Linköping Studies in Science and Technology
Thesis No. 1310

Algorithms for a Partially Regularized Least Squares Problem

Ingegerd Skoglund

Linköpings universitet
INSTITUTE OF TECHNOLOGY

Scientific Computing
Department of Mathematics
Linköpings universitet, SE–581 83 Linköping, Sweden

Linköping 2007

This is a Swedish Licentiate Thesis.
The Swedish Licentiate degree is an intermediate degree, acquired at a level higher than the Master's Degree but below the Doctor's degree.

Algorithms for a Partially Regularized Least Squares Problem

Department of Mathematics
Linköpings universitet
SE–581 83 Linköping
Sweden

ISBN 978-91-85831-93-7 ISSN 0280-7971 LIU-TEK-LIC-2007:19

Printed by LiU-Tryck, Linköping 2007

Abstract

Statistical analysis of data from rivers deals with time series which are dependent, e.g., on climatic and seasonal factors. For example, it is a well-known fact that the load of substances in rivers can be strongly dependent on the runoff. It is of interest to find out whether observed changes in riverine loads are due only to natural variation or caused by other factors. Semi-parametric models have been proposed for estimation of time-varying linear relationships between runoff and riverine loads of substances. The aim of this work is to study some numerical methods for solving the linear least squares problem which arises.

The model gives a linear system of the form $A_1x_1 + A_2x_2 + \eta = b_1$. The vector η consists of identically distributed random variables all with mean zero. The unknowns, x, are split into two groups, x_1 and x_2. In this model, usually there are more unknowns than observations and the resulting linear system is most often consistent having an infinite number of solutions. Hence some constraint on the parameter vector x is needed. One possibility is to avoid rapid variation in, e.g., the parameters x_2. This can be accomplished by regularizing using a matrix A_3, which is a discretization of some norm. The problem is formulated as a partially regularized least squares problem with one or two regularization parameters. The parameter x_2 has here a two-dimensional structure. By using two different regularization parameters it is possible to regularize separately in each dimension.

We first study (for the case of one parameter only) the conjugate gradient method for solution of the problem. To improve rate of convergence block-preconditioners of Schur complement type are suggested, analyzed and tested. Also a direct solution method based on QR decomposition is studied. The idea is to first perform operations independent of the values of the regularization parameters. Here we utilize the special block-structure of the problem. We further discuss the choice of regularization parameters and generalize in particular Reinsch's method to the case with two parameters. Finally the cross-validation technique is treated. Here also a Monte Carlo method is used by which an approximation to the generalized cross-validation function can be computed efficiently.

Sammanfattning

Vid analys av vattenprover tagna från t.ex. ett vattendrag betäms halten av olika ämnen. Dessa halter är ofta beroende av vattenföringen. Det är av intresse att ta reda på om observerade förändringar i halterna beror på naturliga variationer eller är orsakade av andra faktorer. För att undersöka detta har föreslagits en statistisk tidsseriemodell som innehåller okända parametrar. Modellen anpassas till uppmätta data vilket leder till ett underbestämt ekvationssystem. I avhandlingen studeras bl.a. olika sätt att säkerställa en unik och rimlig lösning. Grundidén är att införa vissa tilläggsvillkor på de sökta parametrarna. I den studerade modellen kan man t.ex. kräva att vissa parametrar inte varierar kraftigt med tiden men tillåter årstidsvariationer. Det görs genom att dessa parametrar i modellen regulariseras.

Detta ger upphov till ett minsta kvadratproblem med en eller två regulariseringsparametrar. I och med att inte alla ingående parametrar regulariseras får vi dessutom ett partiellt regulariserat minsta kvadratproblem. I allmänhet känner man inte värden på regulariseringsparametrarna utan problemet kan behöva lösas med flera olika värden på dessa för att få en rimlig lösning. I avhandlingen studeras hur detta problem kan lösas numeriskt med i huvudsak två olika metoder, en iterativ och en direkt metod. Dessutom studeras några sätt att bestämma lämpliga värden på regulariseringsparametrarna.

I en iterativ lösningsmetod förbättras stegvis en given begynnelseapproximation tills ett lämpligt valt stoppkriterium blir uppfyllt. Vi använder här konjugerade gradientmetoden med speciellt konstruerade prekonditionerare. Antalet iterationer som krävs för att lösa problemet utan prekonditionering och med prekonditionering jämförs både teoretiskt och praktiskt. Metoden undersöks här endast med samma värde på de två regulariseringsparametrarna.

I den direkta metoden används QR-faktorisering för att lösa minsta kvadratproblemet. Idén är att först utföra de beräkningar som kan göras oberoende av regulariseringsparametrarna samtidigt som hänsyn tas till problemets speciella struktur.

För att bestämma värden på regulariseringsparametrarna generaliseras Reinsch's metod till fallet med två parametrar. Även generaliserad korsvalidering och en mindre beräkningstung Monte Carlo-metod undersöks.

Acknowledgements

First of all I would like to thank my supervisor Docent Tommy Elfving for his guidance and support during this work.

I would also like to thank Professor Lars Eldén who started up this work once upon a time.

Finally I would like to thank my colleagues in the scientific computing group for providing a good working environment.

Papers

The following manuscripts are appended and will be referred to by their Roman numerals.

(I) Tommy Elfving and Ingegerd Skoglund, A block-preconditioner for a special regularized least squares-problem. Accepted for publication in Numerical Linear Algebra with Applications 2007. A preliminary version was published as Technical Report LiTH-MAT-R-2006-11, November 2006.

(II) Tommy Elfving and Ingegerd Skoglund, A direct method for a special regularized least squares problem, Technical Report LiTH-MAT-R-2007-1, March 2007.

Contents

1 Introduction **1**

2 Solving the linear least squares problem **3**
1 The least squares problem . . . 3
1.1 Overdetermined systems . . . 3
1.2 Underdetermined systems . . . 4
2 Algorithms for computing the least squares solution . . . 4
2.1 Direct methods . . . 4
2.2 Iterative methods . . . 5
3 The QR decomposition and least squares . . . 5
3.1 Overdetermined system . . . 5
3.2 Underdetermined system . . . 6
4 The conjugate gradient method CG and CGLS . . . 6
5 Convergence properties of CGLS . . . 7
5.1 Singular values in two arbitrary intervals . . . 7
6 The conjugate gradient method with preconditioning . . . 8

3 A partially regularized least squares problem **11**
1 An underdetermined system of equations . . . 11
1.1 The model . . . 11
1.2 Matrix formulation . . . 12
1.3 Data . . . 12
1.4 The minimum norm solution . . . 12
2 Partially regularization . . . 12
2.1 The regularized model . . . 12
2.2 Examples . . . 14
3 Choice of regularization parameters . . . 15

4 Summary of papers **17**

References **19**

1

Introduction

Statistical analysis of data from rivers deals with time series which are dependent on climatic factors and therefore have a seasonal variation. For example, it is a well-known fact that the load of substances in rivers can be strongly dependent on the runoff. It is of interest to find out whether observed changes in riverine loads are due only to natural variation or caused by other factors.

Semi-parametric models for estimation of time-varying linear relationships between runoff and riverine loads of substances were studied by Stålnacke and Grimvall [13]. There it was found that the predictivity of such models were superior to that of conventional parametric models. The aim of this work is to study different ways to numerically solve the linear least squares problem which arises.

In Chapter 2 we define the linear least squares problem and describe some standard algorithms for its solution. These are adapted for the specially structured problem treated in Paper I and Paper II.

In Chapter 3 we translate the problem in [13] into matrix formulation. The linear least squares problem which arises includes two regularization parameters. The effect of taking different values of the regularization parameters is shown in some examples.

2

Solving the linear least squares problem

In this chapter first the linear least squares problem is defined. Some standard methods for solving it are discussed and two of them studied in more detail.

1 The least squares problem

1.1 Overdetermined systems In many applications one wants to fit a linear mathematical model to measured data. In order to reduce the influence of errors in the data one can use a larger number of observations than the number of unknown parameters in the model. The problem can be formulated as an overdetermined linear system of equations,

$$Ax = b + \varepsilon, \quad A \in \mathbf{R}^{m\times n}, \quad b \in \mathbf{R}^m, \quad m > n, \tag{2.1.1}$$

where ε is a vector of random errors, and the vector b corresponds to observed data.

The linear least squares problem is then defined as

$$\min_x ||Ax - b||_2, \quad A \in \mathbf{R}^{m\times n}, \quad b \in \mathbf{R}^m, \quad m > n, \tag{2.1.2}$$

where $||.||_2$ denotes the Euklidian vector norm. The least squares solution, x_{LS}, is unique if $\text{rank}(A) = n$.

A statistical motivation for using the least squares criterion is given by the Gauss-Markoff theorem ([3, Theorem 1.1.1]).

Theorem 1.1. *The Gauss-Markoff theorem. Consider the linear model (2.1.1), where $A \in \mathbf{R}^{m\times n}$ is a known matrix of rank n, $\hat{b} = b + \varepsilon$, where ε is a random vector with mean $\mathcal{E}(\varepsilon) = 0$ and variance $\mathcal{V}(\varepsilon) = \sigma^2 I$. Then the best linear unbiased estimator of any linear function $c^T x$ is $c^T x_{LS}$, where x_{LS} is the least squares estimator, obtained by minimizing the sum of squares $||Ax - \hat{b}||_2^2$. Furthermore $\mathcal{E}(s^2) = \sigma^2$, where s^2 is the quadratic form*

$$s^2 = \frac{1}{m-n}(b - Ax_{LS})^T(b - Ax_{LS}) = \frac{1}{m-n}||b - Ax_{LS}||_2^2.$$

The least squares solution satisfies the normal equations

$$A^T Ax = A^T b. \tag{2.1.3}$$

When $\text{rank}(A) = n$ the matrix $A^T A$ is positive definite. In applications often the matrix A is sparse while $A^T A$ may be less sparse or even dense.

1.2 Underdetermined systems If $m < n$ in (2.1.2) the least squares solution is not unique. One possibility is then to compute the minimum norm solution, x_{mn}. When $b \in \mathcal{R}(A)$ then x_{mn} is the solution of the problem

$$\min ||x||_2, \quad Ax = b, \quad A \in \mathbf{R}^{m \times n}, \quad m < n.$$

The solution is given by the normal equations of the second kind, see [3, p. 7]

$$AA^T z = b, \quad x = A^T z.$$

Assuming that $\text{rank}(A) = m$ then $x = A^T (AA^T)^{-1} b$.

2 Algorithms for computing the least squares solution

There are several algorithms for solving the linear least squares problem (2.1.2). A nice historical overview is given by Björck [5], where also references to original works are given.

We distinguish between direct algorithms and iterative algorithms. Using a direct algorithm the solution can be computed in a fixed number of arithmetic operations, for instance cmn^2 or $cm^2 n$, where c is a constant. In an iterative method, given an initial approximation, gradually better approximations to the solution are computed until some stopping criterion is fulfilled.

2.1 Direct methods Before the computer age, the basic tool was to form the normal equations (2.1.3) and use the Cholesky factorization of the matrix $A^T A$ to solve them. However forming $A^T A$ can lead to loss of accuracy. Another disadvantage is that it can be expensive to form $A^T A$. The cost is mn^2 if A is a dense matrix. (Because of symmetry only half of the elements are computed.) We also can lose possible sparseness in A, since if A is sparse $A^T A$ can be much denser.

There are several alternatives to using the normal equations. For instance the singular value decomposition (SVD) is a powerful tool for solving the linear least squares problem. A stable algorithm for computing the SVD was first given by Golub and Kahan in 1965. Although the SVD always include an iterative step, in practice, it converges so fast that it is regarded as a direct method. However, it is an expensive method especially since sparsity of A is usually lost.

Another method for solving the least squares problem uses QR decomposition. It became popular when Golub in 1965 showed how to compute the QR decomposition of A using Householder transformations. The QR decomposition is treated in the next section.

2.2 Iterative methods Iterative methods successively improve an initial approximate solution until an acceptable solution is obtained. They are especially attractive when A is large and sparse while A^TA is almost dense. This is because we never need to form A^TA. Instead the iterative methods work by multiplications of A and A^T on vectors.

The main weakness of iterative methods is that they can perform well on one class of problems but on another perform less well. For methods like the conjugate gradient method, where the convergence rate depends on the condition number, the convergence can be slow because the condition number of A^TA is the square of the condition number of A. To accelerate the convergence some preconditioning is often used. The conjugate gradient method is treated in Section 4.

3 The QR decomposition and least squares

Let $Q \in \mathbf{R}^{m\times m}$ be an orthogonal matrix. Since an orthogonal transformation preserves the Euclidian length it follows that

$$\min_x ||Ax - b||_2 = \min_x ||Q^T(Ax - b)||_2. \tag{2.3.1}$$

This is the basis for using orthogonal matrices when solving the least squares problem.

3.1 Overdetermined system The QR decomposition of A, when $m > n$, is

$$A = Q\begin{pmatrix} R \\ 0 \end{pmatrix},$$

where $R \in \mathbf{R}^{n\times n}$ is upper triangular. R is nonsingular if A has linearly independent columns, i.e. $\text{rank}(A) = n$. Then the solution to (2.3.1) can be computed by

$$x = R^{-1}c_1, \quad \text{where} \quad \begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = Q^Tb,$$

since

$$||Q^TAx - Q^Tb||_2^2 = ||Rx - c_1||_2^2 + ||c_2||_2^2$$

is minimized if $Rx - c_1 = 0$.

The QR decomposition of A can be computed using a sequence of Householder or Givens transformations. In both cases we premultiply A with a product of orthogonal matrices $P_n \cdots P_2P_1 = Q^T$. Here P_k is chosen to zero out the elements below the main diagonal in the kth column in A. With Householder reflections we take

$$P_k = I - 2v_kv_k^T/||v_k||_2^2$$

while with Givens transformations P_k is constructed by using $(m - k)$ elementary rotations.

There is usually no need to form Q explicitly, since the product Q^Tb can be computed simultaneously with doing Q^TA.

Using Householder transformations the factorization requires $2n^2(m - n/3)$ operations while the algorithm using Givens rotations require twice as much. The

cost for the updating of b ($O(mn)$) and for solving $Rx = c_1$ ($O(n^2)$) are not significant compared to the work required to factor A. This is for a dense matrix. When A is sparse and it is possible to zero elements more selectively, Givens rotations usually are the more efficient choice. ([8, p. 215])

3.2 Underdetermined system The QR decomposition of A, when $m < n$, is

$$A = Q \begin{pmatrix} R_1 & R_2 \end{pmatrix},$$

where $R_1 \in \mathbf{R}^{m\times m}$ is upper triangular. Here we premultiply A with a product of orthogonal matrices $P_{m-1} \cdots P_2 P_1 = Q^T$, where P_k is chosen to zero out the elements below the main diagonal in the kth column in A. The orthogonal matrices P_k can be constructed by Householder or Givens transformations.

4 The conjugate gradient method CG and CGLS

The conjugate gradient (CG) method is a special case of Krylov space methods for solving symmetric, positive definite systems of linear equations. It was invented by Hestenes and Stiefel and published in 1952 [11], which also included a version for the normal equations (CGLS). It has the advantage that no estimation of iteration parameters (controlling, e.g. the rate of convergence) is needed. In exact arithmetic it will produce the solution to a $n \times n$ linear system $Hx = c$ in at most n steps. If there are only t distinct nonzero eigenvalues of H it will compute the solution in at most t iterations, see [3]. However the method is most often used as an iterative method, then usually requiring much less then n steps.

For a given starting vector x_0 the CGLS method generates approximations x_k in the affine subspace

$$x_k \in x_0 + \mathcal{K}_k(A^T A, s_0), \qquad s_0 = A^T(b - Ax_0)$$

where $\mathcal{K}_k(A^T A, s_0)$ is the Krylov subspace

$$\operatorname{span}\{s_0, (A^T A)s_0, \ldots, (A^T A)^{k-1} s_0\}.$$

Let $r = b - Ax$ be the residual. Every approximation x_k minimizes the error functional

$$E(x_k) = (x-x_k)^T A^T A(x-x_k) = \|A(x-x_k)\|_2^2 = \|-b+Ax-Ax_k+b\|_2^2 = \|r_k-r\|_2^2.$$

In [11] it was shown that both $\|r - r_k\|_2$ and $\|x - x_k\|_2$ decrease monotonically, but $\|A^T r_k\|_2$ will often exhibit large oscillations when $\kappa(A)$ is large, see [3].

The algorithm for solving the normal equations, originally given by Hestenes and Stiefel [11], is given in Paper I, Section 2. Björck et al. [4] compared this algorithm with several other implementations and found this to be the most accurate. When A is of size $m \times n$ each iteration requires about $4\text{nz}(A)+6n+4m$ operations, where $\text{nz}(A)$ are the number of nonzero elements in A.

5 Convergence properties of CGLS

For CGLS, the following estimate of the rate of convergence is known to hold:

$$\|r - r_k\|_2 < 2\left(\frac{\kappa - 1}{\kappa + 1}\right)^k \|r - r_0\|_2, \tag{2.5.1}$$

where $\kappa = \kappa(A) = \sqrt{\kappa(A^T A)}$, see [3]. From this it follows that an upper bound for the number of iterations, k, to reduce the relative error by a factor of ε so that

$$\|r - r_k\|_2 < \varepsilon \|r - r_0\|_2 \tag{2.5.2}$$

is given by

$$k < \frac{1}{2}\kappa(A) \ln \frac{2}{\varepsilon}. \tag{2.5.3}$$

However, the convergence of the conjugate gradient method also depends on the distribution of the singular values of A. For instance, as mentioned above, if A has only $t \leq n$ distinct singular values, then in exact arithmetic the solution is obtained in t steps. In Axelsson [1], [2], several estimates for the necessary number of iterations needed are given. These utilize specific eigenvalue distribution of the matrix at hand.

5.1 Singular values in two arbitrary intervals Assume that the eigenvalues of a symmetric and positive definite matrix H are situated in two continuous intervals $[a, b] \cup [c, d]$, such that

$$0 < a = \lambda_1(H) < \cdots < \lambda_q(H) = b < c = \lambda_{q+1}(H) < \cdots < \lambda_n(H) = d$$

and q is not too small. When solving $Hx = c$ by the CG method, an upper bound for k to satisfy

$$\|x - x_k\|_{H^{\frac{1}{2}}} < \varepsilon \|x - x_0\|_{H^{\frac{1}{2}}}, \quad \text{where} \quad \|x\|^2_{H^{\frac{1}{2}}} = x^T H x \tag{2.5.4}$$

is given by

$$k = \left(\frac{1}{2}\sqrt{\frac{d}{c}} + \frac{1}{2}\sqrt{\frac{b}{a}} + \frac{1}{4}\sqrt{\frac{b}{a}}\sqrt{\frac{d}{c}} \ln \frac{4d}{b}\right) \ln \frac{2}{\varepsilon}.$$

This result is found in [1].

Let

$$0 < \hat{a} = \sigma_1(A) < \cdots < \sigma_q(A) = \hat{b} < \hat{c} = \sigma_{q+1}(A) < \cdots < \sigma_n(A) = \hat{d},$$

be the singular values of A. To satisfy (2.5.2) we then need to take less than k iterations where

$$k = \left(\frac{\hat{d}}{2\hat{c}} + \frac{\hat{b}}{2\hat{a}} + \frac{\hat{b}}{4\hat{a}}\frac{\hat{d}}{\hat{c}} \ln \frac{4\hat{d}^2}{\hat{b}^2}\right) \ln \frac{2}{\varepsilon}. \tag{2.5.5}$$

Here we utilize that $\lambda_i(A^T A) = \sigma_i^2(A)$. The estimates (2.5.3) and (2.5.5) were used in Paper I.

6 The conjugate gradient method with preconditioning

We have seen that the rate of convergence is dependent on the condition number $\kappa(A)$ and also depends on the spectrum of A. A general technique to improve convergence is by preconditioning. We then transform the problem, to get a matrix with better spectral properties, into the related problem

$$\min_y \|AS^{-1}y - b\|_2, \qquad Sx = y, \tag{2.6.1}$$

where $S \in \mathbf{R}^{n\times n}$ is a nonsingular matrix. If S is chosen so that AS^{-1} has a more favorable spectrum than A, this will improve the convergence of the CGLS method.

A good preconditioner S should have the following properties, see [3]:

(i) AS^{-1} should be better conditioned than A and/or have only a few distinct singular values.

(ii) S should have about the same number of non-zeroes as A.

(iii) It should be cheap to solve equations with S and S^T.

In the iterative method, matrix-vector products of the forms $AS^{-1}y$ and $S^{-T}A^Tr$ will occur. Thus the extra cost of preconditioning will be in solving linear systems of the form $Sx = y$ and $S^Tq = s$ which makes the condition (iii) important.

The normal equations for the problem (2.6.1) are

$$S^{-T}A^TAS^{-1}Sx = S^{-T}A^Tb$$

and if we multiply by S^{-1} we get

$$S^{-1}S^{-T}A^TAx = S^{-1}S^{-T}A^Tb.$$

Apparently, if $M = S^TS$ is an approximation to A^TA, the condition number of $M^{-1}A^TA$ is small.

The preconditioned conjugate method in terms of the original variable x gives the following algorithm, see [3, p. 294]:

PCCGLS

Let x_0 be an initial approximation and set
$r_0 = b - A\,x_0, \quad s_0 = p_0 = S^{-T}A^T r_0, \quad \gamma_0 = \|s_0\|_2^2;$
for $k = 0, 1, 2 \ldots$ while $\gamma_k > tol$ compute

(i)
$t_k = S^{-1}p_k,$
$q_k = A\,t_k,$
$\alpha_k = \gamma_k / \|q_k\|_2^2,$
$x_{k+1} = x_k + \alpha_k t_k,$
$r_{k+1} = r_k - \alpha_k q_k,$
$\theta_{k+1} = A^T r_{k+1},$

(ii)
$s_{k+1} = S^{-T}\theta_{k+1},$
$\gamma_{k+1} = \|s_{k+1}\|_2^2,$
$\beta_k = \gamma_{k+1}/\gamma_k,$
$p_{k+1} = s_{k+1} + \beta_k p_k.$

The preconditioning steps are marked (i) and (ii). One may ask why we do not try to find an appropriate approximation to the matrix $(A^TA)^{-1}$ itself, thus avoiding the need of solving systems of equations. However $(A^TA)^{-1}$ is usually a dense matrix and the sought approximation must then also be a fairly dense matrix. This means that even multiplication with such a matrix usually is too expensive.

In Paper I, we construct special block-preconditioners and analyze their behavior using Algorithm PCCGLS.

3

A partially regularized least squares problem

In Paper I, Section 4 and in Paper II, Section 2 an application in statistics is shortly described. In this chapter this model is studied in somewhat more detail.

1 An underdetermined system of equations

1.1 The model We assume that we have observed a time series L_{ij}, $i = 1, 2, \ldots, n$, $j = 1, 2, \ldots, m$ of riverine loads and a time series q_{ij}, $i = 1, 2, \ldots, n$, $j = 1, 2, \ldots, m$ of runoff data, where n is the length of the study period expressed in years and m is the number of months or samples per year. Then Stålnacke and Grimvall [13] proposed the model

$$L_{ij} = \alpha_{ij} + \beta_{ij} q_{ij} + \eta_{ij}, \quad i = 1, 2, \ldots, n, \ j = 1, 2, \ldots, m, \tag{3.1.1}$$

where α_{ij} and β_{ij} are parameters to be estimated and η_{ij}, $i = 1, 2, \ldots, n$, $j = 1, 2, \ldots, m$ are random variables with mean zero. Mainly the intercepts α_{ij} express changes in point source contributions and the slope parameters β_{ij} changes in diffuse sources.

Since the number of parameters exceeds the number of observations, some kind of constraints on the parameters has to be introduced. The following semi-parametric models were used in [13]. The simplest form is

$$\beta_{ij} = \beta_j, \quad i = 1, 2, \ldots, n, \quad j = 1, 2, \ldots, m,$$

which means that the slope parameters can vary with season but not with year. In the model studied here

$$\beta_{ij} = \beta_j + \gamma_j i, \quad i = 1, 2, \ldots, n, \quad j = 1, 2, \ldots, m, \tag{3.1.2}$$

which allows temporal trends in the slope parameters. In both cases the intercepts α_{ij} are permitted to vary more freely. Now, counting number of parameters we get nm (the α-parameters) and, for model (3.1.2), $2m$ ($\{\beta_j\}$, $\{\gamma_j\}$) whereas the number of data observations $\{L_{ij}\}$ is nm. Hence we have more parameters than unknowns.

1.2 Matrix formulation Order the time series of data L_{ij} in the vector $b_1 = (L_{11}, .., L_{1m}, L_{21}, .., L_{2m}, .., L_{n1}, .., L_{nm})^T$ and in the unknowns as $x_2 = (\alpha_{11}, .., \alpha_{1m}, \alpha_{21}, .., \alpha_{2m}, .., \alpha_{n1}, .., \alpha_{nm})^T$. Let further $x_1 = (\beta_1, \beta_2, .., \beta_m, \gamma_1, \gamma_2, .., \gamma_m)^T$ and $\eta = (\eta_{11}, .., \eta_{1m}, \eta_{21}, .., \eta_{2m}, .., \eta_{n1}, .., \eta_{nm})^T$. Now the model (3.1.1) and (3.1.2) can be formulated

$$A_1 x_1 + A_2 x_2 + \eta = b_1 \tag{3.1.3}$$

where

$$A_1 = \begin{pmatrix} D_1 & D_1 \\ D_2 & 2D_2 \\ \vdots & \vdots \\ D_n & nD_n \end{pmatrix}, \quad D_i = \mathrm{diag}(q_{ij}), \quad i = 1, 2, \ldots, n$$

and

$$A_2 = I_{nm}, \tag{3.1.4}$$

where I_{nm} is the identity matrix of order nm. This is an underdetermined system of equations

$$Ax + \eta = b_1, \quad A = (A_1, A_2), \quad x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \quad A \in \mathbf{R}^{(nm)\times(nm+2m)}, \quad b_1 \in \mathbf{R}^{nm}. \tag{3.1.5}$$

1.3 Data Stålnacke and Grimvall analysed data from the Rhine River in [13]. Biweekly observations of water discharge, q_{ij}, and concentrations of nitrogen, phosphorus or chloride, L_{ij}, were aggregated into monthly runoff and load data. The resulting data are shown in [13, Fig. 1]. Since the aim of this work is only to study how to solve the problem in an efficient way, fictitious data similar to those for chloride were generated using a perturbed sine-function. Instead of data every month, we took data twice a month during 15 years, i.e., $n = 15$ and $m = 24$. This data can be found in Paper II, Figure 4.3.

1.4 The minimum norm solution Due to (3.1.4) the matrix A in (3.1.5) has full rank. Hence the problem (3.1.5) is consistent. Then we may define a unique solution by taking the minimum norm solution, see Chapter 2, Section 1.2. For comparison we have computed this solution, using built in commands in Matlab. The results are displayed in Figure 1.1. It is seen that rapid variations in the intercepts x_2 occur, which is not a required quality. Therefore, following Stålnacke and Grimvall, we consider regularization on the parameter x_2 to avoid such behavior.

2 Partially regularization

2.1 The regularized model In the model (3.1.1) and (3.1.2) Stålnacke and Grimvall added constraints on the intercepts α_{ij}. The basic ideas behind this technique was described by Green and Silverman [9] in 1994. Rapid, irregular variation is penalized during the parameter estimation by minimizing

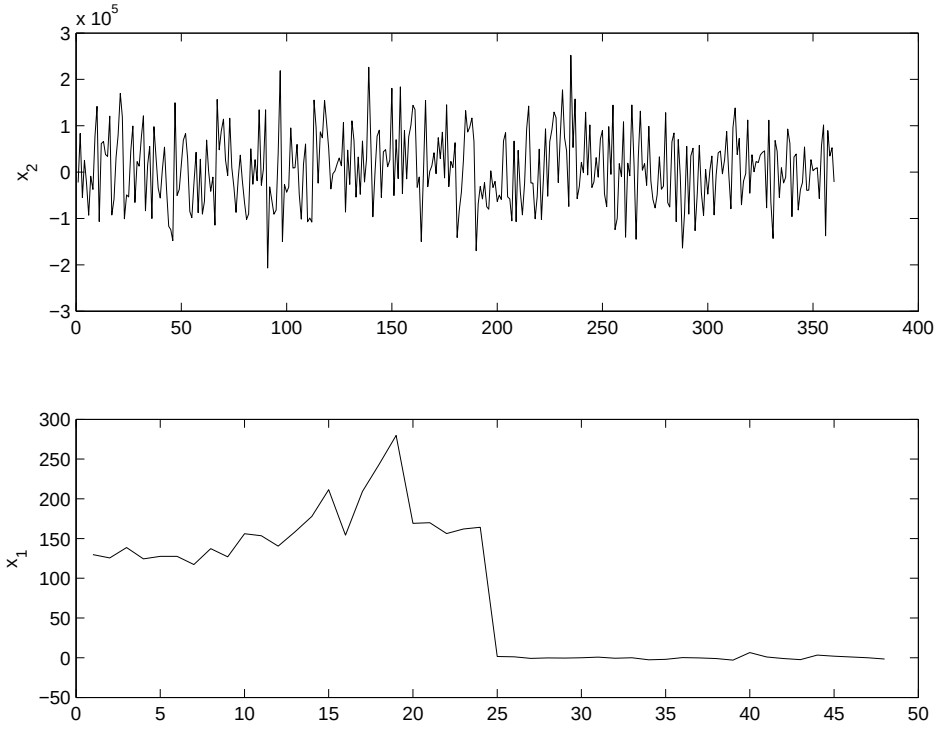

Figure 1.1: *The minimum norm solution, x_2 (top) and x_1 (bottom).*

$$S(\alpha,\beta)=\sum_{i,j}(L_{ij}-\alpha_{ij}-\beta_{ij}q_{ij})^2+\bar{\lambda}_1\sum_{i,j}(\alpha_{ij}-\frac{\alpha_{i,j+1}+\alpha_{i,j-1}}{2})^2+$$

$$\bar{\lambda}_2\sum_{i,j}(\alpha_{ij}-\frac{\alpha_{i+1,j}+\alpha_{i-1,j}}{2})^2 \tag{3.2.1}$$

where $\alpha_{i,m+1}=\alpha_{i+1,1}$, $\alpha_{i,0}=\alpha_{i-1,m}$, and $\bar{\lambda}_1>0,\bar{\lambda}_2>0$ denote penalty coefficients for, respectively, the interseasonal and interannual variation in the intercepts. This means that only a subset $x_2=\{\alpha_{ij}\}$ of the unknowns are regularized. For the other part, x_1, one just uses the submodel (3.1.2).

Let

$$A_3=\begin{pmatrix}A_4\\A_5\end{pmatrix}$$

where

$$A_4=\begin{pmatrix}1&-2&1&&\\&\ddots&\ddots&\ddots&\\&&1&-2&1\end{pmatrix}\text{ and }A_5=\begin{pmatrix}I_m&-2I_m&I_m&&\\&\ddots&\ddots&\ddots&\\&&I_m&-2I_m&I_m\end{pmatrix}.$$

Putting $\bar{\lambda}_i/4 = \lambda_i,\ i = 1, 2$ the expression (3.2.1) becomes

$$S(\alpha, \beta) = \bar{S}(x_1, x_2) = ||b_1 - A_1 x_1 - A_2 x_2||_2^2 + \lambda_1 ||A_4 x_2||_2^2 + \lambda_2 ||A_5 x_2||_2^2 \quad (3.2.2)$$

or, if $\lambda_1 = \lambda_2 = \lambda$,

$$S(\alpha, \beta) = \bar{S}(x_1, x_2) = ||b_1 - A_1 x_1 - A_2 x_2||_2^2 + \lambda ||A_3 x_2||_2^2. \quad (3.2.3)$$

The parameters have a two dimensional structure and A_4, A_5 corresponds to applying the second derivative in each dimension respectively. Because of lack of 'boundary' data the matrix A_4 is of size $(nm - 2) \times nm$ and the matrix A_5 is of size $(nm - 2m) \times nm$.

By defining

$$A = \begin{pmatrix} A_1 & A_2 \\ 0 & \sqrt{\lambda_1} A_4 \\ 0 & \sqrt{\lambda_2} A_5 \end{pmatrix}, \quad A \in \mathbf{R}^{(3nm-2m-2)\times(nm+2m)}$$

$$b = \begin{pmatrix} b_1 \\ 0 \\ 0 \end{pmatrix}, \quad b \in \mathbf{R}^{(3nm-2m-2)}$$

and using (3.2.2) the minimizing problem is

$$\min_x \|Ax - b\|_2 = \min_{x_1, x_2} \left\| \begin{pmatrix} A_1 & A_2 \\ 0 & \sqrt{\lambda_1} A_4 \\ 0 & \sqrt{\lambda_2} A_5 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} - \begin{pmatrix} b_1 \\ 0 \\ 0 \end{pmatrix} \right\|_2. \quad (3.2.4)$$

For the case $\lambda_1 = \lambda_2 = \lambda$ and using (3.2.3) the minimizing problem is

$$\min_x \|Ax - b\|_2 = \min_{x_1, x_2} \left\| \begin{pmatrix} A_1 & A_2 \\ 0 & \sqrt{\lambda} A_3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} - \begin{pmatrix} b_1 \\ 0 \end{pmatrix} \right\|_2. \quad (3.2.5)$$

The sparsity pattern of A can be seen in Paper II, Figure 3.1(a).

2.2 Examples We illustrate the character of the solution for two different λ-values ($\lambda = \lambda_1 = \lambda_2$) in Figure 2.1. The top subplot reflects the fact that the α-parameters have a cyclic behavior (with a period of one year). Note that the scale on the x-axis is in two-week periods, so the time-scale starts at week 1 and ends in week 720 (and one year is for convenience only 48 weeks). The maximum and minimum values respectively, during the considered time period, corresponds to the same season. As seen one can restrict the interseasonal variations by using different values of λ.

The bottom subplot displays the variable x_1, which using (3.1.2) consists of the parameters $\{\beta_j\}, \{\gamma_j\}$. The x_1-parameters with index 1–24 corresponds to $\beta_1, \ldots, \beta_{24}$ and those with index 25–48 to $\gamma_1, \ldots, \gamma_{24}$.

Using $\lambda_1 \neq \lambda_2$ gives the possibility to differently penalize variations in the intercepts, x_2, both interannual (governed by λ_2) and interseasonal (governed by λ_1). For instance, if we take $\lambda_2 > \lambda_1$, then nonlinear interannual variation is more penalized, see Figure 2.2.

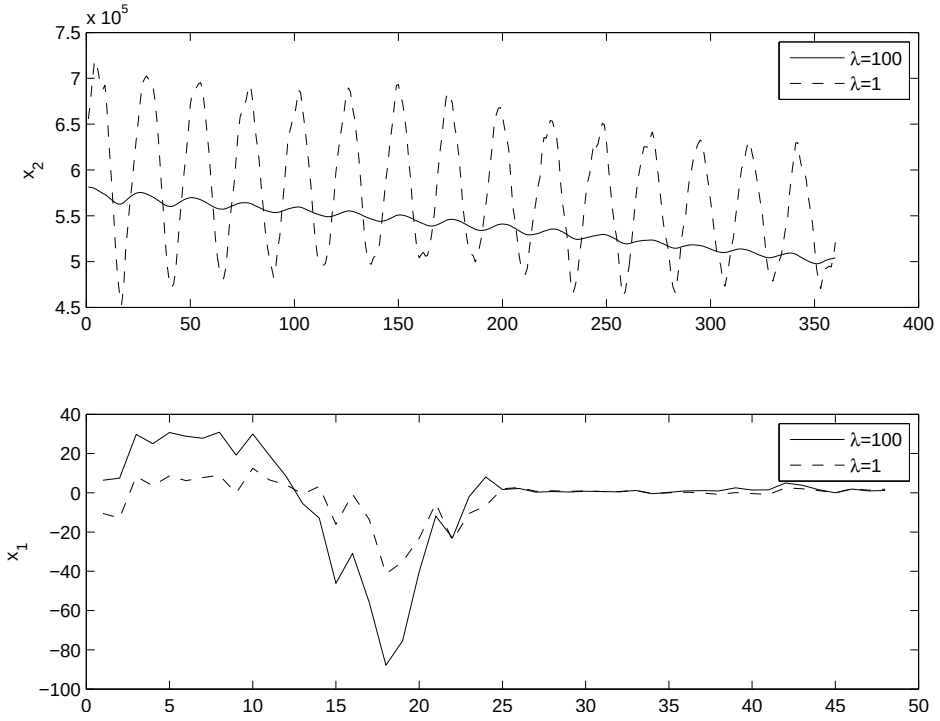

Figure 2.1: *Solution with $\lambda = 1$ and with $\lambda = 100$, x_2 (top) and x_1 (bottom)*

3 Choice of regularization parameters

The parameter λ in (3.2.3) gives the trade-off between data fidelity (measured by $||b_1 - A_1x_1 - A_2x_2||$) and smoothness of x_2 (measured by $||A_3x_2||$). There are many applications where, due to ill-conditioning, it becomes necessary to introduce smoothing on the sought solution, for examples see [10].

There are several techniques available for finding good values of these parameters. The values depend both on the amount of noise in the data, and the degree of ill-posedness in the problem. A well known method is generalized cross-validation [7]. The main idea is that the estimated solution should not change much if one observation is left out during the estimation. This method is implemented in Paper II, Section 4.4 using our QR-based algorithm. There is one computationally demanding step needed when forming the generalized cross-validation function. For this step we therefore investigate, as a cheaper alternative, a Monte Carlo technique due to Girard [6].

In some problems one may know (an estimate) of the norm of the data-error. In this case Reinsch [12] has developed a Newton method for finding the corresponding value of λ. This method is also investigated in Paper II, Section 4.1. We also generalize this method to the case with two regularization parameters.

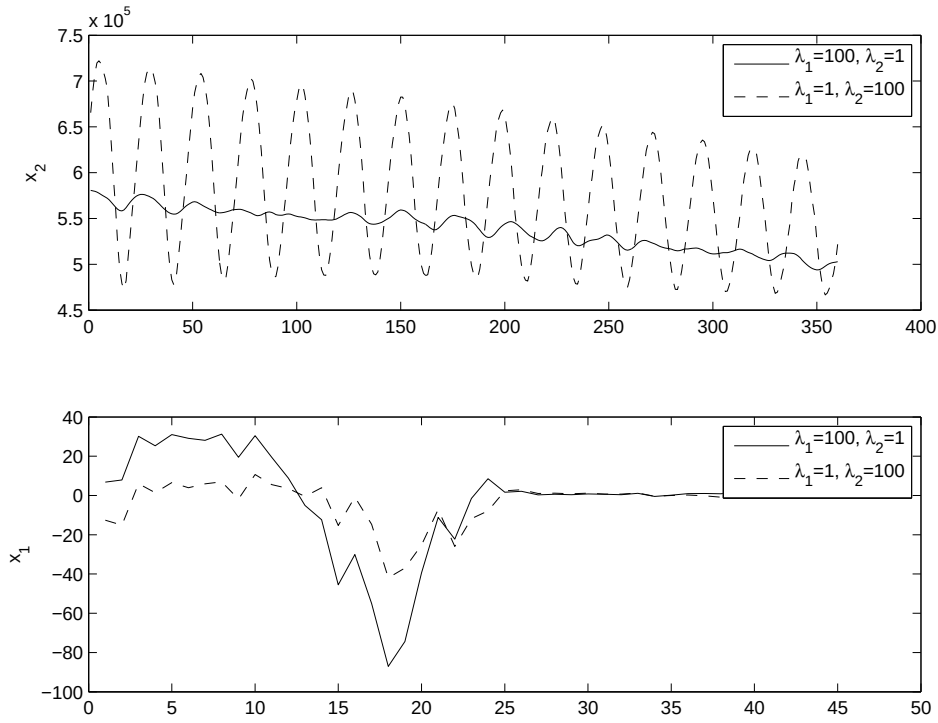

Figure 2.2: *Solution with* $\lambda_1 = 1$, $\lambda_2 = 100$, *and* $\lambda_1 = 100$, $\lambda_2 = 1$, x_2 *(top) and* x_1 *(bottom)*

4

Summary of papers

Paper I
A block-preconditioner for a special regularized least squares-problem

Here the problem (3.2.5) (i.e., using a single regularization parameter) is considered. We propose preconditioners of Schur complement type to be used in Algorithm PCCGLS (given in Chapter 2). We then assume that the matrix $A_1^T A_1$ has full rank and is easy to invert. For one of the suggested preconditioners it is also assumed that $A_3^T A_3$ has full rank. This is then accomplished by adding a multiple of the identity matrix to A_3, which corresponds to regularizing with a Sobolev norm. We also drop the assumption of full rank and derive the corresponding preconditioner.

We derive bounds for the number of iterations and compare these theoretical bounds with the number of iterations obtained in experiments.

Paper II
A direct method for a special regularized least squares-problem

Here both problem (3.2.5) and problem (3.2.4) are treated, i.e., we consider regularizing using one and two regularization parameters respectively.

A direct solution method based on QR decomposition of matrix blocks is studied. The idea is to first perform operations which can be done independently of the regularization parameters, since suitable values of these usually are not known in advance. Therefore the problem needs to be solved for several values.

We next discuss the choice of regularization parameters. In particular we consider the generalized cross-validation method used together with our QR-based algorithm. There is one computationally demanding step needed when forming the generalized cross-validation function. For this step we therefore investigate, as a cheaper alternative, a Monte Carlo technique. In some problems one may know (an estimate) of the norm of the data-error. In this case we study Newton's method for finding the corresponding value of λ. We also generalize this method to the case with two regularization parameters.

Notification

The alphabetic order of the authors of the two papers reflects approximatively equal inputs to the papers. It is of course natural that the adviser mostly inputs ideas and the student works out the details. Many improvements have emerged after the results of numerical experiments. All experimental work was done by the student. Even if most of the work is done jointly we would like to mention that Section 3 of Paper II is mostly the work of the student.

References

[1] O. Axelsson. Solution of linear systems of equations: iterative methods. In V.A. Barker, editor, *Sparse Matrix Techniques*. Springer Verlag, Berlin, Heidelberg, New York, 1976.

[2] O. Axelsson. *Iterative Solution Methods*. Cambridge University Press, Cambridge, 1994.

[3] Å. Björck. *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, 1996.

[4] Å. Björck T. Elfving and Z. Strakos. Stability of conjugate gradient and Lanczos methods for linear least squares problems. *SIAM Journal on Matrix Analysis and Applications*, 19:720–736, 1998.

[5] Å. Björck. The calculation of linear least squares problems. *Acta Numerica*, 13:1–53, 2004.

[6] D.A. Girard. A fast 'Monte-Carlo cross-validation' procedure for large least squares problems with noisy data. *Numerische Mathematik*, 56:1–23, 1989.

[7] G.H. Golub M.T. Heath and G. Wahba. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics*, 21:215–223, 1979.

[8] G.H. Golub and C.F. Van Loan. *Matrix Computaions*. third ed., Johns Hopkins University Press, Baltimore, 1996.

[9] P.J. Green and B.W. Silverman. *Nonparametric regression and generalized linear models - a roughness penalty approach*. Chapman & Hill, London, 1994.

[10] P.C. Hansen. *Rank-Deficient and Discrete Ill-Posed Problems. Numerical Aspects of Linear Inversion*. SIAM, Philadelphia, 1998.

[11] M.R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear stystem. *Journal of Research of the National Bureau of Standards*, B49:409–436, 1952.

[12] C.H. Reinsch. Smoothing by spline functions. *Numerische Mathematik*, 10:177–183, 1967.

[13] P. Stålnacke and A. Grimvall. Semiparametric approaches to flownormalisation and source apportionment of substances transport in rivers. *Environmetrics*, 12:233–250, 2001.