

REGULARIZATION PARAMETER ESTIMATION FOR LARGE SCALE TIKHONOV REGULARIZATION USING A PRIORI INFORMATION

ROSEMARY A RENAUT*, IVETA HNETYNKOVA†, AND JODI MEAD‡

Abstract. This paper is concerned with estimating the solutions of numerically ill-posed least squares problems through Tikhonov regularization. Given *a priori* estimates on the covariance structure of errors in the measurement data \mathbf{b} , and a suitable statistically-chosen σ , the Tikhonov regularized least squares functional $J(\sigma) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_{W_{\mathbf{b}}}^2 + 1/\sigma^2 \|D(\mathbf{x} - \mathbf{x}_0)\|_2^2$, evaluated at its minimizer $\mathbf{x}(\sigma)$, approximately follows a χ^2 distribution with \tilde{m} degrees of freedom. Here $\tilde{m} = m + p - n$ for $A \in \mathcal{R}^{m \times n}$, $D \in \mathcal{R}^{p \times n}$, matrix $W_{\mathbf{b}}$ is the inverse covariance matrix of the mean 0 normally distributed measurement errors \mathbf{e} in \mathbf{b} , and \mathbf{x}_0 is an estimate of the mean value of \mathbf{x} . Using the generalized singular value decomposition of the matrix pair $[W_{\mathbf{b}}^{1/2}AD]$, σ can then be found such that the resulting J follows this χ^2 distribution, Mead and Renaut (2008). Because the algorithm explicitly relies on the direct solution of the problem obtained using the generalized singular value decomposition it is not practical for large scale problems. Here the approach is extended for large scale problems through the use of the Newton iteration in combination with a Golub-Kahan iterative bidiagonalization of the regularized problem. The algorithm is also extended for cases in which \mathbf{x}_0 is not available, but instead a set of measurement data provides an estimate of the mean value of \mathbf{b} . The sensitivity of the Newton algorithm to the number of steps used in the Golub-Kahan iterative bidiagonalization, and the relation between the size of the projected subproblem and σ are discussed. Experiments presented contrast the efficiency and robustness with other standard methods for finding the regularization parameter for a set of test problems and for the restoration of a relatively large real seismic signal. An application for image deblurring also validates the approach for large scale problems. We conclude that the presented approach is robust for both small and large scale discretely ill-posed least squares problems.

Key words. ill-posed problems, Tikhonov regularization, χ^2 -distribution, Golub-Kahan iterative bidiagonalization, hybrid methods, Newton algorithm.

AMS subject classifications. 15A09, 15A29, 62F15, 62G08

1. Introduction. We are concerned with the solution of large-scale linear discrete ill-posed problems such as arise in many physical experiments associated, for example, with the discretization of integral equations [27, 9], and modeled by the ill-posed system of equations $\mathbf{A}\mathbf{x} = \mathbf{b}$. Matrix $A \in \mathcal{R}^{m \times n}$ results from the underlying model discretization and a solution $\mathbf{x} \in \mathcal{R}^n$ is desired for measurements $\mathbf{b} \in \mathcal{R}^m$, which are often noise-contaminated. An approximate solution $\hat{\mathbf{x}}$ may be obtained by solving the weighted regularized least squares problem, with matrix $D \in \mathcal{R}^{p \times n}$, $p \leq n$, chosen dependent on anticipated smoothness properties of the solution \mathbf{x} ,

$$(1.1) \quad \hat{\mathbf{x}} = \operatorname{argmin} J(\mathbf{x}) = \operatorname{argmin} \{ \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_{W_{\mathbf{b}}}^2 + \|D(\mathbf{x} - \mathbf{x}_0)\|_{W_{\mathbf{x}}}^2 \}.$$

*Supported by NSF grants DMS 0513214 and DMS 0652833. Arizona State University, Department of Mathematics and Statistics, Tempe, AZ 85287-1804. Tel: 480-965-3795, Fax: 480-965-4160. Email: renaud@asu.edu

†Supported by the research project MSM0021620839 financed by MSM. Charles University in Prague, Faculty of Mathematics and Physics, and Institute of Computer Science, Academy of Sciences of the Czech Republic. Tel: +420-22191-3362, Fax: +420-22481-1036. Email: hnetynkova@scs.cas.cz

‡Supported by NSF grant EPS 0447689, Boise State University, Department of Mathematics, Boise, ID 83725-1555. Tel: 208426-2432, Fax: 208-426-1354. Email: jmead@boisestate.edu

Notice that we consider the general case in which the fit to data term is measured in a weighted norm with weight matrix $W_{\mathbf{b}}$ which may be available experimentally. For example, when the measurement errors \mathbf{e} in the measurement data \mathbf{b} are assumed to be samples from a multivariate normal distribution, $N_m(0, C_{\mathbf{b}})$ (m variables, mean 0 and covariance $C_{\mathbf{b}}$), the weighting is the inverse covariance $W_{\mathbf{b}} = C_{\mathbf{b}}^{-1}$. For the case of colored noise $C_{\mathbf{b}}$ is diagonal, $C_{\mathbf{b}} = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_m^2)$, whereas for white noise $C_{\mathbf{b}} = \sigma^2 I_m$.

As presented here the regularization term $D(\mathbf{x} - \mathbf{x}_0)$ is also calculated in a weighted norm. If $W_{\mathbf{x}}$ can be assumed to be the inverse covariance matrix for normally distributed errors in the mapped model parameters $D\mathbf{x}$, then it can be shown that the minimum of the functional J follows a χ^2 distribution with $m - n + p$ degrees of freedom, [16, 17]. This extends the standard result on the χ^2 distribution of the unregularized least squares functional [23]. Also, if $D = I_n$ and $W_{\mathbf{x}} = 1/\sigma_{\mathbf{x}}^2 I_n$, where $\sigma_{\mathbf{x}}^2$ is the common white noise variance in model parameters \mathbf{x} , then $\hat{\mathbf{x}}$ is the standard *maximum a posteriori* (MAP) estimate of the solution, [27]. The advantage of assuming that weighting matrices $W_{\mathbf{b}}$ and $W_{\mathbf{x}}$ are inverse covariance matrices is that it permits our focus on the use of the χ^2 property of the minimum functional to efficiently determine the parameter $\sigma_{\mathbf{x}}$ when $W_{\mathbf{x}} = 1/\sigma_{\mathbf{x}}^2 I_n$.

The determination of the optimal $\lambda = 1/\sigma_{\mathbf{x}}$ is a topic of much previous research, and includes methods, amongst others, such as the L-curve, generalized cross validation (GCV), the unbiased predictive risk estimate (UPRE) and the discrepancy principle estimate of the distance between the exact and regularized solution [19], all of which are well-described in the literature, see e.g. [9, 27] for comparisons of the criteria and more references. Other recent approaches analyse the statistical properties of the residual of the least squares functional [24, 10]. Some especially promising efforts for determining λ , particularly for large scale problems, have placed emphasis on regularization methods based on iterative Golub-Kahan bidiagonalization, e.g., LSQR [21, 22] and hybrid methods [4, 9, 12, 13, 2, 20]. In hybrid methods the original problem is projected onto a lower dimensional subspace using the bidiagonalization algorithm, which by itself represents a form of regularization by projection. The projected bidiagonal problem, however, inherits a part of the ill-posedness, and therefore some form of inner regularization is applied to the projected *small* sub-problem. The bidiagonalization may be stopped when the regularized solution of the projected problem matches any of the previously mentioned stopping criteria. Hybrid methods seek to combine the bidiagonalization procedure with determination of an appropriate regularization parameter for solving the projected system. The hybrid method presented in this paper utilizes efficient

iterative bidiagonalization combined with a parameter search algorithm obtained from the χ^2 property of the regularized functional.

Incorporating the χ^2 property of the functional J to find λ is a recent innovation. Provided that the weighting matrix $W_{\mathbf{x}}$ is the appropriate inverse covariance matrix for the regularization term and that the prior information \mathbf{x}_0 approximates the mean of parameters \mathbf{x} , the χ^2 property implies that J lies within an interval centered around its expected value, [17]. Therefore a Newton root-finding algorithm can be designed for determining $\sigma_{\mathbf{x}}$. This algorithm may be implemented for small-scale problems using a direct solve employing the generalized singular value decomposition (GSVD) of the matrix pair $[W_{\mathbf{b}}^{1/2} A | D]$, [17]. Here the GSVD direct solve is replaced by the iterative bidiagonalization for the regularized problem. This algorithm has the benefit of reuse of the bidiagonal system for each step of the Newton iteration, namely for each value of $\sigma = 1/\lambda$ found in the Newton algorithm, and hence solves the regularized problem for optimal λ at almost no overhead as compared to a single solve for one given λ . Note that a general operator D , $D \neq I$, can also be considered, using the conversion to standard form regularization [9, 3]. Another possibility would be to use an algorithm for simultaneous bidiagonalization of A and D as presented in [13].

First, in Section 2 we present extensions of the theory, relevant for solving more general problems. For example, when \mathbf{x}_0 is not the expected value of the model data we show that the minimum functional follows a non-central χ^2 distribution, with centrality parameter related to the choice of \mathbf{x}_0 in relation to the actual mean $\bar{\mathbf{x}}$. Moreover, the use of a truncated (filtered) expansion of the GSVD for providing a filtered direct GSVD solution also modifies the theory, yielding a functional which has fewer degrees of freedom, dependent on the number of terms used in the filtered expansion. The original GSVD-based implementation of the Newton algorithm can be extended for both cases and implementation details are provided in Section 3. For large scale problems the same Newton algorithm is employed but solutions are obtained iteratively using the hybrid algorithm described in Section 3.2. Numerical experiments detailed in Section 4 contrast the performance of the iterative and direct solve algorithms for small-scale problems, for both central and non-central distributions of the underlying functional. These experiments validate both the algorithm for non-central functionals and for the large scale implementation. The dependence of the regularization parameter obtained in relation to the size of the projected problem is also discussed, Section 4.3. As the subproblem size increases, the solution admits higher frequency components and more regularization is needed. Finally, in Sections 4.4-4.5 the hybrid algorithm is shown to yield efficient and robust

results for the deblurring of a relatively large real seismic signal, and the deblurring of a large scale image for which the true image is available. Future work and conclusions are discussed in Section 5.

2. Theoretical Development.

2.1. χ^2 Distribution of the Regularized Functional. The solution of (1.1) with $W_{\mathbf{x}} = \lambda^2 I_n$, assuming invertibility

$$(2.1) \quad \mathcal{N}(A) \cap \mathcal{N}(D) = \emptyset,$$

is given by

$$(2.2) \quad \hat{\mathbf{x}}(\lambda) = (A^T W_{\mathbf{b}} A + \lambda^2 D^T D)^{-1} A^T W_{\mathbf{b}} \mathbf{r} + \mathbf{x}_0 = \mathbf{x}(\lambda) + \mathbf{x}_0,$$

where the residual \mathbf{r} is given by $\mathbf{r} = \mathbf{b} - A\mathbf{x}_0$. This is more compactly written using the **resolution** matrix $R(\lambda) = (A^T W_{\mathbf{b}} A + \lambda^2 D^T D)^{-1} A^T W_{\mathbf{b}}^{1/2}$, or, more generally, $R(W_{\mathbf{D}}) = (A^T W_{\mathbf{b}} A + W_{\mathbf{D}})^{-1} A^T W_{\mathbf{b}}^{1/2}$, where $W_{\mathbf{D}} = D^T W_{\mathbf{x}} D$, yielding $\hat{\mathbf{x}}(W_{\mathbf{D}}) = R(W_{\mathbf{D}}) W_{\mathbf{b}}^{1/2} \mathbf{r} + \mathbf{x}_0$. With this notation, and introducing the **influence** matrix $A(W_{\mathbf{D}}) = W_{\mathbf{b}}^{1/2} A R(W_{\mathbf{D}})$, (1.1) is written as a quadratic form,

$$(2.3) \quad J(\hat{\mathbf{x}}(W_{\mathbf{D}})) = \mathbf{r}^T W_{\mathbf{b}}^{1/2} (I_m - A(W_{\mathbf{D}})) W_{\mathbf{b}}^{-1/2} \mathbf{r}.$$

To obtain the result on the χ^2 distribution of this functional we employ the GSVD, using the notation as given in [17], to reexpress this quadratic form.

LEMMA 2.1. [5] *Assume the invertibility condition (2.1) and $m \geq n \geq p$. There exist unitary matrices $U \in \mathcal{R}^{m \times m}$, $V \in \mathcal{R}^{p \times p}$, and a nonsingular matrix $X \in \mathcal{R}^{n \times n}$ such that*

$$(2.4) \quad A = U \tilde{\Upsilon} X^T, \quad D = V \tilde{M} X^T,$$

where

$$\tilde{\Upsilon} = \begin{bmatrix} \Upsilon & 0 \\ 0 & I_{n-p} \\ 0 & 0 \end{bmatrix}, \quad \Upsilon = \text{diag}(v_1, \dots, v_p) \in \mathcal{R}^{p \times p},$$

$$\tilde{M} = \begin{bmatrix} M, & 0_{p \times (n-p)} \end{bmatrix}, \quad M = \text{diag}(\mu_1, \dots, \mu_p) \in \mathcal{R}^{p \times p},$$

and such that

$$(2.5) \quad 0 \leq v_1 \leq \dots \leq v_p \leq 1, \quad 1 \geq \mu_1 \geq \dots \geq \mu_p > 0,$$

$$v_i^2 + \mu_i^2 = 1, \quad i = 1, \dots, p.$$

Now, using the GSVD of the matrix pair $[W_{\mathbf{b}}^{1/2}A|W_{\mathbf{D}}^{1/2}]$,

$$(2.6) \quad J(\hat{\mathbf{x}}(W_{\mathbf{D}})) = \mathbf{r}^T W_{\mathbf{b}}^{1/2} U (I_m - \tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^T) U^T W_{\mathbf{b}}^{1/2} \mathbf{r}$$

$$(2.7) \quad = \sum_{i=1}^p s_i^2 \mu_i^2 + \sum_{i=n+1}^m s_i^2, \quad \mathbf{s} = U^T W_{\mathbf{b}}^{1/2} \mathbf{r}$$

$$(2.8) \quad = \|\mathbf{k}\|^2 - \sum_{i=p+1}^n k_i^2, \quad \mathbf{k} = Q U^T W_{\mathbf{b}}^{1/2} \mathbf{r}$$

$$(2.9) \quad Q = \text{diag}(\mu_1, \dots, \mu_p, I_{n-p}, I_{m-n}).$$

This is the starting point for showing that J follows a **central** χ^2 distribution with $m + n - p$ degrees of freedom as detailed in Theorem 3.1 [17], provided that \mathbf{x}_0 is the expected value, denoted by $\bar{\mathbf{x}}$, of \mathbf{x} . Typically, $\bar{\mathbf{x}}$ is unknown and $\mathbf{x}_0 = 0$ is chosen. For $\mathbf{x}_0 \neq \bar{\mathbf{x}}$ the following **non-central** generalization is obtained.

THEOREM 2.2 (Non Central χ^2 distribution of the Regularized Functional).

Suppose

- $C_{\mathbf{b}} = W_{\mathbf{b}}^{-1}$ is the symmetric positive definite (SPD) covariance matrix on the mean zero normally-distributed data error, e_i ,
- $C_{\mathbf{D}}$ is the, possibly rank deficient, symmetric positive (semi-)definite covariance matrix for the mean zero normally distributed model errors $\zeta_i = (D(\hat{\mathbf{x}}(W_{\mathbf{D}})_i - \mathbf{x}_0))_i$, with the conditional SPD inverse $W_{\mathbf{D}}$ which satisfies the two Moore-Penrose conditions $W_{\mathbf{D}} C_{\mathbf{D}} W_{\mathbf{D}} = W_{\mathbf{D}}$ and $C_{\mathbf{D}} W_{\mathbf{D}} C_{\mathbf{D}} = C_{\mathbf{D}}$,
- that the invertibility condition (2.1) holds, and
- that the expected value of \mathbf{x} is $\bar{\mathbf{x}}$.

Then in the limit $m - n + p$ sufficiently large, the minimum value of the functional J is a random variable which follows a **non-central** χ^2 distribution with $m - n + p$ degrees of freedom, and **non-centrality** parameter $c = \|\tilde{\mathbf{c}}\|_2^2 = \|\tilde{Q} U^T W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0)\|_2^2$, where

$$(2.10) \quad \tilde{Q} = \text{diag}(\mu_1, \dots, \mu_p, 0_{n-p}, I_{m-n}).$$

Equivalently, $J \sim \chi^2(m - n + p, c)$ has expected value $m - n + p + c$ and variance $2(m - n + p) + 4c$.

Proof. We use the Fisher-Cochran theorem for quadratic forms [23], to show that in the limit as the number of terms increases (2.6) follows a χ^2 distribution. Because (2.8) expresses the quadratic form in terms of the two norm of the vector \mathbf{k} , excepting the components $p + 1 : n$, it is sufficient to consider the distributions of the relevant components k_i of \mathbf{k} . The argument follow as in [17] but the statistical argument is modified when \mathbf{x}_0 is

not the mean. In particular, because the data and model errors are mean zero and normally distributed, the expected value of the random variable \mathbf{b} is $\bar{\mathbf{b}} = A\bar{\mathbf{x}}$, so that $\bar{\mathbf{r}} = A(\bar{\mathbf{x}} - \mathbf{x}_0)$, and \mathbf{k} has mean $\mathbf{c} = QU^T W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0)$. Thus $\mathbf{b} \sim N_m(A\mathbf{x}_1, C_{\mathbf{b}} + AC_{\mathbf{D}}A^T)$ and $W_{\mathbf{b}}^{1/2} \mathbf{r} \sim N_m(W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0), I_m + W_{\mathbf{b}}^{1/2} AC_{\mathbf{D}}A^T W_{\mathbf{b}}^{1/2})$. But, using the GSVD, $C_{\mathbf{D}} = (X^T)^{-1} \text{diag}(M^{-2}, 0_{n-p}) X^{-1}$, and $I_m + \tilde{A}C_{\mathbf{D}}\tilde{A}^T = UQ^{-2}U^T$. Therefore,

$$(2.11) \quad \mathbf{k} \sim N_m(QU^T W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0), I_m).$$

Now, by (2.8), the centrality parameter is calculated excluding the component means of vector \mathbf{k} for components $p+1 : n$ and the result follows, [23]. \square

REMARK 1. *Theoretically, for $\mathbf{b} \in \text{range}(A)$, $\mathbf{r} \in \text{range}(A)$ and the last $n+1 : m$ components of $\mathbf{s} = U^T W_{\mathbf{b}}^{1/2} \mathbf{r}$ are identically zero. This would imply that J has p degrees of freedom instead of $m-n+p$. Practically, \mathbf{b} is error contaminated and $\sum_{i=n+1}^m s_i^2$ in (2.7) is positive and constant with respect to σ . On the other hand, given precise values for $\bar{\mathbf{x}}$ and \mathbf{x}_0 , the last $m-n$ components of $\mathbf{c} = Q\mathbf{q} = QU^T W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0) = Q\tilde{Y}X^T(\bar{\mathbf{x}} - \mathbf{x}_0)$ are identically zero.*

Theorem 2.2 suggests that an approach for finding a suitable regularization parameter in the single variable case $W_{\mathbf{x}} = \lambda^2 I_n$, given sufficient statistical information on the measured data \mathbf{b} , is to find a $\lambda = 1/\sigma_{\mathbf{x}}$ so that as closely as possible J follows a χ^2 distribution, $J(\sigma_{\mathbf{x}}) \sim \chi^2(m-n+p, c(\sigma_{\mathbf{x}}))$. Equivalently, introducing the notation $\tilde{m} = m-n+p$, and $\delta(\sigma_{\mathbf{x}}) = z_{\alpha/2} \sqrt{2\tilde{m} + 4c(\sigma_{\mathbf{x}})}$, we want to determine $\sigma_{\mathbf{x}}$ such that

$$(2.12) \quad \tilde{m} + c(\sigma_{\mathbf{x}}) - \delta(\sigma_{\mathbf{x}}) \leq J \leq \tilde{m} + c(\sigma_{\mathbf{x}}) + \delta(\sigma_{\mathbf{x}}),$$

where $z_{\alpha/2}$ is the relevant z -value for a standard normal distribution, and α defines the $(1-\alpha)$ confidence interval that $J \sim \chi^2(m-n+p, c)$. Taking $c = 0$ this is equivalent to the condition used in [17] for the case when $\mathbf{x}_0 \approx \bar{\mathbf{x}}$. Here, because c depends on σ , the design of an algorithm to find $\sigma_{\mathbf{x}}$ satisfying (2.12) becomes more challenging. The algorithm design is discussed in Section 3, but first we address a modification of the result when the numerical rank is reduced.

2.2. The Truncated GSVD. We are also interested in the case when the numerical rank of the resolution matrix is reduced, in which case the number of degrees of freedom of the random variable J is also reduced. We illustrate this observation through the use of a truncated GSVD expansion for the solution (2.2).

Suppose that the regularized solution $\hat{\mathbf{x}}(\lambda)$ is written in terms of the GSVD expansion.

$$\begin{aligned}
 (2.13) \quad \mathbf{x}(\lambda) &= \sum_{i=1}^p \frac{v_i}{v_i^2 + \lambda^2 \mu_i^2} s_i \tilde{\mathbf{x}}_i + \sum_{i=p+1}^n s_i \tilde{\mathbf{x}}_i, \\
 &= \sum_{i=1}^p \frac{\gamma_i^2}{v_i(\gamma_i^2 + \lambda^2)} s_i \tilde{\mathbf{x}}_i + \sum_{i=p+1}^n s_i \tilde{\mathbf{x}}_i, \\
 &= \sum_{i=1}^p \frac{f_i}{v_i} s_i \tilde{\mathbf{x}}_i + \sum_{i=p+1}^n s_i \tilde{\mathbf{x}}_i, \quad f_i = \frac{\gamma_i^2}{\gamma_i^2 + \lambda^2},
 \end{aligned}$$

using the notation f_i given in [7] and where $\tilde{\mathbf{x}}_i, i = 1 \dots n$ are the columns of matrix $(X^T)^{-1}$. It is well known, however, that the stable numerical calculation of the GSVD relies on $\|D\| \approx \|\tilde{A}\|$, [9, 7]. If C_b is not well conditioned, this property carries through to \tilde{A} , ill-conditioning will be reflected in v_i and $\tilde{\mathbf{x}}_i$ for small i , and the full expansion in (2.13) will lead to a solution that is noise-contaminated. Thus, as with the use of the truncated singular value decomposition for ill-conditioned problems, a truncated GSVD expansion, [7], of the solution has previously been suggested. Setting $f_i = 0$ for small components $v_i, i \leq p - r$ and 1 otherwise, yields a truncated expansion solution

$$(2.14) \quad \mathbf{x}_{\text{TGSVD}} = \sum_{i=p+1-r}^p \frac{1}{v_i} s_i \tilde{\mathbf{x}}_i + \sum_{i=p+1}^n s_i \tilde{\mathbf{x}}_i.$$

More generally, we may consider the filtered solution, following the suggestion in [14] for the regularized TSVD, in which the filter function is defined by $f_i = 0$ for $v_i < \tau$ for some tolerance τ , but for the other terms we maintain the dependence on the regularization parameter λ

$$(2.15) \quad \mathbf{x}_{\text{FILT}}(\lambda) = \sum_{i=p+1-r}^p \frac{\gamma_i^2}{v_i(\gamma_i^2 + \lambda^2)} s_i \tilde{\mathbf{x}}_i + \sum_{i=p+1}^n s_i \tilde{\mathbf{x}}_i = \sum_{i=1}^p \frac{f_i}{v_i} s_i \tilde{\mathbf{x}}_i + \sum_{i=p+1}^n s_i \tilde{\mathbf{x}}_i.$$

This amounts to setting $v_i = 0, \mu_i = 1$, so that $\gamma_i = 0, i \leq p - r$. Consequently, we obtain a new expression for the quadratic form (2.7)

$$(2.16) \quad J(\mathbf{x}_{\text{FILT}}(\lambda)) = \sum_{i=1}^{p-r} s_i^2 + \sum_{i=n+1}^m s_i^2 + \sum_{i=p-r+1}^p \frac{\lambda^2}{\gamma_i^2} f_i s_i^2.$$

Theorem 2.2 is modified appropriately and the distribution applies not for the original $J(\lambda)$ but for a new functional $\tilde{J}(\lambda)$ with the constant terms in (2.16) removed

$$(2.17) \quad \tilde{J}(\lambda) = J(\lambda) - \sum_{i=1}^{p-r} s_i^2.$$

THEOREM 2.3. *Under the same conditions as Theorem 2.2, but with the solution given by (2.15), such that $f_i = 0$, $i = 1 : p - r$, $f_i = \gamma_i^2 / (\gamma_i^2 + \lambda^2)$, $i = p - r + 1 : p$, then the function \tilde{J} is a χ^2 random variable with $m - n + r$ degrees of freedom, and **centrality** parameter*

$$(2.18) \quad c = \|\tilde{\mathbf{c}}\|_2^2 = \|\tilde{Q}U^T W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0)\|_2^2,$$

where, generalizing (2.10),

$$(2.19) \quad \tilde{Q} = \text{diag}(0_{p-r}, \mu_{p-r+1}, \dots, \mu_p, 0_{n-p}, I_{m-n}).$$

Proof. The proof proceeds as before for Theorem 2.2 but replacing \mathbf{k} in (2.8) with $\tilde{\mathbf{k}} = \tilde{Q}U^T W_{\mathbf{b}}^{1/2} \mathbf{r}$ with the new definition (2.19) for \tilde{Q} . The distribution for $W_{\mathbf{b}}^{1/2} \mathbf{r}$ is unchanged, but

$$\tilde{\mathbf{k}} \sim N_m(\tilde{Q}U^T W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0), \text{diag}(0_{p-r}, I_{m-(p-r)})).$$

Whereas the mean and variance of the components of \tilde{k}_i are unchanged from those for k_i for $i > p - r$, the first $p - r$ components of $\tilde{\mathbf{k}}$ are constant, mean and variance 0. Therefore

$$(2.20) \quad \tilde{J} = J - \sum_{i=1}^{p-r} s_i^2$$

$$(2.21) \quad = \sum_{i=p-r+1}^p \tilde{k}_i^2 + \sum_{i=n+1}^m \tilde{k}_i^2,$$

is a sum of $m - n + r$ normally distributed random variables each with variance one and respective mean c_i . \square

Inequality (2.12) applies as before but with r replacing p in the definition of \tilde{m} and c calculated over the relevant components $p - r + 1 : p$. Notice that the distribution of the functional arises only from the components in the expansion for the solution \mathbf{x}_{FILT} which are filtered, the other components are constants independent of the regularization parameter λ . Equivalently the unregularized components do not contribute to the statistical properties of the functional. It appears that the degrees of freedom of the functional is determined by overall numerical rank of the subspace that defines the solution. This observation is a topic for future research.

3. Implementation.

3.1. The Newton Algorithm. In [17] a Newton algorithm to find σ using the original formulation of the theory, without the centrality parameter, was presented. There it was based on the use of the GSVD to find the root of $F(\sigma) = J(\sigma) - \tilde{m} = 0$. The basic Newton iteration, with line search parameter $\alpha^{(k)}$, may be written generally as

$$(3.1) \quad \sigma^{(k+1)} = \sigma^{(k)} \left(1 + \alpha^{(k)} \frac{1}{2} \left(\frac{\sigma^{(k)}}{\|D\mathbf{x}(\sigma^{(k)})\|} \right)^2 (J(\sigma^{(k)}) - \tilde{m}) \right).$$

The derivative is given by

$$(3.2) \quad J'(\sigma) = -\frac{2}{\sigma^3} \|D\mathbf{x}(\sigma)\|^2.$$

This can be determined by implicit differentiation of $J(\sigma)$ in (2.3), but also follows more easily from the expression for J in terms of the GSVD for the pair $[\tilde{A}|D]$,

$$(3.3) \quad J(\sigma) = \|\mathbf{k}(\sigma)\|^2 = \sum_{i=1}^p \frac{s_i^2}{(\sigma^2 \gamma_i^2 + 1)} + \sum_{i=n+1}^m s_i^2, \quad \gamma_i = \frac{\nu_i}{\mu_i}, \quad \sigma = \frac{1}{\lambda},$$

combined with using the expansion (2.13) for the solution $\mathbf{x}(\lambda)$. Observe, as mentioned in Remark 1, that the term $\sum_{i=n+1}^m s_i^2$ is constant and so practically in the GSVD implementation this term is calculated only once. In this case we adjust $\tilde{m} = \tilde{m} - \sum_{i=n+1}^m s_i^2$ and only calculate the update of J for the other relevant components.

While, the algorithm for the truncated functional given by (2.17) is designed similarly with function F defined by

$$(3.4) \quad F(\sigma) = \sum_{i=p-r+1}^p \frac{s_i^2}{(\sigma^2 \gamma_i^2 + 1)} + \sum_{i=n+1}^m s_i^2 - (m - n + r) \approx 0.$$

it is of greater interest to consider an approach to make the calculation efficient for large scale problems. In the following we introduce a method to obtain $\mathbf{x}(\sigma)$ and hence $J(\sigma)$ iteratively.

3.2. The Hybrid LSQR Algorithm. The algorithm is based on Golub-Kahan iterative bidiagonalization [21, 22]. Here, we only describe the essential components of the algorithm, and refer to references for more details. Given initial vectors $\mathbf{g}_0 \equiv 0$, $\mathbf{h}_1 \equiv \mathbf{b}/\beta_1$, where $\beta_1 \equiv \|\mathbf{b}\| \neq 0$, the Golub-Kahan iterative bidiagonalization computes, using two-term recurrences requiring only matrix-vector multiplications with the matrices A and A^T , orthonormal vectors $\mathbf{g}_i, \mathbf{h}_i, i = 1, 2, \dots$

$$(3.5) \quad \alpha_i \mathbf{g}_i = A^T \mathbf{h}_i - \beta_i \mathbf{g}_{i-1}, \quad \|\mathbf{g}_i\| = 1,$$

$$(3.6) \quad \beta_{i+1} \mathbf{h}_{i+1} = A \mathbf{g}_i - \alpha_i \mathbf{h}_i, \quad \|\mathbf{h}_{i+1}\| = 1.$$

Let $H_j \equiv [\mathbf{h}_1, \dots, \mathbf{h}_j] \in \mathcal{R}^{m \times j}$, $G_j \equiv [\mathbf{g}_1, \dots, \mathbf{g}_j] \in \mathcal{R}^{n \times j}$ and

$$L_j \equiv \begin{bmatrix} \alpha_1 & & & & & \\ \beta_2 & \alpha_2 & & & & \\ & \ddots & \ddots & & & \\ & & & \beta_j & \alpha_j & \\ & & & & & \end{bmatrix}, \quad L_{j+} \equiv \begin{bmatrix} L_j \\ \beta_{j+1} \mathbf{e}_j^T \end{bmatrix}, \quad L_{j+} \in \mathcal{R}^{(j+1) \times j}.$$

Then the recurrences (3.5-3.6) can be written in the matrix notation

$$(3.7) \quad A^T H_j = G_j L_j^T, \quad A G_j = [H_j, \mathbf{h}_{j+1}] L_{j+}.$$

With the observation $[H_j, \mathbf{h}_{j+1}]^T \mathbf{b} = \beta \mathbf{e}_1$, the j steps of the bidiagonalization yield a subproblem

$$(3.8) \quad L_{j+} \mathbf{y}_j \approx \mathbf{e}_1 \beta_1.$$

Then, when $D = I$, the least squares solution $\mathbf{y}_j(\sigma)$ of the augmented system

$$(3.9) \quad \begin{bmatrix} L_{j+} \\ 1/\sigma I_j \end{bmatrix} \mathbf{y}_j(\sigma) \approx \mathbf{e}_1 \beta_1,$$

transformed to the original variables through

$$(3.10) \quad \mathbf{x}_j(\sigma) \equiv G_j \mathbf{y}_j(\sigma).$$

represents an LSQR approximation to the solution of the original problem (1.1) for one particular σ , see [21, 22]. Note that $\|\mathbf{x}_j(\sigma)\|_2^2 = \|G_j \mathbf{y}_j(\sigma)\|^2 = \|\mathbf{y}_j(\sigma)\|_2^2$. For the general case when $D \neq I$, the problem (1.1) can be transformed to the so-called *standard* form, and then the LSQR algorithm in the basic form can be applied, details of this transformation are provided for example in [9].

In order to use the root finding algorithm described in Section 3.1 to determine optimal σ , it is important that the updated values for $\mathbf{x}(\sigma^{(k)})$ and $J(\sigma^{(k)})$ can be obtained efficiently for each iteration. The solution $\mathbf{x}(\sigma^{(k)})$ for each $\sigma^{(k)}$ is computed by solving (3.9) with the appropriate number of bidiagonalization steps $j = j(\sigma^{(k)})$, and takes advantage of the fact that the matrix L_{j+} and the right-hand side $\mathbf{e}_1 \beta_1$ do not depend on $\sigma^{(k)}$. If $j(\sigma^{(k)})$ is smaller than $j(\sigma^{(i)})$ for some $1 \leq i < k$, the corresponding subblock of L_{j+} can be used. Otherwise the matrix L_{j+} is augmented by computing additional steps in (3.5-3.6). Note that construction of the matrix L_{j+} is the most expensive part of the LSQR algorithm. Thus reusing L_{j+} makes the computation of $\mathbf{x}(\sigma^{(k)})$ efficient. The calculation of $J(\sigma^{(k)})$ is clear from the update $\mathbf{x}(\sigma^{(k)})$: it is approximated by the augmented residual for the approximate solution $\mathbf{x}_j(\sigma^{(k)})$.

3.3. The Algorithm with Centrality Parameter. To account for the centrality parameter c the algorithm needs modification. For ease of explanation we explicitly write functional J as a function of both \mathbf{b} and σ , $J = J(\mathbf{b}, \sigma)$. We have shown that $J(\mathbf{b}, \sigma) = \|\mathbf{k}(\mathbf{b}, \sigma)\|^2$, and that when σ is chosen appropriately, $J(\mathbf{b}, \sigma)$ will follow a χ^2 distribution with expected value that depends on its expected mean value. But by Theorem 2.2 this mean value is explicitly given by $c = \|\mathbf{c}\|^2 = \|\mathbf{k}(A\bar{\mathbf{x}}, \sigma)\|^2$. We thus need to solve

$$(3.11) \quad F_C(\sigma) := J(\mathbf{b}, \sigma) - J(A\bar{\mathbf{x}}, \sigma) - \tilde{m} \approx 0.$$

Apparently we need to know $\bar{\mathbf{x}}$. But if $\bar{\mathbf{x}}$ were known, we would actually solve using $\mathbf{x}_0 = \bar{\mathbf{x}}$ in (1.1) and then $J(A\bar{\mathbf{x}}, \sigma) = 0$. On the other hand, suppose that indeed $\bar{\mathbf{x}}$ is not known, but we can find an estimate for $\bar{\mathbf{b}}$ from the set of repeated measurements of the experiment which are used to provide the covariance matrix $C_{\mathbf{b}}$. Then we can replace $A\bar{\mathbf{x}}$ in (3.11) by $\bar{\mathbf{b}}$ and seek to solve $F_C(\sigma) \approx 0$. However, from (3.2)

$$(3.12) \quad \frac{\partial F_C}{\partial \sigma} = -\frac{2}{\sigma^3} (\|D\mathbf{x}(\mathbf{b}, \sigma)\|^2 - \|D(\mathbf{x}(\bar{\mathbf{b}}, \sigma))\|^2),$$

and the functional F_C need not be monotonic. Indeed, the root of $F_C(\sigma) = 0$ need not exist.

An example illustrating this is shown in Figure 3.1.

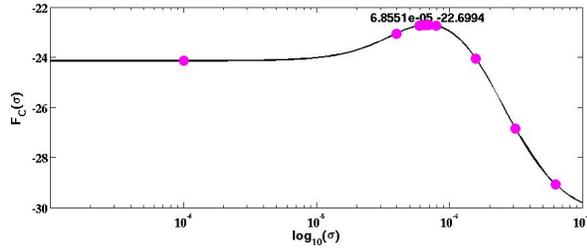


FIG. 3.1. Example of the non-monotonicity of function F_C in (3.11). $F_C(\sigma)$ is plotted against $\log_{10}(\sigma)$. The circles denote the values of the pairs $(\log_{10}(\sigma^{(k)}), F_C(\sigma^{(k)}))$ during the iteration to find the optimal choice of σ , found at $\sigma = 6.9e - 05$ with value $F_C(\sigma) = -22.7$.

The potential for the nondefinite behavior of F'_C , when $\mathbf{q} \neq 0$, illustrated in Figure 3.1 is most apparent if we look at the GSVD formulation. In particular, the equivalent equations for the GSVD implementation are, for (3.11, 3.12), resp.

$$(3.13) \quad \mathcal{B}(\sigma) = \sum_{i=1}^p \frac{s_i^2 - q_i^2}{(\sigma^2 \gamma_i^2 + 1)} + \sum_{i=n+1}^m (s_i^2 - q_i^2) - \tilde{m} = 0, \quad \mathbf{q} = U^T W_{\mathbf{b}}^{1/2} A(\bar{\mathbf{x}} - \mathbf{x}_0), \text{ and}$$

$$(3.14) \quad \frac{\partial F_C}{\partial \sigma} = -2\sigma \sum_{i=1}^p \frac{z_i \gamma_i^2}{(\sigma^2 \gamma_i^2 + 1)^2}, \quad \text{where } z_i = s_i^2 - q_i^2.$$

3.3.1. Newton with Bisection to Minimize F_C . We seek σ such that $F_C(\sigma)$ is close to zero, equivalently, such that $F_C^2(\sigma)$ is minimum. If there is a root such that $F_C(\sigma) = 0$, it will solve the minimization, otherwise we find σ such that $F_C'(\sigma) = 0$, and $F_C^2(\sigma)$ is minimum. As for the original Newton algorithm to solve $F(\sigma) = 0$, the algorithm is made more robust by some basic observations. Both $J(\mathbf{b}, \sigma)$ and $J(\bar{\mathbf{b}}, \sigma)$ which occur in $F_C(\sigma)$ are positive. Indeed,

$$F_C(\sigma) = \|\mathbf{k}(\mathbf{b}, \sigma)\|^2 - \|\mathbf{k}(\bar{\mathbf{b}}, \sigma)\|^2 - \tilde{m},$$

and both norms are monotonically decreasing with σ . Therefore,

$$\|\mathbf{k}(\mathbf{b}, \sigma)\|^2 - \|\mathbf{k}(\bar{\mathbf{b}}, 0)\|^2 - \tilde{m} \leq F_C(\sigma) \leq \|\mathbf{k}(\mathbf{b}, \sigma)\|^2 - \tilde{m}.$$

Because $\|\mathbf{k}(\mathbf{b}, \sigma)\|^2$ is itself monotonically decreasing with σ , for any $\sigma > \sigma_{\max}$, where $\|\mathbf{k}(\mathbf{b}, \sigma_{\max})\|^2 = \tilde{m}$, necessarily $F_C(\sigma) < 0$, and $F_C'(\sigma_{\max}) < 0$. Likewise, $\|\mathbf{k}(\mathbf{b}, \sigma_{\min})\|^2 = \tilde{m} + \|\mathbf{k}(\bar{\mathbf{b}}, 0)\|^2$, provides a lower bound for σ . The algorithm to minimize $F_C^2(\sigma)$ therefore first solves for both σ_{\min} and σ_{\max} , so as to find a bracket on the optimal σ . Because these two values are found using only the functional $J(\mathbf{b}, \sigma)$, the original fast convergent Newton algorithm can be used, [17]. Indeed, determination of σ_{\max} solves the original central distribution problem under the assumption that the given $\mathbf{x}_0 = \bar{\mathbf{x}}$, and can itself be used to give an estimate of the solution, $\mathbf{x}(\mathbf{b}, \sigma_{\max})$. In the rare case that the vector \mathbf{z} , defined from its components z_i , is itself definite, these bounds on σ effectively bracket the root of a monotonically increasing function, and the original Newton algorithm can now be applied to function F_C to find its root. Otherwise we use simple bisection to find the minimum of $F_C'(\sigma)$ within the determined bracket.

REMARK 2. It might appear that another approach to solve in the case that $\mathbf{x}_0 \neq \bar{\mathbf{x}}$ but $\bar{\mathbf{b}}$ is available, would be to use $\bar{\mathbf{b}}$ to find an estimate of $\bar{\mathbf{x}}$, and then to solve again using the obtained value for the estimate of the expected value $\bar{\mathbf{x}}$ as \mathbf{x}_0 in $J(\mathbf{b}, \sigma)$. However, in this case we would need to find the estimate for $\bar{\mathbf{x}}$ without regularization. We know that for ill-conditioned problems, even without noise, the estimate of the solution $\bar{\mathbf{x}}$ is unlikely to be useful. While such an approach would therefore avoid the problem of finding the minimum for F_C , its success would be limited to problems which are well-conditioned. Experiments, not reported here, confirm this observation.

4. Numerical Experiments. The major goal of the presented numerical experiments is that they validate the hypothesis that the hybrid LSQR algorithm can be used to efficiently

obtain the regularization parameter using the χ^2 properties of the regularized functional for large scale problems. Several experiments are presented. First, in Section 4.1, we contrast the basic Newton algorithm implemented using direct GSVD solves and iterative hybrid LSQR solutions for a set of small scale test problems. The algorithm is implemented as described in Section 3.1 and is based on the original work in [17]. In all cases it is implemented in exactly the same way for both direct GSVD and iterative hybrid LSQR versions, namely line search and bracketing are performed equivalently. A comparison, in Section 4.2, is also provided with two other standard methods for estimating the regularization parameter, the L-curve and unbiased predictive risk estimator (UPRE), see for example [9, 27]. Because the hybrid LSQR is iterative, its performance depends on the number of steps of the bidiagonalization algorithm (3.5-3.6). As proposed in the original papers [21, 22], two stopping criteria for the generation of the bidiagonal system are used

$$(4.1) \quad \|\mathbf{r}_j\| \leq btol \|\mathbf{b}\| + atol \|A\| \|\mathbf{x}_j\| \quad \text{and} \quad \frac{\|A^T \mathbf{r}_j\|}{\|A\| \|\mathbf{r}_j\|} \leq atol.$$

The quantities $\|\mathbf{r}_j\|$, $\|\mathbf{x}_j\|$, $\|A\|$, and $\|A^T \mathbf{r}_j\|$ can be estimated at minimal cost in the LSQR iterations. The quantities $atol$, $btol$ are user specified and should reflect the expected accuracy of the data, see [21] for more details. We therefore also examine the relationship between the estimate of σ and the size of the bidiagonal system used at each inner iteration, together with the choices for $atol$ and $btol$, see Section 4.3. Finally, we demonstrate that the new approach can be used for the solution of large scale problems. Two examples are presented, one the deconvolution of real one dimensional seismic signals, Section 4.4 and the second a standard large scale image deblurring test problem, Section 4.5.

4.1. Comparing hybrid LSQR and GSVD. Benchmark problems, `shaw`, `ilaplace`, `heat` and `phillips`, from the Regularization Tool Box [8] are implemented for different noise levels. The parameter σ obtained using the iterative hybrid LSQR solution is contrasted with that obtained by the direct GSVD solution. In examples using the toolbox the *true* solutions \mathbf{x}_{true} , and matrices A defining the models, were obtained using the relevant function calls to the tool box and *true* measurements found from $\mathbf{b}_{\text{true}} = A\mathbf{x}_{\text{true}}$. To obtain noisy data sets the Matlab® function `randn` [15], which yields standard normal data, was used to obtain a perturbation matrix E of size $m \times 500$, for \mathbf{b} of length m . Each column \mathbf{e}^l , $l = 1 : 500$, of E was normalized to two norm length $\|\mathbf{b}\|$ by $\mathbf{e}^l = \mathbf{e}^l \|\mathbf{b}\| / \|\mathbf{e}^l\|$ and the 500 noisy right hand side vectors \mathbf{b}^l obtained as $\mathbf{b}^l = \mathbf{b} + \epsilon \mathbf{e}^l$. Data sets were generated for noise levels $\epsilon = .001, .005, .01, .05$ and $.1$. Given the 500 samples we can then directly calculate

the covariance $C_{\mathbf{b}}$ of the measurements. This method generates a covariance matrix which is nearly diagonal, hence in this first set of experiments we used the diagonal weighting matrix $C_{\mathbf{b}}^{1/2} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_m)$. Also, all experiments, other than those with the problem `shaw` which uses $D = I$, used D to approximate the first derivative operator.

In this first set of experiments we use very small problems, problem size $n = 64$, and impose a very tight confidence interval (2.12) using $\alpha = .9999$, for tolerance $.0014$ on convergence of the Newton iteration. For the LSQR iterations we use $atol = 100 * btol$ and $btol = 10^{-9}$, and allow the subproblem size to reach full size 64 . At each outer Newton step the LSQR algorithm adjusts the size of the projected problem, $j(\sigma)$ such that the LSQR algorithm meets the convergence criteria (4.1), which may be different for each chosen $\sigma^{(k)}$. Therefore, for each noise level we track the average of $j(\sigma^{(k)})$ over outer steps $k = 1, 2, \dots$, where the maximum number allowed is 15 . For both direct and iterative solves we report the total number of calculations $K \leq 15$ of σ , equivalently of problem solves, where the count includes the bracketing and subsequent Newton iterations. P-values for a paired t-test between GSVD and hybrid LSQR results for both K and for σ are also given.

The initial experiments are designed to assess the consistency of the σ update strategy when implemented for the hybrid LSQR algorithm. Therefore the results given in Table 4.1 are for the solution of $F(\sigma) = 0$, using the *unrealistic* estimate $\bar{\mathbf{x}} = \mathbf{x}_0 = 0$.

The results in Table 4.1, demonstrate generally a very high correlation, p-value near 1, between GSVD and hybrid LSQR algorithms for both the average number of solution solves K and for σ . Overall the average size of the projected problem is found to be much smaller than the actual problem size, and decreases with increasing noise level as anticipated. The exceptions, those where the correlation is not high, namely low noise for the problem `heat` and high noise for `ilaplace`, can be explained by the regularizing properties of the hybrid LSQR algorithm. Specifically, for the former case with high noise, the LSQR generates a small bidiagonal system, hence introducing significant regularization itself by excluding noisy components of the solution. Consequently, it is likely that the obtained optimal regularization parameter imposes less regularization than is needed for the GSVD algorithm. On the other hand, in the latter case, the LSQR algorithm iterates almost to the actual full problem size of $n = 64$ and more outer regularization is needed. These characteristics will be examined in more detail in Section 4.3.

Table 4.2 provides equivalent experiments to those in Table 4.1, but for the minimization of the discrepancy $F_C^2(\sigma)$. Therefore, the algorithm first finds the minimum and maximum

shaw					
Noise Level ϵ	Average $j(\sigma^{(k)})$	Iterations K		P value	
		GSVD	LSQR	Iteration	σ
$1e-03$	13.1	7.9	7.9	1.0	1.0
$5e-03$	11.5	7.7	7.6	0.6	0.3
$1e-02$	10.3	7.0	7.0	1.0	1.0
$5e-02$	7.2	7.1	7.1	1.0	1.0
$1e-01$	6.2	7.0	7.0	0.9	0.8
ilaplace					
$1e-03$	16.0	7.3	7.3	1.0	1.0
$5e-03$	12.2	6.3	6.3	1.0	1.0
$1e-02$	10.7	6.4	6.4	0.9	1.0
$5e-02$	8.0	6.5	6.5	0.9	1.0
$1e-01$	7.4	6.5	6.5	0.9	0.5
heat					
$1e-03$	63.8	5.9	5.8	0.9	0.0
$5e-03$	63.9	5.5	5.5	1.0	1.0
$1e-02$	47.6	5.7	5.7	1.0	1.0
$5e-02$	29.4	5.8	5.8	1.0	1.0
$1e-01$	22.4	6.0	6.0	1.0	1.0
phillips					
$1e-03$	37.4	10.7	10.7	1.0	1.0
$5e-03$	28.6	7.4	7.4	1.0	1.0
$1e-02$	25.0	6.8	6.8	1.0	1.0
$5e-02$	15.4	6.1	6.1	1.0	1.0
$1e-01$	12.8	6.1	6.1	1.0	1.0

TABLE 4.1

The last two columns are the P -values for paired t -tests between GSVD and LSQR results, for both the obtained σ and the number of σ updates. The first column is the noise level ϵ used in generating the noisy data. These results are for problem size 64, over 500 tests of each problem, and use $\mathbf{x}_0 = 0$ in the definition of $F(\sigma)$.

values for σ as described in Section 3.3.1, and then carries out bisection to find the minimal discrepancy. The iteration count includes all three stages. To simulate the use of average measurement data, we form the average of the \mathbf{b}^l for the given noise level and use this as $\bar{\mathbf{b}}$. Here the average relative errors in the solutions are also calculated, and the failure count is given, where a failure is indicated by a relative error greater than 50%. The relative errors and failures increase with the noise level, but in general the results indicate that the algorithm is robust at finding good regularization parameters. Also these results demonstrate that a high correlation for the obtained σ is not necessary for achieving low relative error, see for example noise level .01 for problem `ilaplace`. These results reinforce the conclusions about the robustness of the χ^2 method for small scale problems when used with prior information $\bar{\mathbf{b}}$ instead of $\bar{\mathbf{x}}$ as presented in [17].

shaw									
Noise Level ϵ	Average $j(\sigma^{(k)})$	Iterations K		P value		Relative Error		Failures	
		GSVD	LSQR	Iteration	σ	GSVD	LSQR	GSVD	LSQR
$1e-03$	10.5	31.1	31.2	0.8	1.0	0.067	0.067	0	0
$5e-03$	9.7	28.4	28.1	0.9	1.0	0.110	0.110	0	0
$1e-02$	8.7	29.4	29.2	0.7	1.0	0.135	0.135	0	0
$5e-02$	7.4	26.7	26.5	0.9	1.0	0.189	0.189	2	2
$1e-01$	6.9	26.5	25.8	0.8	1.0	0.215	0.215	0	0
ilaplace									
$1e-03$	8.5	28.2	31.2	1.0	1.0	0.023	0.023	0	0
$5e-03$	6.5	26.8	30.6	1.0	1.0	0.053	0.053	1	1
$1e-02$	5.9	26.5	30.6	0.0	0.4	0.084	0.075	1	1
$5e-02$	4.7	25.6	30.8	0.9	1.0	0.161	0.162	16	16
$1e-01$	4.3	24.1	30.9	0.7	1.0	0.211	0.211	48	47
heat									
$1e-03$	34.1	27.9	29.6	0.2	0.8	0.073	0.073	0	1
$5e-03$	28.8	27.3	29.6	0.9	1.0	0.133	0.133	0	0
$1e-02$	23.6	26.6	30.1	1.0	1.0	0.179	0.179	1	1
$5e-02$	14.8	26.1	30.3	1.0	1.0	0.304	0.304	48	49
$1e-01$	11.1	24.5	30.3	0.8	1.0	0.375	0.375	158	158
phillips									
$1e-03$	22.0	33.4	34.5	1.0	1.0	0.017	0.017	0	0
$5e-03$	14.7	28.3	31.6	0.6	1.0	0.029	0.029	0	0
$1e-02$	12.1	27.5	30.9	1.0	1.0	0.039	0.039	0	0
$5e-02$	8.0	24.6	30.2	0.9	1.0	0.079	0.079	0	0
$1e-01$	7.0	25.4	30.5	0.9	1.0	0.112	0.112	0	0

TABLE 4.2

Comparison of the results using the hybrid LSQR and the GSVD based algorithm with centrality parameter obtained using the average of the right hand side vectors. The last two columns are the P-values for paired t-tests between GSVD and LSQR results, for both the obtained σ and the number of σ updates. The first column is the noise level ϵ used in generating the noisy data. These results are for problem size 64, over 500 tests of each problem. The mean relative errors are measured in the least squares norm as compared to the known exact solutions and the failure count is the number of problems which did not achieve relative error less than 50%.

4.2. Other Regularization Techniques. Extensive results contrasting the χ^2 method with *a priori* information $\mathbf{x}_0 = \bar{\mathbf{x}}$ for the L-curve and UPRE techniques for finding the regularization parameter were presented in [17]. In Table 4.3 we thus present just one set of experiments to demonstrate that the new algorithm with the *a priori* information $\bar{\mathbf{b}}$ is competitive.

The problem size is 128 but all other settings are the same as for the experiments reviewed in Table 4.2, except that the size of the projected problem is limited to 128, and the total sample size is 250 instead of 500. The failure counts and average relative least squared errors are given. The entries in bold face in each row indicate the result with minimum average

relative error. Each method is competitive for some problem set, but less competitive for another problem set. It can also now be seen, compare with Table 4.2, that the hybrid LSQR algorithm has the potential to outperform the direct solve using the GSVD, even for this relatively small problem size of 128. We anticipate the GSVD to become less reliable as the problem size increases. Moreover, the results for the hybrid LSQR are not tuned in any regard to its convergence parameters. This tuning is the topic of the next set of results.

shaw								
ϵ	Least Squares Error				Failure Count			
	GSVD	LSQR	Lcurve	UPRE	GSVD	LSQR	Lcurve	UPRE
$1e-03$	0.098	0.098	0.076	0.099	1	1	0	33
$5e-03$	0.120	0.120	0.097	0.121	0	0	0	44
$1e-02$	0.199	0.180	0.183	0.216	0	0	0	44
$5e-02$	0.229	0.206	0.200	0.239	1	1	0	63
$1e-01$	0.276	0.236	0.241	0.259	5	4	0	71
ilaplace								
$1e-03$	0.064	0.052	0.065	0.069	0	1	0	38
$5e-03$	0.091	0.075	0.077	0.088	0	0	0	34
$1e-02$	0.146	0.151	0.125	0.160	3	6	0	42
$5e-02$	0.200	0.204	0.166	0.195	22	18	0	71
$1e-01$	0.247	0.230	0.231	0.230	67	53	63	112
heat								
$1e-03$	0.118	0.118	0.080	0.093	0	0	128	0
$5e-03$	0.152	0.152	0.114	0.127	0	0	157	0
$1e-02$	0.272	0.272	0.245	0.242	16	16	245	15
$5e-02$	0.327	0.324	<i>NaN</i>	0.305	46	40	250	21
$1e-01$	0.408	0.404	<i>NaN</i>	0.396	119	109	250	61
phillips								
$1e-03$	0.025	0.025	0.025	0.040	0	0	0	4
$5e-03$	0.033	0.033	0.028	0.044	0	0	0	4
$1e-02$	0.065	0.065	0.047	0.095	0	0	0	18
$5e-02$	0.100	0.099	0.080	0.109	0	0	0	34
$1e-01$	0.137	0.136	0.123	0.139	2	2	0	28

TABLE 4.3

Results are reported for problem size 128 over datasets each of size 250 for each noise level. The errors are relative least squares errors with respect to the exact solution, and the failure counts are the number of solutions in each case which did not achieve a relative error less than 50%. The parameter settings are set as in Table 4.2, except that the problem size of the projected problem is allowed to reach size 128, consistent with the given problem size. *NaN* indicates that the relative error could not be calculated because the test failed for all samples. Minimum values are in bold face for each problem set.

4.3. Characteristics of the hybrid LSQR algorithm. To verify the robustness of the hybrid LSQR algorithm with respect to the maximum size $j(\sigma)$ of the projected subproblem, as well as to illustrate the dependence of σ on $j(\sigma)$, we present the following experiments.

Problems `shaw` and `phillips` are solved for problem size $n = 512$ with an error level of 10%, the constraint $j(\sigma) \leq 40$ and different choices for $btol$. All other parameters are the same as in the previous section. The legend of each figure, in Figure 4.1 indicates the value of $btol$ and the resulting least squares error in the solution. Plotted are the subproblem size $j(\sigma^{(k)})$ against $\sigma^{(k)}$, for increasing values of $btol$. In all cases the projected problem size is far smaller than the maximum limit of 40, which supports the hypothesis that the hybrid LSQR will be cost effective for larger problems. The actual values of σ obtained in each case for increasing $btol$ are 3.1756, 3.1756, 3.2540, 28.5962, 750000 and are .0016, .0016, .0207, 250000, 750000, for `shaw` and `phillips`, resp. The large values here indicate that no regularization is needed, λ is very small, and the actual value is determined by the upper maximum imposed by the algorithm. Equivalently, when $btol$ is large, the hybrid iteration stops early, with a small projected problem size j , and no noise is introduced into the approximate solution. Therefore very little or no regularization is required but the obtained solution is less accurate because insufficient information was actually included in the projected problem. For smaller $btol$, the projected problem increases in size and the converged value of σ is also smaller, ie more regularization is required.

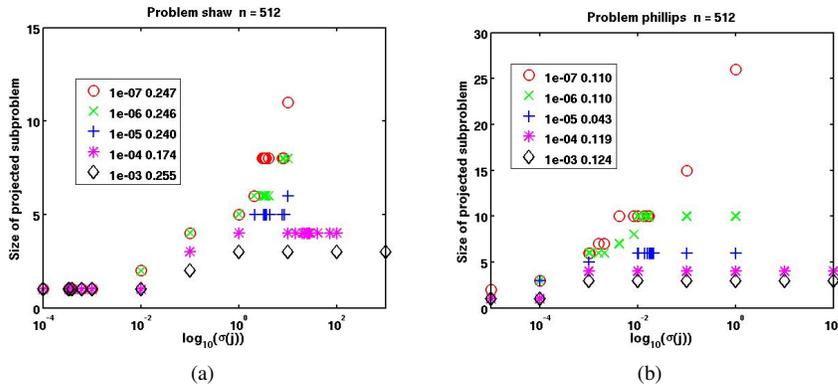


FIG. 4.1. Illustrating the dependence of σ on $j(\sigma)$ for problem with noise level 10% for increasing $btol$. The values in the legend are the noise level and resulting relative least squares error for the solution. Subfigure (a) for problem `shaw` and (b) for `phillips`. Problem size 512.

The optimal solutions in each case, for $btol = 10^{-4}$ and 10^{-5} for `shaw` and `phillips`, resp. are compared with solutions using the L-curve and UPRE in Figure 4.2, for which the errors are .197 and 1.198, and .076 and .501 for each problem resp. Good solutions are obtained by the hybrid LSQR algorithm.

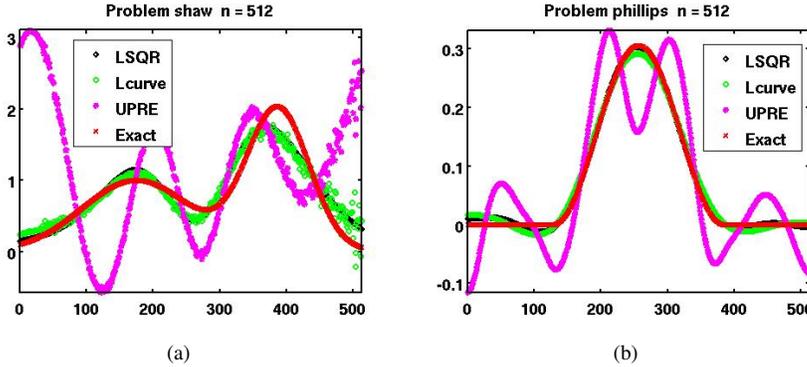


FIG. 4.2. Illustrating the solution for the different methods, for the optimal choice of $btol$ from the data given in Figure 4.1 compared with L-curve and UPRE solutions. Subfigure (a) for problem *shaw* and (b) for *phillips*. Problem size 512. Solution \mathbf{x} is plotted against its index.

4.4. Seismic Signal Restoration. Here we present the result of deblurring real seismic data. A real data set of 48 seismic signals of length 3000 is considered. The signal variance pointwise can therefore be calculated over all 48 signals. For this data the underlying point spread function, ie the signal blur described by matrix A , has been estimated numerically by an approach described in [25], and is not the focus of the analysis here. Instead our purpose is to show that the algorithms can be successfully used on real data. Because there is no known true solution we estimate the reliability of the deblurring result by downsampling the signals and restoring the signals at different resolutions, $2 : 1$, $5 : 1$, $10 : 1$, $20 : 1$, and $100 : 1$, corresponding to using $m = 1500, 600, 300, 150$, and 30 , resp. The signals are restored using UPRE and both versions of the χ^2 algorithm, i.e. with $\mathbf{x}_0 = 0$ and root finding for $F(\sigma)$, as well as using the average $\bar{\mathbf{b}}$ in the minimization of $F_C^2(\sigma)$. The average signal $\bar{\mathbf{b}}$ is formed from the average over all 48 signals. White noise weighting is used, calculated from the average covariance pointwise of all 48 signals.

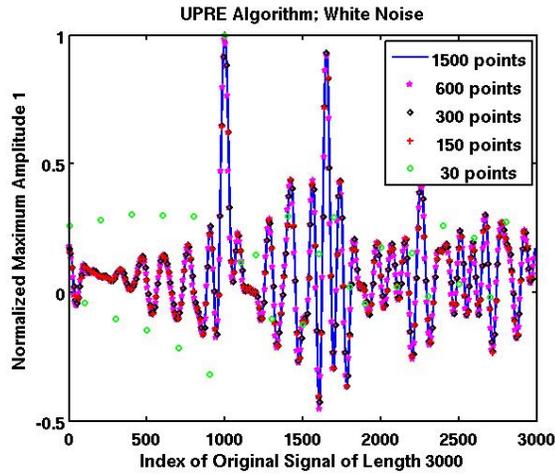
The results indicate, that except for very low resolution, the results of all implementations are consistent across resolutions. Indeed, the calculated values for σ are also consistent across resolutions for all three algorithms. It is apparent that the χ^2 methods are more successful at removing noise from the underlying signal. For this application it is important that spurious oscillations are removed but that the signals are not oversmoothed. The intent is to identify major signal arrival times, and to distinguish the arrivals of different features in the signal, which correspond to different geophysical reflections at the level of the core-mantel boundary in the Earth, and hence assist with interpretation of the structure of the core-mantel boundary. With this in mind, we deduce that the solutions obtained through minimization of

the discrepancy $F_C^2(\sigma)$ oversmooth the solution. Possibly this is due to the use of the overly smoothed prior information $\bar{\mathbf{b}}$. Moreover, this suggests that it may indeed be preferable to in general use the simpler algorithm which seeks to solve $F(\sigma) = 0$ with prior information $\mathbf{x}_0 = 0$. The image deblurring example in the next section therefore uses only the root finding algorithm.

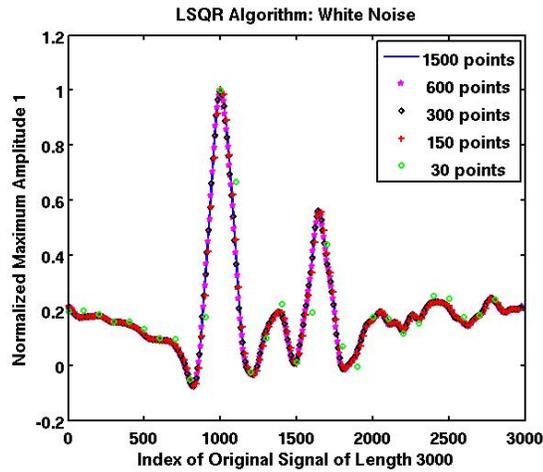
4.5. Image Deblurring. We now turn to our ultimate goal of demonstrating the robustness of the hybrid LSQR method for solving large scale ill-posed problems. To implement the code for large scale problems we take note that calculations in the bidiagonalization require only matrix-vector multiplications. To accomplish this we use the object-oriented approach for evaluating operations with matrices describing point spread functions provided in the `RESTORE TOOLS` [18]. We use problem `Grain` which is provided as one of the examples, with blurring matrix A of effective size $256^2 \times 256^2$, for a stacked 2D image of size 256×256 . A noisy version of `Grain` is obtained by adding noise to \mathbf{b}_{true} in the same way as for the earlier test problems in Section 4.1. Because no prior information is actually available, and would not likely be available for large scale deblurring, we run the basic Newton algorithm for function $F(\sigma)$ with $\mathbf{x}_0 = 0$. In Figure 4.4 we illustrate the original image, its blurred noisy version and the best solution obtained with the basic LSQR algorithm and with the hybrid LSQR algorithm, i.e. with additional regularization. This is for a case with 15% noise, normalized with $\|(\mathbf{b}_{\text{true}} - \mathbf{b})\|/\|\mathbf{b}\| = .15$. Because of the problem size we use $\text{btol} = 10^{-6}$ and $\text{atol} = 10^{-4}$. The best solution was in both cases obtained using the projected problem of the size 15.

The results are comparable but the signal to noise ratios calculated for solving with increasing projected problem size show that the regularization improves the basic LSQR solution, see Figure 4.5. Indeed, these results confirm the semi-convergence behavior of the LSQR iteration, and that the regularization stabilizes the process when the LSQR iteration converges to a solution for which regularization is still required. This is illustrated clearly with the relative error plot in Figure 4.5 and supports the similar observation in [2]. Figure 4.5 (c) illustrates the decrease of σ with the size of the projected problem, corresponding to increase in regularization as greater noise is built into the projected problem. The results presented for this one problem are illustrative of experiments with different noise levels, and additional problems also taken from the `RESTORE TOOLS` set.

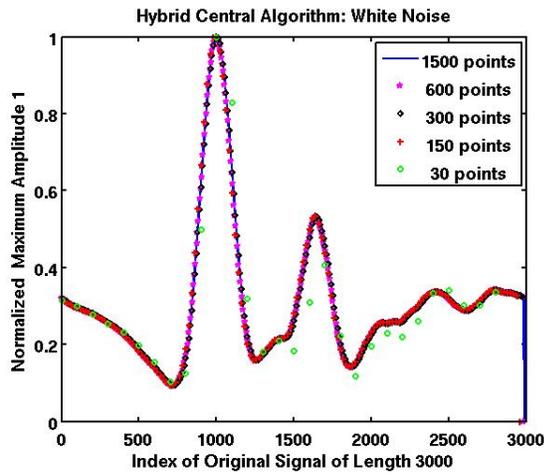
In these experiments with large scale problems we found that the estimates of $F(\sigma)$ obtained from the projected problem were not close enough for use in the χ^2 parameter choice



(a)



(b)



(c)

FIG. 4.3. (a) the UPRE solution, (b) the solution using root of $F(\sigma)$ (c) the solution using the minimum of $F_C^2(\sigma)$. Regularization parameters are consistent across resolutions: $\sigma = 0.01005$, $\sigma = 0.00058$ and $\sigma = 0.000001$ for each method, resp. Each solution is normalized to maximum height 1 and shifted to align at the position of the maximum, as is standard in the seismic literature. The solutions are plotted against the index of the original signal of length 3000.

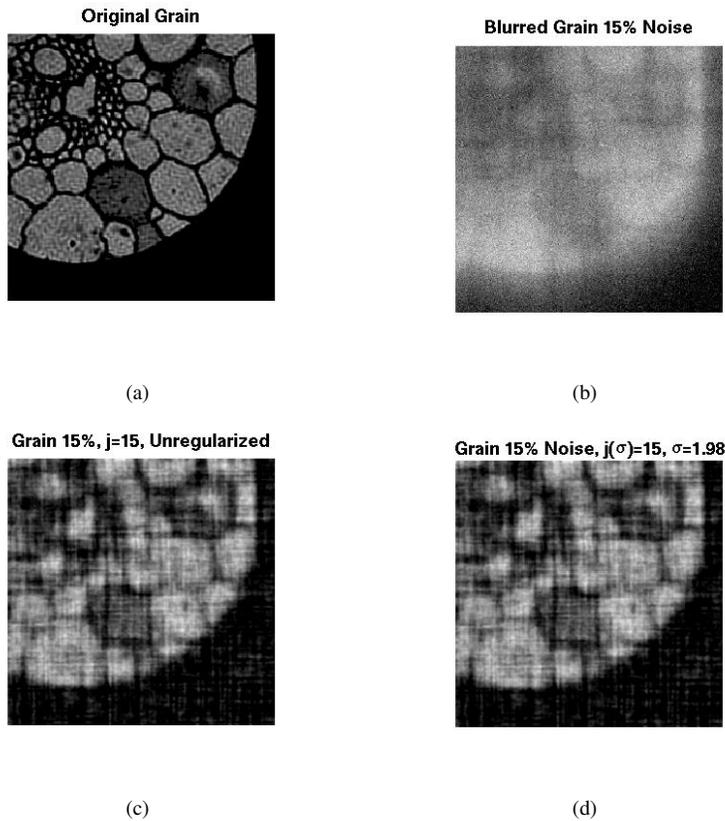


FIG. 4.4. Problem *Grain* with noise of 15% added. (a) the true solution, (b) the noisy blurred right hand side, (c) the basic LSQR solution (d) the hybrid LSQR solution.

algorithm, which is predicated on accurate estimates of this underlying functional. Therefore we actually updated $F(\sigma)$ directly through use of the updated solution $\mathbf{x}(\sigma)$ rather than using the estimate based on the residual of the projected problem. The cost of this additional step is minimal in relation to the overall bidiagonalization step.

5. Conclusions and Future Work. The χ^2 approach for estimation of the regularization parameter arising in the solution of numerically ill-posed problems by generalized Tikhonov regularization has been extended for large-scale problems through its combination with an iterative LSQR algorithm. Numerical results validate the method as compared to the direct GSVD algorithm for a series of test problems available in the literature. The utilization of the theory does rely on the use of some prior information, such as an estimate of the expected value of the model parameters, $\bar{\mathbf{x}}$ and an estimate of the covariance $C_{\mathbf{b}}$ for the measurement variables \mathbf{b} . On the other hand, the image deblurring example presented in

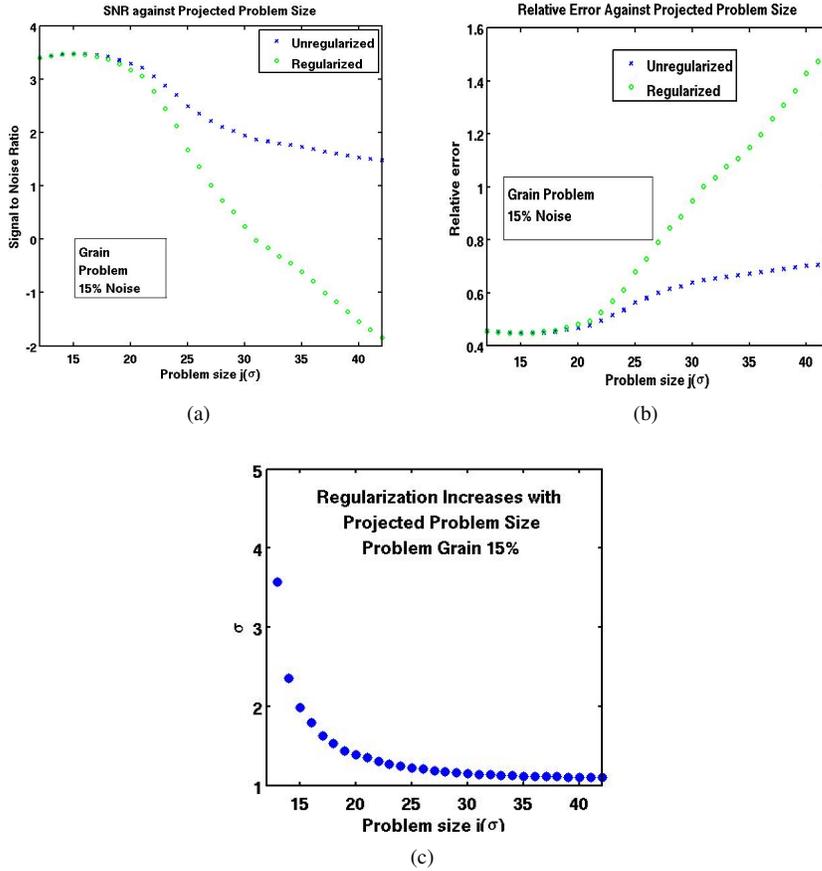


FIG. 4.5. (a) the signal to noise ratio for problem *Grain* with noise of 15% added for increasing subproblem size. Signal to noise ratio is calculated as $10 \log_{10}(1/e)$, where e is the relative error $\|\mathbf{x}_{\text{true}} - \mathbf{x}(\sigma)\| / \|\mathbf{x}_{\text{true}}\|$ illustrated in (b). (c) shows σ in each case. (Recall the regularization parameter $\lambda = 1/\sigma$.)

Section 4.5 used only an estimate of the signal covariance and no estimate of either $\bar{\mathbf{x}}$ or $\bar{\mathbf{b}}$. This indicates that the technique can be useful even without this additional prior information.

For the situation in which $\mathbf{x}_0 \neq \bar{\mathbf{x}}$, for example $\mathbf{x}_0 = 0$, the theory was extended. In particular, modifying the basic theory presented in [17] yields the general noncentral χ^2 distribution of the underlying functional. A new algorithm combining Newton with bisection search for obtaining the regularization parameter in this case was also developed and validated. While the numerical results with simulated data support the use of this more complicated algorithm, the results for the seismic signal restoration and the image deblurring suggest that the algorithm could actually be detrimental and lead to oversmoothing.

The theory has been modified when the underlying resolution matrix is ill-conditioned so that the resulting functional is still a χ^2 random variable at optimum, but with reduced degrees

of freedom. Utilization of this result for severely ill-conditioned problems, and its possible extension to explain results in a basis other than the basis given by the GSVD expansion, is a topic for future work.

There is considerable work in the statistics literature on the estimate of variance in measurement data without repeat measurements [26]. A topic of future study is thus to utilize this theory so as to make the χ^2 approach useful for data with limited experimental data. Additional study of the stabilizing effect of the regularization combined with the LSQR solution is also needed. The hybrid LSQR approach presented here does stabilize the LSQR solution, but as can be seen from Figure 4.5, the stabilization for large scale problems is limited. Further modification of the method may be related to the number of degrees of freedom in the subproblem and the choice of the stopping criteria for the bidiagonalization process. Future work will also consider the impact of preconditioning for improving the algorithm.

Acknowledgement. The first author completed portions of this work while visiting the Institute of Computer Science at the Academy of Sciences of the Czech Republic and the Department of Informatics and Mathematical Modeling at the Technical University of Denmark.

REFERENCES

- [1] Bennett, A., 2005, *Inverse Modeling of the Ocean and Atmosphere*, (Cambridge University Press), 234.
- [2] Chung, J., Nagy, J., and O'Leary, D. P., 2008, *A weighted GCV method for Lanczos hybrid regularization*, ETNA, **28**, 149-167.
- [3] Eldén, L., 1982, *A weighted pseudoinverse, generalized singular values, and constrained least squares problems*, BIT, **22**, 487-502.
- [4] R. D. Fierro, G. H. Golub, P. C. Hansen, and D. P. O'Leary, 1997, *Regularization by truncated total least squares*, SIAM J. Sci. Stat. Comp., **18**, 1225-1241.
- [5] Golub, G. H. and van Loan, C., 1996, *Matrix Computations*, (John Hopkins Press, Baltimore), 3rd ed.
- [6] Hanke, M., and Hansen, P. C., 1993, *Regularization methods for large-scale problems*, Surveys Math. Indust. **3**, 253-315.
- [7] Hansen, P. C., 1989, *Regularization, GSVD and truncated GSVD*, BIT, **6**, 491-504.
- [8] Hansen, P. C., 1994, *Regularization tools: A Matlab package for analysis and solution of discrete ill-posed problems*, Numerical Algorithms, **6**, 1-35.
- [9] Hansen, P. C., 1998, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, (SIAM Monographs on Mathematical Modeling and Computation 4).
- [10] Hansen, P. C., Kilmer, M. E., and Kjeldsen, R. H., 2006, *Exploiting residual information in the parameter choice for discrete ill-posed problems*, BIT, **46**, 41-59.
- [11] Hansen, P. C., Nagy, J. G., and O'Leary, D. P., 2006, *Deblurring Images, Matrices, Spectra and Filtering*, (SIAM Fundamentals of Algorithms).
- [12] Kilmer, M. E., and O'Leary, D. P., 2001, *Choosing regularization parameters in iterative methods for ill-posed problems*, SIAM J. Numer. Anal. Appl., **22**, 1204-1221.
- [13] Kilmer, M. E., Hansen, P. C., and Español, M. I., 2007, *A Projection-based approach to general-form Tikhonov regularization*, SIAM J. Sci. Comput., **29** (1), 315-330.
- [14] Marquardt, D. W., 1970, *Generalized inverses, ridge regression, biased linear estimation, and nonlinear estimation*, Technometrics, **12** (3), 591-612.
- [15] MATLAB is a registered mark of MathWorks, Inc., MathWorks Web Site, <http://.mathworks.com>.
- [16] Mead, J., 2008, *Parameter estimation: A new approach to weighting a priori information*, J. Inv. Ill-posed Problems, **16**, 2, 175-194.
- [17] Mead, J., and Renaut, R. A., *A Newton root-finding algorithm for estimating the regularization parameter for solving ill-conditioned least squares problems*, to appear, Inverse Problems.

- [18] Nagy, J. G., Palmer, K, and Perrone, L., 2004, *Iterative methods for image deblurring: a Matlab object-oriented approach*, Numer. Algorithms, **36**, 73-93.
- [19] O’Leary, D. P. , 2001, *Near-optimal parameters for Tikhonov and other regularization methods*, SIAM J. on Sci. Comput., **23**, 1161-1171.
- [20] O’Leary, D. P., and Simmons, J. A., 1989, *A bidiagonalization-regularization procedure for large scale discretizations of ill-posed problems*, SIAM J. Sci. Stat. Comp., **2** (4), 474-489.
- [21] Paige, C. C., and Saunders, M. A., 1982, *LSQR: An algorithm for sparse linear equations and sparse least squares*, ACM Trans. Math. Software, **8**, 43-71.
- [22] Paige, C. C., and Saunders, M. A., 1982, *ALGORITHM 583 LSQR: Sparse linear equations and least squares problems*, ACM Trans. Math. Software, **8**, 195-209.
- [23] Rao, C. R., 1973, *Linear Statistical Inference and its applications*, (Wiley, New York).
- [24] Rust, B. W., and O’Leary, D. P., 2008, *Residual periodograms for choosing regularization parameters for ill-posed problems*, Inverse Problems, **24**, 034005.
- [25] Stefan, W., 2008, *Total Variation Regularization for Linear Ill-Posed Inverse Problems: Extensions and Applications*, Ph. D. Dissertation, Department of Mathematics and Statistics, Arizona State University, <http://mathpost.la.asu.edu/~stefan>.
- [26] Thompson, A. M., Kay, J. W., and Titterton, D. M., 1991, *Noise Estimation in Signal Restoration using Regularization*, Biometrika, **78**, 3, 475-488.
- [27] Vogel, C. R., 2002, *Computational Methods for Inverse Problems*, (SIAM Frontiers in Applied Mathematics).