

1-1-2013

α^2 Tests for the Choice of the Regularization Parameter in Nonlinear Inverse Problems

J. L. Mead
Boise State University

C. C. Hammerquist
Boise State University

χ^2 TESTS FOR THE CHOICE OF THE REGULARIZATION PARAMETER IN NONLINEAR INVERSE PROBLEMS*

J. L. MEAD[†] AND C. C. HAMMERQUIST[†]

Abstract. We address discrete nonlinear inverse problems with weighted least squares and Tikhonov regularization. Regularization is a way to add more information to the problem when it is ill-posed or ill-conditioned. However, it is still an open question as to how to weight this information. The discrepancy principle considers the residual norm to determine the regularization weight or parameter, while the χ^2 method [J. Mead, *J. Inverse Ill-Posed Probl.*, 16 (2008), pp. 175–194; J. Mead and R. A. Renaut, *Inverse Problems*, 25 (2009), 025002; J. Mead, *Appl. Math. Comput.*, 219 (2013), pp. 5210–5223; R. A. Renaut, I. Hnetynkova, and J. L. Mead, *Comput. Statist. Data Anal.*, 54 (2010), pp. 3430–3445] uses the regularized residual. Using the regularized residual has the benefit of giving a clear χ^2 test with a fixed noise level when the number of parameters is equal to or greater than the number of data. Previous work with the χ^2 method has been for linear problems, and here we extend it to nonlinear problems. In particular, we determine the appropriate χ^2 tests for Gauss–Newton and Levenberg–Marquardt algorithms, and these tests are used to find a regularization parameter or weights on initial parameter estimate errors. This algorithm is applied to a two-dimensional cross-well tomography problem and a one-dimensional electromagnetic problem from [R. C. Aster, B. Borchers, and C. Thurber, *Parameter Estimation and Inverse Problems*, Academic Press, New York, 2005].

Key words. least squares, regularization, nonlinear, covariance

AMS subject classifications. 92E24, 65F22, 65H10, 62J02

DOI. 10.1137/12088447X

1. Introduction. We address nonlinear inverse problems of the form $\mathbf{F}(\mathbf{x}) = \mathbf{d}$, where $\mathbf{d} \in \mathbb{R}^m$ represents measurements and $\mathbf{F} \in \mathbb{R}^{m \times n}$ represents a nonlinear model that depends on parameters $\mathbf{x} \in \mathbb{R}^n$. It is fairly straightforward, although possibly computationally demanding, to recover parameter estimates \mathbf{x} from data \mathbf{d} for a linear operator $\mathbf{F} = \mathbf{A}$ when \mathbf{A} is well-conditioned. If \mathbf{A} is ill-conditioned, typically constraints or penalties are added to the problem, and the problem is regularized. The most common regularization technique is Tikhonov regularization, where least squares is used to minimize

$$(1.1) \quad \|\mathbf{d} - \mathbf{A}\mathbf{x}\|_2^2 + \alpha^2 \|\mathbf{L}\mathbf{x}\|_2^2.$$

The minimum occurs at $\hat{\mathbf{x}}$,

$$\hat{\mathbf{x}} = (\mathbf{A}^T \mathbf{A} + \alpha^2 \mathbf{L}^T \mathbf{L})^{-1} \mathbf{A}^T \mathbf{d},$$

if the invertibility condition

$$\mathcal{N}(\mathbf{A}) \cap \mathcal{N}(\mathbf{L}) \neq \emptyset$$

is satisfied. When $\hat{\mathbf{x}}$ is written explicitly in this manner, it is clear how appropriate choice of α creates a well-conditioned problem. Popular methods for choosing α are

*Received by the editors July 13, 2012; accepted for publication (in revised form) by J. G. Nagy May 30, 2013; published electronically August 6, 2013. This work was supported by NSF grant DMS 1043107.

<http://www.siam.org/journals/simax/34-3/88447.html>

[†]Department of Mathematics, Boise State University, Boise, ID 83725-1555 (jmead@boisestate.edu, chammerquist@gmail.com).

the discrepancy principle, L-curve, and generalized cross validation (GCV) [5], while \mathbf{L} can be chosen to be the identity matrix, or a discrete approximation to a derivative operator. The problem with this approach is that large values of α may significantly change the problem or create undesired smooth solutions.

Nonlinear problems have the same issues but have the added burden of linear approximations [10]. The parameter estimate of the regularized nonlinear problem minimizes

$$(1.2) \quad \|\mathbf{d} - \mathbf{F}(\mathbf{x})\|_2^2 + \alpha^2 \|\mathbf{L}\mathbf{x}\|_2^2.$$

In [3] it is shown that this regularized problem is stable and that there is a weak connection between the ill-posedness of the nonlinear problem and the linearization. The authors of [3] and [12] prove the well-posedness of the regularized problem and convergence of regularized solutions.

The least squares estimate can be viewed as the maximum a posteriori (MAP) estimate [13] when the data and the initial parameter estimate \mathbf{x}_p have errors that are normally distributed with error covariances \mathbf{C}_ϵ and \mathbf{C}_f , respectively. Regularization in this case takes the form of adding a priori information about the parameters. The regularized maximum likelihood estimate minimizes the following objective function:

$$(1.3) \quad \mathcal{J} = (\mathbf{d} - \mathbf{F}(\mathbf{x}))^T \mathbf{C}_\epsilon^{-1} (\mathbf{d} - \mathbf{F}(\mathbf{x})) + (\mathbf{x} - \mathbf{x}_p)^T \mathbf{C}_f^{-1} (\mathbf{x} - \mathbf{x}_p).$$

The difficulty with this viewpoint is that the distribution of the data and initial parameter errors may be unknown, or not normal. In addition, good estimates of \mathbf{C}_ϵ and \mathbf{C}_f can be difficult to acquire. However, different norms or functionals can be sampled which represent error distributions other than normal, and scientific knowledge can be used to form prior error covariance matrices. Empirical Bayes methods [2] can also be used to find hyperparameters, i.e., parameters which define the prior distribution.

In this work we use the least squares estimator and weight the squared errors with inverse covariances matrices but do not assume the errors are normally distributed. Matrix weights alleviate the problem of undesired smooth solutions because they give piecewise smooth solutions; thus discontinuous parameters can be obtained. In preliminary work [7, 9] diagonal matrix weights \mathbf{C}_f were found, and future work involves more robust methods for finding \mathbf{C}_f . In this initial work on nonlinear problems, we approximate the unknown error covariance matrix with a scalar times the identity matrix, but future work involves finding more dense matrices. To find the scalar we use the χ^2 method, which is similar to the discrepancy principle.

The discrepancy principle can be viewed as a χ^2 test on the data residual, but this is not possible when the number of parameters is greater than or equal to the number of data because the degrees of freedom in the χ^2 test are zero or negative. This issue is typically not recognized because the degrees of freedom in the discrepancy principle are often taken to be m , the number of data, while they should be reduced by the number of parameters, i.e., reduced to $m - n$. The reduction is necessary because parameter estimates are found using the data, and hence their dependency should be subtracted from the degrees of freedom. This difference can be significant when the number of parameters is significantly different from the number of data. Alternatively we suggest applying the χ^2 test to the regularized residual (1.1), in which case the degrees of freedom are equal to the number of data [1]. This is called the χ^2 method, and previous work with the method applied to linear problems has been done [7, 8, 9, 11].

The χ^2 method is described in section 2, where we also develop it for nonlinear problems. In this work χ^2 tests are developed for Gauss–Newton and Levenberg–Marquardt algorithms, rather than for the nonlinear functional, because we have found that the statistical tests on the theoretical functional are not satisfied when the optimum value is approximated. We will explain how iterations in these nonlinear algorithms form a sequence of linear least squares problems and show how to apply the χ^2 principle at each iterate. From these tests we determine a regularization parameter, and this can be viewed as an estimate of the error covariance matrices. In section 3 we give results on benchmark problems from [1] and compare the nonlinear χ^2 method to other methods. In section 4 we give conclusions and future work.

2. Nonlinear χ^2 method. We formulate the problem as minimizing weighted errors $\boldsymbol{\epsilon}$ and \mathbf{f} in a least squares sense where the errors are defined as

$$(2.1) \quad \mathbf{d} = \mathbf{F}(\mathbf{x}) + \boldsymbol{\epsilon},$$

$$(2.2) \quad \mathbf{x} = \mathbf{x}_p + \mathbf{f}.$$

The weighted least squares fit of the measurements \mathbf{d} and initial parameter estimate \mathbf{x}_p is found by minimizing the objective function in (1.3). This means the weights are chosen to be the inverse covariance matrices for the respective errors. Statistical information comes in the form of error covariance matrices, but they are used only to weight the errors or residuals and are not used in the sense of Bayesian inference because the errors may not be Gaussian.

An initial estimate \mathbf{x}_p is required for this method, but this often comes from some understanding of the problem. The goal is then to find a way to weight this initial estimate. The elements in \mathbf{C}_f should reflect whether \mathbf{x}_p is a good or poor estimate and ideally quantify the correlation in the parameter errors. Estimation of \mathbf{C}_f is thus a challenge with limited understanding of the parameters, but regularization methods are an approach to its approximation. In this work we use a modified form of the discrepancy principle, called the χ^2 method [7], to estimate it.

It is well known that when the objective function (1.3) is applied to a linear model, it follows a χ^2 distribution [1, 7], and this fact is often used to also check the validity of estimates. This property is given in the following theorem.

THEOREM 1. *If $\mathbf{F}(\mathbf{x}) = \mathbf{A}\mathbf{x}$, \mathbf{d} , and \mathbf{x} are independent and identically distributed random variables from unknown distributions with $\bar{\boldsymbol{\epsilon}} = \bar{\mathbf{f}} = \mathbf{0}$, $\overline{\boldsymbol{\epsilon}\boldsymbol{\epsilon}^T} = \mathbf{C}_\epsilon$, $\overline{\mathbf{f}\mathbf{f}^T} = \mathbf{C}_f$, and $\overline{\boldsymbol{\epsilon}^T\mathbf{f}} = 0$, and the number of data m is large, then the minimum value of (1.3) asymptotically follows a χ^2 distribution with m degrees of freedom.*

A proof of this is given in [7]. In the case that \mathbf{x}_p is not the mean of \mathbf{x} , then (1.3) follows a noncentral χ^2 distribution as shown in [11]. The requirement that $\overline{\boldsymbol{\epsilon}^T\mathbf{f}} = 0$ assumes that the errors in \mathbf{d} and \mathbf{x}_p are uncorrelated, and hence we do not get \mathbf{x}_p from \mathbf{d} .

The χ^2 method is used to find weights for misfits in the initial parameter estimates, \mathbf{C}_f , or data, \mathbf{C}_ϵ , and is based on the χ^2 tests given in Theorem 1. The method can also be used to find weights for the regularization term when it contains a first or second derivative operator \mathbf{L} . In this case the inverse covariance matrix is viewed as weighting the error in an initial estimate of the appropriate derivative. The minimum value of (1.3) is

$$\mathcal{J}(\hat{\mathbf{x}}) = \mathbf{r}^T(\mathbf{A}\mathbf{C}_f\mathbf{A}^T + \mathbf{C}_\epsilon)^{-1}\mathbf{r}$$

with $\mathbf{r} = \mathbf{d} - \mathbf{A}\mathbf{x}_p$ [7]. Its expected value is the number of degrees of freedom, which is the number of measurements in \mathbf{d} . The χ^2 method thus estimates scalar weight \mathbf{C}_f

or \mathbf{C}_ϵ by solving the nonlinear equation

$$(2.3) \quad \mathbf{r}^T (\mathbf{A} \mathbf{C}_f \mathbf{A}^T + \mathbf{C}_\epsilon)^{-1} \mathbf{r} = m.$$

For example, with \mathbf{d} given by the data and \mathbf{x}_p an initial parameter estimate (possibly $\mathbf{0}$), if we have estimates of data uncertainty \mathbf{C}_ϵ , we can solve for \mathbf{C}_f (alternatively we can solve for \mathbf{C}_ϵ given \mathbf{C}_f). This differs from the discrepancy principle in that the regularized residual is used for the χ^2 test rather than just the data residual. One advantage of this approach over the discrepancy principle is that it can be applied when the number of parameters is greater than or equal to the number of data. Results from the χ^2 method for the case when $\mathbf{C}_f = \sigma_f^2 \mathbf{I}$, i.e., the scalar χ^2 method, are given in [8, 11]. It is shown there that this is an attractive alternative to the L-curve, GCV, and the discrepancy principle among other methods. Preliminary work for more dense estimates of \mathbf{C}_ϵ or \mathbf{C}_f has been done in [7, 9], and future work involves efficient solution of nonlinear systems similar to (2.3) to estimate more dense \mathbf{C}_f or \mathbf{C}_ϵ .

Here we extend the estimation of \mathbf{C}_f or \mathbf{C}_ϵ by χ^2 tests to nonlinear problems. It has been shown that the nonlinear data residual has properties similar to those of the linear data residual [10]. In particular, it behaves like a χ^2 random variable with $m - n$ degrees of freedom. Here we show that when Newton-type methods are used to find the minimum of (1.3), the corresponding cost functions at each iterate are also χ^2 , and we give the χ^2 tests for both the Gauss–Newton and Levenberg–Marquardt methods.

2.1. Gauss–Newton method. The Gauss–Newton method uses Newton’s method to estimate the minimum value of a function. When used to find parameters that minimize (1.3), the following estimate is obtained at each iterate:

$$(2.4) \quad \mathbf{x}_{k+1} = \mathbf{x}_k + (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \mathbf{J}_k + \mathbf{C}_f^{-1})^{-1} (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \mathbf{r}_k - \mathbf{C}_f^{-1} \Delta \mathbf{x}_k),$$

where \mathbf{J}_k is the Jacobian of $\mathbf{F}(\mathbf{x})$ about the previous estimate \mathbf{x}_k , $\mathbf{r}_k = \mathbf{d} - \mathbf{F}(\mathbf{x}_k)$, and $\Delta \mathbf{x}_k = \mathbf{x}_k - \mathbf{x}_p$. Iteratively regularized Gauss–Newton updates the regularization parameter at each estimate so that $\mathbf{C}_f^{-1} = \alpha_k \mathbf{I}$. At each iterate, this can be viewed as an estimate that minimizes the following functional:

$$(2.5) \quad \mathcal{J}(\mathbf{x}) \approx \tilde{\mathcal{J}}_k(\mathbf{x}) = (\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x})^T \mathbf{C}_\epsilon^{-1} (\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}) + \alpha_k (\mathbf{x} - \mathbf{x}_p)^T (\mathbf{x} - \mathbf{x}_p),$$

with $\tilde{\mathbf{d}}_k = \mathbf{d} - \mathbf{F}(\mathbf{x}_k) + \mathbf{J}_k \mathbf{x}_k$. With this view, at each iterate the Gauss–Newton functional minimizes weighted errors $\boldsymbol{\epsilon}_k$ and \mathbf{f}_k defined by

$$(2.6) \quad \tilde{\mathbf{d}}_k = \mathbf{J}_k \mathbf{x} + \boldsymbol{\epsilon}_k,$$

$$(2.7) \quad \mathbf{x} = \mathbf{x}_p + \mathbf{f}_k,$$

where $\boldsymbol{\epsilon}_k = \mathbf{v}_k + \boldsymbol{\epsilon}$, \mathbf{v}_k represents error in linear approximation at step k , i.e.,

$$\mathbf{v}_k = \mathbf{F}(\mathbf{x}) - \mathbf{F}(\mathbf{x}_k) - \mathbf{J}_k (\mathbf{x} - \mathbf{x}_k),$$

and $\overline{\mathbf{f}_k \mathbf{f}_k^T}^{-1} = \alpha_k \mathbf{I}$.

Inverse methods typically identify error in data, but it is less common to account for error in the numerical approximation of the model. In most applications, bias in data error is subtracted out so that $\bar{\boldsymbol{\epsilon}} = \mathbf{0}$ is a reasonable assumption. In the following theorem we will assume that $\bar{\boldsymbol{\epsilon}}_k = \mathbf{0}$, which is not valid unless we remove the bias in

the linearization error. Estimating bias in the linearization error is a topic for future work; however, we continue with the theoretical development using this assumption. The fact that the regularization parameter is iteratively adjusted helps to alleviate the effects of this linearization error.

The linear cost function (2.5) is the basis for the linear χ^2 method, which states that

$$\|\tilde{\mathbf{d}}_k - \mathbf{J}_k \hat{\mathbf{x}}\|_{\mathbf{C}_{\epsilon_k}^{-1}} + \|\hat{\mathbf{x}} - \mathbf{x}_p\|_{\mathbf{C}_{f_k}^{-1}} = \mathbf{r}^T \mathbf{P}_k^{-1} \mathbf{r} \sim \chi_m^2$$

with $\mathbf{P}_k = \mathbf{J}_k \mathbf{C}_{f_k} \mathbf{J}_k^T + \mathbf{C}_{\epsilon_k}$ and $\mathbf{r} = \tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}_p$. The Gauss–Newton method can be viewed as optimizing the linear cost function at each iterate, but the algorithm is not completely linear in that the optimal estimate at each iterate is \mathbf{x}_{k+1} in (2.4) rather than $\hat{\mathbf{x}}$. The difference is that \mathbf{x}_{k+1} depends on the estimate at the previous iteration; i.e., $\hat{\mathbf{x}}$ is (2.4) with $\Delta \mathbf{x}_k = \mathbf{0}$. We will show in the following theorem that under specific assumptions, the linear cost function at the Gauss–Newton iterate also follows a χ^2 distribution with m degrees of freedom:

$$\|\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}_{k+1}\|_{\mathbf{C}_{\epsilon_k}^{-1}} + \|\mathbf{x}_{k+1} - \mathbf{x}_p\|_{\mathbf{C}_{f_k}^{-1}} = \mathbf{r}^T \mathbf{P}_k^{-1} \mathbf{r} \sim \chi_m^2.$$

THEOREM 2. *Let $\overline{\epsilon_k} = \overline{\mathbf{f}_k} = \mathbf{0}$, $\overline{\epsilon_k^T \mathbf{f}_k} = 0$, $\overline{\epsilon_k \epsilon_k^T} = \mathbf{C}_{\epsilon_k}$, and $\overline{\mathbf{f}_k \mathbf{f}_k^T} = \mathbf{C}_{f_k}$; then the value of (2.5) at each iterate is*

$$(2.8) \quad \tilde{\mathcal{J}}_k(\mathbf{x}_{k+1}) = (\mathbf{r}_k + \mathbf{J}_k \Delta \mathbf{x}_k)^T \mathbf{P}_k^{-1} (\mathbf{r}_k + \mathbf{J}_k \Delta \mathbf{x}_k)$$

with $\mathbf{P}_k = \mathbf{J}_k \mathbf{C}_{f_k} \mathbf{J}_k^T + \mathbf{C}_{\epsilon_k}$, and $\tilde{\mathcal{J}}_k(\mathbf{x}_{k+1})$ follows a χ^2 distribution with m degrees of freedom.

Proof. The estimate \mathbf{x}_{k+1} and is given by (2.4).

Define $\mathbf{Q}_k = \mathbf{J}_k^T \mathbf{C}_{\epsilon_k}^{-1} \mathbf{J}_k + \mathbf{C}_f^{-1}$ so that

$$(2.9) \quad \mathbf{x}_{k+1} = \mathbf{x}_k + (\mathbf{J}_k \mathbf{C}_f)^T \mathbf{P}_k^{-1} \mathbf{r}_k - \mathbf{Q}_k^{-1} \mathbf{C}_f^{-1} \Delta \mathbf{x}_k.$$

The linear cost function (2.5) at this estimate is

$$(2.10) \quad \tilde{\mathcal{J}}_k(\mathbf{x}_{k+1}) = \mathbf{r}_k^T \mathbf{P}_k^{-1} \mathbf{r}_k + 2 \Delta \mathbf{x}_k^T \mathbf{J}_k^T \mathbf{P}_k^{-1} \mathbf{r}_k + \Delta \mathbf{x}_k^T \mathbf{C}_f^{-1} (\mathbf{C}_f - \mathbf{Q}_k^{-1}) \mathbf{C}_f^{-1} \Delta \mathbf{x}_k.$$

Simplifying further, we note that

$$\begin{aligned} \mathbf{C}_f^{-1} (\mathbf{C}_f - \mathbf{Q}_k^{-1}) \mathbf{C}_f^{-1} &= \mathbf{C}_f^{-1} (\mathbf{C}_f \mathbf{Q}_k - \mathbf{I}) \mathbf{Q}_k^{-1} \mathbf{C}_f^{-1} \\ &= \mathbf{C}_f^{-1} (\mathbf{C}_f \mathbf{J}_k^T \mathbf{C}_{\epsilon_k}^{-1} \mathbf{J}_k \mathbf{Q}_k^{-1}) \mathbf{C}_f^{-1} \\ &= \mathbf{J}_k^T (\mathbf{C}_f^{-1} \mathbf{Q}_k^{-1} \mathbf{J}_k^T \mathbf{C}_{\epsilon_k}^{-1})^T \\ &= \mathbf{J}_k^T (\mathbf{J}_k^T \mathbf{P}_k^{-1})^T, \end{aligned}$$

where the last equality uses the fact that $\mathbf{C}_f \mathbf{J}_k^T \mathbf{P}_k^{-1} = \mathbf{Q}_k^{-1} \mathbf{J}_k^T \mathbf{C}_{\epsilon_k}^{-1}$. Now we have the quadratic form

$$\begin{aligned} \tilde{\mathcal{J}}_k(\mathbf{x}_{k+1}) &= (\mathbf{r}_k + \mathbf{J}_k \Delta \mathbf{x}_k)^T \mathbf{P}_k^{-1} (\mathbf{r}_k + \mathbf{J}_k \Delta \mathbf{x}_k) \\ &= \mathbf{k}^T \mathbf{k}, \end{aligned}$$

where $\mathbf{r}_k + \mathbf{J}_k \Delta \mathbf{x}_k = \tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}_p = \mathbf{P}_k^{1/2} \mathbf{k}$. The square root of \mathbf{P}_k is defined since \mathbf{C}_{ϵ_k} and \mathbf{C}_{f_k} are covariance matrices, and hence \mathbf{P}_k is symmetric positive definite.

The mean of \mathbf{k} is

$$\bar{\mathbf{k}} = \mathbf{P}_k^{-1/2}(\bar{\mathbf{f}}_k + \bar{\boldsymbol{\epsilon}}_k),$$

which is zero under our assumptions. In addition, the covariance of \mathbf{k} is

$$\begin{aligned} \overline{\mathbf{k}\mathbf{k}^T} &= \mathbf{P}_k^{-1/2} (\mathbf{J}_k \mathbf{C}_{f_k} \mathbf{J}_k^T + \mathbf{C}_{\epsilon_k}) \mathbf{P}_k^{-1/2} \\ &= \mathbf{I}. \quad \square \end{aligned}$$

Note that $\mathbf{C}_{\epsilon_k} = \mathbf{C}_{\nu_k} + \mathbf{C}_\epsilon$ and \mathbf{C}_{ν_k} can be approximated using the Hessian \mathbf{H}_k at each iterate; i.e.,

$$(2.11) \quad \mathbf{C}_{\nu_k} \approx \frac{\mathbf{H}_k}{2} \text{cov}((\mathbf{x} - \mathbf{x}_k)^2) \frac{\mathbf{H}_k^T}{2}.$$

2.2. Occam’s method. Occam’s method estimates parameters that minimize the nonlinear functional (1.3) by updating \mathbf{x}_p with \mathbf{x}_k found by minimizing the linear functional (2.5). The regularization parameter is chosen according to the discrepancy principle at each step, and $\mathbf{C}_f^{-1} = \alpha_k^2 \mathbf{L}\mathbf{L}^T$ with the operator \mathbf{L} typically chosen to represent the second derivative [1]. The discrepancy principle finds the value of α_k for which $\|\mathbf{d} - \mathbf{F}(\mathbf{x}_{k+1})\|_{\mathbf{C}_\epsilon^{-1}}^2 \leq \|\boldsymbol{\epsilon}\|_2^2$.

The goal in Occam’s inversion is to find smooth parameters that fit the data [4]. When \mathbf{x}_p , and hence the second derivate estimate, is chosen to be $\mathbf{0}$, smooth results are obtained. Guessing the second derivative of the parameter values to be zero may be a good choice for prior information, if little information is known about \mathbf{x}_p . The discrepancy principle is then applied where a χ^2 test for the data residual is used to find weights α_k . The χ^2 test in Theorem 2 differs from Occam’s inversion in that the regularized residual is used to find the weights α_k , and the Gauss–Newton update rather than the optimum for the linearized cost function is used for the test. When the operator \mathbf{L} is not full rank, the degrees of freedom in the χ^2 test change [8]. This leads to the following Corollary to Theorem 2.

COROLLARY 3. *Let*

$$(2.12) \quad \tilde{\mathcal{J}}_L(\mathbf{x}) = (\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x})^T \mathbf{C}_{\epsilon_k}^{-1} (\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}) + (\mathbf{x} - \mathbf{x}_p)^T \mathbf{L}^T (\mathbf{C}_f^L)^{-1} \mathbf{L} (\mathbf{x} - \mathbf{x}_p),$$

and assume the invertibility condition

$$(2.13) \quad \mathcal{N}(\mathbf{J}_k) \cap \mathcal{N}(\mathbf{L}) \neq \emptyset,$$

where $\mathcal{N}(\mathbf{J}_k)$ is the null space of matrix \mathbf{J}_k . In addition, let the assumptions in Theorem 2 hold except that $\overline{\mathbf{L}\mathbf{f}(\mathbf{L}\mathbf{f})^T} = \mathbf{C}_f^L$. If \mathbf{L} has rank q , the minimum value of the functional (1.2) follows a χ^2 distribution with $m - n + q$ degrees of freedom.

Proof. The proof for linear problems is given in [8]. \square

2.3. Levenberg–Marquardt method. With the Gauss–Newton method it may be the case that the objective function does not decrease, or it may decrease slowly at each iteration. The solution can be incremented in the direction of steepest descent by introducing the damping parameter λ_k . The iterate in this case at each step is

$$(2.14) \quad \mathbf{x}_{k+1} = \mathbf{x}_k + (\mathbf{J}_k^T \mathbf{C}_{\epsilon_k}^{-1} \mathbf{J}_k + \mathbf{C}_f^{-1} + \lambda_k \mathbf{I})^{-1} (\mathbf{J}_k^T \mathbf{C}_{\epsilon_k}^{-1} \mathbf{r}_k - \mathbf{C}_f^{-1} \Delta \mathbf{x}_k).$$

In [1] it is stated that this approach differs from Tikhonov regularization because it does not alter the objective function (1.3). However, at each iterate it does minimize an altered form of $\tilde{\mathcal{J}}_k(\mathbf{x})$ in (2.5); i.e., it minimizes

$$(2.15) \quad \tilde{\mathcal{J}}_k^{LM}(\mathbf{x}) = \tilde{\mathcal{J}}_k(\mathbf{x}) + \lambda_k(\mathbf{x} - \mathbf{x}_k)^T(\mathbf{x} - \mathbf{x}_k).$$

This cost function minimizes the weighted errors ϵ_k , \mathbf{f} , and δ_k defined by

$$\begin{aligned} \tilde{\mathbf{d}}_k &= \mathbf{J}_k\mathbf{x} + \epsilon_k, \\ \mathbf{x} &= \mathbf{x}_p + \mathbf{f}, \\ \mathbf{x} &= \mathbf{x}_k + \delta_k, \end{aligned}$$

where δ_k represents the error in the parameter estimate at step k .

THEOREM 4. *With the same assumptions as in Theorem 2, and if $\bar{\delta}_k = \mathbf{0}$, $\overline{\delta_k \delta_k^T} = \lambda^{-1}\mathbf{I}$, $\overline{\epsilon \delta_k^T} = \mathbf{0}$, and $\overline{\mathbf{f} \delta_k^T} = \mathbf{0}$, then the minimum value of (2.15) at each iterate is*

$$\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1}) = (\mathbf{r}_k + \mathbf{J}_k\mathbf{B}_k\Delta\mathbf{x}_k)^T (\mathbf{P}_k^\lambda)^{-1} (\mathbf{r}_k + \mathbf{J}_k\mathbf{B}_k\Delta\mathbf{x}_k) - s_k$$

with $\mathbf{B}_k = (\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1}\mathbf{C}_f^{-1}$, $s_k = \Delta\mathbf{x}_k^T\mathbf{C}_f^{-1}(\mathbf{B}_k - \mathbf{I})\Delta\mathbf{x}_k$, and $\mathbf{P}_k^\lambda = \mathbf{J}_k\mathbf{B}_k\mathbf{C}_f\mathbf{J}_k^T + \mathbf{C}_{\epsilon_k}$. Here $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1}) + s_k$ follows a χ^2 distribution with m degrees of freedom.

Proof. The proof follows similarly as that of Theorem 2; however, we define $\mathbf{Q}_k^\lambda = \mathbf{Q}_k + \lambda_k\mathbf{I}$ so that

$$(2.16) \quad \mathbf{x}_{k+1} = \mathbf{x}_k + (\mathbf{J}_k\mathbf{B}_k\mathbf{C}_f)^T(\mathbf{P}_k^\lambda)^{-1}\mathbf{r}_k - \mathbf{Q}_k^{-1}\mathbf{C}_f^{-1}\Delta\mathbf{x}_k.$$

The minimum value is

$$\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1}) = \mathbf{r}_k^T(\mathbf{P}_k^\lambda)^{-1}\mathbf{r}_k + 2\Delta\mathbf{x}_k^T(\mathbf{J}_k\mathbf{B}_k)^T(\mathbf{P}_k^\lambda)^{-1}\mathbf{r}_k + \Delta\mathbf{x}_k^T\mathbf{C}_f^{-1}(\mathbf{C}_f - \mathbf{Q}_k^{-1})\mathbf{C}_f^{-1}\Delta\mathbf{x}_k.$$

Now here the last term in $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1})$ does not lead to a quadratic functional. However, the first two terms suggest the quadratic form $\mathbf{k}_\lambda^T\mathbf{k}_\lambda$ with

$$\mathbf{k}_\lambda = (\mathbf{P}_k^\lambda)^{-1/2}(\mathbf{r}_k + \mathbf{J}_k\mathbf{B}_k\Delta\mathbf{x}_k).$$

To determine s_k we note that

$$\begin{aligned} (\mathbf{J}_k\mathbf{B}_k)^T(\mathbf{P}_k^\lambda)^{-1}\mathbf{J}_k\mathbf{B}_k &= (\mathbf{J}_k\mathbf{B}_k)^T(\mathbf{C}_f^{-1}\mathbf{Q}^{-1}\mathbf{J}^T\mathbf{C}_{\epsilon_k}^{-1})^T \\ &= (\mathbf{J}_k(\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1}\mathbf{C}_f^{-1})^T\mathbf{C}_{\epsilon_k}^{-1}\mathbf{J}_k\mathbf{Q}_k^{-1}\mathbf{C}_f^{-1} \\ &= \mathbf{C}_f^{-1}\left((\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1}\mathbf{J}_k^T\mathbf{C}_{\epsilon_k}^{-1}\mathbf{J}\mathbf{Q}^{-1}\right)\mathbf{C}_f^{-1} \\ &= \mathbf{C}_f^{-1}\left((\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1}\mathbf{Q} - \mathbf{I}\right)\mathbf{Q}^{-1}\mathbf{C}_f^{-1} \\ &= \mathbf{C}_f^{-1}\left((\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1} - \mathbf{Q}^{-1}\right)\mathbf{C}_f^{-1}, \end{aligned}$$

where the first equality uses the fact that $\mathbf{Q}_k^{-1}\mathbf{J}_k^T\mathbf{C}_{\epsilon_k}^{-1} = (\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1}\mathbf{J}_k^T\mathbf{P}_k^{-1}$. Unfortunately, this is not $\mathbf{C}_f^{-1}(\mathbf{C}_f - \mathbf{Q}_k^{-1})\mathbf{C}_f^{-1}$, which would give $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1})$ quadratic form, but we can use it to find s_k . Since

$$\mathbf{C}_f - \mathbf{Q}_k^{-1} = \left((\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1} - \mathbf{Q}^{-1}\right) - \left((\mathbf{C}_f^{-1} + \lambda_k\mathbf{I})^{-1} - \mathbf{C}_f\right),$$

we have that

$$\begin{aligned} s_k &= \Delta \mathbf{x}_k^T \mathbf{C}_f^{-1} \left((\mathbf{C}_f^{-1} + \lambda_k \mathbf{I})^{-1} - \mathbf{C}_f \right) \mathbf{C}_f^{-1} \Delta \mathbf{x}_k \\ &= \Delta \mathbf{x}_k^T \mathbf{C}_f^{-1} (\mathbf{B}_k - \mathbf{I}) \Delta \mathbf{x}_k. \end{aligned}$$

Finally, we now show that the mean and covariance of the quadratic part of the functional $\mathbf{k}_\lambda^T \mathbf{k}_\lambda$ are zero and the identity matrix, respectively. First, we note that if $\bar{\mathbf{f}} = \bar{\boldsymbol{\delta}}_k = \mathbf{0}$, then $\overline{\Delta \mathbf{x}_k} = \bar{\boldsymbol{\delta}}_k - \bar{\mathbf{f}} = \mathbf{0}$, and hence $\bar{\mathbf{k}}_\lambda = \mathbf{0}$. For the covariance, note that

$$\begin{aligned} \mathbf{r}_k + \mathbf{J}_k \mathbf{B}_k \Delta \mathbf{x}_k &= \tilde{\mathbf{d}} - \mathbf{J} \mathbf{x}_k + \mathbf{J}_k \mathbf{B}_k \Delta \mathbf{x}_k \\ &= \boldsymbol{\epsilon} + \mathbf{J}_k \boldsymbol{\delta} + \mathbf{J}_k \mathbf{B}_k (\mathbf{f} - \boldsymbol{\delta}). \end{aligned}$$

With the assumption that the errors are not correlated, we thus have

$$\begin{aligned} E \left((\mathbf{r}_k + \mathbf{J}_k \mathbf{B}_k \Delta \mathbf{x}_k) (\mathbf{r}_k + \mathbf{J}_k \mathbf{B}_k \Delta \mathbf{x}_k)^T \right) \\ = \mathbf{C}_{\epsilon_k} + \mathbf{J}_k (1/\lambda \mathbf{I} - 1/\lambda \mathbf{B}_k^T + \mathbf{B}_k \mathbf{C}_f \mathbf{B}_k^T - 1/\lambda \mathbf{B}_k + 1/\lambda \mathbf{B}_k \mathbf{B}_k^T) \mathbf{J}^T. \end{aligned}$$

Simplifying further, we get

$$\begin{aligned} 1/\lambda \mathbf{I} - 1/\lambda \mathbf{B}_k^T + \mathbf{B}_k \mathbf{C}_f \mathbf{B}_k^T - 1/\lambda \mathbf{B}_k + 1/\lambda \mathbf{B}_k \mathbf{B}_k^T \\ = \mathbf{B}_k (1/\lambda \mathbf{B}_k^{-1} - 1/\lambda \mathbf{B}_k^{-1} \mathbf{B}_k^T + \mathbf{C}_f \mathbf{B}_k^T - 1/\lambda \mathbf{I} + 1/\lambda \mathbf{B}_k^T) \\ = \mathbf{B}_k (\mathbf{C}_f + (\mathbf{C}_f + 1/\lambda \mathbf{I} - 1/\lambda \mathbf{B}_k^{-1}) \mathbf{B}_k^T) \\ = \mathbf{B}_k \mathbf{C}_f, \end{aligned}$$

where the second equality uses the fact that $\mathbf{B}_k^{-1} \mathbf{B}_k^T = \mathbf{I}$ both in the second and fourth terms in the previous equality. Thus

$$E(\mathbf{k}_\lambda \mathbf{k}_\lambda^T) = \mathbf{P}^{-1/2} \left(\mathbf{C}_{\epsilon_k} + \mathbf{J} (\mathbf{C}_f^{-1} + \lambda \mathbf{I})^{-1} \mathbf{J}^T \right) \mathbf{P}^{-1/2} = \mathbf{I}. \quad \square$$

We note here that by setting $\lambda_k = 0$ we get the same result as in Theorem 2.

The assumptions in Theorem 4 may be difficult to verify or may not hold exactly. For example, we assume that the damping parameter λ_k represents the inverse of the error covariance for the parameter estimate at step k . The damping parameter is not chosen to estimate the covariance but rather to shorten the step in the Gauss–Newton iteration. There is a problem with its interpretation as the inverse of the standard deviation for the error because as λ_k decreases to zero the covariance is undefined. Alternatively, we take the view that when $\lambda_k \approx 0$, infinite weight is given to $\mathbf{x} - \mathbf{x}_k$ in the objective function since \mathbf{x}_k is considered a good estimate of \mathbf{x} . In section 3 we give histograms of $\tilde{\mathcal{J}}_k(\mathbf{x})$ and $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x})$ and show experimentally that these theorems hold.

3. Numerical experiments. We present two problems to illustrate both the validity of the χ^2 tests and the applicability of these tests to solving nonlinear inverse problems. The nonlinear problems are from Chapter 10 of [1], where the authors not only describe and illustrate solutions to these problems, but also provide corresponding MATLAB codes that both set up the forward problems and find solutions to the inverse problems. In particular, we used the m-files available with the text to compute approximations using Gauss–Newton with the discrepancy principle as described in Algorithm 1 and Occam’s method as described in Algorithm 2.

ALGORITHM 1 (Gauss–Newton with discrepancy principle).

Input \mathbf{L} , \mathbf{C}_ϵ , \mathbf{x}_p
for $i = 1, 2, 3, \dots$ **do**
 Generate logarithmically spaced α_i
 for $k = 1, 2, 3, \dots$ **do**
 Calculate $\mathbf{x}_{k+1} = (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \mathbf{J}_k + \alpha_i^2 \mathbf{L}^T \mathbf{L})^{-1} (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \tilde{\mathbf{d}}_k - \alpha_i^2 \mathbf{L}^T \mathbf{L} \mathbf{x}_p)$
 if $|\mathbf{x}_{k+1} - \mathbf{x}_k| < \text{tol}$ **then**
 $\mathbf{x}_i = \mathbf{x}_{k+1}$
 break
 end if
 end for
 Calculate $\mathcal{J}_{data}^i = \|\mathbf{d} - \mathbf{F}(\mathbf{x}_i)\|_{\mathbf{C}_\epsilon^{-1}}^2$
end for
Choose i for smallest value of $|\mathcal{J}_{data}^i - m|$
 $\mathbf{x} \approx \mathbf{x}_i$.

ALGORITHM 2 (Occam’s inversion [1]).

Input \mathbf{L} , \mathbf{C}_ϵ , \mathbf{x}_0 , δ
for $k = 1, 2, 3, \dots, 30$ **do**
 Calculate \mathbf{J}_k and $\tilde{\mathbf{d}}_k$
 Define $\mathbf{x}_{k+1} = (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \mathbf{J}_k + \alpha_k^2 \mathbf{L}^T \mathbf{L})^{-1} \mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \tilde{\mathbf{d}}_k$
 while $\frac{\|\mathbf{x}_k - \mathbf{x}_{k-1}\|}{\|\mathbf{x}_k\|} > 5e - 3$ or $\|\mathbf{d} - \mathbf{F}(\mathbf{x}_{k+1})\|_{\mathbf{C}_\epsilon^{-1}}^2 > 1.01\delta^2$ **do**
 Choose largest value of α_k such $\|\mathbf{d} - \mathbf{F}(\mathbf{x}_{k+1})\|_{\mathbf{C}_\epsilon^{-1}}^2 \leq m$
 If no such α_k exists, then chose a α_k that minimizes $\|\mathbf{d} - \mathbf{F}(\mathbf{x}_{k+1})\|_{\mathbf{C}_\epsilon^{-1}}^2$
 end while
end for

3.1. Nonlinear inverse problem descriptions. The first problem is an implementation of nonlinear cross-well tomography. The forward model includes ray path refraction, and the refracted rays tend to travel through high-velocity regions and avoid low-velocity regions, adding nonlinearity to the problem. It is set up with two wells spaced 1600 m apart, with 150 sources and receivers equally spaced at 100 m depths down the wells. The travel time between each pair of opposing sources and receivers is recorded, and the objective is to recover the two-dimensional velocity structure between the two wells. The true velocity structure has a background of 2.9 km/s with an embedded Gaussian-shaped region that is about 10% faster and another Gaussian-shaped region that is about 15% slower than the background. The observations for this particular problem consist of 64 travel times between each pair of opposing sources and receivers. The true velocity model along with the 64 ray paths are plotted in Figure 1. The region between the two wells is discretized into 64 square blocks, so there are 64 model parameters (the slowness of each block) and 64 observations (the ray path travel times). The second problem considered is the estimation of the soil electrical conductivity profile from above-ground electromagnetic induction measurements, as illustrated in Figure 2. The forward problem models a ground conductivity meter which has two coils on a one meter long bar. Alternat-




FIG. 1. The setup of the cross-well tomography problem. Shown here is the true velocity model (m/s) and the corresponding ray paths.




FIG. 2. A representation of the soil conductivity estimation. The instrument depicted in the top of image represents a ground conductivity meter creating a time-varying electromagnetic field in the layered earth beneath.

ing current is sent in one of the coils which induces currents in soil, and both coils measure the magnetic field that is created by the subsurface currents. For a complete treatment of the instrument and corresponding mathematical model, see [6]. There are a total of 18 observations, and the subsurface electrical conductivity of the ground is discretized into 10 layers, 20 cm thick, with a semi-infinite layer below 2m, resulting in 11 conductivities to be estimated. As noted in [1], in solving this inverse problem, the Gauss–Newton method does not always converge. Therefore, finding the solution necessitated the use of the Levenberg–Marquardt algorithm.

3.2. Numerical validation of χ^2 tests. We numerically tested Theorem 2 on the cross-well tomography problem by applying the Gauss–Newton method to minimize (1.3) for 1000 different realizations of ϵ and \mathbf{f} which were sampled from $N(\vec{0}, (.001)^2\mathbf{I})$ and $N(\vec{0}, (.00001)^2\mathbf{I})$, respectively. For perspective, the values for \mathbf{d} are $\mathcal{O}(.1)$, while the values for \mathbf{x} are $\mathcal{O}(.0001)$, so that 1% noise is added to the data while the initial parameter estimate has 10% added noise. We then used the Gauss–Newton method to solve the nonlinear inverse problem 1000 times, once for each realization of noise, which is essentially equivalent to sampling $\tilde{\mathcal{J}}_k$ 1000 times. Each of these converged in six iterations. Histograms of samples of $\tilde{\mathcal{J}}_k$ at each iteration are shown




FIG. 3. Histogram showing the sample distribution of $\tilde{\mathcal{J}}_k(\mathbf{x})$ at each Gauss–Newton iteration for the cross-well tomography problem. The mean of the sample is shown as the middle red tick.

in Figure 3, and the mean is labeled in red. Since there are 64 observations and 64 model parameters, the theory says that $\tilde{\mathcal{J}}_k \sim \chi_{64}^2$ with $E(\tilde{\mathcal{J}}_k) = 64$. In the beginning iterations, C_ε slightly underestimates C_{ε_k} , so the mean of $\tilde{\mathcal{J}}_k$ is somewhat larger than the theoretical mean.

We also tested Corollary 3 using this tomography problem when \mathbf{L} does not have full rank. We followed [1], where a discrete approximation of the Laplacian for \mathbf{L} is used in order to regularize this problem. Noise was generated in the same manner as in Figure 3, and histograms of $\tilde{\mathcal{J}}_L$ from (1.2) are plotted in Figure 4. The rank of this operator is 63 so that $E(\mathcal{J}_{Lk}) = 63$. The histograms in Figures 3 and 4 and their respective means at the final iterations show that the sampled distributions are good approximations of the theoretical χ^2 distributions in Theorem 2 and Corollary 3.

The electrical conductivity problem was used to numerically test Theorem 4 since the Levenberg–Marquardt algorithm was used to estimate the model parameters when minimizing (1.3). As in the previous simulations, this inversion was also repeated for 1000 different realizations of $\boldsymbol{\epsilon}$ and \mathbf{f} which were sampled from $N(\vec{0}, (.1)^2 \mathbf{I})$ and $N(\vec{0}, (100)^2 \mathbf{I})$, respectively. The values for \mathbf{d} are $\mathcal{O}(100)$, and the values for \mathbf{x} are $\mathcal{O}(100)$, so that 1% noise is added to data while the initial parameter estimate has 100% added noise. Almost all of these trials converged in six iterations or fewer with the majority converging in five iterations. Histograms for $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1})$ are shown in Figure 5. There were 18 observations for this problem and 11 parameters, so $E(\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1})) = 18$. Theorem 4 states that $E(\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1}) + s_k) = 18$, while at convergence when $\lambda_k = 0$ we have $s_k = 0$. From the histograms, we see that even at early iterations s_k is negligible, and $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1})$ behaves like χ_{18}^2 distribution. Therefore, in the inversion results we neglect the effects of s_k when calculating the regularization parameter.

This experiment was also repeated when \mathbf{L} does not have full rank, and hence there are reduced degrees of freedom in the χ^2 distribution. We used a discrete




FIG. 4. Histogram showing the sample distribution of $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x})$ at each Gauss–Newton iteration for the cross-well tomography problem. The mean of the sample is shown as the middle red tick.

second order differential operator with rank 9 to regularize the inversion; therefore, $E(\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1}) + s_k) = 16$. In this experiment f was sampled from $N(\vec{0}, (10)^2\mathbf{I})$ since $\mathbf{L}\mathbf{x}$ are $\mathcal{O}(10)$. All of these Levenberg–Marquardt (LM) iterations converged in five or fewer iterations with the majority converging in four iterations. These results are also shown in Figure 6, where we see that the sample mean is 17 at convergence. In addition, $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1})$ behaves like χ_{16}^2 at the early iterations, so we also neglect the effects of s_k when \mathbf{L} does not have full rank.

3.3. Inverse problem results.

3.3.1. Cross-well tomography. In [1] the authors solve the nonlinear tomography problem for a range of values of α and choose \mathbf{L} to represent the two-dimensional Laplacian operator. These results are from Algorithm 1, where a value of α that satisfies the discrepancy principle is ultimately the one that is chosen to give the best parameter estimate. A reproduction of those results is given on the left in Figure 7. Here we also found parameter estimates with $\mathbf{L} = \mathbf{I}$, which is a good choice for regularization if there is a good initial estimate \mathbf{x}_p , and those results are on the left in Figure 8. Given that the results with $\mathbf{L} = \mathbf{I}$ are worse than those with the Laplacian operator, we are left to conclude that the initial estimate of constant \mathbf{x}_p equal to 2900 m/s across the entire grid is not as effective as assuming that \mathbf{x} is smoothly varying.

In Figures 7 and 8 the Gauss–Newton estimate with the discrepancy principle is compared to the estimate from the iteratively regularized Gauss–Newton method with the χ^2 test in Theorem 2, as described in Algorithm 3. The iterations were stopped when the change in the parameter estimate was less than the tolerance. In Figure 7 we see that the estimate found with \mathbf{L} representing the Laplacian operator appears the same with both methods. However, in Figure 8 the estimate from Algorithm 3 with $\mathbf{L} = \mathbf{I}$ appears slightly more clear than that from Algorithm 1. It is evident from these figures that estimates found with the Laplacian operator are smoother than the




FIG. 5. Histograms of $\tilde{J}_k^{LM}(\mathbf{x})$ at each Levenberg–Marquardt iteration for the electrical conductivities problem. $C_f = \sigma_f^2 \mathbf{I}$. The sample mean is shown as the middle tick.

solutions found with $\mathbf{L} = \mathbf{I}$.

ALGORITHM 3 (iteratively regularized Gauss–Newton with χ^2 test).

Input $\mathbf{L}, \mathbf{C}_\epsilon, \mathbf{x}_p$
 $\mathbf{x}_1 = \mathbf{x}_p$
for $k = 1, 2, 3, \dots$ **do**
 Calculate \mathbf{J}_k and $\tilde{\mathbf{d}}_k$
 Define $\tilde{J}_k(\mathbf{x}, \alpha) = \|\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}\|_{\mathbf{C}_\epsilon^{-1}}^2 + \alpha^2 \|\mathbf{L}(\mathbf{x} - \mathbf{x}_p)\|_2^2$
 Choose α_k such that
 $\tilde{J}_k(\mathbf{x}_{k+1}, \alpha_k) \approx \Phi_{m-n+q}^{-1}(95\%)$, where Φ_{m-n+q}^{-1} is the inverse CDF of χ_{m-n+q}^2 .
 Calculate $\mathbf{x}_{k+1} = (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \mathbf{J}_k + \alpha_k^2 \mathbf{L}^T \mathbf{L})^{-1} (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \tilde{\mathbf{d}}_k + \alpha_k^2 \mathbf{L}^T \mathbf{L} \mathbf{x}_p)$
 if $|\mathbf{x}_{k+1} - \mathbf{x}_k| < tol$ **then**
 $\mathbf{x} \approx \mathbf{x}_{k+1}$
 return
 end if
end for

Figures 7 and 8 are the results of only one realization of ϵ . In order to establish a good comparison, the above procedure for both Algorithms 1 and 3 were repeated for 200 different realizations of ϵ . The mean and standard deviation of $\|\hat{\mathbf{x}} - \mathbf{x}_{true}\| / \|\mathbf{x}_{true}\|$ for the 200 trials for each method are given in Table 1. The χ^2 method gave better results on average when $\mathbf{L} = \mathbf{I}$, but the discrepancy principle did better on average with the second derivative operator. While these differences are only incremental, the χ^2 method was faster computationally because it only solves the inverse problem once and dynamically estimates α at each iteration. This implementation of the dis-




FIG. 6. Histograms of $\tilde{\mathcal{J}}_k^{LM}(\mathbf{x})$ at each Levenberg–Marquardt iteration for the electrical conductivities problem. $C_f = \sigma_f^2 \mathbf{L}'\mathbf{L}$. The sample mean is shown as the middle tick.




FIG. 7. Solutions found for the tomography problem with \mathbf{L} discrete approximation to the Laplacian. Left: Solution found using the discrepancy principle. Right: Solution found with Algorithm 1.

crepancy principle requires the inverse problem to be solved multiple times, incurring computational cost, as is evident by the outer loop in Algorithm 1.

3.3.2. Subsurface conductivity estimation. The subsurface electrical conductivities inverse problem is in some ways more difficult than the previous problem. The Gauss–Newton step does not always lead to a reduction in the nonlinear cost function, and it is not always possible to find a regularizing parameter for which the solution satisfies the discrepancy principle. In [1] the Levenberg–Marquardt algorithm was implemented to minimize the unregularized least squares problem to demonstrate the ill-posedness of this problem. This estimate, plotted in Figure 9, is wildly oscillating, has extreme values, and is not even close to being a physically possible solution. However, this is not evident from just looking at the data misfit, as this solution




FIG. 8. Solutions found for the tomography problem when $\mathbf{L} = \mathbf{I}$. Left: Solution found using the discrepancy principle. Right: Solution found with Algorithm 1.

TABLE 1
Comparison of discrepancy principle to the χ^2 method for the cross-well tomography problem.

Method	$\mathbf{C}_f^{-1} = \alpha^2 \mathbf{I}$	$\mathbf{C}_f^{-1} = \alpha^2 \mathbf{L}^T \mathbf{L}$
χ^2 method	$\mu = 0.01628$ $\sigma = 0.0006$	$\mu = 0.0206$ $\sigma = 0.00456$
Discrepancy principle	$\mu = 0.01672$ $\sigma = 0.00050$	$\mu = 0.018$ $\sigma = 0.0021$

actually fits the data quite well.

The inverse problem was also solved in [1] using Occam’s inversion as given in Algorithm 2, and this estimate is also shown in Figure 9. Occam’s inversion uses a discrete second derivative approximation for \mathbf{L} . We also implemented it with $\mathbf{L} = \mathbf{I}$, but the algorithm diverged.

Theorem 4 lends itself to a χ^2 method similar to Algorithm 3, but with $\tilde{\mathcal{J}}_k$ replaced with the regularized Levenberg–Marquardt functional $\tilde{\mathcal{J}}_k^{LM}$. The regularized Levenberg–Marquardt method with the regularization parameter found by the χ^2 test from Theorem 4 is given in Algorithm 4. The Levenberg–Marquardt parameter λ was initialized at 0.01 and then decreased by a factor of 10 at each iteration. However, if the functional was decreasing rapidly with that choice, λ decreased, while if it was not decreasing, another choice was found that decreased the functional. We applied this algorithm with both $\mathbf{L} = \mathbf{I}$ and a discrete second derivative operator for \mathbf{L} to find estimates for the subsurface conductivity problem. The resulting parameter estimates are plotted in Figure 10. Comparing the estimates found using the discrete second derivative for \mathbf{L} in Algorithms 2 and 4, in Figures 9 and 10, it is apparent that both estimate the true solution fairly well for this realization of ϵ . While the χ^2 method was still able to find a solution with $\mathbf{L} = \mathbf{I}$, it can be seen in Figure 10 that this choice does not estimate the true solution as well.

ALGORITHM 4 (regularized Levenberg–Marquardt with χ^2 test).

Input $\mathbf{L}, \mathbf{C}_\epsilon, \mathbf{x}_p, \lambda_1$
for $k = 1, 2, 3 \dots$ **do**
 Calculate \mathbf{J}_k and $\tilde{\mathbf{d}}_k$
 if $k > 1$ **then**




FIG. 9. Left: The solution found using a Levenberg–Marquardt algorithm. Right: The solution found with Occam’s inversion with $\mathbf{C}_f^{-1} = \alpha^2 \mathbf{L}^T \mathbf{L}$.

Define:

$$\begin{aligned} \tilde{\mathcal{J}}_k^{LM}(\mathbf{x}, \lambda) &= \|\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}\|_{\mathbf{C}_\epsilon^{-1}}^2 + \alpha_k^2 \|\mathbf{L}(\mathbf{x} - \mathbf{x}_p)\|_2^2 + \lambda^2 \|\mathbf{x} - \mathbf{x}_k\|_2^2, \\ \mathbf{x}_{k+1} &= (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \mathbf{J}_k + \alpha_k^2 \mathbf{L}^T \mathbf{L} + \lambda^2 \mathbf{I})^{-1} (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \tilde{\mathbf{d}}_k + \alpha_k^2 \mathbf{L}^T \mathbf{L} \mathbf{x}_p + \lambda_k^2 \mathbf{x}_k) \end{aligned}$$

Update LM parameter by finding a small λ_k that ensures

$$\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1}, \lambda_k) < \tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_k, \lambda_k)$$

end if

Define:

$$\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}, \alpha) = \|\tilde{\mathbf{d}}_k - \mathbf{J}_k \mathbf{x}\|_{\mathbf{C}_\epsilon^{-1}}^2 + \alpha^2 \|\mathbf{L}(\mathbf{x} - \mathbf{x}_p)\|_2^2 + \lambda_k^2 \|\mathbf{x} - \mathbf{x}_k\|_2^2$$

Choose α_{k+1} such that

$$\tilde{\mathcal{J}}_k^{LM}(\mathbf{x}_{k+1}, \alpha_{k+1}) \approx \Phi_{m-n+q}^{-1}(95\%) \text{ where } \Phi_{m-n+q}^{-1} \text{ is inverse CDF of } \chi_{m-n+q}^2.$$

Calculate:

$$\mathbf{x}_{k+1} = (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \mathbf{J}_k + \alpha_{k+1}^2 \mathbf{L}^T \mathbf{L} + \lambda_k^2 \mathbf{I})^{-1} (\mathbf{J}_k^T \mathbf{C}_\epsilon^{-1} \tilde{\mathbf{d}}_k + \alpha_{k+1}^2 \mathbf{L}^T \mathbf{L} \mathbf{x}_p + \lambda_k^2 \mathbf{x}_k)$$

if $|\mathbf{x}_{k+1} - \mathbf{x}_k| < \text{tol}$ then

$$\mathbf{x} \approx \mathbf{x}_{k+1}$$

return

end if

end for

Once again, to establish a good comparison, each of these methods was run for 200 different realizations of ϵ . The mean and standard deviation of $\|\hat{\mathbf{x}} - \mathbf{x}_{true}\| / \|\mathbf{x}_{true}\|$ for the 200 trials for each method are given in Table 2. While Occam’s inversion with $\mathbf{C}_f^{-1} = \alpha^2 \mathbf{L}^T \mathbf{L}$ was able to find good solutions for some realization of ϵ , such as the solution plotted in Figure 9, the standard deviation for this case given in Table 2 indicates that sometimes it found poor estimates. Occam’s inversion consistently diverged with the choice $\mathbf{C}_f^{-1} = \alpha^2 \mathbf{I}$ over multiple realizations of ϵ . These results show that the χ^2 method consistently gave much better answers than Occam’s inversion for this problem. Even the nonsmoothed χ^2 method with $\mathbf{L} = \mathbf{I}$ found better solutions on average than Occam’s with the smoothing choice $\mathbf{C}_f^{-1} = \alpha^2 \mathbf{L}^T \mathbf{L}$. The relatively small values for σ in Table 2 for the χ^2 method suggest that the solutions found were fairly consistent with each other.




FIG. 10. Left: The parameters found using the χ^2 method and $\mathbf{C}_f^{-1} = \alpha^2 \mathbf{L}^T \mathbf{L}$. Right: The parameters found using the χ^2 method and $\mathbf{C}_f^{-1} = \alpha^2 \mathbf{I}$.

TABLE 2

Comparison of the χ^2 method to Occam's inversion for the estimation of subsurface conductivities.

Method	$\mathbf{C}_f^{-1} = \alpha^2 \mathbf{I}$	$\mathbf{C}_f^{-1} = \alpha^2 \mathbf{L}^T \mathbf{L}$
χ^2 method	$\mu = 0.1827$ $\sigma = 0.0295$	$\mu = 0.0308$ $\sigma = 0.0281$
Occam's inversion	<i>diverges</i>	$\mu = 0.4376$ $\sigma = 0.6615$

4. Summary and conclusions. Weighted least squares with Tikhonov regularization is a popular approach to solving ill-posed inverse problems. For nonlinear problems, the resulting parameter estimates minimize a nonlinear functional, and Gauss–Newton or Levenberg–Marquardt algorithms are commonly used to find the minimizer. These nonlinear optimization methods iteratively solve a linear form of the objective function, and we have shown for both methods that the minimum value of the functionals at each iterate follows a χ^2 distribution. The minimum values of the Gauss–Newton and Levenberg–Marquardt functionals differ, but both have degrees of freedom equal to the number of data. We illustrate this χ^2 behavior for the Gauss–Newton method on a two-dimensional cross-well tomography problem, and for the Levenberg–Marquardt method we use a one-dimensional electromagnetic problem.

Since the minimum value of the functionals follows a χ^2 distribution at each iterate, this gives χ^2 tests from which a regularization parameter can be found for nonlinear problems. We give the resulting algorithms and test them by estimating parameters for the cross-well tomography and electromagnetic problems. It was shown that Algorithm 1 provided parameter estimates that were of accuracy similar to that of the discrepancy principle in a nonlinear cross-well tomography problem. In a subsurface electrical conductivity problem, the χ^2 method with the Levenberg–Marquardt algorithm proved to be more robust than Occam's inversion, providing parameter estimates without the use of a smoothing operator. This algorithm also provided much better estimates than Occam's inversion on average when the smoothing operator was used.

We conclude that the nonlinear χ^2 method is an attractive alternative to the discrepancy principle and Occam's inversion. However, it does share a disadvantage with these methods in that they all require the covariance of the data to be known.

If an estimate of the data covariance is not known, then the nonlinear χ^2 method will not be appropriate for solving such a problem. Future work includes estimating more complex covariance matrices for the parameter estimates. In [9] Mead shows that it is possible to use multiple χ^2 tests to estimate such a covariance, and it seems likely that this could also be extended to solving nonlinear problems.

REFERENCES

- [1] R. C. ASTER, B. BORCHERS, AND C. THURBER, *Parameter Estimation and Inverse Problems*, Academic Press, New York, 2005, p. 301.
- [2] P. C. CARLINE AND T. A. LOUIS, *Bayes and Empirical Bayes Methods for Data Analysis*, Chapman and Hall, London, 1996, p. 397.
- [3] H. W. ENGL, K. KUNISCH, AND A. NEUBAUER, *Convergence rates for Tikhonov regularisation of nonlinear ill-posed problems*, *Inverse Problems*, 5 (1989), pp. 523–540.
- [4] W. P. GOUVEIA AND J. A. SCALES, *Resolution of seismic waveform inversion: Bayes versus Occam*, *Inverse Problems*, 13 (1997), pp. 323–349.
- [5] P. C. HANSEN, *Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion*, SIAM Monogr. Math. Model. Comput. 4, SIAM, Philadelphia, 1998.
- [6] J. M. H. HENDRICKX, B. BORCHERS, D. L. CORWIN, S. M. LESCH, A. C. HILGENDORF, AND J. SCHULE, *Inversion of soil conductivity profiles from electromagnetic induction measurements: Theory and experimental verification*, *Soil Sci. Soc. Amer. J.*, 66 (2002), pp. 673–685.
- [7] J. MEAD, *Parameter estimation: A new approach to weighting a priori information*, *J. Inverse Ill-Posed Probl.*, 16 (2008), pp. 175–194.
- [8] J. MEAD AND R. A. RENAUT, *A Newton root-finding algorithm for estimating the regularization parameter for solving ill-conditioned least squares problems*, *Inverse Problems*, 25 (2009), 025002.
- [9] J. MEAD, *Discontinuous parameter estimates with least squares estimators*, *Appl. Math. Comput.*, 219 (2013), pp. 5210–5223.
- [10] A. RIEDER, *On the regularization of nonlinear ill-posed problems via inexact Newton iterations*, *Inverse Problems*, 15 (1999), pp. 309–327.
- [11] R. A. RENAUT, I. HNETYNKOVA, AND J. L. MEAD, *Regularization parameter estimation for large scale Tikhonov regularization using a priori information*, *Comput. Statist. Data Anal.*, 54 (2010), pp. 3430–3445.
- [12] T. I. SEIDMAN AND C. R. VOGEL, *Well-posedness and convergence of some regularisation methods for non-linear ill posed problems*, *Inverse Problems*, 5 (1989), pp. 227–238.
- [13] A. TARANTOLA, *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia, 2005.