united nations educational, scientific and cultural organization

international atomic energy agency

the
abdus salam
international centre for theoretical physics

Course on "Inverse Methods in Atmospheric Science"
1 - 12 October 2001

301/1332-2

"General Inverse Methods
or
Discrete Linear Ill-Posed Problems"

Michael JACOBSEN
Technical University of Denmark
Brussels, Belgium

Please note: These are preliminary notes intended for internal distribution only.

strada costiera, 11 - 34014 trieste italy - tel. +39 0402240111 fax +39 040224163 - sci_info@ictp.trieste.it - www.ictp.trieste.it

General Inverse Methods
or
Discrete Linear Ill-Posed Problems

Michael Jacobsen (mj@imm.dtu.dk)
PhD-student at Informatics and Mathematical Modelling
Technical University of Denmark

Discrete Linear Ill-Posed Problems, Part 1 of 4
Setting the Stage

Definition:

1. a square or overdetermined system of linear algebraic equations
2. with a huge condition number
3. coming from the discretization of an inverse/ill-posed problem.

Our generic ill-posed problem:

A Fredholm integral equation of the first kind

$$\int_0^1 K(s,t)\, f(t)\, dt = g(s) \;, \qquad 0 \le s \le 1 \;.$$

Numerical Methods

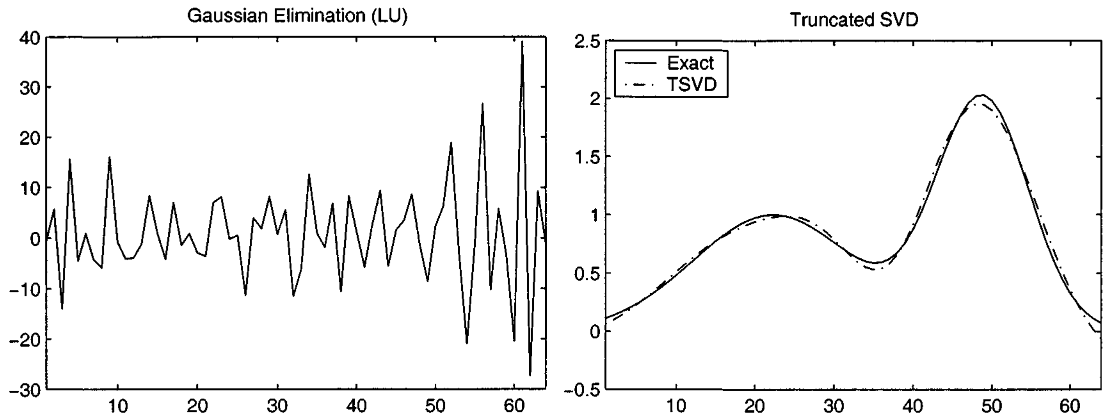

- Standard methods (LU, QR, etc.) produce useless results.
- Specialized methods can produce "reasonable" results.

Some Important Questions

- How to discretize the integral equation?
- Why is the matrix always so ill conditioned?
- Why can we still compute an approximate solution?
- How can we compute it stably and efficiently?
- Is additional information available?
- How can we incorporate it in the solution scheme?
- How should we implement the numerical scheme?

Inverse Problems

Generic model

$$\boxed{\int_{\Omega} \text{input} \times \text{system } d\Omega = \text{output}}$$

Image restoration

scenery $\rightarrow$ lens $\rightarrow$ image

Tomography

X-ray source $\rightarrow$ object $\rightarrow$ damping

Seismology

seismic wave $\rightarrow$ layers $\rightarrow$ reflections

Image Restoration

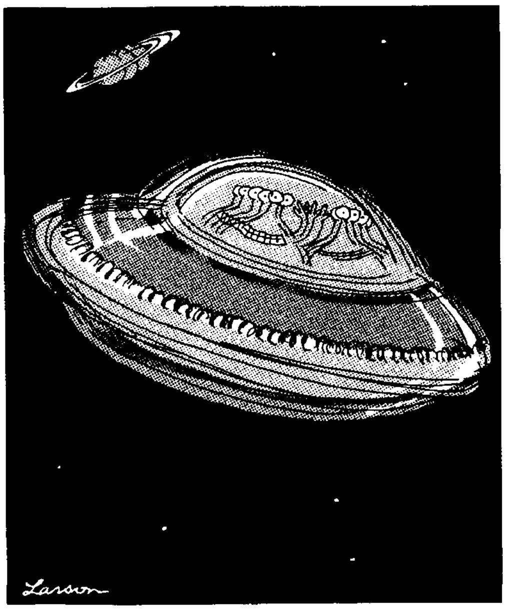

Another photograph from the Hubble telescope

Linear Inverse Problems

Fredholm integral equation of the first kind

$$\int_0^1 K(s,t)\, f(t)\, dt = g(s) \; .$$

Ditto with discrete right-hand side

$$\int_0^1 k_i(t)\, f(t)\, dt = b_i \; , \qquad i = 1, \ldots, m \; .$$

with $k_i(t) = K(s_i, t)$ and $b_i = g(s_i)$.

Integration with K has a smoothing effect on f, i.e., g is smoother than f.

The Riemann-Lebesgue Lemma

Consider the function

$$f(t) = \sin(2\pi p t) \; , \qquad p = 1, 2, \ldots$$

then for $p \to \infty$ and "arbitrary" K we have

$$g(s) = \int_0^1 K(s,t)\, f(t)\, dt \to 0 \; .$$

I.e., high frequencies are damped.

Therefore difficult to reconstruct f from g.

Illustration of Riemann-Lebesgue Lemma

Geomagnetic problem with $f(t) = \sin(2\pi\, pt)$, $p = 1, 2, 4,$ and 8.

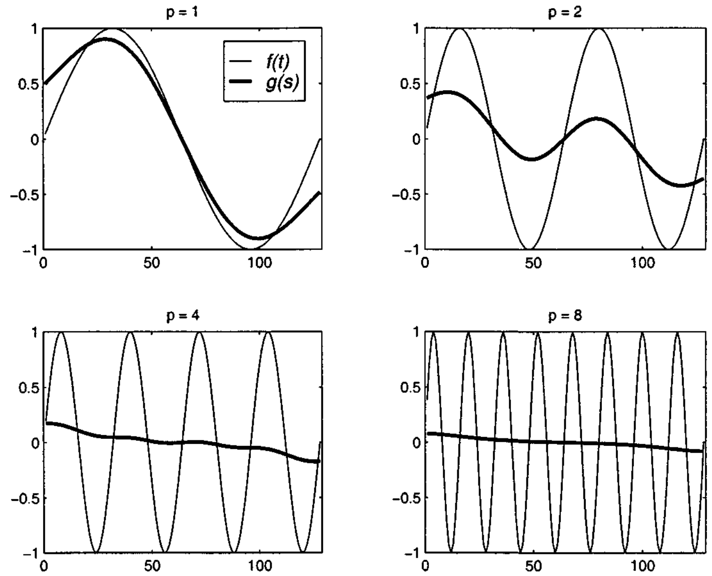

Our Model Problem: Geomagnetic Prospecting

- Iron ore deposit at depth d below surface from 0 to 1 on t axis.
- Measurements of vertical component of magnetic field $g(s)$ at surface, from a to b on the s axis.
- Unknown: the vertical component of the field $f(t)$ at the ore, from 0 to 1 on the t axis.

The Geometry

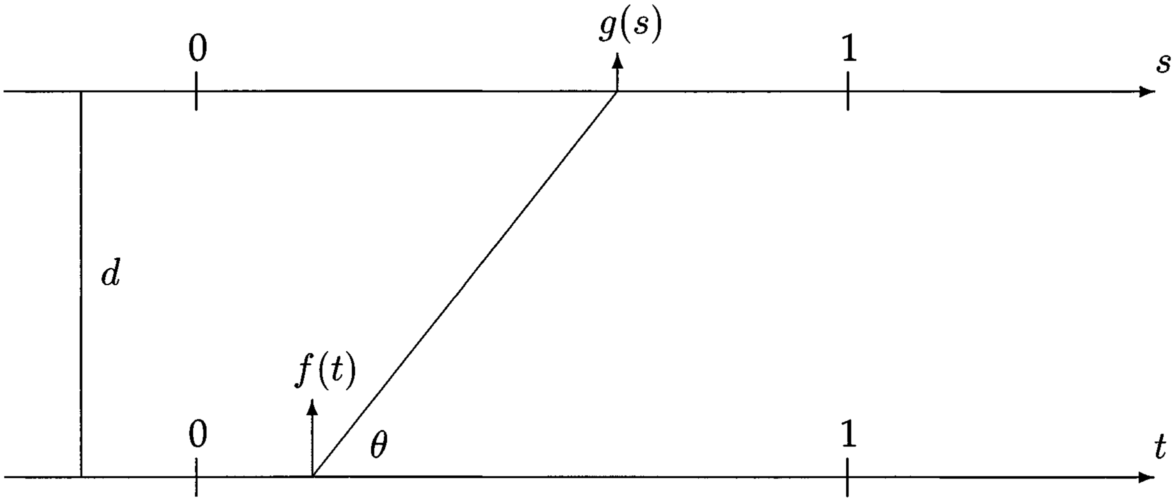

Setting Up the Integral Equation

The value of $g(s)$ due to the part dt on the t axis

$$dg = \frac{\sin\theta}{r^2} f(t)\, dt \;,$$

where $r = \sqrt{d^2 + (s-t)^2}$. Using that $\sin\theta = d/r$, we get

$$\frac{\sin\theta}{r^2} f(t)\, dt = \frac{d}{(d^2 + (s-t)^2)^{3/2}} f(t)\, dt \;.$$

The total value of $g(s)$ for $a \le s \le b$ is therefore

$$g(s) = \int_0^1 \frac{d}{(d^2 + (s-t)^2)^{3/2}} f(t)\, dt \;.$$

Our Integral Equation

Fredholm integral equation of the first kind:

$$\int_0^1 \frac{d}{(d^2 + (s-t)^2)^{3/2}} \, f(t) \, dt = g(s) \, , \qquad a \le s \le b \, .$$

The kernel K, which represents the model, is

$$K(s,t) = h(s-t) = \frac{d}{(d^2 + (s-t)^2)^{3/2}} \, ,$$

and the right-hand side g is what we are able to measure.

From K and g we want to compute f, i.e., an inverse problem.

Discretization: the Quadrature Method

Recall the simple quadrature rule

$$\int_0^1 \phi(t) \, dt \approx \sum_{j=1}^n w_j \, \phi(t_j) \, ,$$

with

$$w_j = \text{weights} \, , \quad t_j = \text{abscissas} \, , \qquad j = 1, \ldots, n \, .$$

Hence, we approximate the integral in our model as follows

$$\int_0^1 K(s,t) \, f(t) \, dt \approx \sum_{j=1}^n w_j \, K(s,t_j) \, \tilde{f}(t_j) \equiv \psi(s) \, .$$

Note that we have replaced f with $\tilde{f}$.

Quadrature Discretization, Cont.

To obtain a linear system of equations, we use collocation.

I.e., we require that ψ equals g at selected points:

$$\psi(s_i) = g(s_i) \ , \qquad i = 1, \ldots, m \ .$$

Here, $g(s_i)$ are really the measured values of the function g.

If $m > n$ we obtain an overdetermined system.

Here we assume $m = n$ for simplicity:

$$\sum_{j=1}^{n} w_j \, K(s_i, t_j) \, \tilde{f}(t_j) = g(s_i) \ , \qquad i, j = 1, \ldots, n \ .$$

The Discrete Problem in Matrix Form

Write out the last equation to obtain

$$\begin{pmatrix} w_1 K(s_1, t_1) & w_2 K(s_1, t_2) & \cdots & w_n K(s_1, t_n) \\ w_1 K(s_2, t_1) & w_2 K(s_2, t_2) & \cdots & w_n K(s_2, t_n) \\ \vdots & \vdots & & \vdots \\ w_1 K(s_n, t_1) & w_2 K(s_n, t_2) & \cdots & w_n K(s_n, t_n) \end{pmatrix} \begin{pmatrix} \tilde{f}(t_1) \\ \tilde{f}(t_2) \\ \vdots \\ \tilde{f}(t_n) \end{pmatrix} = \begin{pmatrix} g(s_1) \\ g(s_2) \\ \vdots \\ g(s_n) \end{pmatrix}$$

or simply $A\,x = b$ (where A is $n \times n$) with

$$\left. \begin{aligned} a_{ij} &= w_j \, K(s_i, t_j) \\ x_j &= \tilde{f}(t_j) \\ b_i &= g(s_i) \end{aligned} \right\} \qquad i, j = 1, \ldots, n \ .$$

A Special Case: the Midpoint Rule

Equidistant abscissas

$$t_j = (j - 0.5)\, n^{-1}\,, \qquad j = 1, \ldots, n$$

with identical weights $w_j = n^{-1}$, $j = 1, \ldots, n$.

Matrix elements:

$$a_{ij} = n^{-1} K(s_i, t_j)\,, \qquad i, j = 1, \ldots, n\,.$$

The Singular Value Decomposition

Assume that A is $m \times n$ and, for simplicity, also that $m \geq n$:

$$A = U\, \Sigma\, V^T = \sum_{i=1}^{n} u_i\, \sigma_i\, v_i^T$$

where U and V consist of *singular vectors*

$$U = (u_1, \ldots, u_n)\,, \qquad V = (v_1, \ldots, v_n)$$

with $U^T U = V^T V = I_n$, and the *singular values* satisfy

$$\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_n)\,, \qquad \sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n \geq 0\,.$$

Then $\|A\|_2 = \sigma_1$ and $\mathrm{cond}(A) = \|A\|_2\, \|A^\dagger\|_2 = \sigma_1/\sigma_n$.

Important SVD Relations

$$\left.\begin{array}{rclrcl} A\,v_i &=& \sigma_i\,u_i & \|A\,v_i\|_2 &=& \sigma_i \\ A^T u_i &=& \sigma_i\,v_i & \|A^T u_i\|_2 &=& \sigma_i \end{array}\right\} \qquad i = 1,\dots,n.$$

These equations are related to the (least squares) solution:

$$\begin{aligned} x &= \sum_{i=1}^{n} (v_i^T x)\, v_i \\ A\,x &= \sum_{i=1}^{n} \sigma_i\,(v_i^T x)\, u_i\ , \quad b = \sum_{i=1}^{n} (u_i^T b)\, u_i \\ A^{-1}b &= \sum_{i=1}^{n} \frac{u_i^T b}{\sigma_i}\, v_i\ . \end{aligned}$$

Discrete Linear Ill-Posed Problems, Part 2 of 4 Regularization

1. The SVD in plots
2. Regularization = stabilization
   - Filtering and/or side constraints
3. Tikhonov's method
   - Formulation and SVD analysis
4. Implementation of Tikhonov's method
5. Related methods
   (a) Least squares with quadratic constraints
   (b) Least squares with inequality constraints

What the Singular Values Look Like

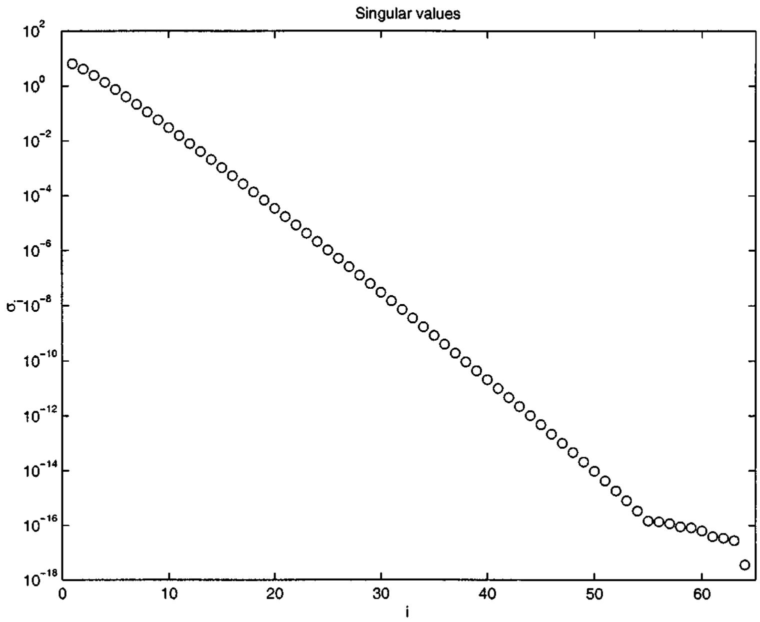

What the Left (and Right) Singular Vectors Look Like

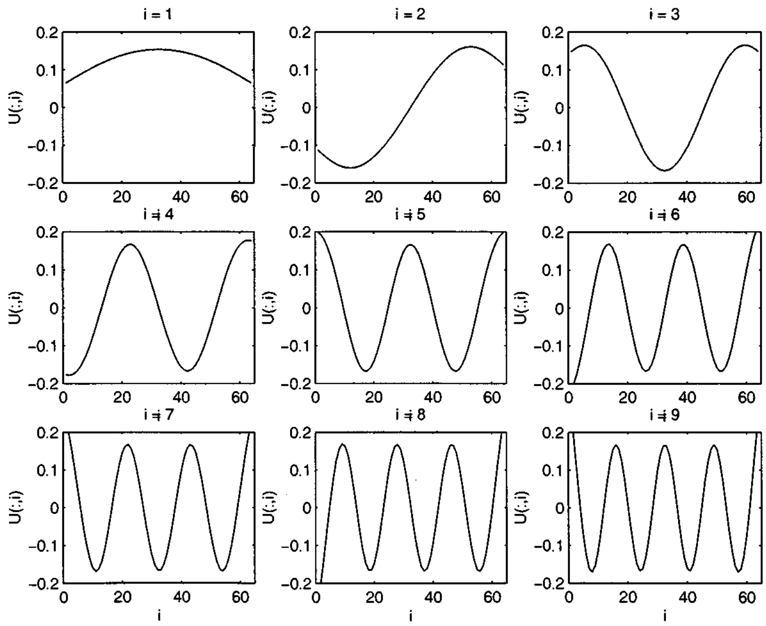

Some Observations

- The singular values decay gradually to zero.
- No gap in the singular value spectrum.
- Condition number cond(A) = "∞."
- Singular vectors have more oscillations as i increases.
- In this problem, # sign changes = $i - 1$.

"Picard" Plots

No Noise

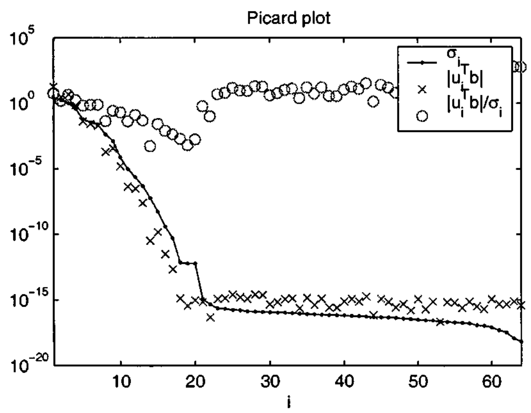

Noise

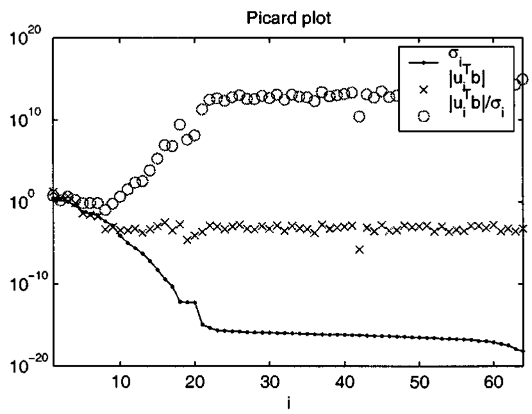

Regularization

Regularization = stabilization: how to deal with solution components corresponding to the small singular values.

Must – somehow – be filtered out or damped.

"Brute force approach": truncate the SVD expansion.

More sophisticated approaches are based on the residual norm

$$\rho(f) = \left\| \int_0^1 K(s,t)\, f(t)\, dt - g(s) \right\| ,$$

with some kind of side constraint(s) to the minimization.

Truncated SVD

Approximate A by the rank-k matrix

$$A_k = \sum_{i=1}^{k} u_i \, \sigma_i \, v_i^T \,, \quad k < n \,.$$

Formulation of the TSVD problem

$$\min \|x\|_2 \quad \text{subject to} \quad \|A_k \, x - b\|_2 = \min .$$

The TSVD solution is

$$x_k = \sum_{i=1}^{k} \frac{u_i^T b}{\sigma_i} \, v_i \,.$$

But minimum 2-norm of x is often undesirable.

The Smoothing Norm

Let the smoothing norm $\omega(f)$ measure the "size" of the solution f.

Example:

$$\omega(f)^2 = \int_0^1 |f^{(p)}(t)|^2 \, dt$$

1. Minimize $\rho(f)$ s.t. $\omega(f) \le \delta$.
2. Minimize $\omega(f)$ s.t. $\rho(f) \le \alpha$.
3. Tikhonov: $\min \left\{ \rho(f)^2 + \lambda^2 \, \omega(f)^2 \right\}$.

SVD Analysis of Discrete Tikhonov Regularization

Can write the discrete Tikhonov solution x_λ in terms of the SVD of A

$$x_\lambda = \sum_{i=1}^{n} \frac{\sigma_i^2}{\sigma_i^2 + \lambda^2} \, \frac{u_i^T b}{\sigma_i} \, v_i$$

Filters components when $\lambda^2 > \sigma_i^2$, i.e., components with small σ_i.

Tikhonov Regularization in General Form

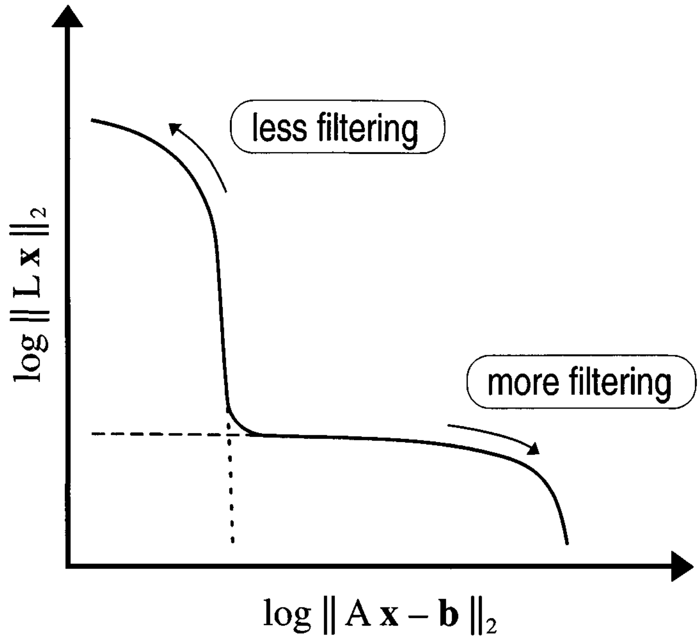

Efficient Implementation The original formulation

$$\min \left\{ \|A\,x - b\|_2^2 + \lambda^2 \, \|x\|_2^2 \right\}.$$

Two alternative formulations

$$(A^T A + \lambda^2 I)\, x = A^T b$$

$$\min \left\| \begin{pmatrix} A \\ \lambda\, I \end{pmatrix} x - \begin{pmatrix} b \\ 0 \end{pmatrix} \right\|_2$$

The first shows that we have a linear problem. The second shows how to solve it stably:

- treat it as a least squares problem
- utilize the sparsity

Least Squares with a Quadratic Constraint

Alternative formulations of Tikhonov regularization

$$\min \|A\,x - b\|_2 \qquad \text{subject to} \qquad \|x - x^*\|_2 \le \alpha$$
$$\min \|x - x^*\|_2 \qquad \text{subject to} \qquad \|A\,x - b\|_2 \le \delta\,,$$

Corresponds to the intersection of the L-curve and the horizontal line $\|x - x^*\|_2 = \alpha$, or the vertical line $\|A\,x - b\|_2 = \delta$.

Requires a root finder, such as Newton's method.

Inequality Constraints Three important constraints to the

solution: nonnegativity, monotonicity, convexity. All three can be put in the general form $G\,x \ge 0$:

$$x \ge 0 \qquad \text{(nonnegativity)}$$
$$L_1\,x \ge 0 \qquad \text{(monotonicity)}$$
$$L_2\,x \ge 0 \qquad \text{(convexity)}$$

where L_1 and L_2 approximate the first and second derivative operators, respectively.

The resulting least squares problem is

$$\min \left\| \begin{pmatrix} A \\ \lambda I \end{pmatrix} x - \begin{pmatrix} b \\ 0 \end{pmatrix} \right\|_2 \qquad \text{subject to} \quad G\,x \ge 0\,.$$

Discrete Linear Ill-Posed Problems, Part 3 of 4
The Regularization Parameter

1. Perturbation and regularization error
2. The Picard condition
3. Parameter choice
   (a) L-curve
   (b) Generalized Cross Validation

Relation to the Regularization Parameter

The regularization parameter (λ or k) determines how many SVD components are included in the regularized solution.
If we write

$$x_{\text{reg}} = A^{\#} b \qquad \text{and} \qquad b = b_{\text{exact}} + e\,,$$

then λ or k should *balance* the perturbation and regularization errors

$$\begin{aligned} x_{\text{exact}} - x_{\text{reg}} &= A^{\dagger} b_{\text{exact}} - A^{\#} b \\ &= (A^{\dagger} - A^{\#})\, b_{\text{exact}} - A^{\#} e\,. \end{aligned}$$

A typical situation in practice:

- The norm $\|e\|_2$ is not known.
- The errors are fixed (not practical to repeat measurements).

The Discrete Picard Condition

The relative decay of the singular values and the Fourier coefficients plays a major role!

The Discrete Picard Condition. Let τ_A denote the level at which the singular values of A level off. Then the discrete Picard condition is satisfied if, for all singular values $\sigma_i > \tau_A$, the corresponding coefficients $|u_i^T b_{\mathrm{exact}}|$, on the average, decay to zero faster than the σ_i.

Can base the analysis on the moving geometric mean

$$\rho_i = \sigma_i^{-1} \left(\prod_{j=i-q}^{i+q} |u_i^T b| \right)^{1/(2q+1)} , \quad i = 1+q, \ldots, n-q \, .$$

Properties of the L-Curve

Theorem 4.5.1. The semi-norm $\|x_\lambda\|_2$ is a monotonically decreasing convex function of the norm $\|A\, x_\lambda - b\|_2$.

Define x_{ls} = least squares solution and

$$\delta_0 = \|(I_m - U\, U^T)\, b\|_2 \qquad \text{(inconsistency measure)}$$

Then

$$\delta_0 \le \|A\, x_{L,\lambda} - b\|_2 \le \|b\|_2$$
$$0 \le \|x_\lambda\|_2 \le \|x_{\mathrm{ls}}\|_2 \, .$$

More Properties of the L-curve

Any point (δ, η) on the L-curve is a solution to the following two inequality-constrained least squares problems:

$$\delta = \min \|A\,x - b\|_2 \qquad \text{subject to} \qquad \|x\|_2 \le \eta$$

$$\eta = \min \|x\|_2 \qquad \text{subject to} \qquad \|A\,x - b\|_2 \le \delta \,.$$

Can study the L-curve by means of the expressions

$$\begin{aligned} \|x_{\text{reg}}\|_2^2 &= \sum_{i=1}^{p} \left(f_i \, \frac{u_i^T b}{\sigma_i} \right)^2 \\ \|A\,x_{\text{reg}} - b\|_2^2 &= \sum_{i=1}^{p} \left((1 - f_i)\, u_i^T b \right)^2 + \delta_0^2 \,. \end{aligned}$$

The L-Shaped Appearance of the L-curve

Analysis: study L-curves for b_{exact} and e.

Result: the L-curve has two distinctly different parts.

- The horizontal part where the regularization errors dominate.
- The vertical part where the perturbation errors dominate.

The optimal regularization parameter must lie somewhere near the L-curve's corner.

The corner is located approximately at

$$\left(\|A\,x_\lambda - b\|_2 \,,\, \|x_\lambda\|_2 \right) \approx \left(\sqrt{\sigma_0^2 n} \,,\, \|x_{\text{exact}}\|_2 \right)$$

Analysis of the L-Curve

Assume that b lies in the range of A, such that

$$u_i^T b = 0 \,, \qquad i = n+1, \ldots, m \,.$$

Can analyze the L-curve by means of the expressions

$$\begin{aligned} \|x_\lambda\|_2^2 &= \sum_{i=1}^{m} \left(\frac{\sigma_i^2}{\sigma_i^2 + \lambda^2} \frac{u_i^T b}{\sigma_i} \right)^2 \\ \|b - A\, x_\lambda\|_2^2 &= \sum_{i=1}^{n} \left(\frac{\lambda^2}{\sigma_i^2 + \lambda^2} \, u_i^T b \right)^2 \,. \end{aligned}$$

Recall that $b = b_{\text{exact}} + e$.

The Flat and Steep Parts

The component b_{exact} dominates when λ is small:

$$\begin{aligned} \|x_\lambda\|_2 &\approx \|x_{\text{exact}}\|_2 \\ \|b - A\, x_\lambda\|_2^2 &\approx \lambda^4 \sum_{i=1}^{n} \left(\frac{u_i^T b}{\sigma_i^2} \right)^2 \end{aligned}$$

The error e dominates when λ is large ($u_i^T e \approx \pm\epsilon_0$):

$$\begin{aligned} \|A^\# e\|_2^2 &\approx \lambda^{-4} \sum_{i=1}^{n} \left(\sigma_i \, u_i^T e \right)^2 \approx \lambda^{-4} \, \epsilon_0^2 \|A\|_F^2 \\ \|b - A\, A^\# e\|_2^2 &\approx \epsilon_0^2 \sum_{i=1}^{n} \left(\frac{\lambda^2}{\sigma_i^2 + \lambda^2} \right)^2 \approx n\, \epsilon_0^2 \,. \end{aligned}$$

The Key Idea

The flat and the steep parts of the L-curve represent solutions that are dominated by regularization errors and perturbation errors.

The balance between these two errors must occur near the L-curve's corner.

The two parts – as well as the corner – are emphasized in log-log scale.

Log-log scale is insensitive to scalings of A and b.

The Curvature of the L-Curve

Want to derive an analytical expression for the L-curve's curvature κ in log-log scale; define

$$\eta = \|x_\lambda\|_2^2 \,, \qquad \rho = \|A\,x_\lambda - b\|_2^2$$

and

$$\hat{\eta} = \log \eta \,, \qquad \hat{\rho} = \log \rho \,.$$

Then the curvature is given by

$$\kappa = 2\,\frac{\hat{\rho}'\hat{\eta}'' - \hat{\rho}''\hat{\eta}'}{((\hat{\rho}')^2 + (\hat{\eta}')^2)^{3/2}} \,.$$

Generalized Cross Validation (GCV)

Statistical approach: Seeks to minimize the expected value of

$$\|Ax - b^{\text{exact}}\|_2$$

Notice b^{exact}. Another viewpoint: If any measurement is left out then a solution from the remaining should predict the left out measurement.

Minimize the GCV-functional

$$\mathcal{G}(\lambda) = \frac{\|Ax_\lambda - b\|_2^2}{\text{trace}(I - AA\#)}$$

Experiences with GCV and the L-Curve

- The GCV method, on the average, leads to a slight oversmoothing which accounts for the increased average error, compared to the optimal results. Occasionally GCV undersmooths, leading to larger errors.
- The L-curve criterion consistently oversmooths—there is no $\lambda < \lambda_{\text{opt}}$. Hence, the average error is greater than that for GCV.
- The L-curve criterion is more *robust* than GCV, in the sense that the L-curve criterion never leads to large errors while GCV occasionally does.

Discrete Linear Ill-Posed Problems — Part 4 of 4
Iterative Methods

Two different classes of iterative methods.

- Iterative solution of a regularized problem, such as Tikhonov

$$\left(A^T A + \lambda^2 L^T L\right) x = A^T b\,.$$

  Challenge: to construct a good preconditioner!

- Iterate on the unregularized system, e.g., on

$$A^T A\, x = A^T b$$

  and use the iteration number as the regularization parameter!

Advantages of Iterative Methods

- The matrix A is never altered, only "touched" via matrix-vector multiplications $A\,x$ and $A^T y$.
- The matrix A is not explicitly required – we only need a "black box" that computes the action of A or the underlying operator.
- Produces a natural sequence of regularized solutions; stop when the solution is "satisfactory" (parameter choice).
- Atomic operations are easy to parallelize.

Disadvantages

- Convergence may be (very) slow.

ART or Kaczmarz's Method

Kaczmarz's method = algebraic reconstruction technique (ART):

$$x \leftarrow x + \frac{b_i - a_i^T x}{\|a_i\|_2^2} a_i \, , \qquad i = 1, \ldots, m \, ,$$

where b_i is the ith component b.

Mathematically equivalent to Gauss-Seidel's method for the problem

$$x = A^T y \, , \qquad A\,A^T y = b \, .$$

Used successfully in computerized tomography.

In general: fast initial convergence, then slow.

Conjugate Gradients

CGLS: CG applied to the normal equations $A^T A\,x = A^T b$:

$$\begin{aligned}
\alpha_k &= \|A^T r^{(k-1)}\|_2^2 \,/\, \|A\,d^{(k-1)}\|_2^2 \\
x^{(k)} &= x^{(k-1)} + \alpha_k\, d^{(k-1)} \\
r^{(k)} &= r^{(k-1)} - \alpha_k\, A\,d^{(k-1)} \\
\beta_k &= \|A^T r^{(k)}\|_2^2 \,/\, \|A^T r^{(k-1)}\|_2^2 \\
d^{(k)} &= A^T r^{(k)} + \beta_k\, d^{(k-1)}
\end{aligned}$$

where $r^{(k)}$ is the residual vector $r^{(k)} = b - A\,x^{(k)}$, and $d^{(k)}$ is an auxiliary m-vector.

Initialization: starting vector $x^{(0)}$, residual $r^{(0)} = b - A\,x^{(0)}$, and $d^{(0)} = A^T r^{(0)}$.

Slow Convergence

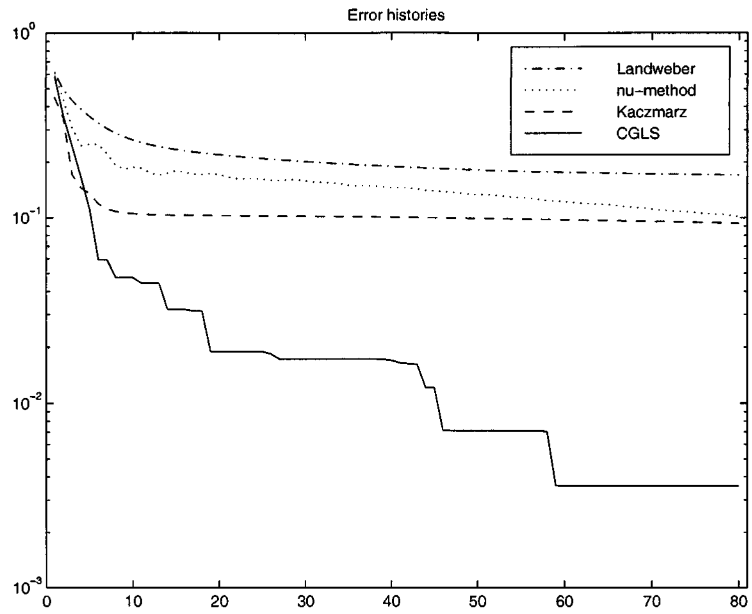

Convergence in the Presence of Data Errors

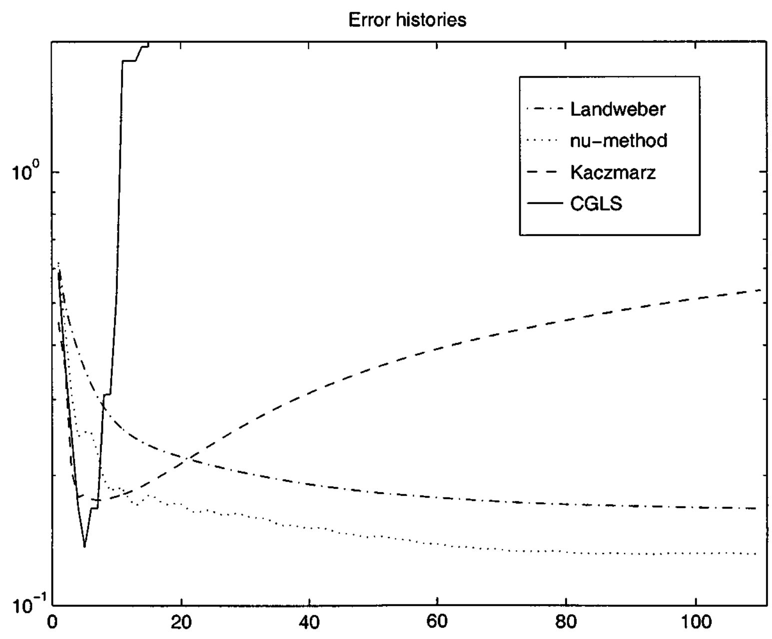

Semi-Convergence

CGLS exhibits *semi-convergence:*

- initial convergence towards x_{exact},
- followed by (slow) convergence to $x_{\text{ls}} = A^{\dagger}b$.

Must stop at the end of the first stage!

A full understanding of this phenomenon is still lacking and is subject of current research.

Pretty Pictures: Filter Factors

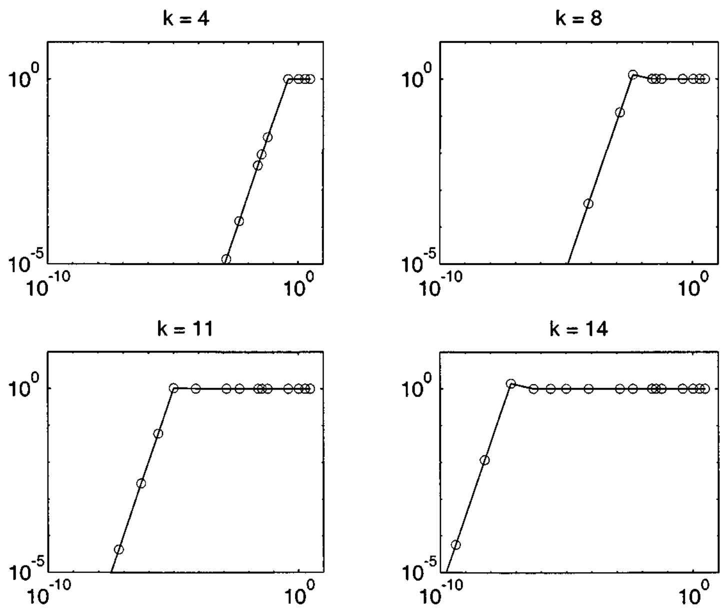

Finite Precision Filter Factors

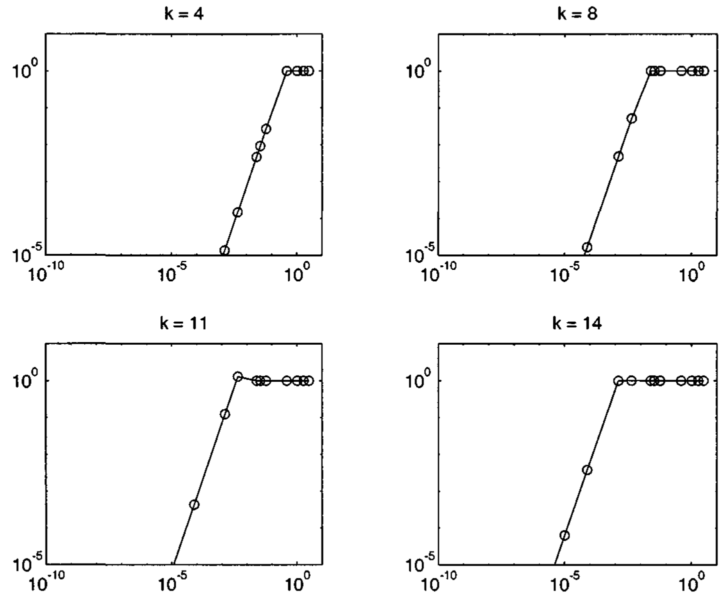

Finite Precision Slow-Down

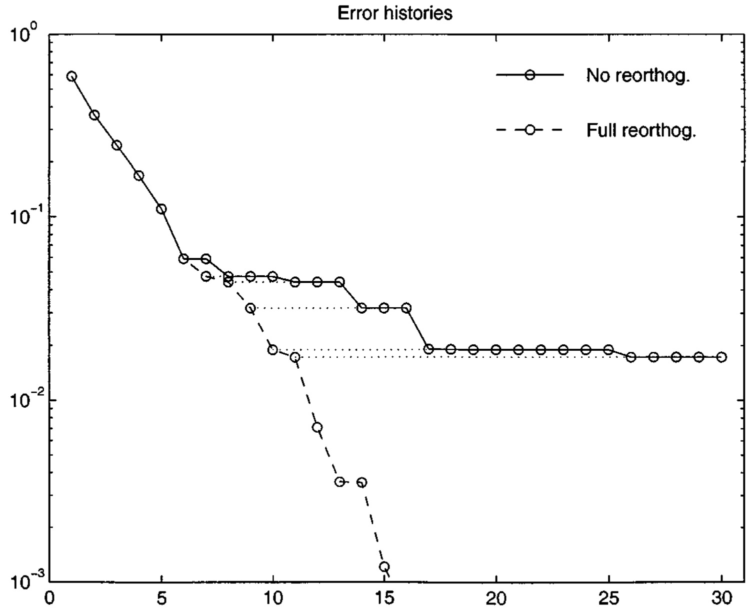

References

Recommendable

- Per Christian Hansen: "Rank-Deficient and Discrete Ill-Posed Problems", SIAM, 1998
- Per Christian Hansen: "Regularization Tools Version 3.1" for Matlab 6.0
- Arnold Neumaier: "Solving Ill-Conditioned and Singular Linear Systems: A Tutorial on Regularization", SIAM Review, 1998.

WWW:

- Per Christian Hansen: http://www.imm.dtu.dk/~pch
- Me: http://www.imm.dtu.dk/~mj