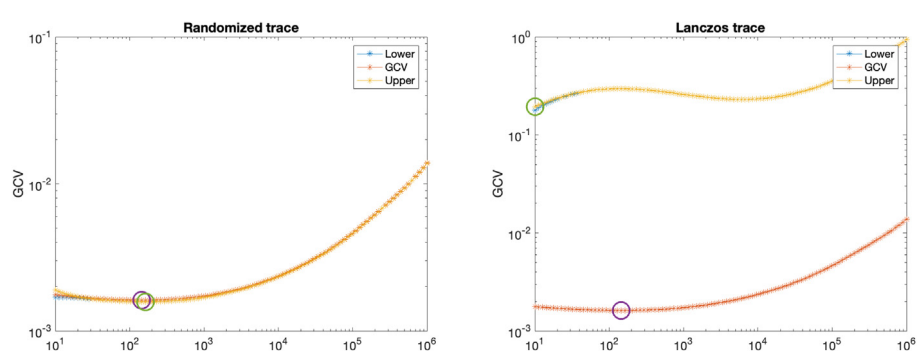

(a) GCV approximation using the trace estimator.

(b) GCV approximation using the Ritz values.

Fig. 9. Comparison of trace estimators for the GCV for $k = 30$. The circles denote the minimizers.

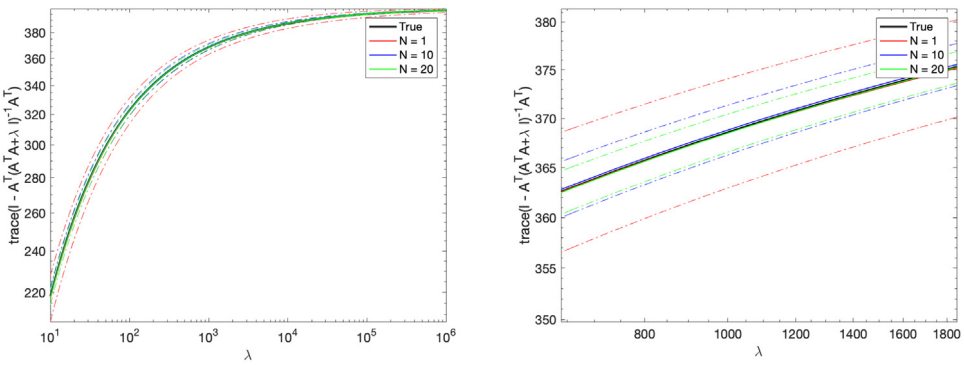

(a) Full view (b) Zoom

Fig. 10. Randomized trace estimator for increasing N for the tomography problem. We show the average for 10 random realizations and the dotted line are plus minus one standard deviation.

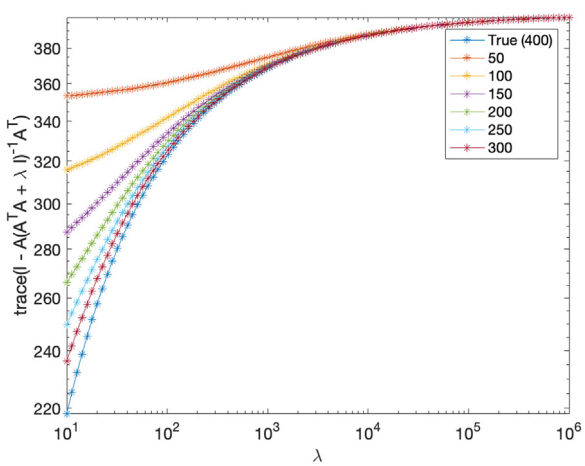

Fig. 11. Approximation of the trace using the Ritz values for increasing k.

Appendix B

In this section we describe the relation (15). Our aim is to give an explanation of how the piece-wise measure works for the reader that has no experience with measure theory, without any mathematical rigour, but simply to give an intuitive idea. To understand the measure in (15), it suffices to think of a measure as a weighted integral. Consider the following integral where $g(x)$ is a continuously differentiable function:

$$\int_{\Omega} f(x)\mathrm{d}g(x) = \int_{\Omega} f(x)\frac{\mathrm{d}g(x)}{\mathrm{d}x}\mathrm{d}x = \int_{\Omega} f(x)g'(x)\mathrm{d}x. \tag{B.1}$$

Hence, if the measure is a continuously differentiable function we can regard the measure as a weighted integral, where $g'(x)$ is the weight. Now if the function $g(x)$ is piecewise constant it is no longer differentiable. Consider the following function:

$$f(x) = \begin{cases} 0 \text{ if } 0 \le x \le 1 \\ 1 \text{ if } 0 < x \le 2 \end{cases} \tag{B.2}$$

The function is everywhere differentiable except at $x = 1$. Now, for $h > 0$

$$\frac{f(1+h) - f(1)}{h} = \frac{1}{h}. \tag{B.3}$$

If $h \to 0$ the function will go to infinity. We can regard the derivative of this function at $x = 1$ as a delta function. Hence, when integrating a function with a piecewise constant measure we can regard this as taking point evaluations. If we let the discontinuities be at the eigenvalues of a matrix, we get (15).

References

Bai, Z., Fahey, M., Golub, G., 1996. Some large-scale matrix computation problems. J. Comput. Appl. Math. 74, 71–89.

Bakushinskii, A., 1981. Remarks on choosing a regularization parameter usnig the quasi-optimality and ratio criterion. USSR Comput. Math. Math. Phys. 24, 181–182.

Bauer, F., Lukas, M.A., 2011. Comparing parameter choice methods for regularization of ill-posed problems. Math. Comput. Simulation 81 (9), 1795–1841.

Calvetti, D., Golub, G.H., Reichel, L., 1999. Estimation of the L-curve via lanczos bidiagonalization. BIT Numer. Math. 39 (4), 603–619.

Chung, J., Nagy, J.G., O'Leary, D.P., 2008. A weighted GCV method for lanczos hybrid regularization. Electron. Trans. Numer. Anal. 28, 149–167.

Engl, H.W., Hanke, M., Neubauer, A., 1996. Regularization of Inverse Problems. Kluwer Academic Publishers, Dordrecht.

Golub, G.H., Heath, M., Wahba, G., 1979. Generalized cross-validation as a method for choosing a good ridge parameter. Technometrics 21 (2), 215–223.

Golub, G.H., von Matt, U., 1995. Generalized cross-validation for large scale problems. J. Comput. Graph. Statist. 6, 1–34.

Golub, G.H., von Matt, U., 1997. Tikhonov regularization for large scale problems. In: Workshop on Scientific Computing. Springer, New York, pp. 3–26.

Golub, G.H., Meurant, G., 2009. Matrices, Moments and Quadrature with Applications. Princeton University Press, Princeton, NJ, USA.

Golub, G., Van Loan, C., 2013. Matrix Computations. In: Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press.

Halko, N., Martinsson, P., Tropp, J., 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. SIAM Rev. 53 (2), 217–288, URL https://doi.org/10.1137/090771806.

Hansen, P.C., 1987. The truncated SVD as a method for regularization. BIT 27, 534–553.

Hansen, P.C., 1992. Analysis of discrete ill-posed problems by means of the L-curve. SIAM Rev. 34, 561–580.

Hansen, P., 1998. Rank-Deficient and Discrete Ill-Posed Problems. Springer-Verlag, Phialdelphia.

Hansen, P.C., 2010. Discrete Inverse Problems: Insight and Algorithms. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.

Hansen, P.C., 2013. Oblique projections and standard-form transformations for discrete inverse problems. Numer. Linear Algebra Appl. 20, 250–258.

Hansen, P.C., O'Leary, D.P., 1993. The use of the L-curve in the regularization of discrete ill-posed problems. SIAM J. Sci. Comput. 14 (6), 1487–1503.

Hansen, P.C., Saxild-Hansen, M., 2012. AIR Tools — A MATLAB package of algebraic iterative reconstruction methods. J. Comput. Appl. Math. 236 (8), 2167–2178, Inverse Problems: Computation and Applications.

Hutchinson, M., 1990. A stochastic estimator for the trace of the influence matrix for Laplacian smoothing splines. Comm. Statist. Simulation Comput. 19, 433–450.

Kilmer, M.E., O'Leary, D.P., 2001. Choosing regularization parameters in iterative methods for ill-posed problems. SIAM J. Matrix Anal. Appl. 22 (4), 1204–1221.

Leonov, A., 1978. On the choice of regularization parameters by means of the quasi-optimality and ratio criteria. Svoiet Math. Dokl. 19, 537–540.

Leonov, A., 1991. On the accuracy of tikhonov regularizing algorithms and quasioptimal selection of a regularization parameter. Soviet Math. Dokl. 44, 711–716.

Lukas, M.A., 2006. Robust generalized cross-validation for choosing the regularization parameter. Inverse Problems 22 (5), 1883.

Lukas, M.A., 2008. Strong robust generalized cross-validation for choosing the regularization parameter. Inverse Problems 24 (3), 034006.

Mallows, C.L., 1973. Some comments on Cp. Technometrics 15, 661–675.

Mejer Hansen, T., Mosegaard, K., 2008. VISIM: Sequential simulation for linear inverse problems. Comput. Geosci. 34, 53–76.

Morozov, V., 1984a. Methods for Solving Incorrectly Posed Problems. Springer-Verlag, New York.

Morozov, V., 1984b. Methods for Solving Incorrectly Posed Problems. Springer Verlag, New York.

Neupauer, R., Borchers, B., 2001. A MATLAB implementation of the minimum relative entropy method for linear inverse problems. Comput. Geosci. 27, 757–762.

Reginska, T., 1996. A regularization parameter in discrete ill-posed problems. SIAM J. Sci. Comput. 17.

Stoer, J., Bulirsch, R., 1983. Introduction to Numerical Analysis. Springer Verlag.

Ubaru, S., Chen, J., Saad, Y., 2017. Fast estimation of $tr(f(a))$ via stochastic lanczos quadrature. SIAM J. Matrix Anal. Appl. 38 (4), 1075–1099, URL https://doi.org/10.1137/16M1104974.

Varah, J., 1983. Pitfalls in the numerical solution of linear ill-posed problems. SIAM J. Stat. Sci. Comput. 4 (2), 164–176.

Vatankhah, S., Renaut, R., Ardestani, V., 2018. A fast algorithm for regularized focused 3D inversion of gravity data using randomized singular-value decomposition. Geophysics 83 (4), 25–34.

Vogel, C.R., 2002. Computational Methods for Inverse Problems. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.

Wahba, G., 1990. Spline Models for Observational Data. Society for Industrial and Applied Mathematics, Philadelphia.

Wei, Y., Xie, P., Zhang, L., 2016. Tikhonov regularization and randomized GSVD. SIAM J. Matrix Anal. Appl. 37, 649–675.

Xiang, H., Zou, J., 2013. Regularization with randomized SVD for large-scale discrete inverse problems. Inverse Problems 29 (8), 085008.

Xiang, H., Zou, J., 2015. Randomized algorithms for large-scale inverse problems with general form Tikhonov regularization. Inverse Problems 29, 085088.

Zunino, A., Mosegaard, K., 2019. An efficient method to solve large linearizable inverse problems under Gaussian and separability assumptions. Comput. Geosci. 122, 77–86.