Numerical Methods for Ill-Posed Problems I

Lothar Reichel

Summer School on Applied Analysis
TU Chemnitz, Oct. 4-8, 2010

Outline of Lecture 1:

- Inverse and ill-posed problems
- The singular value decomposition
- Tikhonov regularization

Inverse problems

- Inverse problems arise when one seeks to determine the cause of an observed effect.
  - Helioseismology: Determine the structure of the sun by measurements from earth or space.
  - Medical imaging, e.g., electrocardiographic imaging, computerized tomography.
  - Image restoration: Determine the unavailable exact image from an available contaminated version.
- Inverse problems often are ill-posed.

Ill-posed problems

A problem is said to be ill-posed if it has at least one of the properties:

- the problem does not have a solution,
- the problem does not have a unique solution,
- the solution does not depend continuously on the data.

Linear discrete ill-posed problems

$$Ax = b$$

arise from the discretization of linear ill-posed problems (Fredholm integral equations of the 1st kid) or, naturally, in discrete form (image restoration).

- The matrix A is of ill-determined rank, possibly singular. System may be inconsistent.
- The right-hand side b represents available data that generally is contaminated by an error.

Available contaminated, possibly inconsistents, linear system

$$Ax = b \qquad (1)$$

Unavailable associated consistent linear system with error-free right-hand side

$$Ax = \hat{b} \qquad (2)$$

Let $\hat{x}$ denote the desired solution of (2), e.g., the minimal-norm solution.

Task: Determine an approximate solution of (1) that is a good approximation of $\hat{x}$.

Define the error

$$e := \hat{b} - b \qquad \text{noise}$$

and let

$$\epsilon := \|e\|$$

The choice of solution method depends on whether an estimate of ϵ is available.

How much damage can a little noise really do?

How much noise requires the use of special techniques?

Example 1: Fredholm integral equation of the 1st kind

$$\int_0^\pi \exp(-st)x(t)dt = 2\frac{\sinh(s)}{s}, \quad 0 \leq s \leq \frac{\pi}{2}.$$

Determine solution $x(t) = \sin(t)$.

Discretize integral by Galerkin method using piecewise constant functions. Code baart from Regularization Tools.

This gives a linear system of equations

$$Ax = \hat{b}, \qquad A \in \mathrm{R}^{200\times 200}, \;\; \hat{b} \in \mathrm{R}^{200}.$$

A is numerically singular.

Let the "noise" vector e in b have normally distributed entries with mean zero and

$\epsilon = \|e\| = 10^{-3}\|b\|$

$b := \hat{b} + e$

i.e., 0.1% relative noise

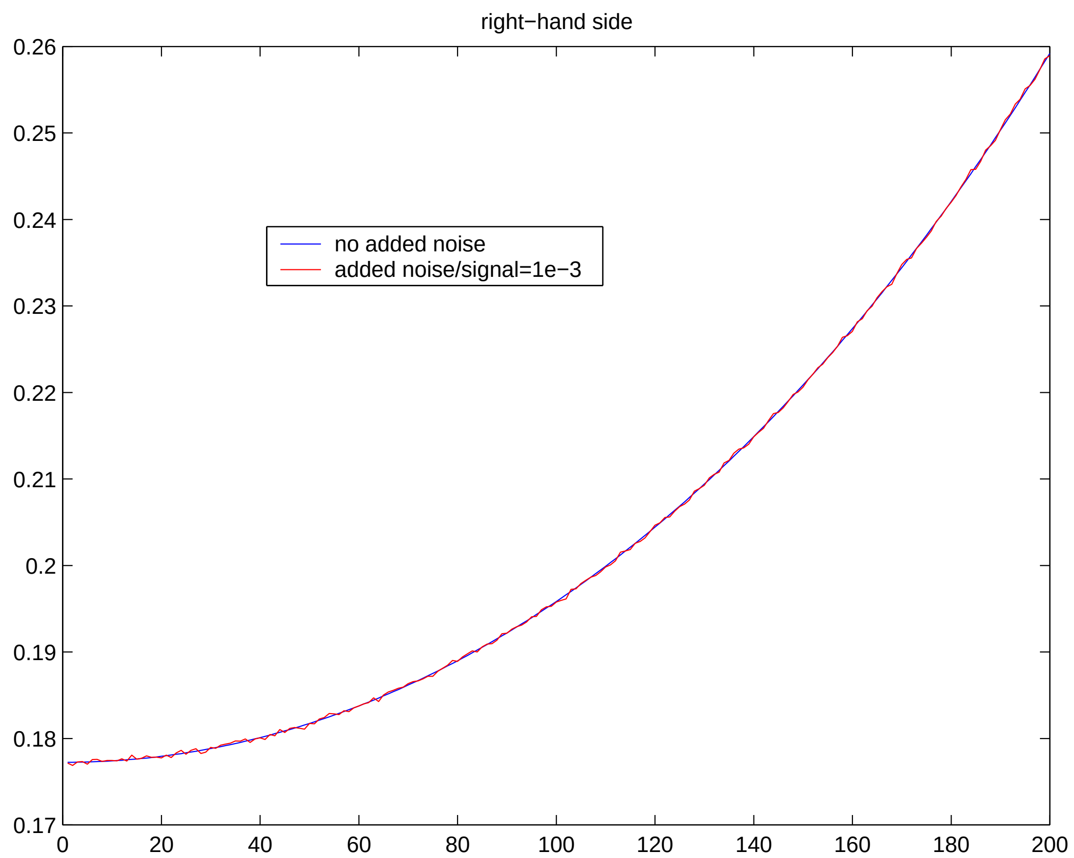
right−hand side
no added noise
added noise/signal=1e−3
0.26
0.25
0.24
0.23
0.22
0.21
0.2
0.19
0.18
0.17
0
20
40
60
80
100
120
140
160
180
200

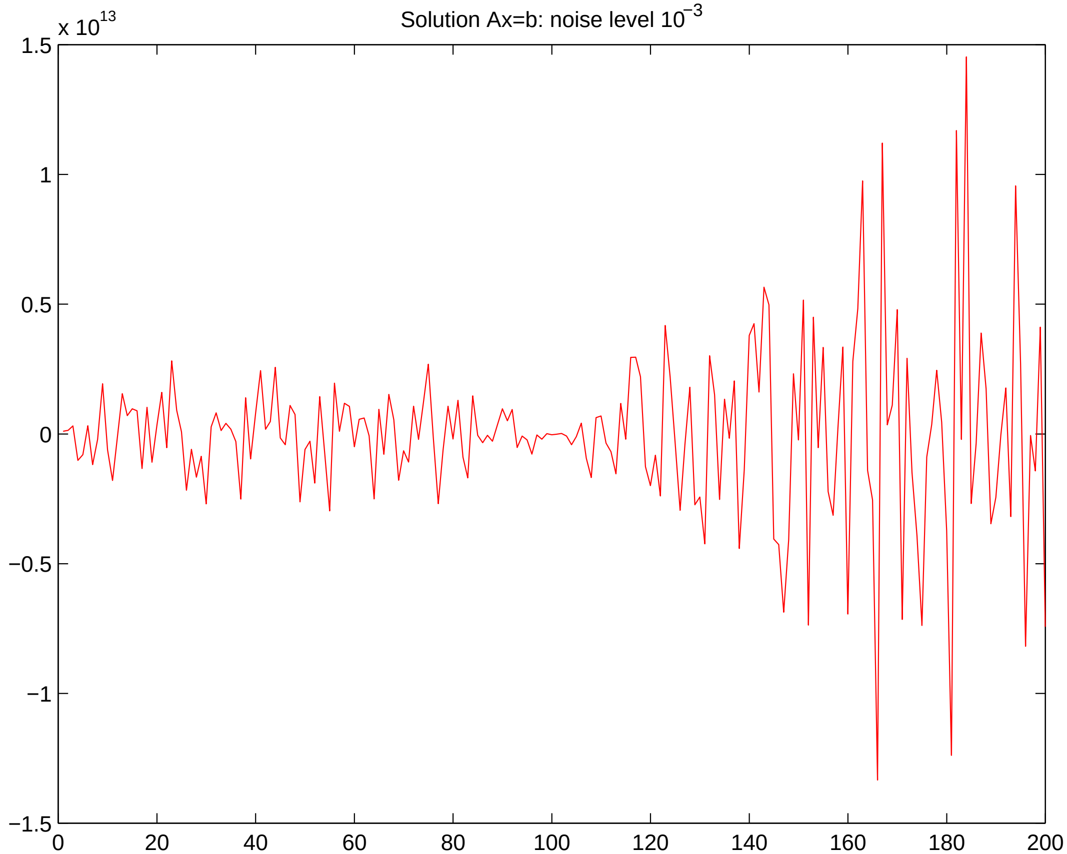
Solution Ax=b: noise level 10^{-3}
x 10^{13}
1.5
1
0.5
0
−0.5
−1
−1.5
0
20
40
60
80
100
120
140
160
180
200

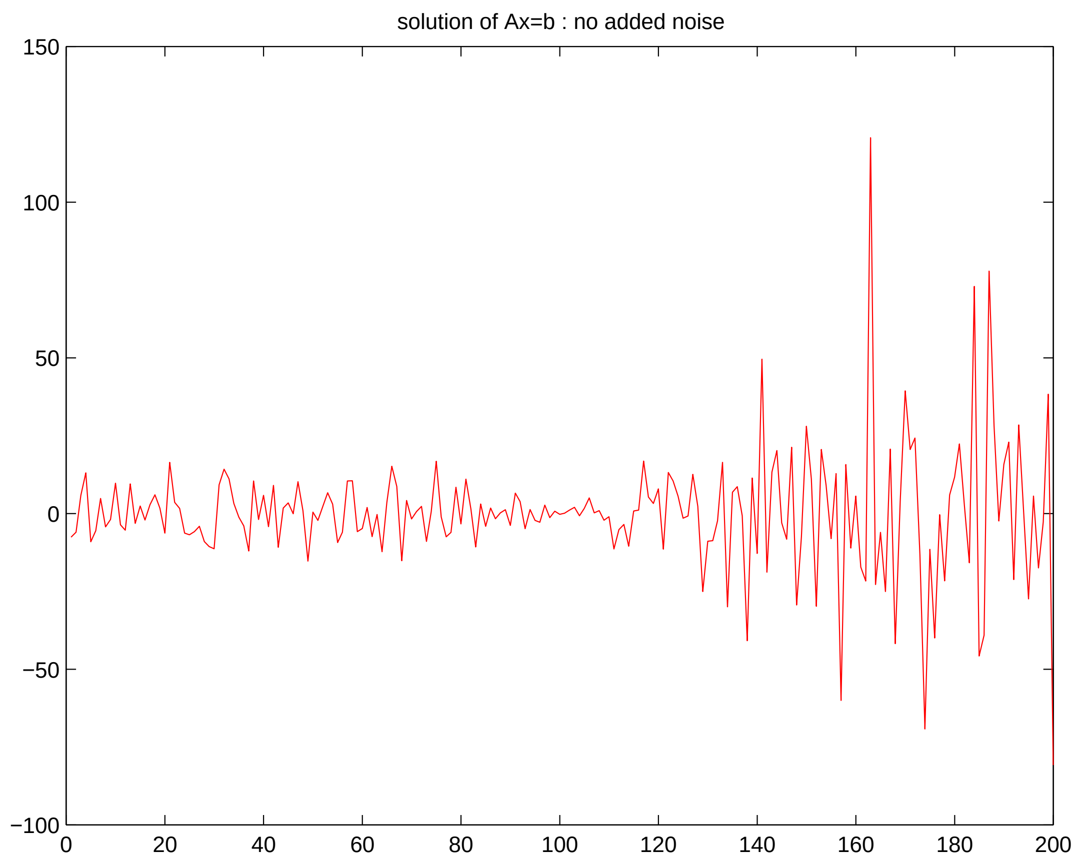
solution of Ax=b : no added noise
150
100
50
0
−50
−100
0
20
40
60
80
100
120
140
160
180
200

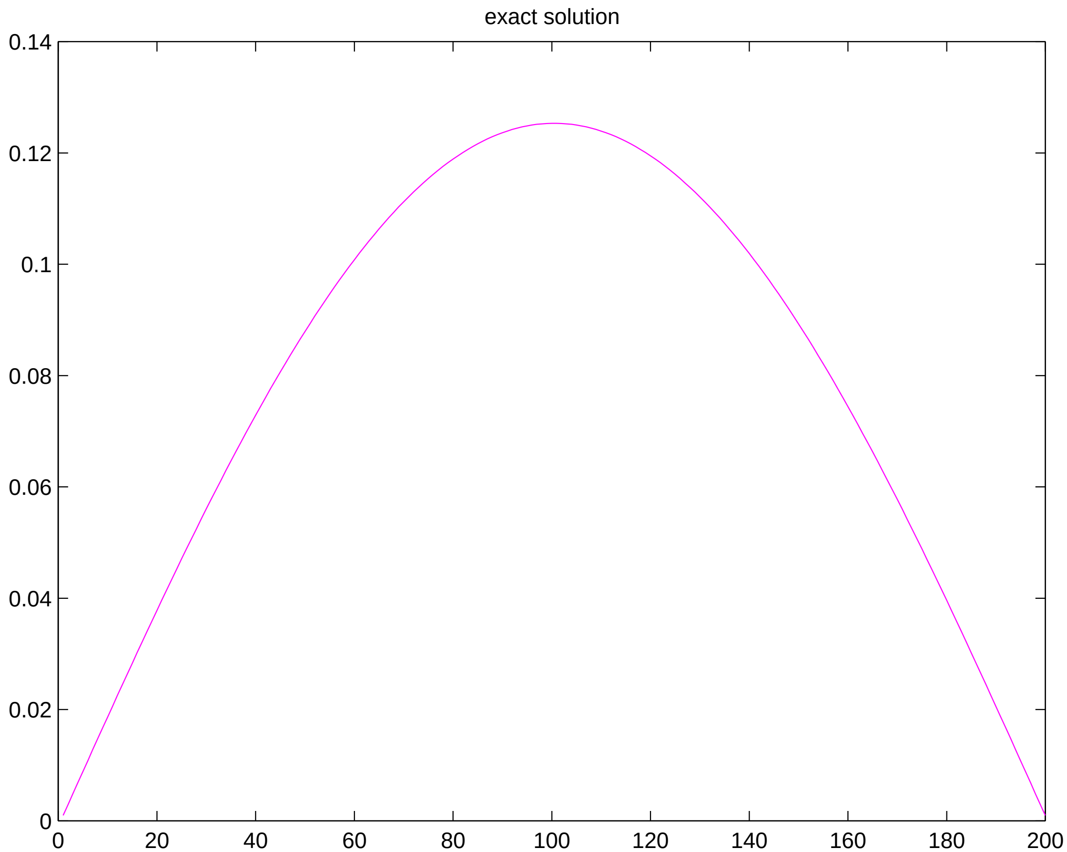

exact solution
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
0
20
40
60
80
100
120
140
160
180
200

The singular value decomposition

The SVD of the $m \times n$ matrix A, $m \geq n$:

$$A = U \Sigma V^T$$

$$\begin{aligned} U &= [u_1, u_2, \ldots, u_m] \qquad \text{orthogonal}, \quad m \times m, \\ V &= [v_1, v_2, \ldots, v_n] \qquad \text{orthogonal}, \quad n \times n, \\ \Sigma &= \text{diag}[\sigma_1, \sigma_2, \ldots, \sigma_n], \quad m \times n \end{aligned}$$

with

$$\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_n \geq 0.$$

Application: Least-squares approximation

Let the matrix $A \in \mathbf{R}^{m\times n}$, $m \geq n$, represent the "model" and the vector $b \in \mathbf{R}^m$ the data. Solve

$$\min_x \|Ax-b\|^2 = \min_x \|U\Sigma V^T x-b\|^2 = \min_x \|\Sigma V^T x-U^T b\|^2.$$

Let $y = [y_1, y_2, \ldots, y_n]^T = V^T x$ and $b' = [b'_1, b'_2, \ldots, b'_m]^T = U^T b$. Then

$$\min_x \|Ax-b\|^2 = \min_y \|\Sigma y-b'\|^2 = \sum_{j=1}^{n}(\sigma_j y_j-b'_j)^2+\sum_{j=n+1}^{m}(b'_j)^2.$$

If A is of full rank, then all $\sigma_j > 0$ and

$$y_j = \frac{b'_j}{\sigma_j}, \qquad 1 \leq j \leq n,$$

yields the solution

$$x = Vy.$$

If some $\sigma_j = 0$, then least-squares solution not unique. Often one is interested in the least-squares solution of minimal norm. Arbitrary components y_j are set to zero.

Assume that

$$\sigma_1 \geq \sigma_2 \ldots \geq \sigma_\ell > \sigma_{\ell+1} = \ldots = \sigma_n = 0.$$

Then A is of rank ℓ. Introduce the diagonal matrix

$$\Sigma^\dagger = \text{diag}[1/\sigma_1, 1/\sigma_2, \ldots, 1/\sigma_\ell, 0, \ldots, 0], \quad n \times m.$$

The matrix

$$A^\dagger = V \Sigma^\dagger U^T$$

is known as the Moore-Penrose pseudoinverse of A.

The solution of the least-squares problem

$$\min_x \|Ax - b\|$$

of minimal Eucliden norm can be expressed as

$$x = A^\dagger b.$$

Moreover,

$$A^\dagger A = I, \qquad AA^\dagger = P_{\mathcal{R}(A)}.$$

Note

$$A = U\Sigma V^T = \sum_{j=1}^{n} \sigma_j u_j v_j^T.$$

Define

$$A_k := \sum_{j=1}^{k} \sigma_j u_j v_j^T, \qquad 1 \le k \le \ell.$$

Then A_k is of rank k; A_k is the sum of k rank-one matrices $\sigma_j u_j v_j^T$.

Moreover,

$$\|A - A_k\| = \min_{\text{rank}(B) \le k} \|A - B\| = \sigma_{k+1},$$

i.e., A_k is the closest matrix of rank $\le k$ to A.

Let $b = \hat{b} + e$, where e denotes an error. Then

$$\begin{aligned} x := A^{\dagger} b &= \sum_{j=1}^{\ell} \frac{u_j^T b}{\sigma_j} v_j \\ &= \sum_{j=1}^{\ell} \frac{u_j^T \hat{b}}{\sigma_j} v_j + \sum_{j=1}^{\ell} \frac{u_j^T e}{\sigma_j} v_j \\ &= \hat{x} + \sum_{j=1}^{\ell} \frac{u_j^T e}{\sigma_j} v_j. \end{aligned}$$

If $\sigma_\ell > 0$ tiny, then

$$\frac{u_\ell^T e}{\sigma_\ell}$$

might be huge and x a meaningless approximation of $\hat{x}$.

Recall

$$A_k = \sum_{j=1}^{k} \sigma_j u_j v_j^T$$

best rank-k approximation of A.

Pseudoinverse of A_k:

$$A_k^\dagger := \sum_{j=1}^{k} \sigma_j^{-1} v_j u_j^T, \quad \sigma_k > 0$$

Approximate $\hat{x}$ by

$$\begin{aligned} x_k &:= A_k^{\dagger} b \\ &= \sum_{j=1}^{k} \frac{u_j^T b}{\sigma_j} v_j \\ &= \sum_{j=1}^{k} \frac{u_j^T \hat{b}}{\sigma_j} v_j + \sum_{j=1}^{k} \frac{u_j^T e}{\sigma_j} v_j. \end{aligned}$$

for some $k \leq \ell$.

How to choose k?

Example 1 cont'd: Singular values of A

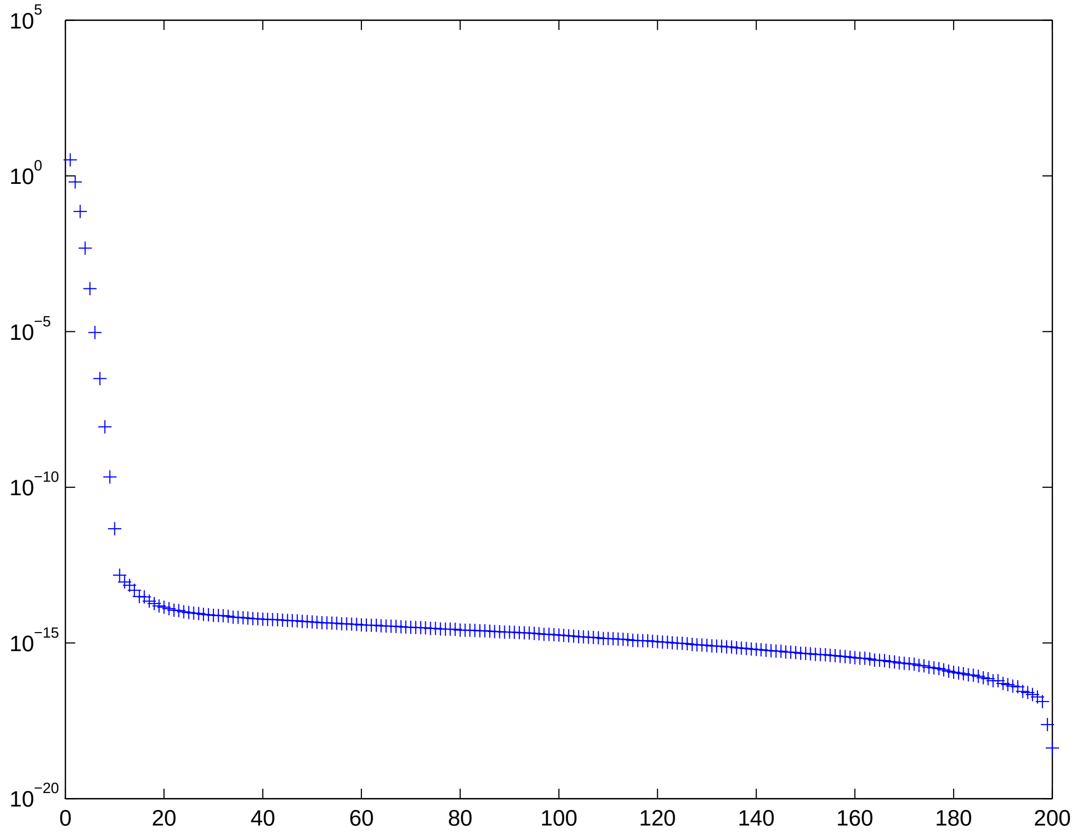

Example 1 cont'd: Right-hand side without noise

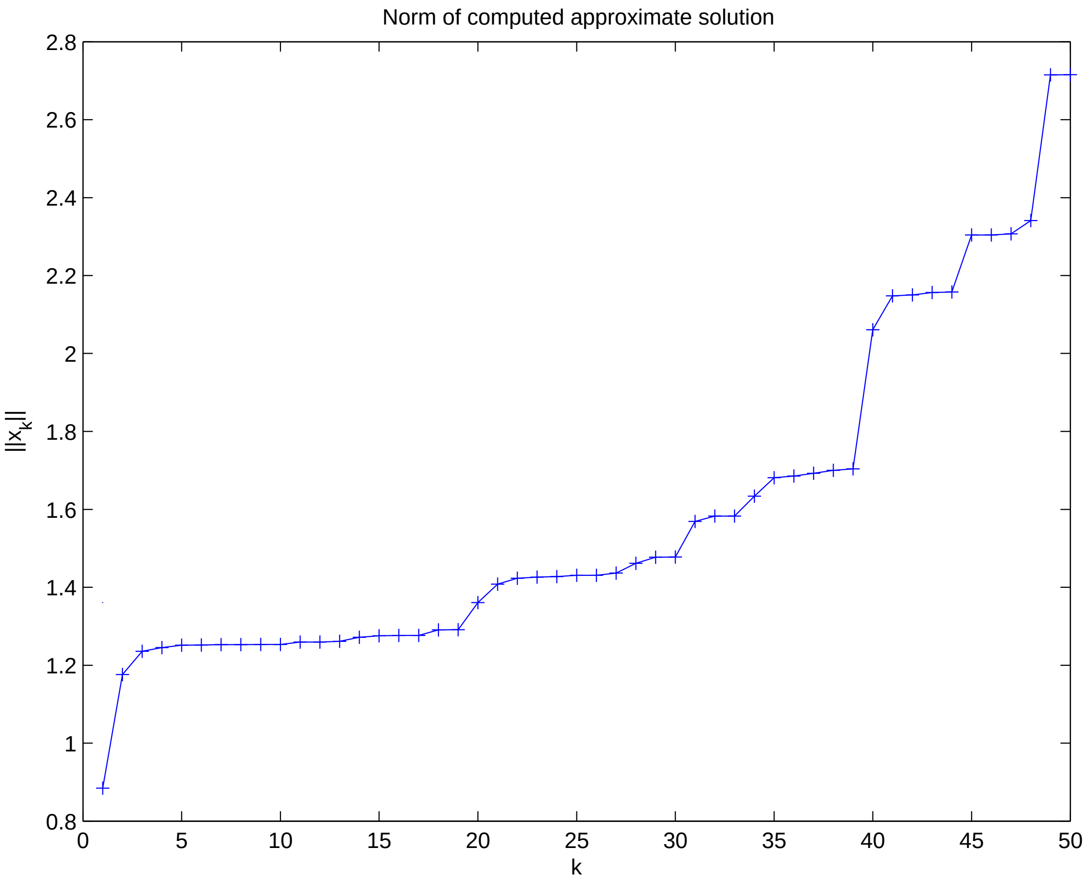

Example 1 cont'd: Right-hand side without noise

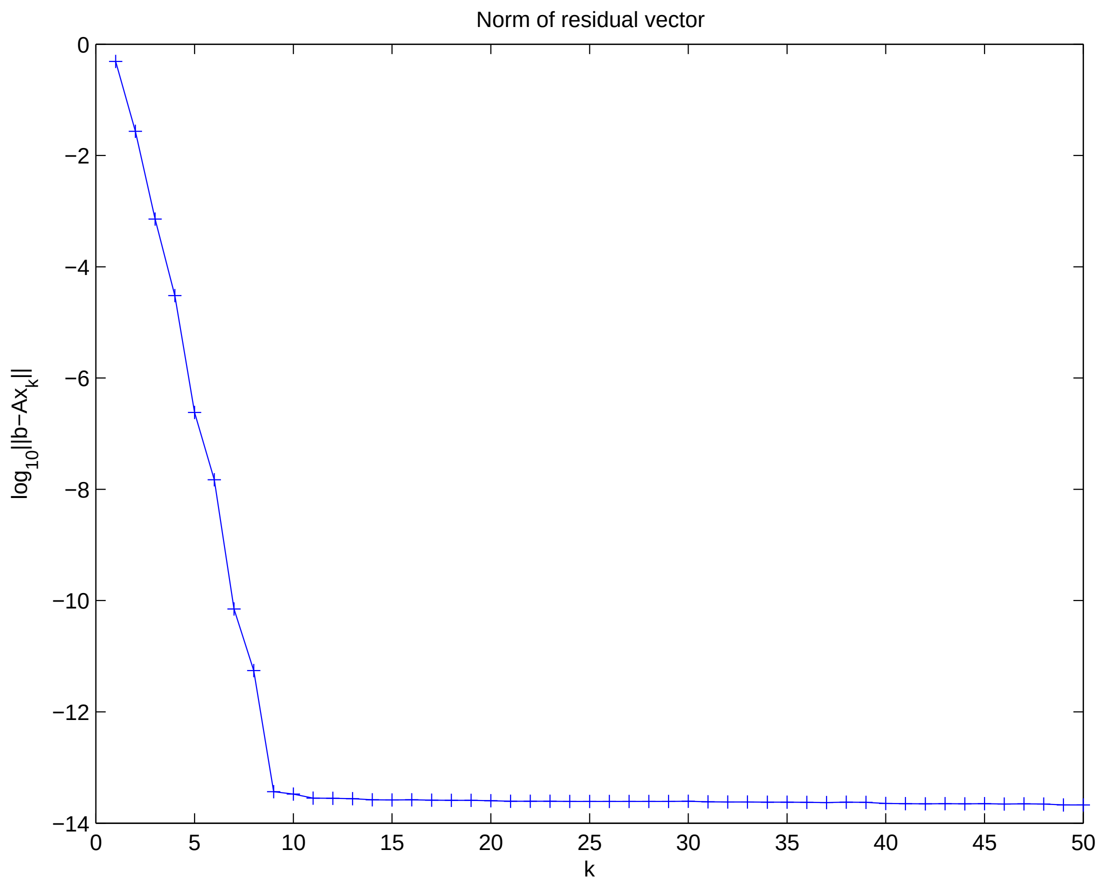

Example 1 cont'd: Right-hand side without noise

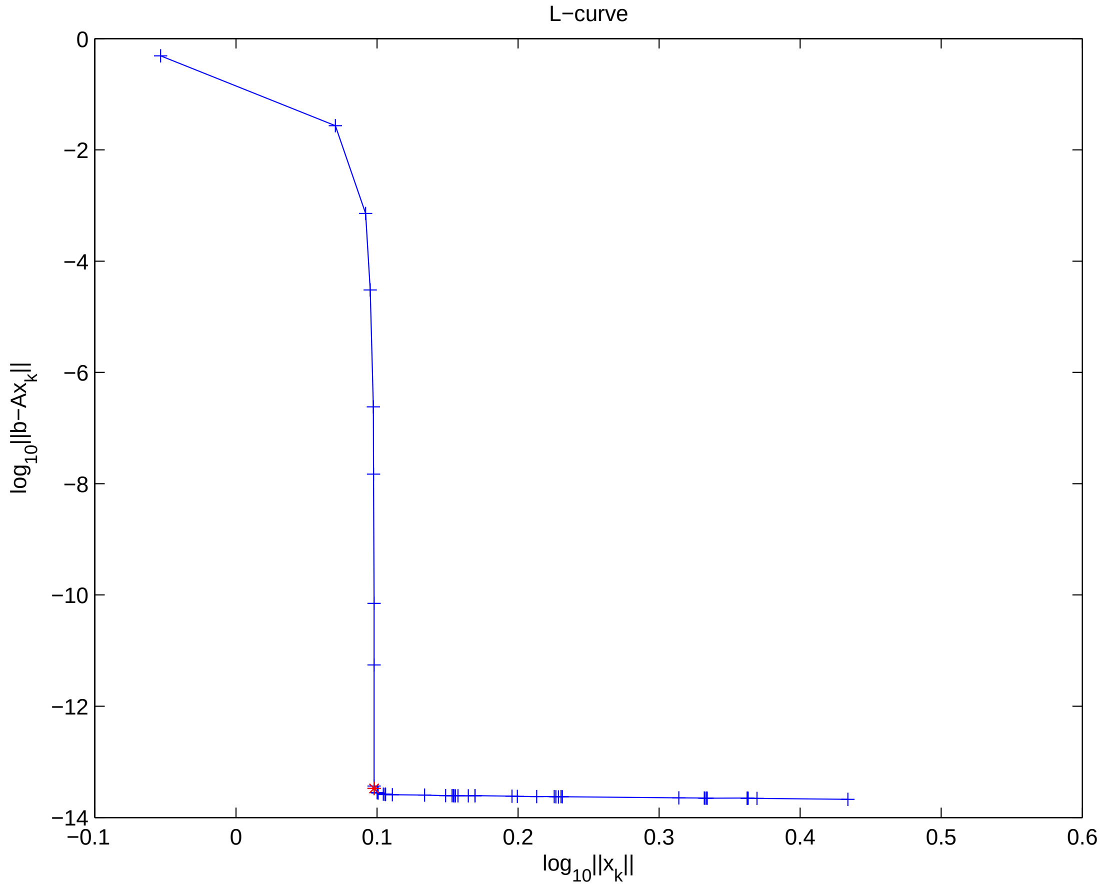

Example 1 cont'd: Right-hand side without noise: Blow-up

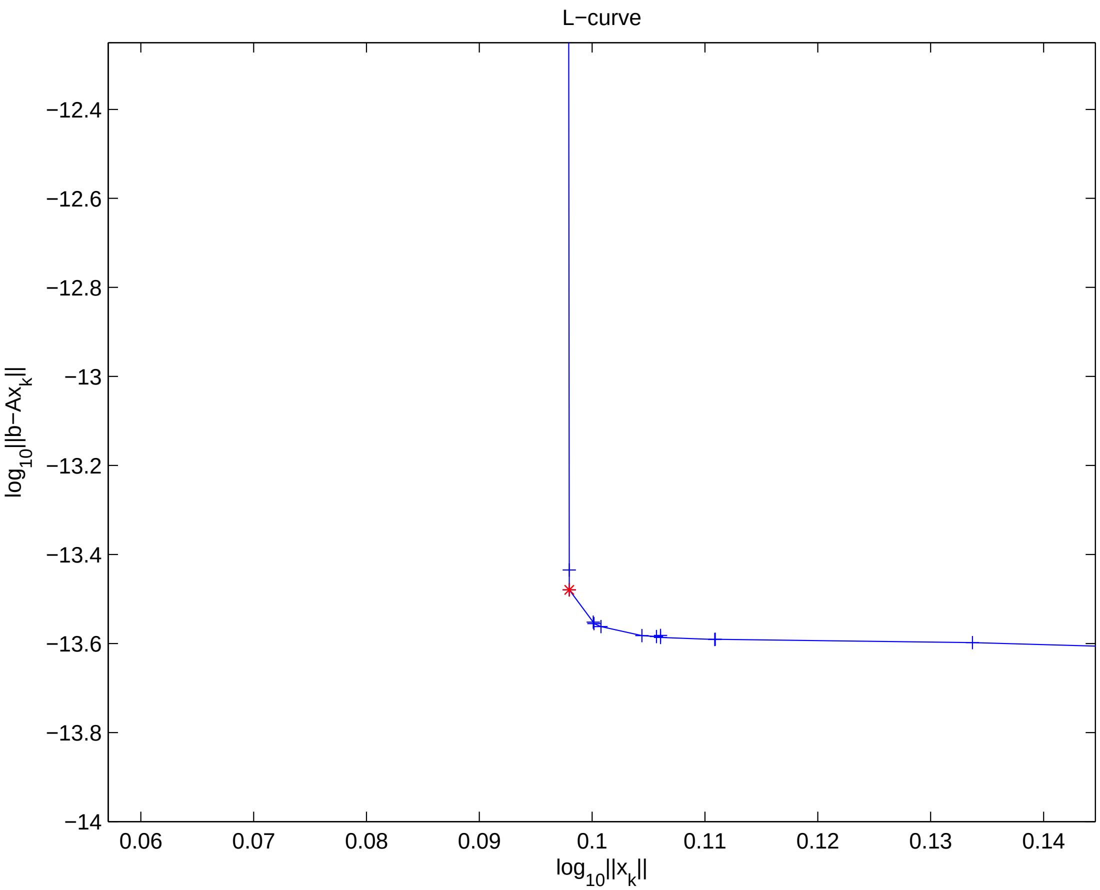

Example 1 cont'd: Exact and computed solutions

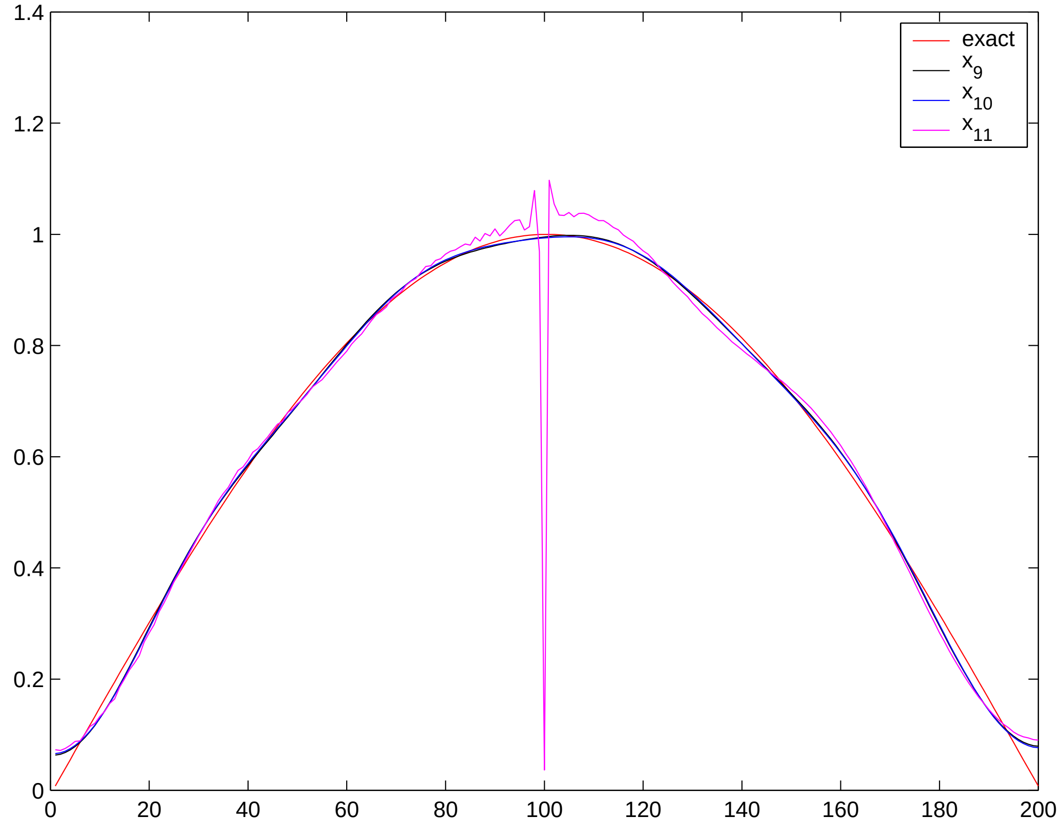

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

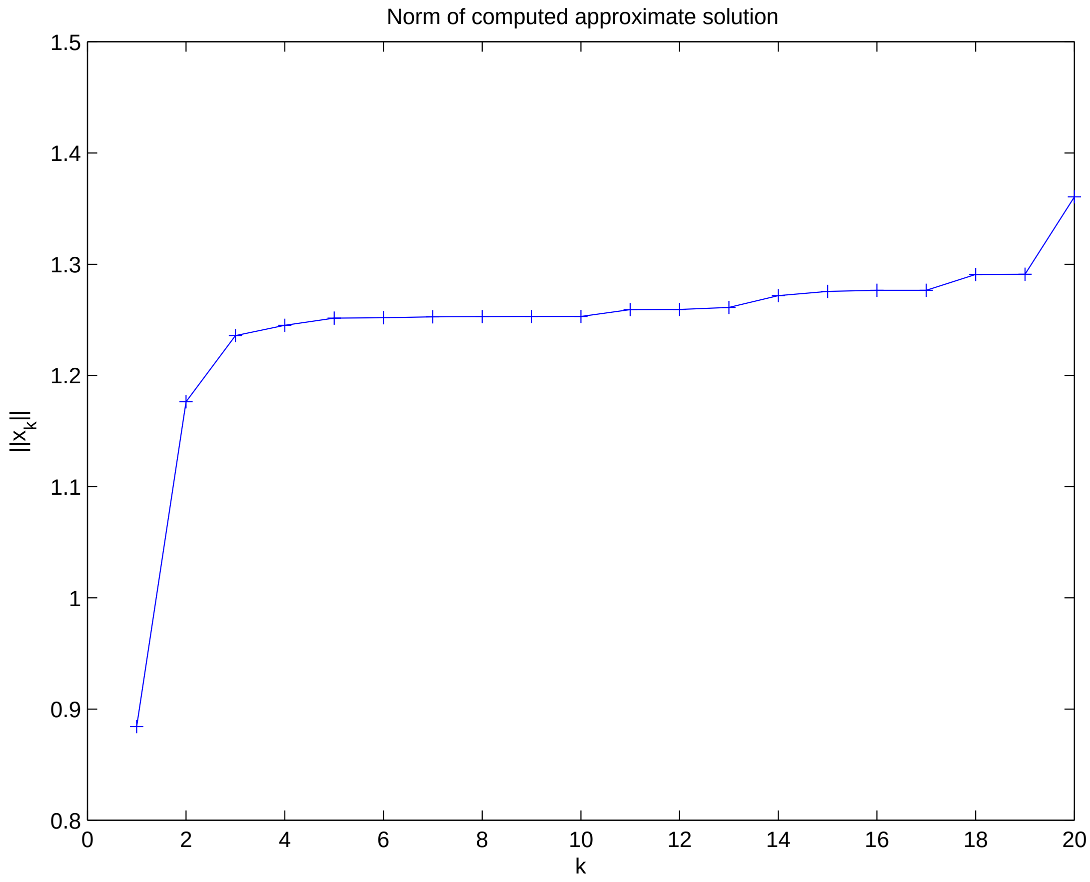

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

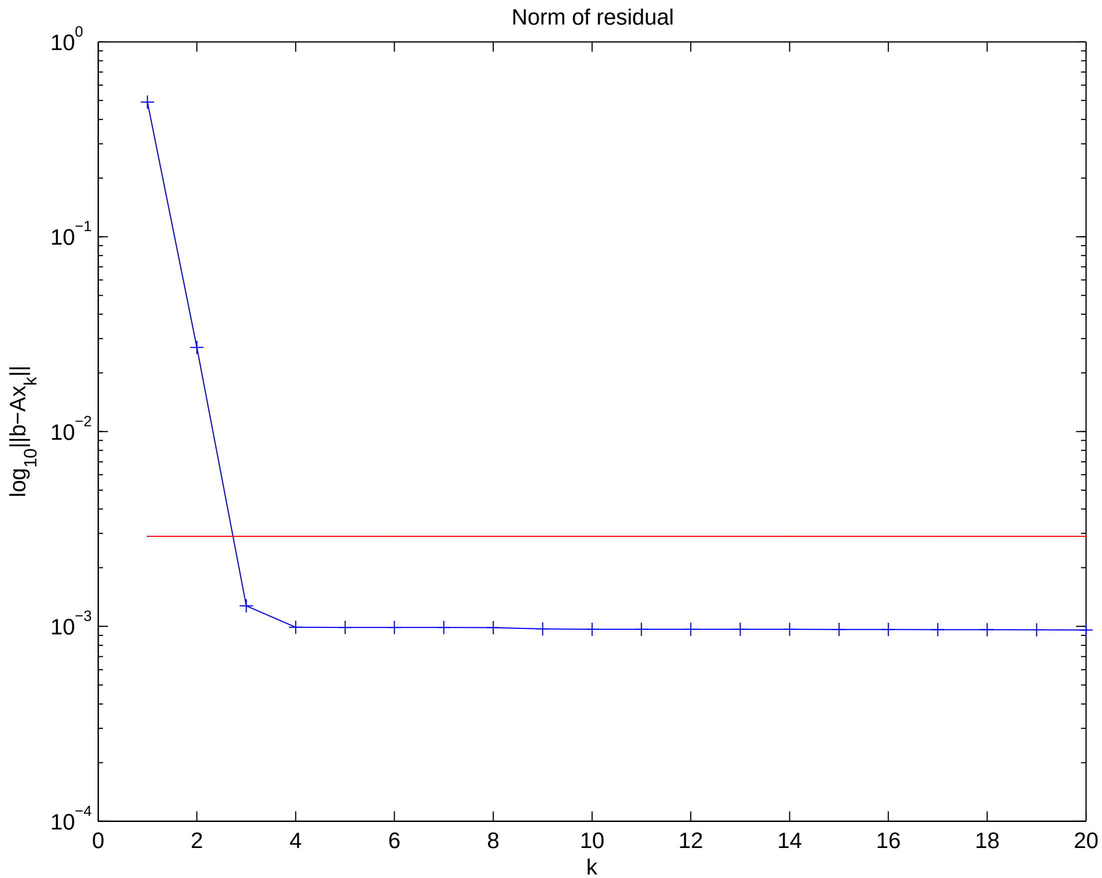

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

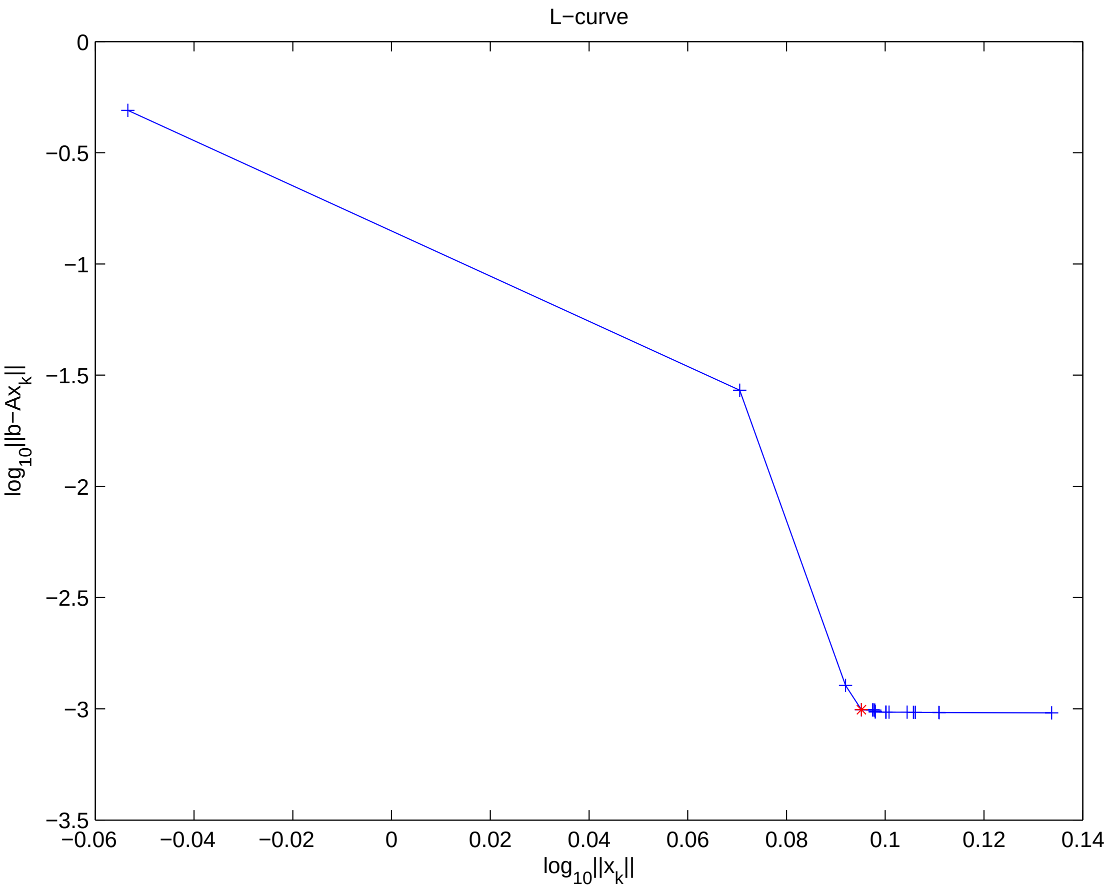

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

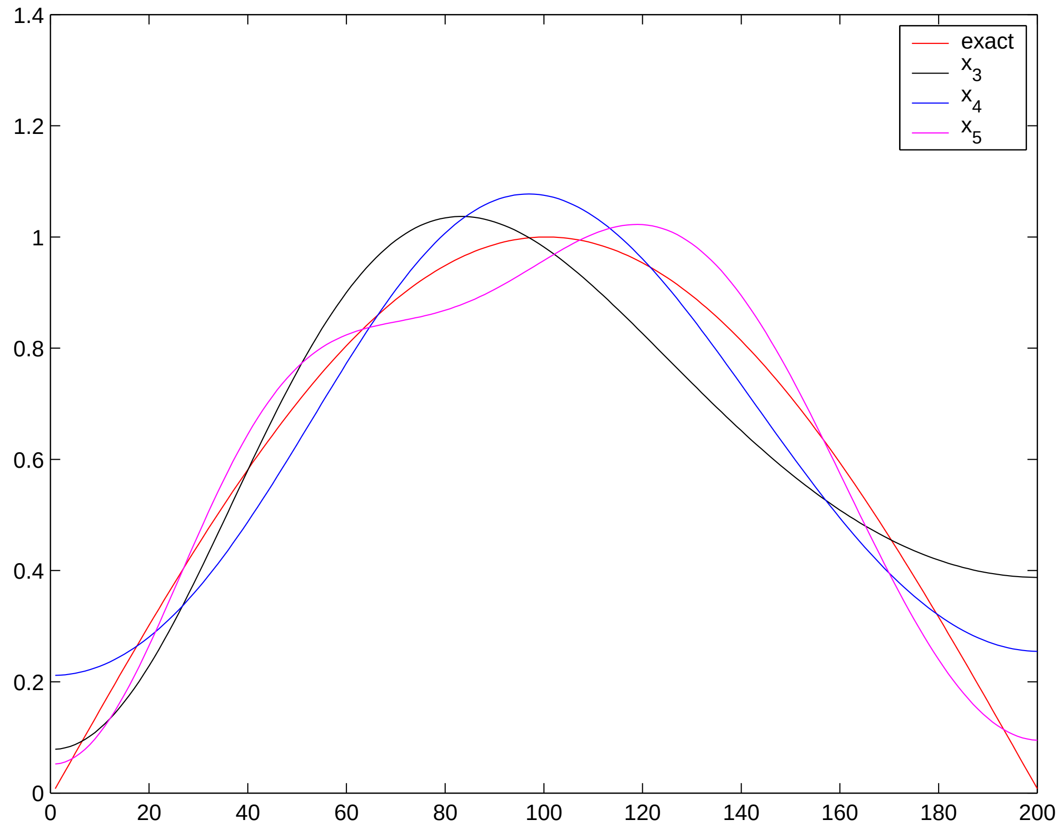

The SVD in Hilbert space

$A : \mathcal{X} \to \mathcal{Y}$ compact linear operator

$\mathcal{X}$, $\mathcal{Y}$ Hilbert spaces,
$\langle \cdot, \cdot \rangle$ inner products in $\mathcal{X}$ and $\mathcal{Y}$,
$\| \cdot \|$ associated norms.

Compute minimal-norm solution $\hat{x} \in \mathcal{X}$ of

$$Ax = \hat{y}, \qquad \hat{y} \in \mathcal{R}(A),$$

by the SVD of A.

Singular triplets $\{\sigma_j, u_j, v_j\}_{j=1}^{\infty}$ of A:

σ_j singular value,

$$\sigma_1 \geq \sigma_2 \geq \sigma_3 \geq \ldots > 0,$$

$u_j \in \mathcal{Y}$ left singular function,
$v_j \in \mathcal{X}$ right singular function, satisfy

$$\langle u_j, u_\ell \rangle = \langle v_j, v_\ell \rangle = \begin{cases} 1, & j = \ell, \\ 0, & j \neq \ell, \end{cases}$$

and

$$
\begin{aligned}
Av_j &= \sigma_j u_j, \qquad A^* u_j = \sigma_j v_j, \qquad j = 1, 2, 3, \ldots\ , \\
Ax &= \sum_{j=1}^{\infty} \sigma_j \langle x, v_j \rangle u_j, \qquad \forall x \in \mathcal{X}, \\
A^* y &= \sum_{j=1}^{\infty} \sigma_j \langle y, u_j \rangle v_j, \qquad \forall y \in \mathcal{Y},
\end{aligned}
$$

where $A^* : \mathcal{Y} \to \mathcal{X}$ the adjoint of A.

A compact $\Rightarrow$ σ_j cluster at zero.

Minimal-norm solution

$$\hat{x} = \sum_{j=1}^{\infty} \frac{\langle \hat{y}, u_j \rangle}{\sigma_j} v_j.$$

$\hat{x} \in \mathcal{X}$ implies Picard condition:

$$\sum_{j=1}^{\infty} \frac{|\langle \hat{y}, u_j \rangle|^2}{\sigma_j^2} < \infty.$$

Fourier coefficients

$$\hat{c}_j = \langle \hat{y}, u_j \rangle, \qquad j = 1, 2, 3, \dots \ ,$$

have to converge to zero rapidly.

$y = \hat{y} + e$ available contaminated rhs

Determine approximation of $\hat{x}$ by TSVD:

$$x_k = \sum_{j=1}^{k} \frac{\langle y, u_j \rangle}{\sigma_j} v_j.$$

$\hat{k} \geq 1$ smallest index, such that

$$\|x_{\hat{k}} - \hat{x}\| = \min_{k \geq 1} \|x_k - \hat{x}\|.$$

Assume norm of error in rhs known:

$$\delta = \|y - \hat{y}\|.$$

z is said to satisfy the discrepancy principle if

$$\|Az - y\| \leq \tau\delta,$$

where $\tau > 1$ constant independent of δ.

The discrepancy principle selects the smallest index $k = k_\delta$, such that

$$\|Ax_{k_\delta} - y\| \leq \tau\delta.$$

Then

- $\|Ax_k - y\|$ decreases monotonically as k increases.
- k_δ increases monotonically as $\delta \searrow 0$.
- $x_{k_\delta} \longrightarrow \hat{x}$ as $\delta \searrow 0$.

Tikhonov regularization

Solve the minimization problem

$$\min_x \{\|Ax - b\|^2 + \mu\|x\|^2\}, \qquad (3)$$

where $\mu > 0$ is the (fixed) regularization parameter and $L \in \mathbf{R}^{p\times n}$, $p \leq n$, the regularization operator.

Normal equations associated with (3):

$$(A^T A + \mu I)x = A^T b.$$

Solution of the Tikhonov minimization problem

$$x_\mu := (A^T A + \mu I)^{-1} A^T b, \qquad \mu > 0.$$

Note that:

$$\lim_{\mu \searrow 0} x_\mu = A^\dagger b, \qquad \lim_{\mu \to \infty} x_\mu = 0.$$

A proper choice of the value of the regularization parameter μ is important.

If $\epsilon := \|e\|$ is known, then we can use the discrepancy principle.

For now, assume that ϵ is not known.

To see how the value of μ affects the solution x_μ and the residual error $b - Ax_\mu$ plot the curve

$$(\log_{10} \|x_\mu\|, \log_{10} \|b - Ax_\mu\|), \qquad \mu > 0$$

known as the L-curve.

Choose the value of μ, denoted μ_L , that corresponds to the *vertex* of the L-curve.

Reasons for choosing $\mu = \mu_L$:

- When μ is too small, the solution is contaminated by errors, hence of large norm.
- When μ is too large, x_μ solves a faraway problems, hence the norm of the residual $\|b - Ax_\mu\|$ is large.
- The solution corresponding to μ_L is balances these errors.

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

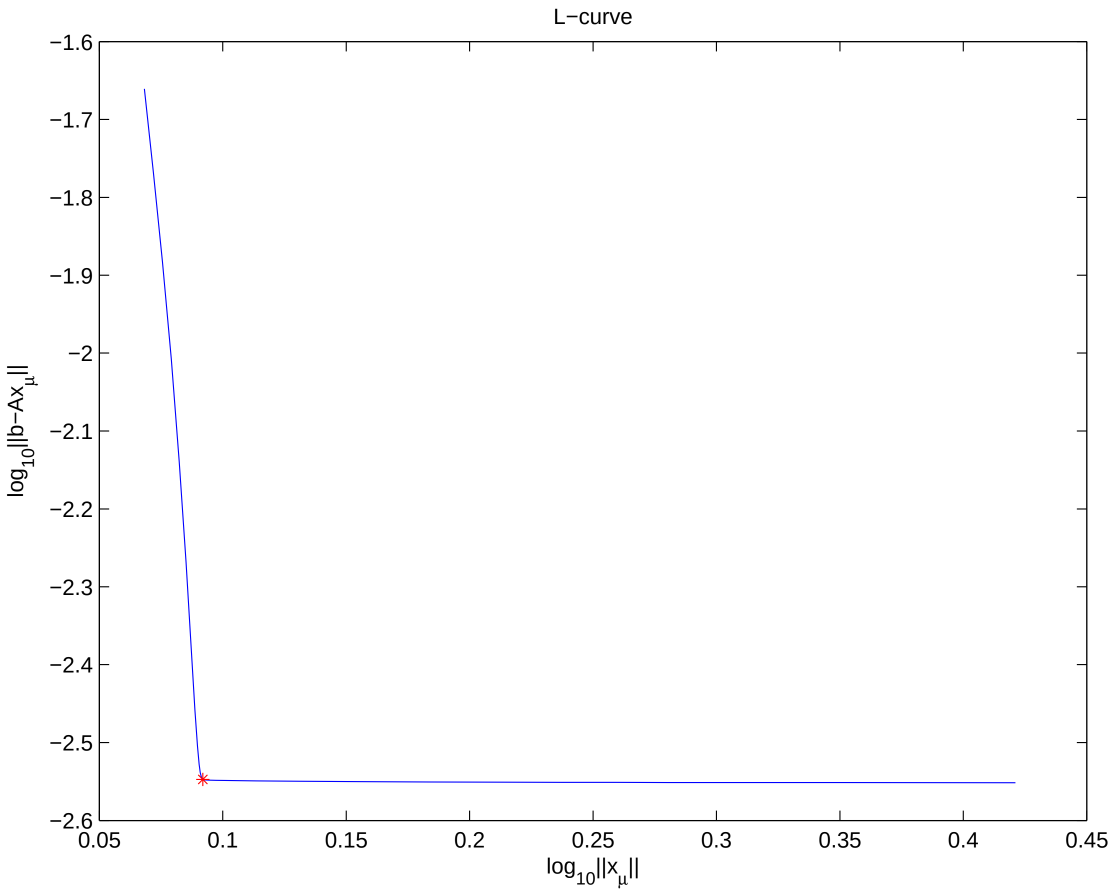

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

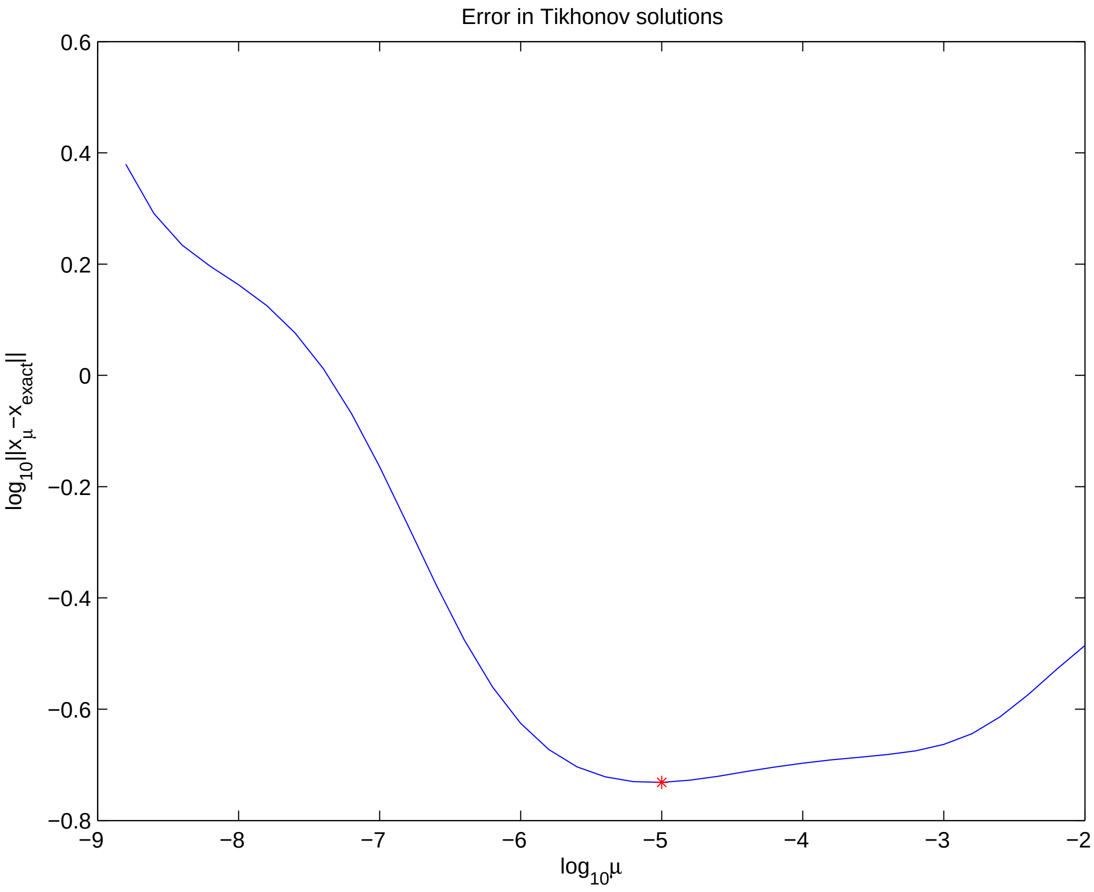

- For small to medium-sized problems, we compute the SVD of A. It is then inexpensive to determine points

$$(\log_{10} \|x_\mu\|, \log_{10} \|b - Ax_\mu\|)$$

on the L-curve for several values of μ.

- For large problems it is expensive to determine points on the L-curve.

Extrapolation enhanced SVD

Let $A = A_1 \otimes A_2$ with $A_1, A_2 \in \mathbf{R}^{1500\times1500}$ defined by the MATLAB functions baart and foxgood. Then A is $2.25 \cdot 10^6 \times 2.25 \cdot 10^6$. The SVD of A can be computed from the SVDs of A_1 and A_2. Extrapolation is equivalent to post-processing of the singular values.

Error in TSVD solutions (dashed graph) and in extrapolated TSVD solutions (solid graph). The errors are measured in the Frobenius norm.

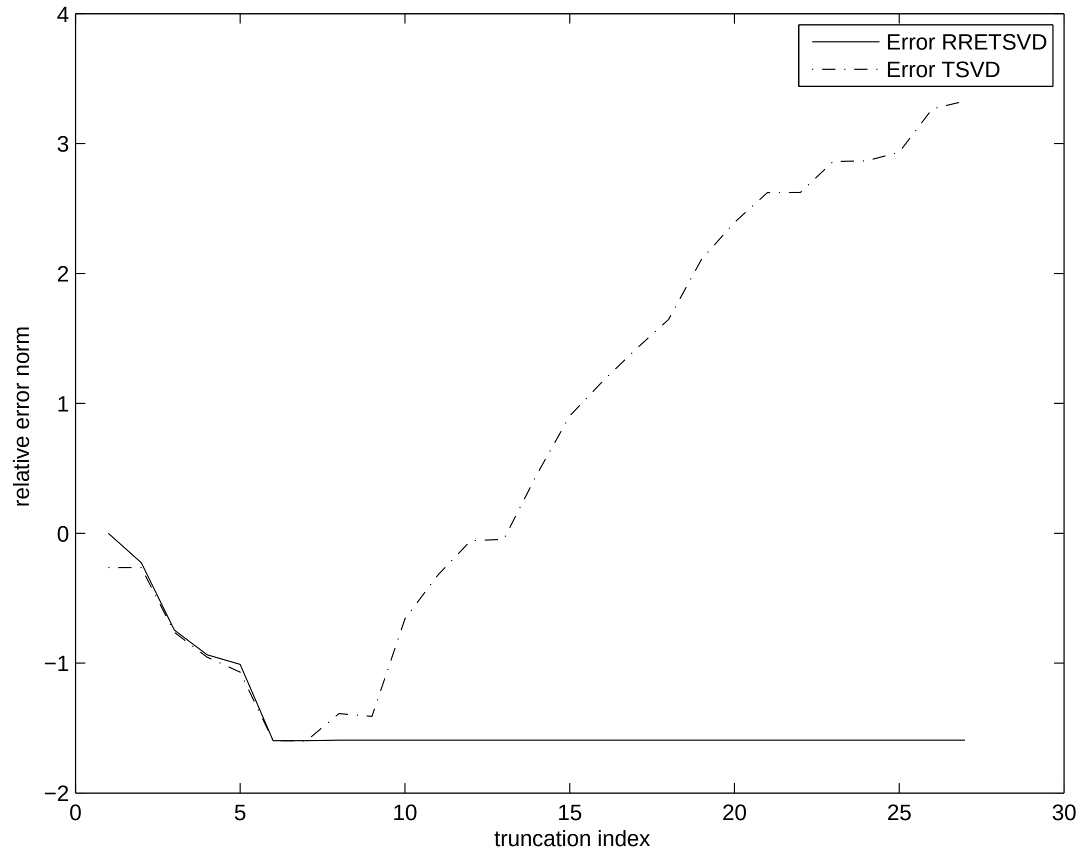

Computed extrapolated TSVD solution x_{23}.

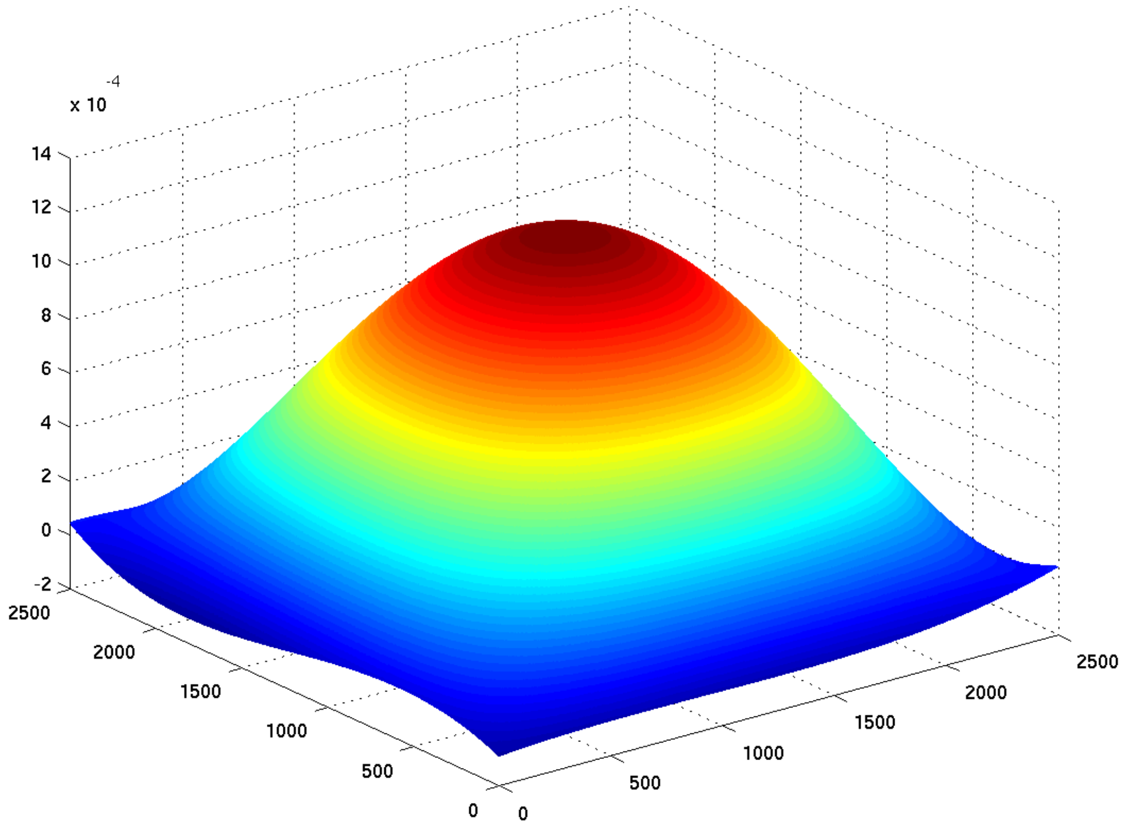

Computed TSVD solution x_{23}.

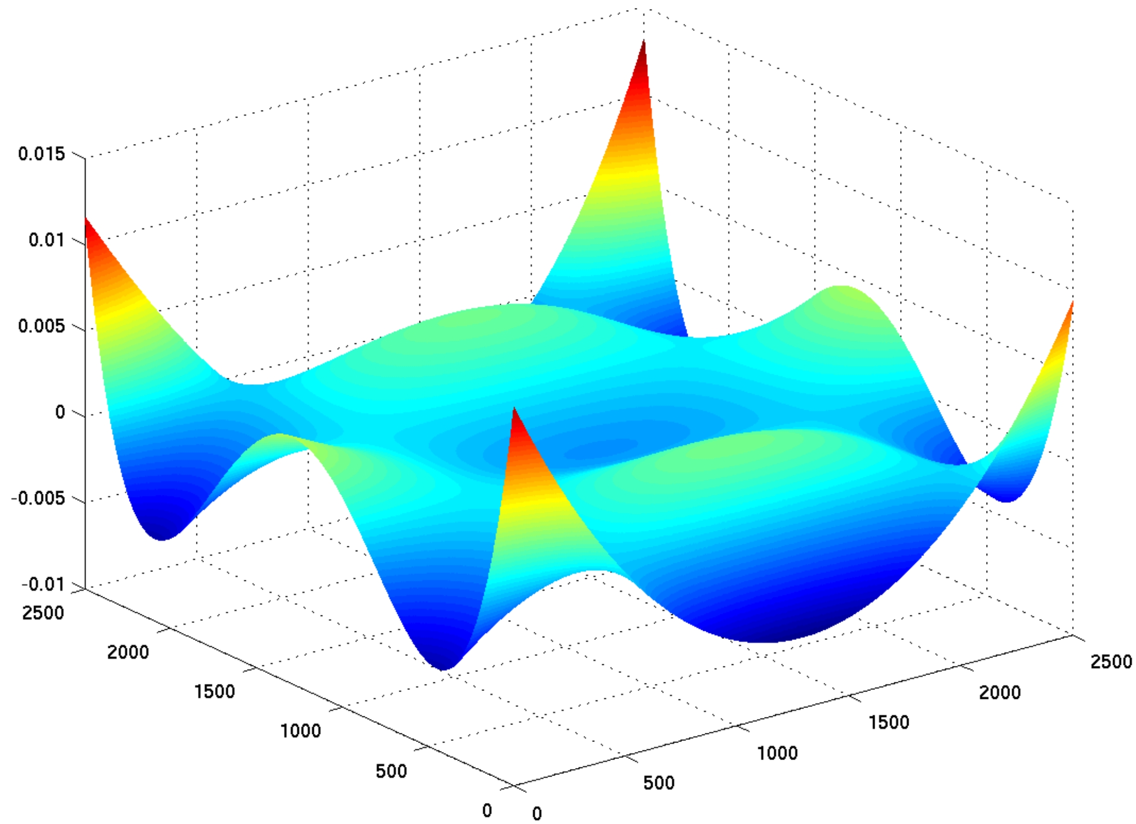

Danke!