Geophys. J. R. astr. Soc. (1972) **28**, 97–109

Interpretation of Inaccurate, Insufficient and Inconsistent Data

D. D. Jackson

(Received 1972 January 27)*

Summary

Many problems in physical science involve the estimation of a number of unknown parameters which bear a linear or quasi-linear relationship to a set of experimental data. The data may be contaminated by random errors, insufficient to determine the unknowns, redundant, or all of the above. This paper presents a method of optimizing the conclusions from such a data set. The problem is formulated as an ill-posed matrix equation, and general criteria are established for constructing an 'inverse' matrix. The 'solution' to the problem is defined in terms of a set of generalized eigenvectors of the matrix, and may be chosen to optimize the resolution provided by the data, the expected error in the solution, the fit to the data, the proximity of the solution to an arbitrary function, or any combination of the above. The classical 'least-squares' solution is discussed as a special case.

1. Formulation of the problem

Suppose that we wish to determine a set of unknown parameters x_j, $j = 1, \ldots, m$ from a set of data y_i, $i = 1, \ldots, n$ where y_i are each functionally related to the x_j in a known way. That is

$$\begin{aligned} y_1 &= A_1(x_1, \ldots, x_m) \\ &\vdots \\ y_n &= A_n(x_1, \ldots, x_m). \end{aligned}$$

Such a set of equations may arise by approximation of a continuous relationship $y(\eta) = A[\eta, x(\xi)]$ by a discrete representation, letting $y_i = y(\eta_i)$, $x_j = x(\xi_j)$; or by expansion of the continuous functions $y(\eta)$ and $x(\xi)$ in terms of appropriate sets of orthogonal functions, in which case y_i and x_j represent expansion coefficients (See Table 1 for notation conventions.)

If the functions $A_i(x_j)$ are linear in x_j, we may write the problem in matrix form

$$y_i = A_{ij} x_j \tag{1a}$$

or

$$\mathbf{y} = \mathbf{A}\mathbf{x}. \tag{1b}$$

* Received in original form 1971 October 29.

Table 1

Notation Conventions

1. Repeated indices imply summation unless otherwise noted.
2. A bold-face block capital letter represents matrix of coefficients (e.g. $\mathbf{A} = (A_{ij})$).
3. A bold-face lower case letter represents a column vector (e.g. $\mathbf{x} = (x_i)$).
4. var (x_k) represents the variance of x_k.

If the functions $A_i(x_j)$ are not strictly linear, but vary smoothly enough, they may be expanded in a Taylor series about some set of initial values of the x_j, say x_j^0

$$y_i = A_i(x_j^0) + \left.\frac{\partial A_i}{\partial x_j}\right|_{x_j^0} \Delta x_j + \dots . \tag{2}$$

Defining $y_i \equiv A_i(x_j^0) + \Delta y_i$ and ignoring second and higher order terms in (2), we have

$$\Delta y_i = \left.\frac{\partial A_i}{\partial x_j}\right|_{x_j^0} \Delta x_j .$$

This is the same form as equation (1a), with the substitution of Δy_i for y_i; Δx_j for x_j, and

$$\left.\frac{\partial A_i}{\partial x_j}\right|_{x_j^0} = A_{ij} .$$

For simplicity, we shall proceed using the notation of equation (1), with the understanding that the above substitution can be made at any stage of the calculations for a system which results from the perturbation of a quasi-linear problem.

Problem (1) may be approached by operating on both sides with an $(m \times n)$ 'inverse' matrix $\mathbf{H}$ and letting the 'solution', or *model*, be

$$\hat{\mathbf{x}} \equiv \mathbf{HAx} = \mathbf{Hy}. \tag{3}$$

The operator $\mathbf{H}$ will be a good inverse if it satisfies the following criteria.

(a) $\mathbf{AH} \approx \mathbf{I}_n$, the $n \times n$ identity matrix. This is a measure of how well the model fits the data, since $\mathbf{Ax} = \mathbf{y}$ if $\mathbf{AH} = \mathbf{I}_n$.

(b) $\mathbf{HA} \approx \mathbf{I}_m$. This is a measure of the uniqueness of the solution, since there may exist only one solution if $\mathbf{HA} = \mathbf{I}_m$.

(c) the uncertainties in $\hat{\mathbf{x}}$ are not too large, i.e. var $(\hat{\mathbf{x}})$ is small. For statistically independent data,

$$\text{var}\,(\hat{x}_k) = \sum_{i=1}^{n} H_{ki}^2 \,\text{var}\,(y_i). \tag{4}$$

Backus & Gilbert (1968) pointed out that, for under-determined systems, the product matrix $\mathbf{HA}$ has a physical meaning related to the establishment of 'unique' properties of the solutions to equation (1). That is,

$$\hat{\mathbf{x}} = \mathbf{Rx}, \tag{5}$$

where $\mathbf{R} \equiv \mathbf{HA}$.

The matrix $\mathbf{R}$ maps the entire set of solutions $\mathbf{x}$ into a single vector $\hat{\mathbf{x}}$. Any element of $\hat{\mathbf{x}}$, say $\hat{x}_k$, may be interpreted as the result of convolving the k^{th} row of $\mathbf{R}$ with any vector which satisfies (1). Thus, $\mathbf{R}$ is a matrix whose rows are 'windows' through which we may view the general solution $\mathbf{x}$ and obtain a unique result. If $\mathbf{R}$ is an

identity matrix, the solution $\hat{\mathbf{x}}$ is unique, and each element is perfectly resolved. If $\mathbf{R}$ is a near diagonal matrix, each element $\hat{x}_k$ is in fact a weighted sum of nearby values x_j, j near k, for *any* solution to (1). Thus, the degree to which $\mathbf{R}$ approximates the identity matrix is a measure of the resolution obtainable from the data. The rows of $\mathbf{R}$ are thus called 'resolving kernels'.

In a similar way, Wiggins (1972) showed that, for over-constrained systems, the product $\mathbf{S} \equiv \mathbf{AH}$ is a measure of the independence of the data. The 'theoretical' data $\hat{\mathbf{y}} \equiv \mathbf{A}\hat{\mathbf{x}} = \mathbf{AHy}$ are a convolution of the matrix $\mathbf{S}$ with the actual data $\mathbf{y}$. The matrix $\mathbf{S}$ is referred to as the information density matrix.

The criteria (a), (b) and (c) are not equally important in all cases, and may be weighted for specific problems. Additional physical constraints may apply in some cases. The procedure below allows these to be incorporated easily.

2. The well-posed case

Consider as an example the case where $\mathbf{A}$ is a square, symmetric, non-singular matrix. Elementary matrix theory (e.g. Hildebrand 1965) tells us that there exists a unique inverse $\mathbf{A}^{-1}$ such that $\mathbf{AA}^{-1} = \mathbf{A}^{-1}\mathbf{A} = \mathbf{I}_n$, the $n \times n$ identity matrix. Thus, criteria (a) and (b) are satisfied exactly and there is a unique solution, $\hat{\mathbf{x}} = \mathbf{A}^{-1}\mathbf{y}$, which satisfies the data exactly. The variance of the model will depend upon how non-singular the matrix $\mathbf{A}$ really is.

The nature of the singularity is best understood in terms of the eigenvalues of the matrix $\mathbf{A}$. The eigenvectors $\mathbf{v}_i$ and eigenvalues λ_i are defined by the equations

$$\mathbf{A}\mathbf{v}_i = \lambda_i \mathbf{v}_i \qquad \text{(no summation)}$$

where the $\mathbf{v}_i$ are orthonormal, i.e.

$$\mathbf{v}_i^T \mathbf{v}_j = \delta_{ij}.$$

The matrix $\mathbf{A}$ may be factored into the product

$$\mathbf{A} = \mathbf{V}\Lambda\mathbf{V}^T$$

where $\mathbf{V}$ is an $n \times n$ matrix called an 'orthonormal modal matrix', whose columns are the eigenvectors $\mathbf{v}_i$, and Λ is an $n \times n$ diagonal matrix whose elements are the eigenvalues λ_i. It may be assumed without loss of generality that the eigenvalues are written in decreasing order of their absolute value, and that the eigenvectors are ordered correspondingly in the matrix $\mathbf{V}$. The matrix $\mathbf{V}$ has the convenient properties

$$\mathbf{V}^T\mathbf{V} = \mathbf{V}\mathbf{V}^T = \mathbf{I}_n.$$

The inverse matrix A^{-1} may then be written

$$\mathbf{A}^{-1} = \mathbf{V}\Lambda^{-1}\mathbf{V}^T \tag{6}$$

where Λ^{-1} is a diagonal matrix whose elements are λ_i^{-1}. Thus, the inverse matrix will cease to exist if any of the eigenvalues λ_i is zero. If λ_i is non-zero but very small, the operator A^{-1} exists but does not satisfy criterion (c); that is, for statistically independent data,

$$\operatorname{var} \hat{x}_k = \sum_{i=1}^{n} \left\{ \sum_{j=1}^{n} V_{kj} \lambda_j^{-1} V_{ij} \right\}^2 \operatorname{var}(y_i)$$

will be very large because of the reciprocal eigenvalue in the bracket.

3. Generalized eigenvector analysis

Eigenvector analysis can be extended to the general real $n \times m$ case as discussed by Lanczos (1961). Two sets of eigenvectors $\mathbf{u}_i$ and $\mathbf{v}_j$ may be found such that

$$\mathbf{A}\mathbf{v}_j = \lambda_j \mathbf{u}_j \quad \text{(no summation)} \tag{7a}$$

$$\mathbf{A}^T \mathbf{u}_i = \lambda_i \mathbf{v}_i \quad \text{(no summation)} \tag{7b}$$

or

$$\mathbf{A}^T \mathbf{A}\mathbf{v}_j = \lambda_j^2 \mathbf{v}_j \quad j = 1, \ldots, m \quad \text{(no summation)} \tag{8a}$$

$$\mathbf{A}\mathbf{A}^T \mathbf{u}_i = \lambda_i^2 \mathbf{u}_i \quad i = 1, \ldots, n \quad \text{(no summation)} \tag{8b}$$

We shall again assume that the eigenvalues are ranked in decreasing order of magnitude It can then be shown that

$$\lambda_i = \lambda_j \quad \text{if} \quad i = j, \quad i \leqslant p$$

$$\lambda_i = 0 \qquad i > p$$

$$\lambda_j = 0 \qquad j > p$$

for some integer p less than or equal to the minimum of m and n. That is, there are p non-zero eigenvalues common to (8a) and (8b), and all other eigenvalues are zero. The integer p may be interpreted as the potential number of degrees of freedom in the data.

The matrix $\mathbf{A}$ can be factored into the product

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^T \tag{9}$$

where $\mathbf{U}$ is an $n \times p$ matrix whose columns are the eigenvectors $\mathbf{u}_i, i = 1, \ldots, p$; $\mathbf{V}$ is the $m \times p$ matrix whose columns are the eigenvectors $\mathbf{v}_i$, $i = 1, \ldots, p$ and $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues. After $\mathbf{U}$ and $\mathbf{V}$ are formed from the eigenvectors corresponding to the p non-zero eigenvalues, there remain $(n-p)$ eigenvectors $\mathbf{u}_i$, and $(m-p)$ eigenvectors $\mathbf{v}_j$, corresponding to zero eigenvalues. It is convenient to assemble these into the columns of matrices also, and to denote these matrices as $\mathbf{U}_0$ [an n by $(n-p)$ matrix] and $\mathbf{V}_0$ [an m by $(m-p)$ matrix]. By the orthonormality of the eigenvectors $\mathbf{u}_i$, we may obtain the following relationships

$$\mathbf{U}^T \mathbf{U} = \mathbf{I}_p \tag{10a}$$

$$\mathbf{U}_0{}^T \mathbf{U}_0 = \mathbf{I}_{n-p} \tag{10b}$$

$$\mathbf{U}^T \mathbf{U}_0 = \mathbf{0} \qquad p \times (n-p) \tag{10c}$$

$$\mathbf{U}_0{}^T \mathbf{U} = \mathbf{0} \qquad (n-p) \times p\,. \tag{10d}$$

If $p = n$, then $\mathbf{U}$ contains the complete set of eigenvectors of the symmetric matrix $[\mathbf{A}\mathbf{A}^T]$, and is therefore an orthonormal modal matrix. It then satisfies

$$\mathbf{U}\mathbf{U}^T = \mathbf{I}_n \tag{10}$$

I shall demonstrate below that the condition $p = n$ guarantees the existence of a solution.

The eigenvectors $\mathbf{v}_i$, $i = 1, \ldots, m$ are similarly a complete set of orthonormal vectors in the 'model' space. We have separated them into two groups: those to

which the model is sensitive (i.e. those in V), and those which are annihilated by A (i.e. those in $\mathbf{V}_0$). The matrices V and $\mathbf{V}_0$ obey the same type of relationship as U and $\mathbf{U}_0$

$$\mathbf{V}^T\mathbf{V} = \mathbf{I}_p \tag{11a}$$

$$\mathbf{V}_0{}^T\mathbf{V}_0 = \mathbf{I}_{m-p} \tag{11b}$$

$$\mathbf{V}^T\mathbf{V}_0 = \mathbf{0} \qquad px(m-p) \tag{11c}$$

$$\mathbf{V}_0{}^T\mathbf{V} = \mathbf{0} \qquad (m-p)\times p. \tag{11d}$$

If $p = m$, then V contains the complete set of eigenvectors of the symmetric matrix $[\mathbf{A}^T\mathbf{A}]$, and is therefore an orthonormal modal matrix. It then satisfies

$$\mathbf{V}\mathbf{V}^T = \mathbf{I}_m. \tag{11}$$

The condition $p = m$ guarantees that, if a solution to equation (1) exists, it will be unique.

Before proceeding, it may be worthwhile to give a simple interpretation to equation (9). In an equation such as $\mathbf{B}\,\mathbf{f} = \mathbf{g}$ the column vector $\mathbf{g}$ may be looked upon as a weighted sum of the columns of $\mathbf{B}$ with the weighting factors given by the elements of the column vector $\mathbf{f}$. The matrix product $\mathbf{BF} = \mathbf{G}$ may be considered column by column: the first column of the matrix $\mathbf{G}$ is a weighted sum of the columns of $\mathbf{B}$, with weighting factors $\mathbf{F}_{11}$, $\mathbf{F}_{21}$, etc., and similarly for the rest of the columns. A similar argument holds for row vectors: in the equation $\mathbf{f}^T\mathbf{B}^T = \mathbf{g}^T$, the row vector $\mathbf{g}^T$ is a sum of the rows of $\mathbf{B}^T$ with the same weighting factors as above. In the product $\mathbf{F}^T\mathbf{B}^T = \mathbf{G}^T$, each row of $\mathbf{G}^T$ is a weighted sum of the rows of $\mathbf{B}^T$, with the weighting factors given in the appropriate rows of $\mathbf{F}^T$. Thus, in a long string of matrix multiplications, each column of the product matrix is a weighted sum of the columns of the *first* matrix in the string, and each row of the product is a sum of the rows of the *last* column in the string. Thus, from equation (9), each column of the matrix A is a weighted sum of the columns of U; that is, the eigenvectors $\mathbf{u}_i$, $i = 1, \ldots, p$. Each row of A is a sum of the eigenvectors $\mathbf{v}_i$, $i = 1, \ldots, p$

Because the n eigenvectors $\mathbf{u}_i$, $i = 1, \ldots, n$ form a complete set, we may express the data vector y as a sum of the $\mathbf{u}_i$; that is,

$$\mathbf{y} = \sum_{i=1}^{n} \beta_i \mathbf{u}_i = \sum_{i=1}^{p} \beta_i \mathbf{u}_i + \sum_{i=p+1}^{n} \beta_i \mathbf{u}_i = \mathbf{U}\boldsymbol{\beta} + \mathbf{U}_0\,\boldsymbol{\beta}_0 \tag{12}$$

where the vector $\boldsymbol{\beta}$ has the p components $\beta_1, \ldots, \beta_p$, and $\boldsymbol{\beta}_0$ has the $(n-p)$ components $\beta_{p+1}, \ldots, \beta_n$. By (10a) and (10b), $\boldsymbol{\beta} = \mathbf{U}^T\mathbf{y}$, and $\boldsymbol{\beta}_0 = \mathbf{U}_0{}^T\mathbf{y}$. In a similar way, we may express the 'unknown' vector $\boldsymbol{\varepsilon}$ in terms of the $\mathbf{v}_i$, $i = 1, \ldots, m$.

$$\mathbf{x} = \mathbf{V}\boldsymbol{\alpha} + \mathbf{V}_0\,\boldsymbol{\alpha}_0 \tag{13}$$

where $\boldsymbol{\alpha} = \mathbf{V}^T\mathbf{x}$ is a vector with p components and $\boldsymbol{\alpha}_0 = \mathbf{V}_0{}^T\mathbf{x}$ is a vector with $m-p$ components.

If we then replace all the quantities in equation (1) with the equivalent expressions in (9), (12) and (13), the 'transformed' equation becomes

$$\mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^T[\mathbf{V}\boldsymbol{\alpha} + \mathbf{V}_0\boldsymbol{\alpha}_0] = \mathbf{U}\boldsymbol{\beta} + \mathbf{U}_0\,\boldsymbol{\beta}_0 \tag{14}$$

and the problem is now to find the unknown vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}_0$. The existence of a solution is equivalent to the vanishing of the 'residual' vector, $\boldsymbol{\varepsilon} \equiv \mathbf{A}\mathbf{x} - \mathbf{y}$. In the transformed notation

$$|\boldsymbol{\varepsilon}|^2 = |\boldsymbol{\varepsilon}^T\boldsymbol{\varepsilon}| = |\boldsymbol{\Lambda}\boldsymbol{\alpha} - \boldsymbol{\beta}|^2 + |\boldsymbol{\beta}_0|^2 \quad . \tag{15}$$

The 'least squares' solution is that which minimizes $|\boldsymbol{\varepsilon}|^2$, and requires that

$$\boldsymbol{\alpha} = \boldsymbol{\Lambda}^{-1}\,\boldsymbol{\beta} \tag{16}$$

leaving the 'least square error' $|\boldsymbol{\beta}_0|^2 = |\mathbf{U}_0{}^T \mathbf{y}|^2$. There will exist an exact solution only if $\mathbf{U}_0{}^T \mathbf{y} = \mathbf{0}$. For $p = n$, this condition must hold, since $\mathbf{U}_0 = \mathbf{0}$. For $p < n$, there may be an exact solution only if the data contain no contribution from the eigenvectors in $\mathbf{U}_0$; this leads to the $n-p$ constraints, $\mathbf{u}_T . \mathbf{y} = \mathbf{0}$, $i = p+1, \ldots, n$. When $p < n$, the system is said to be overconstrained.

The vector $\boldsymbol{\alpha}_0$ does not appear in (15), and thus the α_i may be chosen arbitrarily for $i = p+1, \ldots, m$. The least squares solution will be unique only if $p = m$. When $p < m$, the system (1) is said to be underdetermined.

Notice that, if $p < m$ *and* $p < n$, the system (1) is both overconstrained *and* underdetermined. In this case, an exact solution may not exist. However, there will exist an infinite number of solutions satisfying the least-squares criterion.

4. The Lanczos inverse

To handle linear systems with the problems discussed above, Lanczos (1961) introduced the 'natural' inverse, which is equivalent to the 'generalized inverse' of Penrose (1955),

$$\mathbf{H}_L = \mathbf{V}\boldsymbol{\Lambda}^{-1}\,\mathbf{U}^T. \tag{17}$$

This has a form reminiscent of equation (6). The corresponding model will be

$$\hat{\mathbf{x}}_L = \mathbf{H}_L\,\mathbf{y} = \mathbf{V}\boldsymbol{\Lambda}^{-1}\,\mathbf{U}^T\mathbf{y}. \tag{18}$$

Expressing $\mathbf{x}_L$ and $\mathbf{y}$ in terms of their 'transform' variables as above,

$$\boldsymbol{\alpha} = \mathbf{V}^T\hat{\mathbf{x}}_L = \boldsymbol{\Lambda}^{-1}\,\boldsymbol{\beta} \tag{19a}$$

and

$$\boldsymbol{\alpha}_0 = \mathbf{V}_0{}^T\hat{\mathbf{x}}_L = \mathbf{0}. \tag{19b}$$

The Lanczos inverse has the following desirable properties. (a) It always exists! (b) Comparing (19a) with (16), it is evident that the Lanczos inverse is a 'least squares' inverse, and is thus an exact solution, if any exists. (c) Further, $\hat{\mathbf{x}}_L$ is that least squares solution which minimizes $|\mathbf{x}|^2$ since, for a least squares solution,

$$|\mathbf{x}|^2 = |\boldsymbol{\Lambda}^{-1}\,\boldsymbol{\beta}|^2 + |\boldsymbol{\alpha}_0|^2.$$

The first term is fixed, and the second term is minimized by $\boldsymbol{\alpha}_0 = \mathbf{0}$. The property $|\mathbf{x}|^2 = \min$ is a useful one if $\mathbf{x}$ represents a perturbation to some starting model in a quasi-linear problem. (d) The resolution matrix for the Lanczos inverse is given by

$$\mathbf{R} = \mathbf{HA} = \mathbf{V}\boldsymbol{\Lambda}^{-1}\,\mathbf{U}^T.\mathbf{U}\boldsymbol{\Lambda}\mathbf{V}^T = \mathbf{V}\mathbf{V}^T. \tag{20}$$

This is the optimum resolving matrix in the sense that it minimizes

$$r_k = \sum_{j=1}^{m} (H_{ki}\,A_{ij} - \delta_{kj})^2 \tag{21}$$

for each value of k. Equation (21) may be interpreted thus: each row of $\mathbf{R}$ is the best fit to the corresponding row of the identity matrix, in the least squares sense, which may be formed from the rows of A, and (e) The Lanczos inverse similarity provides the best information density matrix

$$\mathbf{S} = \mathbf{AH} = \mathbf{U}\mathbf{U}^T \tag{22}$$

in the sense that it minimizes

$$s_k = \sum_{j=1}^{n} (A_{ki}\,H_{ij} - \delta_{kj})^2 \tag{23}$$

I have claimed above that the Lanczos inverse is a least squares inverse, and thus for a purely overconstrained system, it must be identical to the inverse provided

by the standard least squares procedure. This procedure may be stated as follows: find the vector $\mathbf{x}$ which minimizes $|\boldsymbol{\varepsilon}|^2 = |\mathbf{A}\mathbf{x}-\mathbf{y}|^2$. We have

$$|\boldsymbol{\varepsilon}|^2 = (\mathbf{A}\mathbf{x}-\mathbf{y})^T(\mathbf{A}\mathbf{x}-\mathbf{y}) = \mathbf{x}^T\mathbf{A}^T\mathbf{A}\mathbf{x}-\mathbf{x}^T\mathbf{A}^T\mathbf{y}-\mathbf{y}^T\mathbf{A}\mathbf{x}+\mathbf{y}^T\mathbf{y}.$$

Differentiating with respect to $\mathbf{x}^T$, and, setting the result equal to 0, we get

$$[\mathbf{A}^T\mathbf{A}]\,\mathbf{x} = \mathbf{A}^T\mathbf{y}. \tag{24}$$

If the matrix product in brackets is non-singular,

$$\hat{\mathbf{x}} = [\mathbf{A}^T\mathbf{A}]^{-1}\mathbf{A}^T\mathbf{y}. \tag{25}$$

Expressing $\mathbf{A}$ in terms of equation (9),

$$\hat{\mathbf{x}} = [\mathbf{V}\boldsymbol{\Lambda}^2\mathbf{V}^T]^{-1}\mathbf{V}\boldsymbol{\Lambda}\mathbf{U}^T\mathbf{y}. \tag{26}$$

As long as $p = m$, (that is, the system is not underdetermined) the inverse of the bracketed quantity is

$$\mathbf{V}\boldsymbol{\Lambda}^{-2}\mathbf{V}^T,$$

and we have

$$\hat{\mathbf{x}} = \mathbf{V}\boldsymbol{\Lambda}^{-1}\mathbf{U}^T\mathbf{y} = \mathbf{H}_L\mathbf{y}$$

as advertized.

It is interesting to observe what happens when $p < m$, that is the system is underdetermined as well as overconstrained. Then $\mathbf{A}^T\mathbf{A}$ will be singular, and the standard least squares procedure will fail. However, direct substitution will demonstrate that $\hat{\mathbf{x}} = \mathbf{H}_L\mathbf{y}$ satisfies equation (24) regardless of the singularity of $\mathbf{A}^T\mathbf{A}$. Thus, the standard least squares procedure is a short-cut to the Lanczos inverse for the strictly overconstrained case, but will not guarantee stability for the underdetermined case.

An analogous procedure to 'least squares' applies to the strictly underdetermined case (Smith & Franklin 1969). Here

$$\hat{\mathbf{x}} = \mathbf{A}^T[\mathbf{A}\mathbf{A}^T]^{-1}\mathbf{y} \tag{27}$$

which again is identical to the Lanczos inverse, provided that $\mathbf{A}\mathbf{A}^T$ is non-singular. This requires that $p = n$, that is, that the system not be overconstrained.

5. Modifications of the Lanczos inverse

Until this point, we have assumed that the data were dimensionless and statistically independent, and we have ignored criterion (c), the variance of the model. In this Section, I shall discuss modifications of the Lanczos inversion procedure which address these considerations.

First consider an overconstrained system for which the data are statistically independent, but have different units. One might then wish to generalize the least squares error criterion to minimize

$$\boldsymbol{\varepsilon}^T\mathbf{D}\boldsymbol{\varepsilon}$$

where $\mathbf{D}$ is a diagonal matrix whose elements are

$$D_i = 1/\mathrm{var}\,(y_i).$$

In this way, the residual for each data point is compared with its expected error. When the data are not statistically independent, it is logical to choose for D the inverse of the covariance matrix for the data (Kaula 1966). The problem may be handled with the same formulation used in the last section, provided that $\mathbf{D}$ is a symmetric, positive definite matrix (a reasonable assumption for an inverse covariance matrix). I shall show below that there will exist a matrix $\mathbf{E}$ such that $\mathbf{E}^T\mathbf{E} = \mathbf{D}$. We have already solved the problem $\boldsymbol{\varepsilon}^T\boldsymbol{\varepsilon} = \min$, where $\boldsymbol{\varepsilon} = \mathbf{A}\hat{\mathbf{x}}-\mathbf{y}$. Now consider the equation

$$\mathbf{E}\boldsymbol{\varepsilon} = \mathbf{E}\mathbf{A}\hat{\mathbf{x}}-\mathbf{E}\mathbf{y} \tag{28}$$

and let $\mathbf{E}\boldsymbol{\varepsilon} = \boldsymbol{\varepsilon}'$, $\mathbf{EA} = \mathbf{A}'$, and $\mathbf{Ey} = \mathbf{y}'$. We now have a problem of the form of equation (1), and we may use the standard least squares procedure or the Lanczos procedure to find $\hat{\mathbf{x}}$ which minimizes

$$\boldsymbol{\varepsilon}'^{T}\boldsymbol{\varepsilon}' = \boldsymbol{\varepsilon}^{T}\mathbf{D}\boldsymbol{\varepsilon}.$$

Thus, the problem is solved once $\mathbf{E}$ is found. This may always be accomplished, because the symmetric, positive definite matrix $\mathbf{D}$ may be decomposed into its eigenvectors and eigenvalues (which will be positive)

$$\mathbf{D} = \mathbf{V}_D\boldsymbol{\Lambda}_D\mathbf{V}_D^{T} \tag{29}$$

and we may set

$$\mathbf{E} = \boldsymbol{\Lambda}_D^{\frac{1}{2}}\mathbf{V}_D^{T} \tag{30}$$

This weighting matrix E puts the data in a dimensionless, statistically independent form. Where the data were already statistically independent, the weighting by $\mathbf{E}$ is equivalent to dividing each equation in the system (1) by the standard deviation of the corresponding data point.

A similar transformation allows the use of the Lancsoz procedure to find that solution of an underdetermined problem which minimizes

$$\mathbf{x}^{T}\mathbf{F}\mathbf{x}$$

where again $\mathbf{F}$ is assumed to be a symmetric positive definite matrix. Where we are able to guess something about the statistics of the unknown parameters (cf. Jordan & Franklin 1971), the matrix $\mathbf{F}$ may be chosen to be the inverse covariance matrix of the model parameters. We then let $\mathbf{x}' = \mathbf{G}\mathbf{x}$ and $\mathbf{A}' = \mathbf{A}\mathbf{G}^{-1}$, where $\mathbf{G}$ is selected such that $\mathbf{G}^{T}\mathbf{G} = \mathbf{F}$. We may then use the Lanczos procedure to find that $\hat{\mathbf{x}}'$ which minimizes $|\mathbf{x}'|^2$, and our desired solution is $\hat{\mathbf{x}} = \mathbf{G}^{-1}\hat{\mathbf{x}}'$.

The above transformations may be combined for the general solution, to find that $\hat{\mathbf{x}}$ which simultaneously minimizes $\boldsymbol{\varepsilon}^{T}\mathbf{D}\boldsymbol{\varepsilon}$ and $\mathbf{x}^{T}\mathbf{F}\mathbf{x}$. This is accomplished by letting $\hat{\mathbf{x}} = \mathbf{G}^{-1}\mathbf{H}_L'\mathbf{y}'$ where $\mathbf{H}_L'$ is the Lanczos inverse for the matrix $\mathbf{A}' = \mathbf{E}\mathbf{A}\mathbf{G}^{-1}$, and $\mathbf{y}' = \mathbf{E}\mathbf{y}$. The effect of these transformations on the resolution and information density matrices is expressed by

$$\mathbf{R} = \mathbf{G}^{-1}\mathbf{R}\,\mathbf{G} \tag{31}$$

$$\mathbf{S} = \mathbf{E}^{-1}\mathbf{S}'\mathbf{F}. \tag{32}$$

Let us now address the problem of the variance of the model parameters. I have shown that we may assume the data to be statistically independent and to have unit variance, if the appropriate transformations are made in advance. By equation (4), we shall then have

$$\operatorname{var}\hat{x}_k = \sum_{j=1}^{p}\left(\frac{V_{kj}}{\lambda_j}\right)^2. \tag{33}$$

This will be finite, but because the smallest non-zero eigenvalue may be pretty small, the variance could be unacceptably large. A sensible way to control the variance (Wiggins 1971) is to construct the inverse $\mathbf{H}$ out of only those eigenvectors corresponding to the q largest eigenvalues, where $q \leqslant p$. This is equivalent to considering an eigenvalue to be zero if it is less than some modest threshold (i.e. assuming p less than it really is). The threshold may be set such that

$$\sum_{j=1}^{q}\left(\frac{V_{kj}}{\lambda_j}\right)^2 < t_k \tag{34}$$

for all k, where t_k is the maximum allowable variance of $\hat{x}_k$. The effect of reducing q is to reduce the number of eigenvectors belonging to $\mathbf{U}$ and $\mathbf{V}$, while increasing by the same amount those belonging to $\mathbf{U}_0$ and $\mathbf{V}_0$. This degrades the resolution and information density, so that we have a tradeoff as indicated symbolically in Fig. 1.

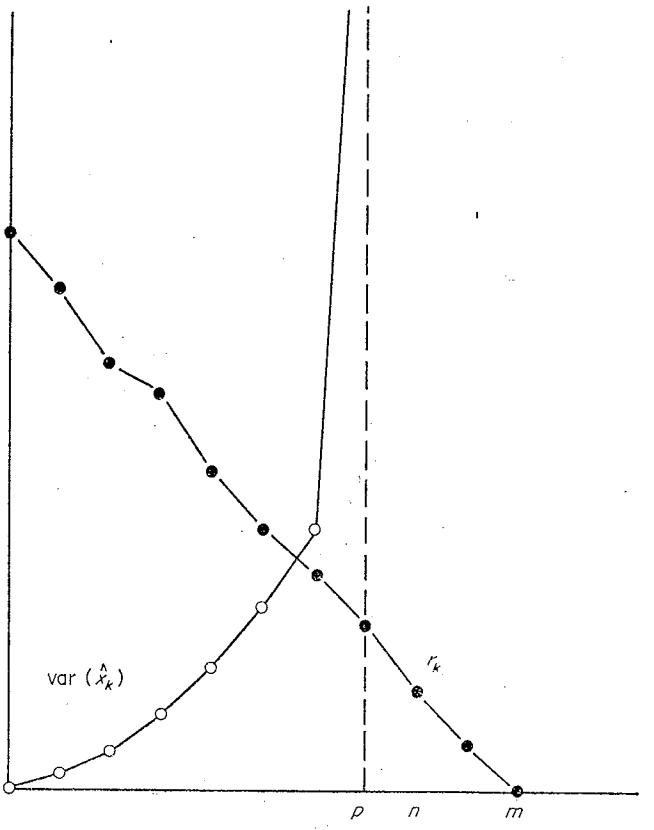

FIG. 1 Schematic diagram showing tradeoff between resolution and variance in the estimation of the k^{th} unknown parameter for a simultaneously underdetermined and overconstrained problem. The effective number of degrees of freedom, q, may be taken to be any integer less than or equal to p. Use of the ' generalized inverse ' procedure, which requires that $p = n$, would lead to numerical instabilities for this case.

The integer q is the effective number of degrees of freedom in the data, and depends on the uncertainties in the data as well as on our need for certainty in the model. A great deal of care should be exercised in choosing the appropriate value of q. For systems which are fundamentally underdetermined, the ability to make any reliable interpretation from the model $\hat{\mathbf{x}}$ may be limited either by lack of resolution, or by large variance in $\hat{\mathbf{x}}$. Surprisingly, the most important conclusions may often be made on the basis of a model which fits the data very poorly, even though exact solutions exist. This is because the exact solutions are not unique, and because the inverse operators which generate them may rely heavily upon poorly determined features of the data.

Similar arguments hold for systems which are basically overconstrained. If care is not exercised in choosing q, the model may be subject to unnecessarily large variance in an effort to satisfy poorly determined features of the data.

6. Special inverses for underdetermined systems

Let us examine in more detail the case of the strongly underdetermined system. This case will include those problems in which equation (1) is the result of discretizing a continuous relationship between a known function and an unknown function, because we may only handle finite amounts of data, yet we would in principle like to know an infinitude of details about the unknown function. A wise procedure is to use more parameters to describe the unknowns than are likely to be uniquely determined

by the data. One may then form a family of inverses, compare the tradeoff between resolution and variance for this family, and select that inverse which is most appropriate for interpreting the solution. This procedure simultaneously formulates intelligent questions (what is the effective number of degrees of freedom in the data, and which unknowns may be independently determined), while providing an answer with an acceptable variance. In this section, I shall describe some alternate techniques for optimizing the resolving kernels; describe a technqiue for finding the closest solution to some preassigned function, and discuss the wisdom of taking more data to obtain better resolution.

I have stated above that the Lanczos inverse provides the closest resolution matrix available, in a least squares sense, to the identity matrix. It is instructive to show that, in addition, it provides the most deltalike matrix which may be constructed from the incomplete set of p eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_p$. That is, assume the resolution matrix to be of the form

$$\mathbf{R} = \mathbf{B}\mathbf{V}^T. \tag{35}$$

Each row of $\mathbf{R}$ will be some linear combination of the eigenvectors of $\mathbf{V}$. Let us denote by $\mathbf{b}_k{}^T$ the k^{th} row of $\mathbf{B}$, and by $\boldsymbol{\delta}_k{}^T$ the k^{th} row of the identity matrix (i.e. the j^{th} element of this vector is zero for $j \neq k$, 1 for $j = k$). We now determine $\mathbf{b}_k{}^T$ as the least squares solution to

$$\mathbf{b}_k{}^T \mathbf{V}^T = \hat{\boldsymbol{\delta}}_k{}^T \tag{36}$$

or, transposing,

$$\mathbf{V}\mathbf{b}_k = \boldsymbol{\delta}_k. \tag{37}$$

By (25), the least squares solution is

$$\hat{\boldsymbol{\delta}}_k = [\mathbf{V}^T \mathbf{V}]^{-1} \mathbf{V}^T \boldsymbol{\delta}_k = \mathbf{V}^T \boldsymbol{\delta}_k \tag{38}$$

and, transposing back, we have

$$\hat{\boldsymbol{\delta}}_k{}^T = \boldsymbol{\delta}_k{}^T \mathbf{V}. \tag{39}$$

Performing this operation for every row, we have

$$\mathbf{B} = \mathbf{V} \tag{40}$$

and the most deltalike resolution matrix is

$$\mathbf{R} = \mathbf{V}\mathbf{V}^T \tag{41}$$

identical to the Lanczos inverse.

The modified Lanczos inverse discussed in the last section controls the variance of the model by using only q of the possible p eigenvectors in the inverse. As phrased above, the same q eigenvectors are used in constructing each row of $\mathbf{H}$. However, it is useful to note in equation (4) that the variance of the k^{th} model parameter depends only on the parameters of the k^{th} row of $\mathbf{H}$. We may construct an inverse row by row, getting a different tradeoff curve of resolution *vs* variance for each row. That is

$$r_k = \sum_{i=1}^{n} \left(\sum_{j=1}^{q} V_{kj} V_{ij} - \delta_{ki} \right)^2 \tag{42}$$

$$\text{var}(\hat{x}_k) = \sum_{j=1}^{q} \left(\frac{V_{kj}}{\lambda_j} \right)^2 \tag{43}$$

where the integer q may now be a function of k. This allows us to examine each model parameter $\hat{x}_k$ individually. In so doing, we implicitly ignore criterion (a), and the fit of the model $\hat{\mathbf{x}}$ to the data. If we form a model by independently estimating each value $\hat{x}_k$ by this scheme, it is very likely that the model will not fit the data. This is

not a serious concern, if indeed the separate values $\hat{x}_k$ have an intrinsic meaning which is useful in the eventual interpretation of this problem. On the other hand, if a true solution (1) is desired, one may increase the values of $q(k)$ at the expense of increased variance for the values $\hat{x}_k$.

There are some cases where it may be desirable to have a particular resolution kernel resemble some function other than the delta-function. For example, attempting to create a delta-like kernel may result in having a sharply peaked kernel with positive and negative 'side lobes', which could be avoided by letting the kernel have a somewhat broader peak. Or, it may be desirable to form an averaging kernel with a 'dipole' behaviour, that is, a sharp negative peak adjacent to a sharp positive peak, to obtain an estimate of the derivatives of the general solution $\mathbf{x}$. Such a problem may be very easily handled by substituting the desired optimum kernel function for $\boldsymbol{\delta}_k$ in equation (39). The corresponding inverse matrix may be constructed by setting

$$\mathbf{H} = \mathbf{B}\boldsymbol{\Lambda}^{-1}\,\mathbf{U}^T. \tag{44}$$

In other cases, it may be desirable to have the resolution kernel resemble a delta function, but proximity to a delta function may be more important for some elements than for others. For example, suppose one wants to estimate $\hat{x}_k$, and has evidence that x_j is very large for some particular value $j \neq k$. Then one would want to assure that the element R_{kj} is very small, at the expense of allowing other off-diagonal elements of $\mathbf{R}$ to be somewhat larger. This may be accomplished by applying a weighting factor matrix to both sides of (25). If the weighting matrix is $\mathbf{W}$, and we let $\mathbf{W}^T\mathbf{W} = \mathbf{W}^2$, then

$$\mathbf{B} = [\mathbf{V}^T\mathbf{W}^2\,\mathbf{V}]^{-1}\,\mathbf{V}^T\mathbf{W}^2. \tag{45}$$

Thus, the weighted resolution kernel, and its corresponding inverse $\mathbf{H}$, may be easily computed from the eigenvectors of the original matrix $\mathbf{A}$, avoiding the need to compute eigenvectors of a new, weighted matrix.

To summarize the above discussion, we may use the fact that each row of the resolution matrix must be composed of the eigenvectors $\mathbf{v}_i$, $i = 1, \ldots, p$, and the standard least squares procedure (equation 25)) to find that resolution kernel nearest to some desired function. The same trick may be used to find that solution of equation (1) which is nearest to some favourite function. The favourite function might be a hypothetical solution to (1), which would have special importance for interpretation. We know that the form of the general solution is

$$\mathbf{x} = \mathbf{V}\boldsymbol{\alpha} + \mathbf{V}_0\,\boldsymbol{\alpha}_0 \tag{46}$$

where $\boldsymbol{\alpha}$ is determined by equation (1), while $\boldsymbol{\alpha}_0$ is arbitrary. We then seek the least squares solution to

$$\mathbf{V}\boldsymbol{\alpha} + \mathbf{V}_0\,\boldsymbol{\alpha}_0 = \mathbf{f} \tag{47}$$

where $\mathbf{f}$ is the favourite function. By equation (25), we have

$$\boldsymbol{\alpha}_0 = \mathbf{V}_0{}^T[\mathbf{f} - \mathbf{V}\boldsymbol{\alpha}] = \mathbf{V}_0{}^T\mathbf{f} \tag{48}$$

so that the solution which best approximates $\mathbf{f}$ is

$$\mathbf{x}_f = \mathbf{V}\boldsymbol{\alpha} + \mathbf{V}_0\,\mathbf{V}_0{}^T\mathbf{f} = \hat{\mathbf{x}}_L + \mathbf{V}_0\,\mathbf{V}_0{}^T\mathbf{f}. \tag{49}$$

Let us now consider the problem of, marginal utility of data'. Suppose we have analysed a large linear system according to the above procedures, and have found that the resolution provided by the data is insufficient. We wish to consider the improvement which would result from adding an $(n+1)$ *st* measurement. Generally, the values of $A_{n+1,\,1}, \ldots, A_{n+1,\,m}$ may be calculated in advance, and the variance of the $(n+1)$ *st* measurement may be at least estimated in advance. We could, of course, add a new row to the matrix $\mathbf{A}$ and start from scratch, finding a new set of eigenvalues

and eigenvectors, etc. However, a first order perturbation to the existing eigenvalues will probably be sufficient to answer the question at hand. Let us denote by $\mathbf{a}$ the $m \times 1$ column vector whose values are $A_{n+1,1}, \ldots, A_{n+1,m}$: then the new $(n+1)\,st$ row of $\mathbf{A}$ is $\mathbf{a}^T$. Before adding the $(n+1)_s^t$ row, we had from equation (8a) that

$$\mathbf{A}^T\mathbf{A}\mathbf{V} = \mathbf{V}\Lambda^2. \tag{50}$$

If we denote by $\mathbf{A}'$ the new matrix with the $(n+1)st$ row added.

$$\mathbf{A}' = \begin{pmatrix} \mathbf{A} \\ \hline \mathbf{a}^T \end{pmatrix} \qquad (n+1) \times m \tag{51}$$

the problem becomes that of finding $\mathbf{V}'$ and Λ' such that

$$\mathbf{A}'^T\mathbf{A}'\mathbf{V}' = \mathbf{V}'\Lambda'^2. \tag{52}$$

First, we note that

$$\mathbf{A}'^T\mathbf{A}' = [\mathbf{A}^T\mathbf{A} + \mathbf{a}\mathbf{a}^T]. \tag{53}$$

Then, to first order, we assume that the eigenvectors are unchanged, that is $\mathbf{V}' = \mathbf{V}$. We then have

$$[\mathbf{A}^T\mathbf{A} + \mathbf{a}\mathbf{a}^T]\mathbf{V} = \mathbf{V}\Lambda'^2 \tag{54}$$

or

$$\Lambda'^2 = \mathbf{V}^T[\mathbf{A}^T\mathbf{A} + \mathbf{a}\mathbf{a}^T]\mathbf{V} = \Lambda^2 + \delta\Lambda^2. \tag{55}$$

Because the matrix $\mathbf{V}$ is only approximately the matrix of eigenvectors for the problem (52), Λ'^2 will not be a truly diagonal matrix, but for large systems the off-diagonal elements will be small compared with the diagonal elements. (If they are not, first order perturbation theory is insufficient.) Ignoring off-diagonal elements, we have

$$\delta\lambda_j^2 = (\mathbf{V}_j^T\mathbf{a})^2.$$

We may now estimate the possible improvement in resolution and variance by recomputing the tradeoff curve, using the same eigenvectors $\mathbf{v}_j$ but augmented eigenvalues

$$\lambda_j' = \sqrt{\lambda_j^2 + \delta\lambda_j^2}.$$

The direct effect of augmenting the λ_j will be to decrease the variance in $\hat{\mathbf{x}}$, which in turn may allow more eigenvectors $\mathbf{V}$ to be used in the inverse without violating the preset variance limit.

7. Special inverses for overconstrained systems

By now, it should be clear that each feature of the underdetermined system has its analog in the overconstrained system. For the underdetermined system, we were primarily interested in the tradeoff between resolution and variance and the completeness of the set of eigenvectors in $\mathbf{V}$. Here, we are most interested in the tradeoff between information density and variance, and the completeness of $\mathbf{U}$. Using the essentials of the least squares technique, we may design an inverse $\mathbf{H}$ which will provide an information density $\mathbf{S}$ closest in the least squares sense to any desired matrix of the proper dimensionality. If we design $\mathbf{S}$ to approximate the unit matrix, we find which data act independently and which are seen through a hazy window. Designing $\mathbf{S}$ to approximate a family of somewhat broader peaks may eliminate side lobes. Where the data represent successive values read from a continuous curve, designing the $\mathbf{S}$ matrix to approximate a series of 'dipoles' will help reveal to what extent the derivative of the data curve represents independent information.

For the underdetermined system, the free eigenvectors in $\mathbf{V}_0$ could be used to search for a solution closest to some hypothetical function with special physical meaning. For overconstrained systems, the same principle may be used to test the

residual vector for components with special physical meaning (e.g. predictable errors with known properties but unknown magnitude). Removal of these components is equivalent to 'filtering' the data.

We may also consider the value of creating an additional unknown to help fit the data. This results in adding a new column to the matrix **A**. First order perturbation theory may tell us quite easily the improvement this would allow in the information density matrix, and thus in the fit to the data.

8. Summary

The matrix equations which occur in real life situations do not need to be well conditioned. They may be simultaneously overconstrained and underdetermined. Even when the matrix relating the unknowns to the data appears non-singular, random errors in the data may reduce the effective number of degrees of freedom in a linear system. However, even the most ill-conditioned system may contain useful information. The matrix equation $\mathbf{Ax} = \mathbf{y}$ may be solved in terms of a linear operator on the data, $\hat{\mathbf{x}} = \mathbf{Hy}$. Useful criteria for constructing H are (a) $\mathbf{R} \equiv \mathbf{HA} \sim \mathbf{I}_m$; (b) $\mathbf{S} \equiv \mathbf{AH} \sim I_n$; and (c) var $(\hat{x}_k) = \sum H_{ki}^2$ var (y_i) is small.

For underconstrained systems, the matrix **R** has a physical interpretation: the rows of **R** represent windows through which the general solution (and thus, the 'real' solution) may be observed. For overconstrained systems, the model $\hat{\mathbf{x}}$ will in fact fit certain combinations of the data, determined by passing the real data through the windows represented by the rows of **S**. In general, the performance of the inverse under (a) and (b) will be inversely related to performance under (c).

Decomposing the matrix **A** into its eigenvalues and eigenvectors, $\mathbf{A} = \mathbf{U}\Lambda\mathbf{V}^T$, allows the investigator to construct easily inverse operators which optimize any combination of the three criteria above, if (a) and (b) are taken in the least squares sense. Furthermore, special inverses may be easily constructed to test hypothetical features of solutions. Moreover, simple perturbation theory allows one to calculate easily the improvement in resolution and information density provided by taking more data or increasing the number of unknown parameters.

References

Backus, George & Gilbert, F., 1968. The resolving power of gross earth data, *Geophys. J. R. astr. Soc.*, **16**, 169–205.

Backus, George & Gilbert, F., 1970. Uniqueness in the inversion of inaccurate gross earth data, *Phil. Trans. R. Soc. Lond.*, **266A**, 123–192.

Hildebrand, F., 1965. *Methods of applied mathematics*, Prentice-Hall, Englewood Cliffs, New Jersey.

Jordan, T. H. & Franklin, J. N., 1971. Optimal solutions to a linear inverse problem in geophysics, *Proc. Nat. Acad. Sci.*, **68**, 291–293.

Kaula, W., 1966. *Theory of satellite geodesy*, Blaisdell, Waltham, Mass.

Lanczos, C., 1961. *Linear differential operators*, D. Van Nostrand Co., Lond.

Penrose, R., 1955. A generalized inverse for matrices, *Proc. Camb. Phil. Soc.*, **51**, 406-413.

Smith, M. L. & Franklin, J. N., Geophysical application of generalized inverse theory, *J. geophys. Res.*, **74**, 2783–2785.

Wiggins, Ralph, 1972. The general linear inverse problem: implications of surface waves and free oscillations on earth structure, *Rev. Geophys. Space Phys.*, **10**, in press.