TOPICAL REVIEW

Separable nonlinear least squares: the variable projection method and its applications

To cite this article: Gene Golub and Victor Pereyra 2003 *Inverse Problems* **19** R1

INSTITUTE OF PHYSICS PUBLISHING INVERSE PROBLEMS

Inverse Problems **19** (2003) R1–R26 PII: S0266-5611(03)52278-X

TOPICAL REVIEW

Separable nonlinear least squares: the variable projection method and its applications

Gene Golub[1] **and Victor Pereyra**[2]

[1] Scientific Computing and Computational Mathematics, Stanford University, Stanford, CA, USA
[2] Weidlinger Associates, 4410 El Camino Real, Los Altos, CA, USA

Received 8 August 2002
Published 14 February 2003
Online at stacks.iop.org/IP/19/R1

Abstract

In this paper we review 30 years of developments and applications of the variable projection method for solving separable nonlinear least-squares problems. These are problems for which the model function is a linear combination of nonlinear functions. Taking advantage of this special structure, the method of variable projections eliminates the linear variables obtaining a somewhat more complicated function that involves only the nonlinear parameters. This procedure not only reduces the dimension of the parameter space but also results in a better-conditioned problem. The same optimization method applied to the original and reduced problems will always converge faster for the latter. We present first a historical account of the basic theoretical work and its various computer implementations, and then report on a variety of applications from electrical engineering, medical and biological imaging, chemistry, robotics, vision, and environmental sciences. An extensive bibliography is included. The method is particularly well suited for solving real and complex exponential model fitting problems, which are pervasive in their applications and are notoriously hard to solve.

I. Introduction and historical background

1. Introduction

We consider in this paper nonlinear data fitting problems which have as their underlying model a linear combination of nonlinear functions. More generally, one can also consider that there are two sets of unknown parameters, where one set is dependent on the other and can be explicitly eliminated. Models of this type are very common and we will show a variety of applications in different fields. Inasmuch as many inverse problems can be viewed as nonlinear data fitting problems, this material will be of interest to a wide cross-section of researchers and practitioners in parameter, material or system identification, signal analysis, the analysis of

spectral data, medical and biological imaging, neural networks, robotics, telecommunications, and model order reduction, to name but a few.

The authors published in April 1973 the paper [45] on 'The differentiation of pseudoinverses and separable nonlinear least squares'. This was work initiated in 1971, motivated by and generalizing earlier work of Guttman *et al* [49], which in turn elaborated on and generalized work in Hugo Scolnik's Doctoral Thesis produced at the University of Zurich [112, 113].

Scolnik's original work dealt with the case in which the nonlinear functions depended on one variable each and were of exponential type (t_i^α). In [49] this was extended to general functions of one variable, while in [45] functions of several variables were considered. In this last paper the authors also developed a differential calculus for projectors and pseudoinverses, and proved that the separation-of-variables approach led to the same solution set as that of the original problem. Then Ruhe and Wedin [101] extended these ideas to the more general case of two arbitrary sets of variables.

The separability aspect comes from the idea of eliminating the linear variables (i.e., either the coefficients in the linear combinations or one set of the nonlinear variables), and then minimizing the resulting *variable projection* (VP) functional that depends only on the remaining variables.

Improvements in the algorithm, most especially a simplification due to Kaufman [57], extended by Ruhe and Wedin [101] to the general nonlinear case, made the cost per iteration for the reduced functional comparable to that for the full functional, as proven by the complexity analysis of Ruhe and Wedin.

At first it was thought that the main benefit of the elimination of one set of parameters was that fewer variables needed to be estimated. Numerical experience showed, however, that in most cases the reduced problem converged in fewer iterations than if the same minimization algorithm were used for the full functional. This was later substantiated theoretically for the Gauss–Newton (GN) method by the work of Ruhe and Wedin [101], where the asymptotic rates of convergence for the Golub–Pereyra (GP), the Kaufman variant, the NIPALS algorithm, and the unreduced functional are compared by studying the eigenvalues of the corresponding fixed point iteration functions.

Since the reduced functional was more complicated, this reduction in the number of iterations did not guarantee that the total computing time was smaller, but there were classes of problems for which the reduction was dramatic, and in fact, it was clearly beneficial to use the VP formulation. Exponential fitting was one such problem. See for instance [45, 102].

The main contributions of the GP paper were to give a clean, matrix-based formulation of the problem, a differential calculus for matrix functions, orthogonal projectors, and generalized inverses, and a modern (for the early 1970s) and detailed algorithm for dealing with the complexities arising from the VP formulation, that included an early efficient implementation of an adaptive Levenberg–Marquardt (LM) algorithm. They also showed that the reduced problem had the same minima as the original one, provided that the linear parameters were recovered by solving an appropriate linear least-squares problem.

We believe that an important part of the impact of this work has come from the fact that a usable, public computer implementation of the algorithm was made available. The 1973 paper is evenly divided between theoretical derivation and implementation details. In a Stanford report with the same name [44], we included a listing of the program VARPRO that actually implemented the VP algorithm and produced the numerical results in the paper. Variants of this original code are still in use in some of the applications that we mention below.

The purpose of this paper is to take the story from where we left it after our second paper in 1976 [46], which already contains details on a number of related contributions, mostly

clustered around the early 1970s. In fact, what we want to stress here is the surprising richness of applications of this idea and its impact in a number of very different fields, with lively developments carrying over to this century.

Thus, we will classify the applications that we have collected through the years roughly by field, giving a more detailed description of the basic problem solved for some selected items and pointing out whatever new insights of general value have been discovered. We hope that this strategy will help different practitioners find clustered the material of most interest to them, while also calling their attention to possible cross-pollination.

We do not attempt to be comprehensive, and we refer the reader to the excellent bibliographies of many of the quoted papers for connected contributions. The selection is more by expediency than by any attempt to single out some contributions and slight others.

Some interesting trends are observed: during the first few years most of the contributions relate to enhancements, modifications, theoretical results, and comparisons. Some of the early applications occur in the area of signal localization, which is still one of the most active fields today. This is not totally surprising given the recent explosion in new telecommunication and mobile applications. Another very active field is that of the modelling and interpretation of nuclear magnetic resonance data, where VARPRO has a stellar position. A more recent interesting field of application is that of neural network training.

We will use repeatedly some acronyms that we define here for further reference: LLS, linear least squares; NLLS, nonlinear least squares; SNLLS, separable nonlinear least squares; VP, variable projection; SVD, singular value decomposition.

2. Separable nonlinear least squares and the variable projection method

Given a set of observations $\{y_i\}$, a separable nonlinear least-squares problem is defined in [45] as one for which the model is a linear combination of nonlinear functions that can depend on multiple parameters, and for which the ith component of the residual vector is written as

$$r_i(\boldsymbol{a}, \boldsymbol{\alpha}) = y_i - \sum_{j=1}^{n} a_j \phi_j(\boldsymbol{\alpha}; t_i).$$

Here the t_i are independent variables associated with the observations y_i, while the a_j and the k-dimensional vector $\boldsymbol{\alpha}$ are the parameters to be determined by minimizing the functional $\|\boldsymbol{r}(\boldsymbol{a}, \boldsymbol{\alpha})\|_2^2$, where the components of $\boldsymbol{r}(\boldsymbol{a}, \boldsymbol{\alpha})$ are $r_i(\boldsymbol{a}, \boldsymbol{\alpha})$, and $\|\cdot\|_2$ stands for the l_2 vector norm. We can write this functional using matrix notation as

$$\|\boldsymbol{r}(\boldsymbol{a}, \boldsymbol{\alpha})\|_2^2 = \|\boldsymbol{y} - \boldsymbol{\Phi}(\boldsymbol{\alpha})\boldsymbol{a}\|_2^2,$$

where the columns of the matrix $\boldsymbol{\Phi}(\boldsymbol{\alpha})$ correspond to the nonlinear functions $\phi_j(\boldsymbol{\alpha}; t_i)$ of the k parameters $\boldsymbol{\alpha}$ evaluated at all the t_i-values, and the vectors $\boldsymbol{a}$ and $\boldsymbol{y}$ represent the linear parameters and the observations respectively.

Now it is easy to see that if we knew the nonlinear parameters $\boldsymbol{\alpha}$, then the linear parameters $\boldsymbol{a}$ could be obtained by solving the linear least-squares problem:

$$\boldsymbol{a} = \boldsymbol{\Phi}(\boldsymbol{\alpha})^{+}\boldsymbol{y}, \tag{1}$$

which stands for the minimum-norm solution of the linear least-squares problem for fixed $\boldsymbol{\alpha}$, where $\boldsymbol{\Phi}(\boldsymbol{\alpha})^{+}$ is the Moore–Penrose generalized inverse of $\boldsymbol{\Phi}(\boldsymbol{\alpha})$. On replacing this $\boldsymbol{a}$ in the original functional, the minimization problem takes the form

$$\min_{\boldsymbol{\alpha}} \tfrac{1}{2}\|(\boldsymbol{I} - \boldsymbol{\Phi}(\boldsymbol{\alpha})\boldsymbol{\Phi}(\boldsymbol{\alpha})^{+})\boldsymbol{y}\|_2^2, \tag{2}$$

where the linear parameters have been eliminated.

We define $\boldsymbol{r}_2(\boldsymbol{\alpha}) = (\boldsymbol{I} - \boldsymbol{\Phi}(\boldsymbol{\alpha})\boldsymbol{\Phi}(\boldsymbol{\alpha})^+)\boldsymbol{y}$, which will be called the VP of $\boldsymbol{y}$. Its name stems from the fact that the matrix in parentheses is the projector on the orthogonal complement of the column space of $\boldsymbol{\Phi}(\boldsymbol{\alpha})$, which we will denote in what follows by $P^{\perp}_{\boldsymbol{\Phi}(\boldsymbol{\alpha})}$. We will also refer to $\frac{1}{2}\|\boldsymbol{r}_2(\boldsymbol{\alpha})\|_2^2$ as the VP functional.

This is a more powerful paradigm than the simple idea of alternating between minimization of the two sets of variables (such as the NIPALS algorithm of Wold and Lyttkens [139]), which can be proven theoretically and practically not to result, in general, in the same enhanced performance.

In summary, the VP algorithm consists of first minimizing (2) and then using the optimal value obtained for $\boldsymbol{\alpha}$ to solve for $\boldsymbol{a}$ in (1). One obvious advantage is that the iterative nonlinear algorithm used to solve the first minimization problem works in a reduced space and, in particular, fewer initial guesses are necessary. However, the main payoff of this algorithm is the fact that it always converges in fewer iterations than the minimization of the full functional, including convergence when the same minimization algorithm for the full functional diverges (see for instance [64]), i.e., the minima for the reduced functional are better defined than those for the full one.

We demonstrated also that the set of stationary points of the original and reduced functionals are the same. This theorem has been reassuring to many practitioners and has been used to derive other theoretical results in similar situations.

A different reason for using the reduced functional is seen by observing from the above results that the linear parameters are determined by the nonlinear ones, and therefore the full problem must become increasingly ill-conditioned as (and if) it converges to the optimal parameters. That is probably one of the reasons that the important and prevalent problem of real or complex exponential fitting is so hard to solve directly. See for instance [117] for a theoretical discussion of this issue and an interesting application to the training of nonlinear neural networks. Further comments on the basic results can also be found in the textbooks of Seber and Wild [114] and Björck [14].

3. Numerical methods for nonlinear least-squares problems

General numerical optimization methods can be used to solve NLLS problems, but it pays to take into consideration the special form of the goal functional (a sum of squares), just as it pays to take advantage of the special form of separable problems. We review briefly some of the basic concepts that lead to the main numerical methods used today in standard and SNLLS problems.

We assume in what follows that the model functions, $\phi_j(\boldsymbol{\alpha}; t)$, are twice differentiable with respect to $\boldsymbol{\alpha}$. A fundamental quantity for any optimization method for NLLS that uses derivatives is the Jacobian matrix $\boldsymbol{J}(\boldsymbol{\alpha})$ of the vector of residuals: $\boldsymbol{r}_2(\boldsymbol{\alpha})$. This appears when calculating the gradient of the VP functional

$$\nabla \tfrac{1}{2}\|\boldsymbol{r}_2(\boldsymbol{\alpha})\|_2^2 = \boldsymbol{J}^{\mathrm{T}}(\boldsymbol{\alpha})\boldsymbol{r}_2(\boldsymbol{\alpha}),$$

while its Hessian is given by

$$\nabla^2 \tfrac{1}{2}\|\boldsymbol{r}_2(\boldsymbol{\alpha})\|_2^2 = \boldsymbol{J}^{\mathrm{T}}(\boldsymbol{\alpha})\boldsymbol{J}(\boldsymbol{\alpha}) + \sum_{i=1}^{n} r_{2i}(\boldsymbol{\alpha})\, \nabla^2 r_{2i}(\boldsymbol{\alpha}).$$

The Gauss–Newton (GN) method for nonlinear least squares can be viewed as a quasi-Newton method in which the Hessian is approximated by $\boldsymbol{J}^{\mathrm{T}}(\boldsymbol{\alpha})\boldsymbol{J}(\boldsymbol{\alpha})$, while the Levenberg–Marquardt (LM) enhancement adds a positive definite matrix and a coefficient λ in order to combat

ill-conditioning. GN is very effective for small-residual problems, since in that case the neglected term is unimportant.

By using a trust-region strategy for step control, needed to stabilize the plain GN and LM iterations and making them more globally convergent, one obtains Moré's implementation [76]. See also [93] for an early proof of convergence of the GN algorithm, [90] for an early proof of convergence for LM, and the GP paper for a detailed implementation of an adaptive LM method for SNLLS problems.

In [91], Pereyra also describes a detailed implementation of an SVD-based trust-region LM algorithm for NLLS problems that appear in large-scale seismic inversion tasks. These are notoriously ill-conditioned problems and often outright singular. It would be worthwhile to consider such an algorithm for the regularization of ill-conditioned SNLLS problems.

Instead of describing the implementation in the GP paper we go directly to that of Kaufman [57], as described in Gay and Kaufman [41]. In this discussion we omit the $\boldsymbol{\alpha}$-dependence in all the matrices in order to lighten the notation. They first observe that GP have proven that the full Moore–Penrose pseudoinverse is not necessary to represent the orthogonal projector, and that a symmetric generalized inverse $\boldsymbol{\Phi}^-$ satisfying only $\boldsymbol{\Phi}\boldsymbol{\Phi}^-\boldsymbol{\Phi} = \boldsymbol{\Phi}$ and $(\boldsymbol{\Phi}\boldsymbol{\Phi}^-)^{\mathrm{T}} = \boldsymbol{\Phi}\boldsymbol{\Phi}^-$ suffices. Thus, the jth column of the Jacobian $\boldsymbol{J}$ can be written as

$$\boldsymbol{J}_{\cdot j} = -\left[\left(\boldsymbol{P}_{\boldsymbol{\Phi}}^{\perp}\frac{\partial\boldsymbol{\Phi}}{\partial\alpha_j}\boldsymbol{\Phi}^-\right) + \left(\boldsymbol{P}_{\boldsymbol{\Phi}}^{\perp}\frac{\partial\boldsymbol{\Phi}}{\partial\alpha_j}\boldsymbol{\Phi}^-\right)^{\mathrm{T}}\right]\boldsymbol{y}.$$

Kaufman's simplification for the VP functional consists in approximating the Hessian in the GN method by using only the first term in the Jacobian formula: $\boldsymbol{L} = -\big(\boldsymbol{P}_{\boldsymbol{\Phi}}^{\perp}\frac{\partial\boldsymbol{\Phi}}{\partial\alpha_j}\boldsymbol{\Phi}^-\big)\boldsymbol{y}$, thus saving in the numerical linear algebra involved at the cost of marginally more function and gradient evaluations. It has been extensively demonstrated, as we indicate below, that savings of up to 25% are achieved by this simplification, making the VPK method as cost efficient per iteration as working with the full functional.

Kaufman's argument to justify the small impact of her simplification is that if one writes $\boldsymbol{J} = \boldsymbol{K} + \boldsymbol{L}$, where $\boldsymbol{K}_{\cdot j} = -\big(\boldsymbol{P}_{\boldsymbol{\Phi}}^{\perp}\frac{\partial\boldsymbol{\Phi}}{\partial\alpha_j}\boldsymbol{\Phi}^-\big)^{\mathrm{T}}$, then:

$$\boldsymbol{J}^{\mathrm{T}}\boldsymbol{J} = \boldsymbol{K}^{\mathrm{T}}\boldsymbol{K} + \boldsymbol{L}^{\mathrm{T}}\boldsymbol{L} + \boldsymbol{K}^{\mathrm{T}}\boldsymbol{L} + \boldsymbol{K}\boldsymbol{L}^{\mathrm{T}} = \boldsymbol{K}^{\mathrm{T}}\boldsymbol{K} + \boldsymbol{L}^{\mathrm{T}}\boldsymbol{L},$$

since $\boldsymbol{K}$ lies in the null space of $\boldsymbol{\Phi}^{\mathrm{T}}$, while $\boldsymbol{L}$ lies in the range of $\boldsymbol{\Phi}$, so only the term $\boldsymbol{L}^{\mathrm{T}}\boldsymbol{L}\boldsymbol{r}(\boldsymbol{\alpha})$ is being dropped from the exact Hessian.

We point to the reference above for the linear algebra aspects involved in the Kaufman simplification, which combine to produce the quoted 25% reduction in time per iteration.

4. Variations and related algorithms

Bates and Lindstrom [7] give an interesting statistical interpretation of both the GP and the Kaufman algorithms. After some analysis and numerical comparisons they conclude that these algorithms are very attractive and provide greater stability than methods for the full functional, besides reducing the dimensionality of the optimization problem.

Later, Ruhe and Wedin [101] demonstrated the asymptotic convergence properties of the two methods, confirming the experimental results that GN always converges in fewer iterations for the reduced functional than for the full one. They also extended VP to more general separable problems, where the set of variables splits into two, and where, presumably, one of the sets can be easily eliminated. They showed that VP, with the Kaufman simplification, had the same cost per iteration as the full functional approach.

Golub and LeVeque [43] extended VP to problems with multiple right-hand sides. See also [59].

Böckmann [10] has considered regularization through a trust-region algorithm combined with separation of variables and has obtained excellent results in comparison with state-of-the-art solvers applied to the unreduced problem.

In a recent MSc Thesis from Dalhousie University (advisor Patrick Keast), Lukeman [72] extends the application of the Shen–Ypma algorithm [115] to overdetermined systems, establishing the connection with the GP approach, and he gives a very lucid and accurate description of the early developments. See also [78].

Osborne and collaborators [85, 86] have studied through the years Prony's method, another reduction technique valid for the exponential fitting problem, and they have derived variations that make it more stable.

5. Constrained problems

Kaufman and Pereyra [58] extended VP to problems with separable equality constraints of the form

$$\boldsymbol{H}(\boldsymbol{\alpha})\boldsymbol{a} = \boldsymbol{g}(\boldsymbol{\alpha}).$$

They show how to reduce this to an unconstrained problem that can be solved by a standard SNLLS solver, such as VARPRO. They also go on to develop a more efficient algorithm that takes into account the special structure of this problem. See also [27] for further refinements.

Parks [87] considered the basic theory for separable nonlinear programming problems in the spirit of Ruhe and Wedin, that she called reducible. Then she went on to consider specific special cases, including the one corresponding to the equality constrained problem studied by Kaufman and Pereyra, that she called semilinear.

More recently, Conn *et al* [25] elaborated on this idea in the context of trust-region algorithms for general nonlinear programming problems, where a subset of variables occur in some predictable functional form that can be used to optimize them more economically.

6. Various implementations

The VARPRO program [44] was written by Pereyra. At Stanford University, John Bolstad streamlined the code and improved the documentation of VARPRO under the guidance of Gene Golub. He also included the Kaufman simplification and the calculation and output of the covariance matrix, a very important addition that is frequently missing in least-squares codes written by numerical analysts. Another graduate student at Stanford, Randy LeVeque, wrote a modified version based on the original VARPRO code called VARP2 which extends the original code to problems with multiple right-hand sides.

Both VARPRO and VARP2 have been publicly available for a long time in netlib, on the port library compiled by David Gay [40]. There one can also find versions of the Gay and Kaufman implementation for unconstrained (nsf.f) and bound-constrained SNLLS problems, which include the option of using finite differences to approximate the derivatives. A FORTRAN 90 version of VARPRO can be found in Alan Miller's repository [75].

Sagara and Fukushima [106] use parameter continuation to solve SNLLS problems and report reasonable success in increasing the domain of convergence for certain complicated exponential–polynomial fitting problems.

An advocate of VARPRO, Bert Rust, brought the code early on to Oak Ridge National Laboratory and the National Bureau of Standards (NBS, now known as NIST (National Institute of Standards and Technology)), where it has been used through the years in many applications, some of which we mention below. Rall and Funderlic [97] wrote an interactive front end at

Oak Ridge National Laboratory, INVAR, to facilitate the use of VARPRO and to add more statistics. They also added a finite-difference evaluation of the derivatives as an option. Later this code was improved at NBS, including graphical output [140], and it is still in use at NIST.

Another group of scientists who adopted VARPRO were those involved in analysing *in vivo* magnetic resonance spectra (MRS); this line of research began with the work of van der Veen *et al* [130] at Delft University in The Netherlands.

A version of VARPRO can be found in Magnetic Resonance User Interface (MRUI) [81], a widely used system for MR imaging maintained by the Universidad Autónoma de Barcelona, Spain (a version in MATLAB [82] is also available). This is mostly a non-profit effort financed through grants of the European Union under the project on advanced signal processing for medical magnetic resonance imaging and spectroscopy (TMR-CT970160). More than 300 research groups in 40 countries have licensed MRUI.

7. Performance and comparisons

In the original paper we showed the main performance characteristics of the VP method in several problems involving exponentials, Gaussians, and a rational model for fitting iron Mössbauer spectra, as compared to using the full functional. Four different methods were considered: PRAXIS, a general optimization code that does not use derivatives produced by Richard Brent [17], a GN method with step control, a Marquardt-type code with adaptive choice of the regularization parameter that we developed, and a variable metric code produced by M Osborne.

The conclusions were in line with what was known at the time and they are still valid today: GN was fastest when it worked, requiring a good initial estimate, while the variable metric code was not competitive. Brent's code is recommended if analytical derivatives are a problem, but otherwise it usually requires more function evaluations. Since the cost per iteration was higher for the VP functional, reduction in the number of iterations alone was not a guarantee of less computer time, as shown in some of the examples. Again, these results were dependent on the method used, and for the same problems with the same initial values and final error we obtained different comparative performances for different methods. In these examples, the VP approach consistently required fewer iterations, including problems in which it converged while the iteration for the full functional diverged.

For the exponential fitting problem, the VP approach was consistently faster, and the results for GN and Marquardt were comparable and best by far. Since the Kaufman simplification would give an additional 25% edge to the VP method, we see that by adjusting the computer times accordingly in [45], in all cases considered the time performance of VPK is better than that corresponding to the minimization of the full functional. As mentioned above, this has also been confirmed theoretically in the paper by Ruhe and Wedin.

That is to say, when VP is combined with the Kaufman modification, the costs per iteration for the full and reduced functionals are similar, debunking the original notion that the VP functional was considerably more complicated and therefore more expensive to calculate than the original one.

Krogh [64], Kaufman [57], Corradi [26], Gay and Kaufman [41], and Nielsen [84], among others, had similar experiences. The most comprehensive independent studies are those of Corradi who also considered problems with noise, and Gay and Kaufman, who were concerned with proving that the Kaufman modification did not alter the robustness or number of iterations necessary. Unfortunately Corradi does not report computer times.

See also the section on medical and biological applications for a discussion of some detailed comparisons within magnetic resonance spectroscopy applications.

II. Applications

8. Applications to numerical analysis, mathematics and optimization

One of the uses of the derivative of projectors and generalized inverses, as indicated in [45, 131], is in the study of sensitivity to errors in solving linear least-squares problems. For instance, the usual perturbation analysis can be extended to the rank-deficient case (for rank-preserving perturbations): if $\boldsymbol{A}(\epsilon) = \boldsymbol{A} + \epsilon\boldsymbol{B}$ and $\|\boldsymbol{A}\| = \|\boldsymbol{B}\| = 1$, we can estimate the error in the least-squares solution of $\boldsymbol{A}(\epsilon)\boldsymbol{x}(\epsilon) = \boldsymbol{b}$, for small ϵ.

By using a Taylor expansion,

$$\boldsymbol{A}^{+}(\epsilon) - \boldsymbol{A}^{+}(0) = \epsilon \boldsymbol{D}\boldsymbol{A}^{+}(0)\boldsymbol{B} + \mathrm{O}(\epsilon^2),$$

and the formula from [45] for the Fréchet derivative of the generalized inverse with respect to ϵ,

$$\boldsymbol{D}\boldsymbol{A}^{+}(0) = -\boldsymbol{A}^{+}\boldsymbol{B}\boldsymbol{A}^{+} + \boldsymbol{A}^{+}\boldsymbol{A}^{+\mathrm{T}}\boldsymbol{B}^{\mathrm{T}}\boldsymbol{P}_{\boldsymbol{A}}^{\perp} + \boldsymbol{Q}\boldsymbol{B}^{\mathrm{T}}\boldsymbol{A}^{+\mathrm{T}}\boldsymbol{A}^{+},$$

where $\boldsymbol{Q}$ stands for the projector on the orthogonal complement of the *row* space of $\boldsymbol{A}$ and we have used the fact that $\boldsymbol{D}\boldsymbol{A}(0) = \boldsymbol{B}$, we obtain the classical estimate,

$$\|\boldsymbol{x}(\epsilon) - \boldsymbol{x}(0)\| \leqslant 2\epsilon\|\boldsymbol{A}^{+}\|\|\boldsymbol{x}(0)\| + \epsilon\|\boldsymbol{A}^{+}\|^2\|\boldsymbol{r}\| + \mathrm{O}(\epsilon^2),$$

where $\boldsymbol{r} = \boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}$ is the residual vector. Additional perturbation analyses of this kind are considered in [33, 47, 48, 118].

Koliha [63] has extended the differentiability results of the Moore–Penrose inverse in [45] to C*-algebras.

In [123] Trosset considers the problem of computing distances between convex sets of positive definite matrices as a reducible programming problem in the sense of Park, and takes good advantage of the separability of the variables.

Some of the results of [45, 46] are used by Bramley and Winnicka [16] when solving linear inequalities and by Byers *et al* [20] when looking for the nearest non-regular pencil of matrices.

Lanzkron and Rose [67] discuss in detail approximate nonlinear elimination. This is the problem of solving large-scale, sparse, nonlinear algebraic equations by a nonlinear Gauss–Seidel method. So, basically there is an external Gauss–Seidel iteration and an internal iterative solver for single nonlinear equations (or systems in the block case), and the old question of how to manage the inner iterations such that the outer iterations preserve some good properties (see [89, section 3] for an early solution) is revisited. They also discuss the problem of slow convergence, associated perhaps with poor scaling of certain sets of variables, and consider the separation of those variables in order to improve the overall convergence. This is then connected to SNLLS problems and the VP approach.

9. Applications to chemistry

The moment-singularity method for the calculation of densities of state in the vicinity of van Hove singularities is considered by Prais *et al* [94, 137]. Because the asymptotic behaviour of the modified moments is related to the singular behaviour of the density, information about the locations and functional forms of the singularities can be determined directly from the moments themselves. In the 1974 paper a least-squares method for calculating the locations and exponents of the singularities is described. This NLLS problem is separable and the authors report good success with the 'very fast' VARPRO code, when using reasonable initial estimates for the location of the singularities. Several numerical results and comparisons are

reported. In the 1986 paper they refine the method and still make good use of VARPRO, reporting that they have found this code leading to accurate determination of the singularities when as few as 20 moments are used. Several new calculations, comparisons, and validations are included.

Lee and Baeck [68] demonstrate, using a highly positive uranium ion as a test case, that the exact relationship between small and large components of a Dirac spinor in relativistic self-consistent field calculations is not fully satisfied by the kinetic balance condition, even for a two-electron system. In order to obtain a basis set for a multiple-electron system, the numerical atomic spinors obtained by Dirac–Hartree–Fock calculations are fitted by a given number of Slater-type functions. This SNLLS problem is solved using VP. The calculations for a uranium ion demonstrate the difference between the exponents for the small and large components, giving numerical evidence that the kinetic balancing is not an exact relationship between small and large components of spinors.

In [11], Beece *et al* use a VP algorithm in an exponential fitting problem associated with the effect of viscosity on the kinetics of the photochemical cycle of bacteriorhodopsin. Marque and Eisenstein [73] extend this work to consider pressure effects on the photocycle of purple membrane. By considering several kinetic data sets taken at the same temperature and pressure but with different monitoring wavelengths and an exponential model, they are able to use VARP2 to separate the variables and efficiently solve a problem with multiple right-hand sides. The first to use this method in such problems was Richard Lozier [71], who motivated the development of the VARP2 extension and became a champion in this field for many years (we thank Randy LeVeque for this insight).

In [51], Holmström considers constrained SNLLS for chemical equilibrium analysis. He uses his own implementation of a VP algorithm that can be found in the LAKE program system [54]. See also the description of the MATLAB-based TOMLAB system in [52]. A recent review paper with Pettersson [53] is also noteworthy. In an earlier publication with Andersson and Ruhe [2] the use of SNLS in chemical applications is described in detail.

10. Mechanical systems

Prells and Friswell [95] consider the application of the VP method to the update of finite-element models of mechanical systems when the forces applied to them are unknown, an inverse problem. Thus, the forces and the model parameters are estimated from observed values of the system response. The VP method leads to the elimination of the unknown forces and results in an extension of the output residual method, which is frequently used to solve such inverse problems. The method is exemplified in three applications that are discussed in detail:

- damage detection of a multi-storey building subject to wind excitation;
- estimation of a foundation model of a rotary machine;
- model estimation of a steel beam tested by hammer excitation.

Prells and Friswell observe that the inclusion of prior knowledge (constraints) helps stabilize some of these ill-conditioned problems, and that in general the use of VP is very successful. They plan to incorporate better regularization techniques and to consider industrial-scale problems in future work.

Westwick *et al* [135] consider the nonlinear identification of automobile dynamics when the car is attached to a vibration test-rig. This time-dependent problem is analysed, and after some manipulations it is recognized that in the nonlinear case a SNLLS problem has to be solved. They use the algorithms of [117] and report excellent results in a benchmark example

of a continuous-time simulation of the vibrations of a quarter-car model. This model consists of two second-order linear ordinary differential equations and a nonlinear algebraic constraint that stands for the suspension spring.

Kaufman and collaborators [59, 60] consider large-scale system identification problems arising in the estimation of modal parameters from multiple-driver, multiple-receiver transfer function data in the frequency domain. These transfer functions describe the vibrations of structurally complex objects and they contain many modes. By judiciously applying the VP reduction they are able to make very large problems amenable to computation. This is a separable problem with multiple right-hand sides, where the nonlinear parameters are common but the linear ones are not, and therefore the separation of variables produces a tremendous reduction in the dimensionality of the problem.

Similarly, Edrissi *et al* [32] have considered the parameter identification problem associated with linear, time-invariant state-space models by using separation of variables. They cite as the main advantages better conditioning and performance than if the full functional approach is used. Similar results are obtained by Bruls *et al* [18] for the linear output error and the innovation model and Wiener system for the nonlinear case. Again they claim better conditioning for the reduced approach and they apply it to two real industrial problems.

11. Neural networks

A very interesting and important application can be found in the training of neural networks. Feedforward neural networks, such as multi-layer perceptrons or radial function networks, are nonlinear parametric models, as clearly explained by Weigl *et al* [132–134] and as discussed in the excellent book by Bishop [9]. In these works, some types of neural networks are conceived as linear combinations of basis functions, for instance, sigmoids, or radial functions, with free parameters on them in the nonlinear adaptive case. Training corresponds to fitting these separable models in the least-squares sense by using pre-classified data (a training set) and an optimization algorithm. Using VARPRO makes the training an order of magnitude faster than traditional back-propagation algorithms.

This speed in the training phase allows the researcher to design rapidly the neural network that best fits its needs. A transparent description of a simple neural network of this type is provided by Weigl *et al*. Given an input node, a hidden layer made of two nodes, and an output node, the various steps are summarized as follows.

- The input node accepts the input value x and sends it to the two hidden layer nodes, which are sub-indexed with i.
- The hidden nodes calculate a nonlinear function $g_i(a_i, \theta_i; x) = \mathrm{sigm}(a_i x + \theta_i)$, where the function $\mathrm{sigm}(z) = \frac{1}{1+e^{-z}}$ is the classic perceptron sigmoid mapping function, a_i is a scalar parameter, and θ_i is a bias.
- Each hidden layer node then sends its function value to the output node, which combines the various values linearly by multiplying them with a coefficient A_i and adding them up:

$$f(x) = \sum_i A_i g_i(a_i, \theta_i; x).$$

Given a training set $(x_j, f(x_j))$, $j = 1, \ldots, J$, we teach or train the neural network by finding the parameters $(\boldsymbol{A}, \boldsymbol{a}, \theta)$ that solve

$$\min_{\boldsymbol{A},\boldsymbol{a},\theta} \sum_j \left(f(x_j) - \sum_i A_i g_i(a_i, \theta_i; x_j) \right)^2,$$

a SNLLS problem. The extension to general multi-input, multi-layer, and multi-output neural networks is now straightforward.

As in other applications, one advantage of this approach is that the coefficients $\boldsymbol{A}$ are computed directly, without a learning step, and thus convergence is much faster. Other basis functions can be used instead of the sigmoids. See for instance [83] and [88]. An interesting extension would be to consider very large networks, such as those used in VLSI design, by using block techniques combined with separability [91].

In [34] the authors consider issues of regularization for NLLS problems that appear in training feedforward neural networks. It is odd that these authors, who are very familiar with the SNLLS technology, seem to have missed the above connection, at least in this contribution. However, some fellow Swedes (Sjöberg and Viberg) have not [117] (although they do not seem to be aware of the work of Weigl *et al*).

More importantly, besides using VP to train neural networks, Sjöberg and Viberg are the only researchers who we have seen mentioning explicitly and demonstrating the fact that, in general, the VP functional is better conditioned than the full functional.

12. Parameter estimation and approximation

This is a rich application area. Schwetlick [110] gives a comprehensive survey, including SNLLS. A traditional application (see [56, 109, 111]) is the fitting of data by splines with free knots. In this problem, a data set $\{x_i, y_i\}$ is given, where the x_i are abscissae and the y_i are noisy measurements of values of an unknown smooth function $f(\boldsymbol{x})$. One wishes to approximate $f(\boldsymbol{x})$ in the least-squares sense by a function $s(\boldsymbol{x})$ belonging to a finite-dimensional space. Splines provide a powerful choice for the approximation functions. They require for their definition a knot sequence $a \leqslant t_1 < \cdots < t_i < \cdots < t_n \leqslant b$, where it is advisable to take $n \ll m$.

Thus, the approximation functions are linear combinations of basis functions,

$$s(x) = \sum_{j=1}^{n} \alpha_j B_j(x),$$

where the knots appear in the definition of the B-splines. See [30, 4.2] for more details.

If the knots are fixed (usually uniformly distributed), then the fitting problem,

$$\min_{\boldsymbol{\alpha}} \sum_{i=1}^{m} \left(y_i - \sum_{j=1}^{n} \alpha_j B_j(x_i) \right)^2,$$

is a simple LLS one. However, if we consider the basis function knots as unknowns, thus potentially leading to better approximation properties for functions with high variability in the interval of interest, then this is a SNLLS problem, which has been successfully solved using VP. In fact, for practical applications, constraints need to be imposed to avoid a change in the linear ordering of the knots, but the resulting problem is still separable in the sense of Kaufman and Pereyra. There is an unexploited connection here with unequally spaced fast Fourier, wavelet, and Radon transforms [12].

Francos *et al* [37] consider the problem of estimating discrete homogeneous random fields. Again, the separation paradigm helps significantly in solving this problem which involves models with discrete complex exponentials.

An interesting paper that connects VP with the manipulation of tensor products [92] is [124], where the approximation in the Frobenius norm of a rectangular matrix by a Kronecker product is discussed.

Bates and Watts [6] consider the problem of multi-response nonlinear least-squares estimation $\boldsymbol{Z}(\theta) = \boldsymbol{Y} - \boldsymbol{H}(\theta)$, where the $N \times K$ matrix $\boldsymbol{Y}$ contains the N observations of K responses, and $\boldsymbol{H}(\theta)$ is a nonlinear model, where the P-dimensional vector of parameters θ is to be estimated. For this estimation, they chose to minimize with respect to θ the determinant of the matrix $\boldsymbol{Z}^{\mathrm{T}}\boldsymbol{Z}$

$$\min_{\theta} |\boldsymbol{Z}^{\mathrm{T}}\boldsymbol{Z}|,$$

where $|\cdot|$ stands for determinant. In order to apply standard optimization methods they derive formulae for the derivatives of this matrix. In fact, the elements g_{θ_p} of the gradient vector of the logarithm of the determinant

$$\nabla \tfrac{1}{2} \log |Z^{\mathrm{T}}(\boldsymbol{\alpha}) Z(\boldsymbol{\alpha})| \tag{3}$$

can be written as $g_{\theta_p} = \mathrm{trace}(Z^{+} Z_{\theta_p})$ [5], and the Hessian can then be obtained from the formulae for the differentiation of the pseudoinverse.

13. Telecommunications; electrical and electronic engineering

A number of applications of VP are related to classical and modern telecommunication problems, many of which can be cast as SNLLS fitting problems for linear combinations of complex exponentials, where the linear coefficients represent amplitude, while the nonlinear ones are the phases of the signals (plane waves).

Roy and Kailath [100] describe in detail applications to practical signal processing problems. The objective there is to estimate from measurements a set of constant (time-independent) parameters upon which the received signal depends. Among these, high-resolution direction-of-arrival (DOA) estimation is important in many sensor systems such as radar, sonar, electronic surveillance, and seismic exploration. High-resolution frequency estimation is important in numerous applications, such as the design and control of robots and large flexible space structures. In such problems, the functional form of the underlying signals can often be assumed (e.g., narrow-band plane waves, cisoids). The quantities to be estimated are parameters in these functional descriptions, such as frequencies and directions of arrival for plane waves, or cisoid frequencies.

Several approaches have been developed through the years for solving these problems, including Capon's [23] maximum-likelihood and Burg's [19] maximum-entropy methods. These methods have significant limitations and Pisarenko was one of the first to consider the structure of the data model to estimate the parameters of cisoids in additive noise using a covariance approach. Schmidt [108] and Bienvenu [13] were the first to exploit correctly the measurement model in the case of a sensor array of arbitrary form. Schmidt's algorithm, MUSIC (MUltiple SIgnal Classification), which according to that author was inspired by the separation-of-variables technique, has been widely studied and was considered in an MIT study of that time as the most promising high-resolution algorithm. However, MUSIC's success came at a high computational cost that involved a search in parameter space and the storage of array calibration data.

Roy and Kailath developed a new algorithm, called ESPRIT, that dramatically reduced the computational cost and storage for sensor arrays that show what they call displacement invariance. These are arrays where the sensors come in matched pairs with identical displacement vectors.

We will concentrate now on the direction-of-arrival problem for plane waves (i.e., the sensors are in the far field of the source of energy), in an isotropic, homogeneous, and non-

dispersive medium, so that energy propagates in straight lines. Under those assumptions, the complex signal output of the kth sensor at time t can be written as

$$\boldsymbol{x}(t) = \sum_{i=1}^{d} s_i(t)\boldsymbol{a}(\theta_i),$$

where $\boldsymbol{a}(\theta_i) = [a_1(\theta_i \mathrm{e}^{-\mathrm{j}\omega_0\tau_1(\theta_i)}), \ldots, a_m(\theta_i \mathrm{e}^{-\mathrm{j}\omega_0\tau_m(\theta_i)})]$ is the array steering vector for direction θ_i. Observe that this is a separable model.

Unfortunately, many of the earlier simplified algorithms are ineffective when some of the sources are coherent. This can stem from multi-path effects or it can be introduced artificially to impede detection. Kumaresan and Shaw [65] and Cadzow [21] have studied in detail the application of separation of variables to this classical problem. More recently, a number of new algorithms have been developed to consider the more challenging problem of multiple broadband source location. A variety of least-squares modelling methods provide viable means for overcoming the difficulties of coherent sources. As we have seen above, in the standard LLS approach, the measured sensor signals are modelled as linear combinations of the source steering vectors.

Cadzow [21] presents a method that models the signal eigenvectors. These are linear combinations of the steering vectors instead of the sensor signals, which introduces a smoothing effect and decreases the computational cost, while the use of the VP approach produces significant additional computational savings. As Roy and Kailath [100] indicate, VP-type algorithms were considered too expensive until fairly recently, thus justifying the use of the simplified SVD-based ones. However, the increasing power of modern computers has rendered some of those arguments and simplified methods obsolete, especially in low-signal-to-noise-ratio situations, where they do not work well.

Friedlander [38] has analysed the sensitivity of the maximum-likelihood method for the problem above. This is a separable problem and the sensitivity study involves the differentiation formulae of [45]. This analysis is valuable because the fast algorithms require a knowledge of the antenna array that is hard to come by in real situations, and thus have not been used as often as they deserve.

Talwar *et al* [120, 121] have considered the problem of estimating co-channel digital signals using an antenna array when the spatial response of the array is unknown. Traditional techniques, such as MUSIC or ESPRIT, are dependent on the reliability of the array manifold. In the application the authors envision (mobile communications), the array manifold is poorly determined because of a highly variable propagation environment. They consider instead a block SNLLS approach, which is both fast and reliable.

Rao and Arun [98] discuss the problem of estimating closely spaced frequencies of multiple, superimposed sinusoids from noisy measurements as a NLLS problem. This variant of the problem discussed earlier has wide applications to radio-astronomy, interference spectroscopy, seismic data processing, and MR spectroscopy (which we discuss in detail later). Because of the cost of the computation, as compared to the simplified methods, SNLLS is only advisable at low signal-to-noise ratios.

Abel [1] applies separation-of-variables techniques to problems in underwater acoustic testing and global positioning system (GPS) navigation. The GPS problem involves using pseudo-range measures to determine a user's position. GPS consists of a set of satellites transmitting time-stamped signals; pseudo-ranges are formed by comparing satellite signal arrival times according to the user clock, to their transmission times according to GPS time, and then scaling by the propagation velocity. This can be stated as a SNLLS problem that is solved by a VP-like algorithm.

Zhou *et al* [144] consider the DOA problem for multiple moving targets by means of a passive array of sensors, a problem of great interest in communications, air traffic control, and tactical and strategical defence operations. In satellite and personal communication systems it is also advantageous to deploy sensor arrays to reject undesired signals. The classical techniques mentioned above deteriorate rapidly in the presence of moving targets, since they provide poor resolution because of the spread array spatial spectrum caused by the target motion. This deterioration increases with the number of sensors. Zhou *et al* propose a maximum-likelihood algorithm, where the target motion is assumed to be locally linear, which helps eliminate the spread spectrum effects and provides accurate target dynamical state estimates. Since they use the array signal model for an array of omnidirectional sensors, their approach leads to a separable problem that is solved by a VP method.

Liang *et al* [69] consider a two-stage hybrid approach for separate co-channel interference reduction and intersymbol interference equalization in a slow Raleigh fading channel, by using a space–time filter. This design involves an optimization problem that is recognized as a (novel) separable problem and it is dealt with by VP techniques.

Lilleberg *et al* [70] from Nokia Mobile Phones consider a near–far resistant iterative algorithm for multi-user signature sequence delay estimation. VP is used to separate the delay and data to be estimated, obtaining a so-called blind maximum-likelihood estimator that does not require any knowledge of user amplitudes and data.

Heredia and Arce [50] have considered the splitting of a signal into a set of multi-level components as a SNLLS problem. They use as a comparative example a system identification problem for wave propagation through a nonlinear multi-layer channel, where they test the new concepts against linear, Volterra, and neural network alternatives. They show that the realization of piecewise linear filters with unknown thresholds leads to a SNLLS problem. In the test problem they verify that the new approach can cope with the difficulties of the problem that trip the Volterra and neural network approaches.

Borden [15] considers the problem of obtaining accurate and high-resolution images of complex targets from radar scattering. Simply matching templates is not adequate when the radar images contain erroneous or surplus information. The remedy here is to try to locate some salient features of the scattering volume. In noisy and data-limited environments, the best that can be done is to try to locate the position and strength of scattering centres. Even that process may be difficult in real-life situations. Borden derives a model that includes noise and eventually reduces to a fairly complicated separable model which he solves by VP. He develops his own implementation of the reduction and optimization using Householder transformations and gives numerical results for some synthetic data. He concludes that this problem, previously intractable, can now be efficiently and robustly solved by taking advantage of separability. The method is robust against noise contamination and displays good super-resolution capabilities. He also indicates that by using truncated SVDs one could also take care of the problem of automatically determining the correct number of scatterers, in the case where that number is overestimated in the model.

Baum *et al* [8] review the singularity expansion method (SEM) for quantifying the transient electromagnetic scattering from targets illuminated by pulsed EM radiation. The SEM theory suggests that the late-time scattered field of a target, interrogated by pulsed EM radiation, can be represented as a sum of natural-resonance modes. Since the excitation-independent natural frequencies depend upon the detailed size and shape of the target, then the full complement of those frequencies is unique to a specified target and provides a potential basis for its identification. The first efforts to extract such natural frequencies from measured target pulse responses were based on Prony's method. However, in the practical low-signal-to-noise-ratio environment in which this inverse problem occurs, only one or a few modes could be extracted

reliably using that inherently unstable algorithm. Although several efforts have improved the reliability of Prony-based methods, realistic problems require a nonlinear approach, and since the problem is separable, VP has found another good application in the radar cross-section identification business.

Robertazzi and Schwartz [99] consider the problem of applying Kalman filters to nonlinear regression models. Their idea is to process the data offline in a random instead of a causal manner. They argue that the standard approach processes the data as they arrive and only provides sub-optimal solutions. They use as an example a separable regression problem and its solution through VP to show how bad such a sub-optimal solution can be.

In some recent publications, Escovar and Suaya [35, 36] discuss some difficult problems that arise in the design of cutting-edge microcircuits when the clock speed runs over the 1 GHz regime. Critical signals, such as clocks and high-speed data-buses, have lengths and transition times comparable to the transit time of light propagation in the medium over those distances. Optimization methods applied to the layout problem based on rigorous solution to Maxwell's equations in the media provide an appealing approach to the design of high-speed data lines, which are the critical components on present-day microprocessors. At certain stages of this work, separable nonlinear models appear and they are handled with VARPRO. Suaya [119] also considers the role of resistance, inductance and capacitance coupling, and finds again a good use for VARPRO.

14. Differential equations and dynamical systems

Wikström [138] has considered in detail the problem of parameter estimation associated with ordinary differential equations arising in chemical kinetics, theoretical biology, and population dynamics in ecology. She shows how some of these problems can be profitably attacked via constrained SNLLS techniques. A MATLAB toolbox has also been developed.

The parameter estimation problem considered by Wikström is the following:

$$\min_{\boldsymbol{y}(t,\boldsymbol{k}),\boldsymbol{k}} \tfrac{1}{2}\sum_{i=1}^{M}(\tilde{\boldsymbol{y}}_i - \boldsymbol{y}(t_i,\boldsymbol{k}))^{\mathrm{T}}\boldsymbol{W}_i(\tilde{\boldsymbol{y}}_i - \boldsymbol{y}(t_i,\boldsymbol{k})), \tag{4}$$

$$\text{s.t.} \qquad \dot{\boldsymbol{y}}(t,\boldsymbol{k}) = \boldsymbol{G}(t,\boldsymbol{y}(t))\boldsymbol{k}, \qquad \boldsymbol{y}(t_0) = \boldsymbol{y}_0, \tag{5}$$

where $\boldsymbol{W}_i$ are diagonal weight matrices, $\tilde{\boldsymbol{y}}_i$ are observed vector values of the solution of the differential equation $\boldsymbol{y}(t_i,\boldsymbol{k})$ at the times t_i, $\boldsymbol{G}(t,\boldsymbol{y}(t))$ is a matrix-valued function, and the elements of the vector $\boldsymbol{k}$ are the parameters to be determined. By numerically integrating the differential equations and considering the discrete system in matrix form, Wikström obtains the following equality-constrained separable NLLS problem:

$$\min_{\boldsymbol{y},\boldsymbol{k}} \tfrac{1}{2}\|\boldsymbol{W}^{1/2}(\tilde{\boldsymbol{y}} - \boldsymbol{y})\|^2, \tag{6}$$

$$\text{s.t.} \qquad \boldsymbol{y}_i = \boldsymbol{y}_0 + \boldsymbol{H}_i(\boldsymbol{y})\boldsymbol{k}, \qquad \boldsymbol{y}(t_0) = \boldsymbol{y}_0, \tag{7}$$

where $\boldsymbol{H}_i(\boldsymbol{y})$ is the result of numerically integrating $\boldsymbol{G}(t,\boldsymbol{y}(t))$ from t_0 to t_i.

15. Environmental sciences and time series analysis

Bert Rust and collaborators have used VARPRO in a number of applications related to modelling time series in environmental applications [62, 103–105]. In their most recent and most comprehensive contribution, they consider the inverse modulation of global fossil fuel production $P(t)$ by variations in northern hemispheric temperatures $T(t)$. They propose a

number of new models using some additional data. As an example, a delay differential equation that they propose is

$$\mathrm{d}P/\mathrm{d}t = (\alpha - \beta[T(t) - T(t-\tau)]/\tau)P,$$

where α, β, and the lag parameter τ are to be determined by fitting. The function $T(t)$ is obtained as an optimal smoothing cubic spline fit of the temperature data provided by Jones *et al* [55], where the smoothing parameter was chosen to minimize the generalized cross-validation statistics [42]. All the fittings in the paper are done with INVAR, the interactive version of VARPRO at NIST.

See also [61] for an application that strips a time series of its trend, which can be a combination of polynomial, exponential, autoregressive, and sinusoidal terms. As an illustration, a purely mathematical fit to the Mauna Loa carbon dioxide monthly averages from 1958 to 1976 is performed. The following model is fitted using VARPRO:

$$\begin{aligned} y(t) = a + d\mathrm{e}^{\alpha t} &+ f_1 \sin[2\pi/12(t+\gamma_i)] + f_2 \sin[2\pi/6(t+\gamma_2)] \\ &+ \rho_1 \sin[2\pi/\tau_1(t+\theta_1)] + \rho_2 \sin[2\pi/\tau_2(t+\theta_2)]. \end{aligned} \tag{8}$$

16. Robotics and vision

In [122] Tomasi and Shi consider the problem of determining the direction of heading from image deformations, which is important in applications to robotics and vision. They show how to eliminate the unknown depth values by using VP. As observed earlier in other applications, they report better-defined minima for the reduced functional, leading to a more reliable solution. The method also degrades gracefully with increasing uncertainty.

In related work, Chiuso *et al* [24] consider the problem of estimating spatial properties of a three-dimensional scene from the motion of its projection onto a two-dimensional surface, such as the retina. They address this problem, which has been studied for more than 20 years, from the point of view of an engineer trying to use conventional algorithms that have been tested on controlled sequences of images and do not behave as advertised in real-life applications. They indicate that the problem is noise and show how to handle it in a robust manner. Another problem is the trade-off between optical flow/feature tracking, which requires image motion of a few pixels to work well. This makes the problem of getting structure from motion very hard, since then the average image motion is comparable with the localization error. Again, a separation-of-variables argument allows them to simplify the solution and provides a more robust alternative in the presence of noise.

Radke *et al* [96] consider the problem of estimating projective transformations associated with standard problems in image processing and computer vision. The estimation problem leads to the minimization of a nonlinear functional of eight parameters, which, through a separation argument, can be reduced to a problem in only two variables.

Gardner and Milanfar [39] consider several algorithms for reconstructing convex bodies from brightness functions. One of the algorithms results in a bound-constrained SNLLS problem. Their current implementation uses the optimization toolbox from MATLAB, without taking advantage of the separability, and results in an algorithm that is too slow to be used in three dimensions. The authors conjecture that they could do much better by taking into account the separability.

Yen and Petzold have considered in a series of papers [141–143] a coordinate-splitting formulation of the equations of motion for multi-body systems, which is very effective when solving certain nonlinear, highly oscillatory problems. The problem leads to an index-3

differential–algebraic system, involving generalized coordinates $\boldsymbol{q}$ and Lagrange multipliers λ:

$$\begin{aligned} &\boldsymbol{M}(\boldsymbol{q})\ddot{\boldsymbol{q}} - \boldsymbol{f}(\boldsymbol{q}, \dot{\boldsymbol{q}}, t) + \boldsymbol{G}(\boldsymbol{q})^{\mathrm{T}}\lambda = 0, \\ &\boldsymbol{g}(\boldsymbol{q}) = \boldsymbol{0}. \end{aligned} \tag{9}$$

It is possible to eliminate the Lagrange multipliers by choosing an appropriate annihilation matrix $\boldsymbol{P}(\boldsymbol{q})$, such that $\boldsymbol{P}(\boldsymbol{q})\boldsymbol{G}^{\mathrm{T}}(\boldsymbol{q}) = \boldsymbol{0}$. Writing (9) in first-order form and performing some additional manipulations yields a stabilized index-2 system. Premultiplying this system by the matrix $\boldsymbol{P}(\boldsymbol{q})$ results in the system

$$\begin{aligned} &\boldsymbol{P}(\boldsymbol{q})(\dot{\boldsymbol{q}} - \boldsymbol{v}) = \boldsymbol{0}, \\ &\boldsymbol{P}(\boldsymbol{q})(\boldsymbol{M}(\boldsymbol{q})\dot{\boldsymbol{v}} - \boldsymbol{f}(\boldsymbol{v}, \boldsymbol{q}, t)) = \boldsymbol{0}, \\ &\boldsymbol{G}(\boldsymbol{q})\boldsymbol{v} = \boldsymbol{0}, \\ &\boldsymbol{g}(\boldsymbol{q}) = \boldsymbol{0}. \end{aligned} \tag{10}$$

The results on differentiation of projectors in [45] are now combined with discretization via a kth-order BDF formula to yield a system of nonlinear difference equations.

Dutre *et al* [31] discuss parallel kinematic structures encountered in the design of robotic manipulators. They consider an analytical description of the velocity closure equations and their Jacobian matrix, i.e., the linear mapping from driving joint velocities to end effector velocities or twist. Since available numerical methods are always more efficient than this analytic approach, the justification for the study is to provide more geometrical insights on the problem. The analytic expressions of the Jacobian and higher derivatives also open up new approaches in dynamics, acceleration analysis, calibration, and manipulability assessment of parallel robotic manipulators. After deriving the analytical velocity closure and the dependence matrix, the authors proceed to the calculation of the Jacobian matrix. The time derivative of the dependence matrix requires the differentiation of a pseudoinverse and the results of [45] are used to obtain closed analytical expressions.

17. Medical and biological imaging

Magnetic resonance (MR) in liquids and solids was discovered over 50 years ago by Ed Purcell (Harvard) and Felix Bloch (Stanford), who shared the Nobel Prize in Physics in 1952 for their work. MR is now a fundamental analytical tool in synthetic chemistry, plays an important role in biomedical research [22], and it has revolutionized modern radiology and neurology. Its applications outside the laboratory or medical clinic are as diverse as oil-well imaging, food analysis, and the detection of explosives.

A typical MR experiment involves placing the sample under study in a strong magnetic field, which forces the magnetic moments or spins of all the nuclei in the sample to line up along the main applied field and precess around this direction. The spins precess at the same frequency but with random phases. Pulses of radio-frequency (RF) magnetic fields are then applied that disturb the spin alignments but make the phases coherent and detectable. As this state precesses in the magnetic field, the spins emit radio-frequency radiation that can be analysed to reveal the structural, chemical, and dynamical properties of the sample. The idea of applying strong RF pulses is due to H C Torrey and E L Hahn. Higher and higher magnetic fields of the order of teslas have been used in order to increase the sensitivity of the method until just recently, when researchers in the USA and Germany showed that MR spectroscopy can be performed with fields in the microtesla range [74], by pre-polarizing the nuclei and using a superconducting quantum interference device (SQUID) to detect individual flux quanta. In a magnetic field of 1.8 μT the researchers observe proton MR in a liquid sample at about 100 Hz with an astonishing degree of sensitivity.

In vivo MR spectroscopy has the strongest connection and owes most to the VP methodology. In 1988, van der Veen *et al* [130], a group of researchers at Delft University and Phillips Medical Systems in The Netherlands, published a very influential paper on the accurate quantification of *in vivo* MRS data, using the VP method and prior knowledge (constraints). According to one of the authors, as of August 2002 this paper had had more than 190 citations. The problem here is to fit MR spectra in the time domain, using models whose parameters have physical significance. An example would be a linear combination of exponentially damped sinusoids, but other types of nonlinear function are also considered.

van der Veen *et al* consider MRS measurements of human calf muscle and human brain tissue and compare the FFTs of the original data with a linear prediction and SVD decomposition, and VARPRO with and without prior knowledge. The best results are obtained with VARPRO plus prior knowledge, which was difficult or impossible to impose on previous simpler approaches. They also list as desirable features the fact that starting values of the amplitudes of the spectral components are not required and that there are no restrictions on the form of the model functions.

Some years later, a group in Leuven, Belgium [125–127], made a comparative study for this problem, including artificial noise and also using prior knowledge. They examined one of the datasets considered earlier by van der Veen *et al*, added different levels of white Gaussian noise: small (5%), medium (15%), and large (25%), and considered 300 runs per level in a Monte Carlo simulation.

The model, without prior knowledge, involves 256 complex data points and 11 exponentials, for a total of 44 parameters divided equally between linear and nonlinear. In order to obtain good initial values, the time signal was Fourier transformed and displayed. Then, an interactive peak-picking was performed. This provided good initial values for the frequency and the damping of each peak. When the full functional is minimized, one also needs initial values for the linear parameters, and these are obtained by solving the LLS problem obtained by evaluating the nonlinear part at the chosen initial values. As observed in our original paper [45], this is much better than taking arbitrary values.

They first consider three different optimization methods for the full and the reduced problem, using VP with the Kaufman improvement: the original VARPRO with the LM implementation, a secant-type code NL2SOL [29], and LMDER, a modern LM implementation from MINPACK [77]. According to their results, NL2SOL with the separated functional seems to be the most reliable, even for large levels of noise. MINPACK's routine is almost as reliable and systematically faster. We should indicate here that the reported average times for this problem are about one minute on a SUN ULTRA2 (200 MHz), so in current and future platforms the differences in performance are negligible.

The most important conclusion drawn here is that if a VARPRO-type code is to be used, it should include the Kaufman simplification and should take advantage of the advances in numerical optimization. The MINPACK NLLS solver LMDER or the Gay–Kaufman implementation are good candidates for replacement codes. This has been confirmed in [129], where the author, in the process of developing an object-oriented system for the analysis of *in vivo* MR signals, made similar comparisons and came to the conclusion that VP did not help if used with an external NLLS code, but it was quite efficient if properly implemented within a modern solver such as Gay and Kaufman's NSG code.

The Vanhamme *et al* study continues to consider three different methods for obtaining starting values: HSVD [66], a fully automatic parameter estimation method that combines a state-space approach with SVDs; pick 1, the peak-picking method described above; and pick 2, a more careful (and time-consuming) version of pick 1. Since the influence of these different initial value choices seems to be method independent, only results for MINPACK are

presented. The conclusions are that the desirable automatic procedure works very well for low and medium noise levels, but not as well for high levels of noise. The procedures pick 1 and pick 2 are similar in performance and reliability, with a slight edge for pick 2.

Finally, they consider the effect of using prior knowledge about the problem. In prior 1, the number of linear variables is reduced to 11 by noting that all the peaks have a phase of 135°. In prior 2 all the known prior knowledge is used to eliminate variables, obtaining a problem with 5 linear and 12 nonlinear parameters. Results are given for MINPACK only, and here with prior 1, VP has a slight edge in performance for small and medium noise levels, although for medium and high levels of noise there is a deterioration in the reliability. For prior 2, VP is consistently more reliable than the full functional approach, with a slight penalty in performance (under 8%,) which has now come down to under 10 s of CPU time.

As a conclusion to this study, the authors suggested using the full functional instead of the reduced one and they proceeded to write their own solver AMARES to do that. This new solver also includes some other features special to the problem that were not present in [130]. The study of van Leeuwen indicates that with the Gay and Kaufman implementation of VP, the balance would clearly tilt in favour of the reduced functional, both in speed and reliability.

Both VARPRO and AMARES are currently offered in the MRUI system [81]. See also [128] for a recent review article.

Recently, de Beer *et al* [28], in a multi-centre study (done in the context of the European Union BIOMED 1 Concerted Action No PL 920432) with the purpose of reducing the variations in quantitative results of the participating *in vivo* MR groups, worked with the same data-analysis protocol. A central part of that protocol was the VARPRO procedure.

Sala *et al* [107] study respiratory rehabilitation, including lower-limb exercise training, as part of the management for patients with chronic obstructive pulmonary disease (COPD). In order to understand better the physiologic mechanisms underlying the beneficial consequences of training, they made MRS measurements in the quadriceps of 13 patients with COPD and 8 healthy patients, before and during a period of exercises. Free induction decays were analysed in the time domain using the VARPRO software, by fitting single Lorentzian functions for the resonance frequencies and amplitudes. The individual half-time of phosphocreatine recovery was fitted in the time domain to an exponential function. The authors then present a detailed discussion of the physiological conclusions stemming from this study.

Barone [4] has considered the problem of fast deconvolution in the case of noisy data. This is a general problem that occurs in many applications, such as the recovery of images observed in noise through a linear system representing a physical measuring device. Speed of computation is the driving factor for this study, especially when large amounts of data are present. The problem is shown to be separable, and using a VP approach is considerably more efficient than a previous simulated annealing method used on the full functional. Three examples are considered, including the restoration of a MR image of a human brain that includes Gibbs oscillations.

Mosher *et al* [3, 79, 80] have considered the problem of modelling the spatio-temporal neuromagnetic field or magneto-encephalogram (MEG) produced by neural activity of the brain. A popular model for the neural activity produced in response to a given sensory stimulus is a set of current dipoles, where each dipole represents the primary current associated with the combined activation of a large number of neurons located in a small volume of the brain. An important problem in the interpretation of MEG data is the localization of these neural current dipoles. The key concept here is that given any arbitrary static current distribution, the magnetic field can be obtained by using the Biot–Savart law. The general model of p_r rotating dipoles and p_f fixed dipoles results in a problem with $3r$ unknown location parameters ($r = 2p_r + p_f$) and p_f unknown constrained moment parameters. Using m SQUID biomagnetometers, n samples

are collected to form a spatio-temporal data matrix $\boldsymbol{F}$. The problem is then

$$\min_{\boldsymbol{L},\boldsymbol{M},\boldsymbol{S}} \|\boldsymbol{F} - \boldsymbol{H}(\boldsymbol{L}, \boldsymbol{M})\boldsymbol{S}\|_F^2,$$

where the model $\boldsymbol{H}(\boldsymbol{L}, \boldsymbol{M})\boldsymbol{S}$ involves the gain matrix $\boldsymbol{H}$, the time series $\boldsymbol{S}$, the moment orientation matrix $\boldsymbol{M}$, and $\|\cdot\|_F$ is the Frobenius norm. This separable model is one of the most ambitious and large-scale applications of the VP approach that we have seen to date and it is much too complicated to discuss here. We refer the reader to the original papers for more details.

Solving this problem directly would require finding five parameters for each of p dipoles at each point of n time snapshots, for an overall total of $5pn$ parameters. By considering the separability of the problem and using the VP algorithm one can reduce the dimensionality of the problem by factoring out the linear moments. Given the complexity of the model, these authors recommend the use of a derivative-free approach, such as a Nelder–Mead-type algorithm. The case of rotating dipoles is most favourable, since there the nonlinear parameters are not time dependent, and therefore the problem can be reduced to finding only $3p$ location parameters, while deferring the calculation of the $2pn$ linear parameters, to be obtained *a posteriori* with a single LLS solve. SVDs are recommended to optimize the operation count and, more importantly, to regularize the problem in the ill-conditioned or even rank-deficient case.

Silva *et al* [116] also consider realistic head models, instead of spherical ones, applied to the neural source localization problem associated with epileptic foci from scalp EEG data. Usually, accurate head models are constructed from MR images of the brain, skull, and scalp. Since this is an expensive and time-consuming procedure, the authors have developed a parametrized realistic head model that can be adjusted to a particular patient by using external distances between anatomic landmarks. To localize the neural sources, an inverse problem, one needs to be able to solve the direct problem of determining the scalp potentials at the electrodes, generated by a given source. The authors use the boundary element method (BEM) as their forward solver and the VP method for the inverse problem. Extensive validation studies are reported, together with application to three clinical cases.

Westwick and Kearney [136] consider the identification of a Hammerstein model of a biological system using VP. Again they cite the simplification arising from the reduced dimensionality and indicate that this work can be extended to other commonly used model structures, such as the Wiener cascade.

18. Conclusions and further developments

We have reviewed the ideas behind the VP method as an efficient solution to SNLLS problems. We have then briefly visited a significant number of very different applications that indicate how prevalent these problems are and how successful the method has been during the last thirty years, including very recent results.

One of the most exciting new applications is to the training of some types of neural network. It would be interesting to extend the method to large-scale problems, i.e., those with too many parameters to be solved directly by a VP-type code. By their nature, those problems will be sparse, although sparsity may be destroyed by the elimination of the linear variables.

What we have learned from the original results and the accumulated experience of a large number of researchers in very diverse disciplines is that the VP approach is chosen not only because it is more efficient than solving the unreduced problem, but also because it is a good preconditioner and leads to problems with better-defined minima.

A public domain version of VARPRO with the Kaufman modification and using modern optimization technology would also be welcome, together with one that uses a truncated SVD approach to regularize ill-conditioned or over-parametrized models.

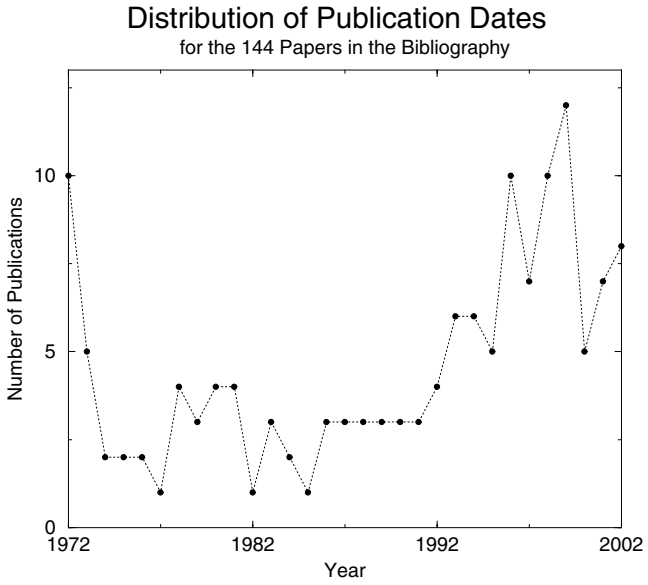

Figure 1. Distribution of cited references per year.

19. Distribution of references

We show in a graph (figure 1) the distribution of cited references per year. We observe an increase of interest in the last decade, although that observation may be tainted by the fact that there are more items of recent vintage available on the Internet.

Acknowledgments

We thank Dr Burt Rust for his valuable comments, for promoting the use of VARPRO through the institutions in which he worked, and for providing a number of interesting papers for our collection.

We also thank Drs Linda Kaufman and Dirk van Ormondt and his colleagues for their role in developing VARPRO through the years.

Special thanks for Professor Michael Saunders, for his careful reading of the manuscript and perceptive comments that have helped improve an earlier version quite considerably.

The work of GG was partially supported by NSF Grant CCR9971010. The work of VP was partially supported by NSF SBIR Grant 0124874.

References

[1] Abel J S 1994 A variable projection method for additive components with application to GPS *IEEE Trans. Aerosp. Electron. Syst.* **30** 928–30

[2] Andersson L, Holmström K and Ruhe A 1987 Complex formation constants. A challenging data fitting problem from solution chemistry *Algorithms for Approximation* vol 10 (Oxford: Clarendon) pp 557–72

[3] Baillet S, Mosher J C and Leahy R M 2001 Electromagnetic brain mapping *IEEE Signal Process. Magn.* 14–30

[4] Barone P 1999 Fast deconvolution by a two-step method *SIAM J. Sci. Comput.* **21** 882–99

[5] Bates D M 1983 The derivative of $|\boldsymbol{X}^{\mathrm{T}}\boldsymbol{X}|$ and its uses *Technometrics* **25** 373–6

[6] Bates D M and Watts D G 1987 A generalized Gauss–Newton procedure for multi-response parameter estimation *SIAM J. Sci. Stat. Comput.* **8** 49–55

[7] Bates D M and Lindstrom M J 1986 Nonlinear least squares with conditionally linear parameters *Proc. Annual Mtg* (Chicago, IL: American Statistical Association) pp 152–7
[8] Baum C E, Rothwell E J, Chen K-M and Nyquist D P 1991 The singularity expansion method and its application to target identification *Proc. IEEE* **79** 1481–92
[9] Bishop C M 1995 *Neural Networks for Pattern Recognition* (Oxford: Oxford University Press)
[10] Böckmann C 1995 A modification of the trust-region Gauss–Newton method for separable nonlinear least squares problems *J. Math. Syst. Estim. Control* **5** 1–16
[11] Beece D, Bowne S F, Czege J, Eisenstein L, Frauenfelder H, Good D, Marden M C, Marque J, Ormos P, Reinisch L and Yue K T 1981 The effect of viscosity on the photocycle of bacteriorhodopsin *Photochem. Photobiol.* **33** 517–22
[12] Beylkin G 1998 On applications of unequally spaced fast Fourier transforms *Mathematical Geophysics Summer School (Stanford, CA)*
[13] Bienvenu G and Kopp L 1980 Adaptivity to background noise spatial coherence for high resolution passive methods *Proc. IEEE Acoust. Speech Signal Process.* 307–10
[14] Björck Å 1996 *Numerical Methods for Least Squares Problems* (Philadelphia, PA: SIAM)
[15] Borden B 2001 Radar scattering centre localization by subspace fitting *Inverse Problems* **17** 1483–91
[16] Bramley R and Winnicka B 1996 Solving linear inequalities in a least squares sense *SIAM J. Sci. Comput.* **17** 275–86
[17] Brent R 1973 *Algorithms for Finding Zeros and Extrema of Functions Without Calculating Derivatives* (Englewood Cliffs, NJ: Prentice-Hall)
[18] Bruls J, Chou C T, Haverkamp B R J and Verhaegen M 1999 Linear and non-linear system identification using separable least squares *Eur. J. Control* **5** 116–28
[19] Burg J P 1967 Maximum entropy spectral analysis *Soc. Explor. Geophys.* (reprinted 1978 *Modern Spectral Analysis* ed D G Childers (New York: IEEE) pp 34–41)
[20] Byers R, He C and Mehrmann V 1998 Where is the nearest non-regular pencil? *Linear Algebr. Appl.* **285** 81–105
[21] Cadzow J A 1990 Multiple source location—the signal subspace approach *IEEE Trans. Acoust. Speech Signal Process.* **38** 1110–25
[22] Callaghan P 2002 Spectroscopy scales new peaks *Phys. World* **15** 23–5
[23] Capon J 1969 High-resolution frequency–wavenumber spectrum analysis *Proc. IEEE* **57** 1408–18
[24] Chiuso A, Brockett R and Soatto S 2000 Optimal structure from motions: local ambiguities and global estimates *Int. J. Comput. Vision* **39** 195–228
[25] Conn A R, Vicente L N and Visweswariah C 1999 Two-step algorithm for nonlinear optimization with structured applications *SIAM J. Opt.* **9**
[26] Corradi C and Stefanini L 1978 Computational experience with algorithms for separable nonlinear least squares problems *Calcolo* **15** 317–30
[27] Corradi C 1981 A note on the solution of separable nonlinear least-squares problems with separable nonlinear equality constraints *SIAM J. Numer. Anal.* **18** 1134–8
[28] de Beer R, Barbiroli B, Gobbi G, Knijn A, Kugel H, Langenberger K W, Tkac I and Topp S 1998 Multicentre 1H MRS of the human brain addressed by one and the same data-analysis protocol *Magn. Reson. Imag.* **16** 1107–11
[29] Dennis J E, Gay D M and Welsch R E 1981 An adaptive nonlinear least-squares algorithm *ACM Trans. Math. Softw.* **7** 348–68
[30] Dierckx P 1993 *Curve and Surface Fitting with Splines* (Oxford: Clarendon)
[31] Dutre S, Bruyninckx H and de Schutter J 1997 The analytical Jacobian and its derivative for a parallel manipulator *IEEE Conf. on Robotics and Automation (Albuquerque, NM, 1997)* pp 2961–6
[32] Edrissi H E, Verhaegen M H G, Haverkamp B R J and Chou C T 1998 Off- and on-line identification of discrete time LTI systems using separable least squares *Proc. Conf. Decision and Control (Tampa, FL)* (New York: IEEE) (cdc98 WM07)
[33] Eriksson J, Wedin P-Å and Gulliksson M 1997 Regularization methods for nonlinear least squares problems: Part I. Exact rank-deficiency *Technical Report* UMINF 96.04 (Umea University, Sweden)
[34] Eriksson J, Gulliksson M, Lindström P and Wedin P-Å 1998 Regularization tools for training feed-forward neural networks *J. Optim. Methods Softw.* **10** 49–69
[35] Escovar R 2001 Optimal design of leading edge microprocessors. ENCIMAG INPG *MSc Thesis* University of Grenoble, France
[36] Escovar R and Suaya R 2002 Transmission line design of clock trees *Proc. IEEE ACM Int. Conf. on CAD* pp 334–40
[37] Francos J M, Narasimhan A and Woods J W 1996 Maximum likelihood parameter estimation of the harmonic

evanescent and purely indeterministic components of discrete homogeneous random fields *IEEE Trans. Inf. Theory* **42** 916–30

[38] Friedlander B 1990 Sensitivity analysis of the maximum likelihood direction-finding algorithm *IEEE Trans. Aerosp. Electron. Syst.* **26** 953–68

[39] Gardner R J and Milanfar P 2003 Reconstruction of convex bodies from brightness functions *Discrete Comput. Geom.* **29** 279–303

[40] Gay D M 2002 Port Library: http://netlib.bell-labs/netlib/master/readme.html

[41] Gay D M and Kaufman L 1991 Tradeoffs in algorithms for separable nonlinear least squares *Proc. World Congress on Computations in Applied Mathematics* ed R Vichnevetsky and J J H Miller (Dublin: Criterion) pp 157–8

[42] Golub G H, Heath M and Wahba G 1979 Generalized cross-validation as a method for choosing a good ridge parameter *Technometrics* **21** 215–23

[43] Golub G H and LeVeque R 1979 Extensions and uses of the variable projection algorithm for solving nonlinear least squares problems *Proc. Army Numerical Analysis and Computers Conf.* (Washington, DC: US Army Research Office) pp 1–12

[44] Golub G H and Pereyra V 1972 The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate *Stanford University Computer Science Department Report* STAN-CS-72-261

[45] Golub G H and Pereyra V 1973 The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate *SIAM J. Numer. Anal.* **10** 413–32

[46] Golub G H and Pereyra V 1976 Differentiation of pseudoinverses, separable nonlinear least squares and other tales *Generalized Inverses and Applications* ed M Z Nashed (New York: Academic) pp 303–24

[47] Golub G H and Van Loan C F 1989 *Matrix Computations* 3rd edn (Baltimore, MD: John Hopkins University Press)

[48] Gulliksson M and Wedin P-Å 2000 The use and properties of Tikhonov filter matrices *SIAM J. Matrix Anal. Appl.* **22** 276–81

[49] Guttman I, Pereyra V and Scolnik H D 1973 Least squares estimation for a class of nonlinear models *Technometrics* **15** 209–18

[50] Heredia E A and Arce G R 1996 Piecewise linear systems modeling based on a continuous threshold decomposition *IEEE Trans. Signal Process.* **44**

[51] Holmström K 1998 Constrained separable NLLS algorithms for chemical equilibrium analysis *Proc. 5th Mtg Nordic Section of Mathematics Programming Society (Molde University, Norway)* ed A Lökketangen

[52] Holmström K 1999 The TOMLAB optimization environment in MATLAB *Adv. Model. Optim.* **1** 47–69

[53] Holmström K and Pettersson J 2002 A review of the parameter estimation problem of fitting positive exponential sums to empirical data *Appl. Math. Comput.* **126** 31–61

[54] Ingri N, Andersson I, Pettersson L, Andersson L, Yagasaki A and Holmström K 1996 LAKE—a program system for equilibrium analytical treatment of multimethod data, especially combined potentiometric and NMR data *Acta Chem. Scand.* **50** 717–34

[55] Jones P D 1988 Hemispheric surface air temperature variations: recent trends and update to 1987 *J. Clim.* **1** 654–9

[56] Jupp D L 1978 Approximation to data by splines with free knots *SIAM J. Numer. Anal.* **15** 328–43

[57] Kaufman L 1975 A variable projection method for solving separable nonlinear least squares problems *BIT* **15** 49–57

[58] Kaufman L and Pereyra V 1978 A method for separable nonlinear least squares problems with separable equality constraints *SIAM J. Numer. Anal.* **15** 12–20

[59] Kaufman L and Sylvester G 1992 Separable nonlinear least squares with multiple right-hand sides *SIAM J. Matrix Anal. Appl.* **13** 68–89

[60] Kaufman L, Sylvester G and Wright M H 1994 Structured linear least-squares problems in system identification and separable nonlinear data fitting *SIAM J. Optim.* **4** 847–71

[61] Kirk B L and Rust B W 1981 Inductive modeling of time series. A detrending approach *Environmetrics* vol 81 (Philadelphia, PA: SIAM)

[62] Kirk B L and Rust B W 1983 The solar cycle effect on atmospheric carbon dioxide levels *Weather and Climate Responses to Solar Variations* (Boulder, CO: Colorado Associated University Press)

[63] Koliha J J 2001 Continuity and differentiability of the Moore–Penrose inverse in C*-algebras *Math. Scand.* **88** 154–60

[64] Krogh F T 1974 Efficient implementation of a variable projection algorithm for nonlinear least squares problems *Commun. ACM* **17** 167–9

[65] Kumaresan R and Shaw A K 1988 Superresolution by structured matrix approximation *IEEE Trans. Antennas Propag.* **36** 34–44

[66] Kung S Y, Arun K S and Bhaskar Rao D V 1983 State-space and singular value decomposition based approximation methods for the harmonic retrieval problem *J. Opt. Soc. Am.* **73** 1799–811
[67] Lanzkron P J, Rose D J and Wilkes J T 1996 An analysis of approximate nonlinear elimination *SIAM J. Sci. Comput.* **17** 538–59
[68] Lee Y S and Baeck K K 1986 Basic set requirement for small components besides kinetic balance in relativistic self-consistent-field calculations of many electron systems *Bull. Korean Chem. Soc.* **7** 428–33
[69] Liang Jen-Wei, Chen J-T and Paulraj A J 1997 A two-stage hybrid approach for CCI-ISI reduction with space–time processing *IEEE Commun. Lett.* **1** 163–5
[70] Lilleberg J, Nieminen E and Latva-aho M 1996 Blind iterative multiuser delay estimator for CDMA *Proc. IEEE Int. Symp. on Personal Indoor and Mobile Radio Communications (PIMRC) (Taipei, Taiwan)* (New York: IEEE) pp 565–8
[71] Lozier R H, Bogomolni R A and Stoeckenius W 1975 *Biophys. J.* **15** 955–62
[72] Lukeman G G 1999 Application of the Shen–Ypma algorithm for separable overdetermined nonlinear systems *MSc Thesis* Dalhousie University, Halifax, NS, Canada
[73] Marque J and Eisenstein L 1984 Pressure effects on the photocycle of purple membrane *Biochemistry* **23** 5556–63
[74] McDermott R, Trabesinger A H, Muck M, Hahn E L, Pines A and Clarke J 2002 Liquid-state NMR and scalar couplings in microtesla magnetic fields *Science* **295** 2247
[75] Miller A J 2002 http://users.bigpond.net.au/amiller
[76] Moré J J 1978 The Levenberg–Marquardt algorithm: implementation and theory *Numerical Analysis* ed G A Watson (New York: Springer) pp 105–16
[77] Moré J J, Garbow B S and Hillstrom K E 1980 User Guide to MINPACK *Argonne National Laboratory Report* ANL-80-74 (Argonne, IL)
[78] Morrison D B 1998 Application of the Shen–Ypma algorithm for nonlinear systems with some linear unknowns *MSc Thesis* Dalhousie University, Halifax, NS, Canada
[79] Mosher J J, Lewis P S and Leahy R M 1992 Multiple dipole modeling and localization from spatio-temporal MEG data *IEEE Trans. Biomed. Eng.* **39** 541–57
[80] Mosher J J and Leahy R M 1998 Recursive MUSIC: a framework for EEG and MEG source localization *IEEE Trans. Biomed. Eng.* **4511** 1342–54
[81] MRUI 2002 Magnetic Resonance User Interface at http://carbon.usb.es/mrui
[82] Naressi A, Couturier C, Devos J M, Castang I, de Beer R and Graveron-Demilly D 2001 Java-based graphical user interface for the MRUI quantitation package *Comput. Biol. Med.* **31** 269–86
[83] Ngia L 2000 System modeling using basis functions and applications to echo cancellation *PhD Thesis* Chalmers Institute of Technology, Sweden (NNSP2000)
[84] Nielsen H B 2000 Separable nonlinear least squares *Technical Report* IMM-REP-2000-01 (Technical University of Denmark)
[85] Osborne M R 1975 Some special nonlinear least squares problems *SIAM J. Numer. Anal.* **12** 571–92
[86] Osborne M R and Smyth G K 1995 A modified Prony algorithm for exponential function fitting *SIAM J. Sci. Comput.* **16** 119–38
[87] Parks T A 1985 Reducible nonlinear programming problems *PhD Thesis* Rice University, Houston, TX
[88] Pati Y C and Krishnaprasad P S 1993 Analysis and synthesis of feedforward neural networks using discrete affine wavelet transformations *IEEE Trans. Neural Netw.* **4** 73–85
[89] Pereyra V 1967 Accelerating the convergence of discretization algorithms *SIAM J. Numer. Anal.* **4** 508–33
[90] Pereyra V 1967 Iterative methods for solving nonlinear least squares problems *SIAM J. Numer. Anal.* **4** 27–36
[91] Pereyra V 2000 Ray tracing methods for inverse problems *Inverse Problems* **16** R1–35
[92] Pereyra V and Scherer G 1973 Efficient computer manipulation of tensor products with applications to multidimensional approximation *Math. Comput.* **27** 595–605
[93] Pereyra V and Zadunaisky P 1965 On the convergence and precision of a process of successive differential corrections *Proc. IFIP* **65**
[94] Prais M G and Wheeler J C 1986 Improved modified-moment-singularity method *Phys. Rev.* A **33** 1233–45
[95] Prells U and Friswell M I 1999 Application of the variable projection method for updating models of mechanical systems *J. Sound Vib.* **225** 307–25
[96] Radke R, Ramadge P and Echigo T 2001 Efficiently estimating projective transformations *Proc. IEEE Int. Conf. on Image Processing (2000); IEEE Trans. Pattern Anal. Mach. Intell.* submitted
[97] Rall J E and Funderlic R E 1980 Interactive VARPRO (INVAR) A nonlinear least squares program *Report* ORNL/CSD-55 (Oak Ridge National Laboratory, TN)
[98] Rao B D and Arun K S 1992 Model based processing of signals: a state space approach *Proc. IEEE* **80** 283–309

[99] Robertazzi T G and Schwartz S C 1989 On applying the extended Kalman filter to nonlinear regression models *IEEE Trans. Aerosp. Electron. Syst.* **25** 433–8
[100] Roy R and Kailath T 1989 ESPRIT—estimation of signal parameters via rotational invariance techniques *IEEE Trans. Acoust. Speech Signal Process.* **37** 984–95
[101] Ruhe A and Wedin P-Å 1980 Algorithms for separable nonlinear least squares problems *SIAM Rev.* **22** 318–37
[102] Rust B W, Leventhal M and McCall S L 1976 Evidence for a radioactive decay hypothesis for supernova luminosity *Nature* **262** 118–20
[103] Rust B W and Kirk B L 1978 Inductive modelling of population time series *Time Series and Ecological Processes* ed H H Shugart (Philadelphia, PA: SIAM) pp 154–92
[104] Rust B W and Kirk B L 1982 Modulation of fossil fuel production by global temperature variations *Environ. Int.* **7** 419–22
[105] Rust B W and Crosby F J 1994 Further studies on the modulation of fossil fuel production by global temperature variations *Environ. Int.* **20** 429–56
[106] Sagara N and Fukushima M 1984 A continuation method for solving separable nonlinear least squares problems *J. Comput. Appl. Math.* **10** 157–61
[107] Sala E, Roca J, Marrades R M, Alonso J, Gonzalez de Suso J M, Moreno A, Barbera J A, Nadal J, Luis de Jover, Rodriguez-Roisin R and Wagner P D 1999 Effects of endurance training on skeletal muscle bioenergetics in chronic obstructive pulmonary disease *Am. J. Respir. Crit. Care Med.* **159** 1726–34
[108] Schmidt R O 1986 Multiple emitter location and signal parameter estimation *IEEE Trans. Antennas Propag.* **34** 276–80
[109] Schutze T and Schwetlick H 1997 Constrained approximation by splines with free knots *BIT* **37** 105–37
[110] Schwetlick H 1992 Nonlinear parameter estimation: models, criteria, and algorithms *Numerical Analysis* ed D F Griffiths and G A Watson (New York: Longman–Wiley) pp 164–93
[111] Schwetlick H and Schutze T 1995 Least squares approximation by splines with free knots *BIT* **35** 1–23
[112] Scolnik H D 1971 On the solution of non-linear least squares problems *PhD Thesis* University of Zurich, Switzerland
[113] Scolnik H D 1972 On the solution of non-linear least squares problems *Proc. IFIP 71* (Amsterdam: North-Holland) pp 1258–65
[114] Seber G A F and Wild C J 1989 *Nonlinear Regression* (New York: Wiley–Interscience)
[115] Shen Y-Q and Ypma T J 1990 Solving $N+m$ equations with only m nonlinear variables *Computing* **44** 259–70
[116] Silva C, Almeida R, Oostendorp T, Ducla-Soares E, Foreid J P and Pimentel T 1999 Interictal spike localization using standard realistic head model: simulations and analysis of clinical data *Clin. Neurophysiol.* **110** 846–55
[117] Sjöberg J and Viberg M 1997 Separable non-linear least squares minimization—possible improvements for neural net fitting *IEEE Workshop in Neural Networks for Signal Processing (Amelia Island Plantation, FL)* (New York: IEEE)
[118] Stewart G W 1977 On the perturbation of pseudo-inverses, projections and linear least squares problems *SIAM Rev.* **19** 634–62
[119] Suaya R 2002 Interconnect modeling for high speed digital circuits. The role of RLC coupling *Proc. Int. Conf. on Modeling and Simulation of Microsystems* pp 722–5
[120] Talwar S 1996 Blind space-time algorithms for wireless communication systems *PhD Thesis* SCCM, Stanford University, CA
[121] Talwar S, Viberg M and Paulraj A 1994 Blind estimation of multiple co-channel digital signals arriving at an antenna array *IEEE SP Lett.* **1** 29–31
[122] Tomasi C and Shi J 1993 Direction of heading from image deformations *IEEE Conf. on Computer Vision and Pattern Recognition*
[123] Trosset M W 1997 Computing distances between convex sets and subsets of the positive semidefinite matrices *Technical Report* 97-3 (Rice University, Houston, TX)
[124] Van Loan C and Pitsianis N 1993 Approximation with Kronecker products *Linear Algebra for Large Scale and Real Time Applications* ed M S Moonen and G H Golub (Dordrecht: Kluwer) pp 293–314
[125] Vanhamme L, van den Boogaart A and Van Huffel S 1996 Fast and accurate parameter estimation of noisy complex exponentials with use of prior knowledge *Proc. EUPSICO-96 (Trieste, Italy)* (Amsterdam: Elsevier)
[126] Vanhamme L, van den Boogaart A and Van Huffel S 1997 Improved method for accurate and efficient quantification of MRS data with use of prior knowledge *J. Magn. Reson.* **129** 35–43
[127] Vanhamme L 1999 Advanced time-domain methods for nuclear magnetic resonance spectroscopy data analysis *PhD Thesis* University of Leuven
[128] Vanhamme L, Sundin T, Van Hecke P and Van Huffel S 2001 MR spectroscopic quantitation: a review of time domain methods *NMR Biomed.* **14** 233–46

[129] Van Leeuwen J P 1994 Object georiënteerd parameterschatten *MSc Thesis* Technical University of Delft
[130] van der Veen J W C, de Beer R, Luyten P R and van Ormondt D 1988 Accurate quantification of *in vivo* PNMR signals using the variable projection method and prior knowledge *Magn. Reson. Med.* **6** 92–8
[131] Wedin P Å 1973 Perturbation theory for pseudo-inverses *BIT* **13** 217–32
[132] Weigl K and Berthod M 1993 Neural networks as dynamical bases in function space *Report* No 2124 (INRIA, Sophia-Antipolis, France)
[133] Weigl K, Giraudon G and Berthod M 1993 Application of projection learning to the detection of urban areas in SPOT satellite images *Report* No 2143 (INRIA, Sophia-Antipolis, France)
[134] Weigl K and Berthod M 1994 Projection learning: alternative approach to the computation of the projection *Proc. European Symp. on Artificial Neural Networks (Brussels, Belgium)* pp 19–24
[135] Westwick D T, George K and Verhaegen M 1999 Nonlinear identification of automobile vibration dynamics *Proc. MED99: 7th Mediterranean Conf. on Control and Automation (Haifa, Israel)* pp 724–37
[136] Westwick D T and Kearney R E 2000 Identification of a Hammerstein model of the stretch reflex EMG using separable least squares *EMBS22; Proc. IEEE* (New York: IEEE)
[137] Wheeler J C, Prais M G and Blumstein C 1974 Analysis of spectral densities using modified moments *Phys. Rev.* B **10** 2429–47
[138] Wikstrom G 1997 Algorithms and software for the computation of parameters occurring in ODE-models *Thesis* Umeå University, Sweden
[139] Wold H and Lyttkens E 1969 Nonlinear iterative partial least squares (NIPALS) estimation procedures *Bull. ISI* **43** 29–51
[140] Wolfe C M, Rust B W, Dunn J H and Brown I E 1987 An interactive nonlinear least squares program *NBS Technical Note* 1238 (Gaithersburg, MD)
[141] Yen J and Petzold L R 1995 Convergence of the iterative methods for coordinate-splitting formulation in multibody dynamics *Technical Report* TR95-052 (Minnesota University, MN)
[142] Yen J and Petzold L R 1996 Computational challenges in the solution of nonlinear oscillatory multibody dynamical systems *Proc. Biennial Conf. on Numerical Analysis (Pittman Research Notes in Mathematics No 344)* ed D F Griffiths and G A Watson (Reading, MA: Addison-Wesley)
[143] Yen J and Petzold L R 1998 An efficient Newton-like iteration for the numerical solution of highly oscillatory constrained multibody dynamical systems *SIAM J. Sci. Comput.* **19** 1513–34
[144] Zhou Y, Yip P C and Leung H 1999 Tracking the direction-of-arrival of multiple moving targets by passive arrays: algorithm *IEEE Trans. Signal Process.* **47** 2655–66