Numerical Methods for Linear Discrete Ill-Posed Problems

Lothar Reichel

Como, May 2018

Part 1: The SVD and Krylov subspace methods

Outline of Part 1:

- Inverse and ill-posed problems
- Solution of small to moderately sized problems: Direct methods based on the SVD.
- Solution of large scale problems: Iterative methods

Inverse problems

Inverse problems arise when one seeks to determine the cause of an observed effect.

- Inverse heat conduction problems: What was the temperature of a rod an hour ago?
- 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.

Example of an ill-posed problem:

Fredholm integral equation of the first kind,

$$\int_0^1 k(s,t)x(t)dt = f(t), \qquad 0 \le s \le 1,$$

with a continuous kernel k.

By the Riemann–Lebesgue lemma, small perturbations in f may correspond to large perturbations in x:

$$\max_{0\le s\le 1} \left| \int_0^1 k(s,t)\cos(2\pi\ell t)dt \right|$$

can be made "tiny" by choosing $|\ell|$ large.

Linear discrete ill-posed problems

Let $A \in \mathbf{R}^{m \times n}$ and $b \in \mathbf{R}^m$ with $m \geq n$. When $m > n$, consider the least-squares problems

$$\min_{x \in \mathbf{R}^n} \|Ax - b\|_2$$

or, when $m = n$, consider the linear system of equations

$$Ax = b.$$

Matrices that arise in inverse problems, such as problems of remote sensing or image restoration problems, are of ill-determined rank, possibly rank deficient.

Least-squares problems or linear systems of equations with a matrix of this kind are referred to as linear discrete ill-posed problems.

The vector b contains available data and is not required to be in $\mathcal{R}(A)$.

Linear discrete ill-posed problems arise from the discretization of linear ill-posed problems, such as Fredholm integral equations of the 1st kind or, in discrete form, such as in image restoration.

The vector b in linear discrete ill-posed problems that arise in applications is generally determined by measurement and therefore is contaiminated by error (noise).

In image restoration problems b represents an observed image.

Example 1: Consider the 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 by Hansen.

The code baart gives

- the matrix $A \in \mathbf{R}^{200 \times 200}$, which is numerically singular,
- the desired solution $x_{\text{exact}} \in \mathbf{R}^{200}$, and
- the error-free right-hand side $b_{\text{exact}} \in \mathbf{R}^{200}$.

Then

$$Ax_{\text{exact}} = b_{\text{exact}}.$$

Assume that b_{exact} is not available. Instead a noise-contaminated vector

$$b = b_{\text{exact}} + e$$

is known. Here e represents white Gaussian noise scaled to correspond to 0.1% relative noise, i.e.,

$$\|e\|_2 = 10^{-3}\|b_{\text{exact}}\|_2$$

We would like to determine an approximation of x_{exact} by solving

$$Ax = b.$$

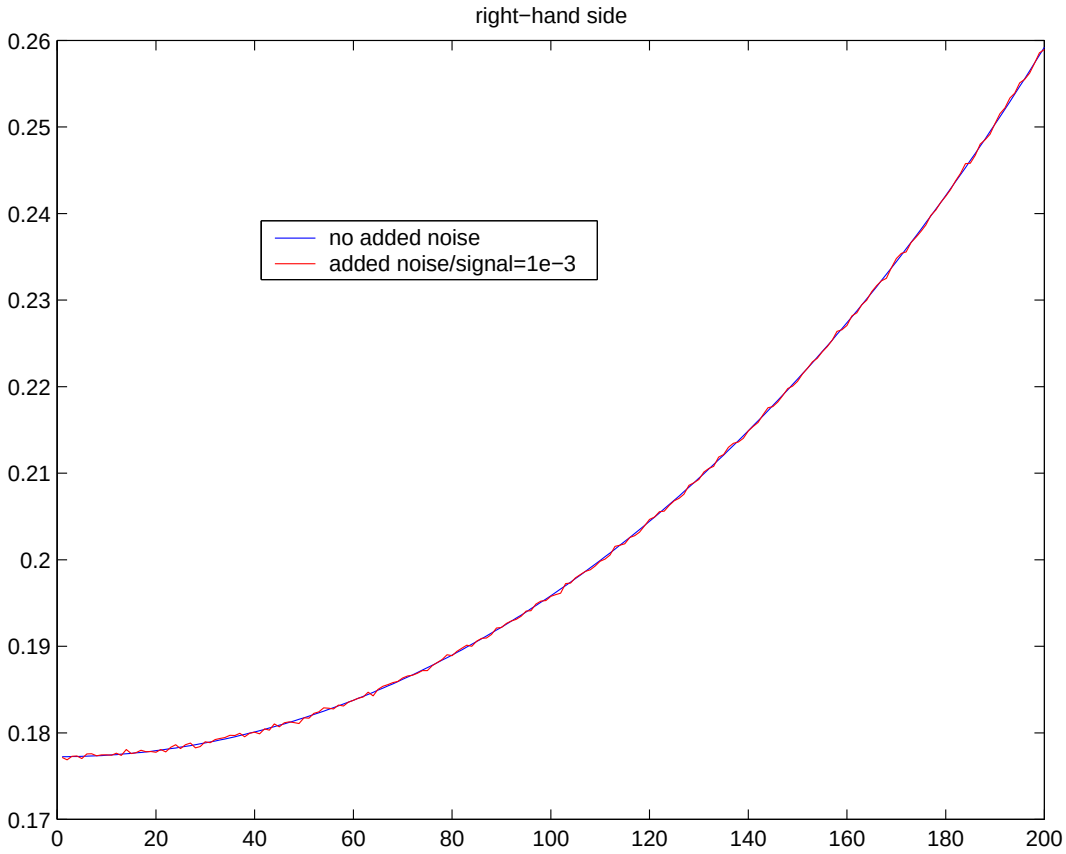
right−hand side
no added noise
added noise/signal=1e−3
0.26
0.25
0.24
0.23
0.22
0.21
0.2
0.19
0.18
0.17
0
20
40
60
80
100
120
140
160
180
200

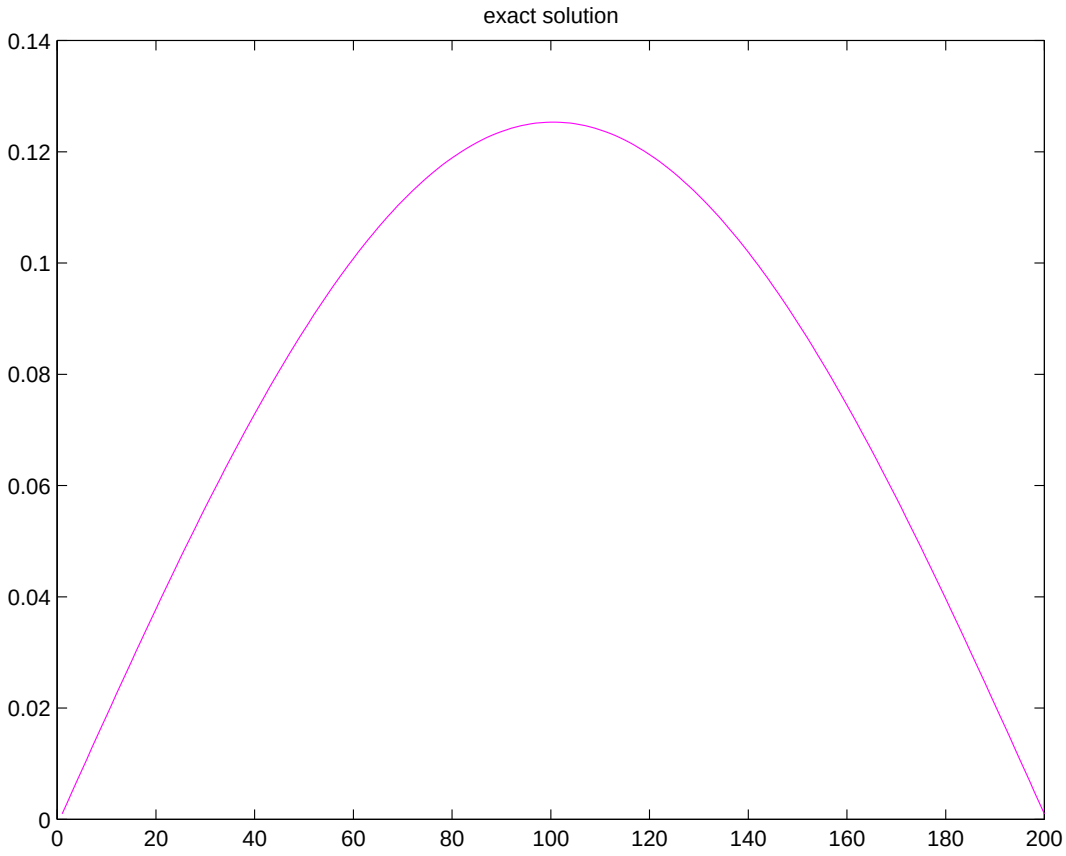
exact solution
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
0
20
40
60
80
100
120
140
160
180
200

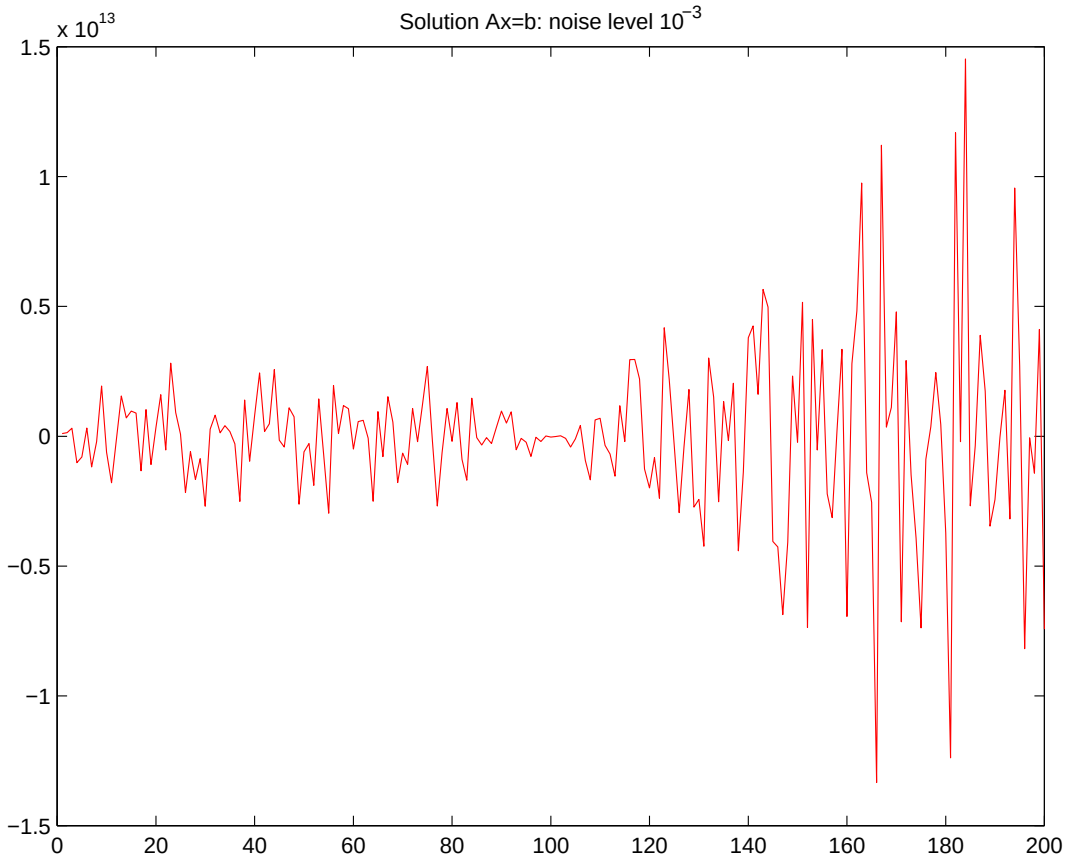
Solution Ax=b: noise level 10^-3
x 10^13
1.5
1
0.5
0
−0.5
−1
−1.5
0
20
40
60
80
100
120
140
160
180
200

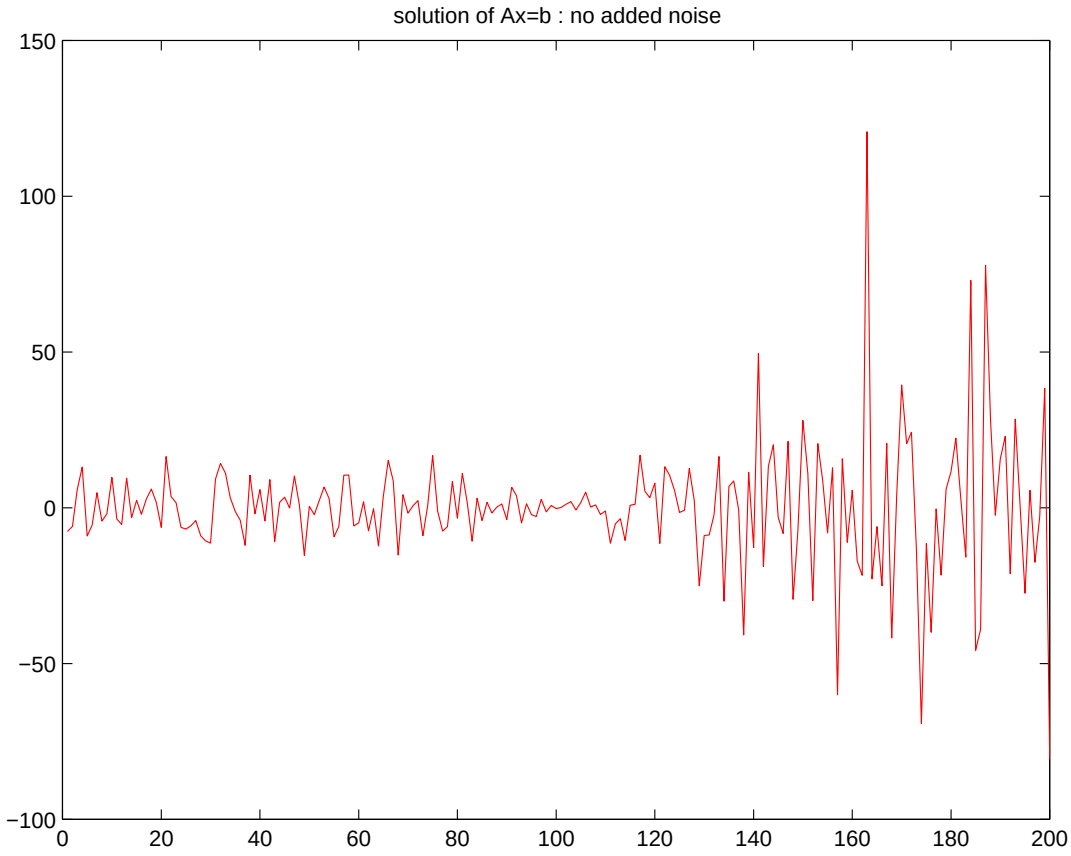
solution of Ax=b : no added noise
150
100
50
0
−50
−100
0
20
40
60
80
100
120
140
160
180
200

The singular value decomposition

The SVD of $A \in \mathbf{R}^{m \times n}$, $m \geq n$:

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

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

with

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

Application to least-squares approximation

$$\min_x \|Ax-b\|_2^2 = \min_x \|U\Sigma V^T x-b\|_2^2 = \min_x \|\Sigma V^T x-U^T b\|_2^2.$$

Let

$$\begin{aligned} y &= [y_1, y_2, \ldots, y_n]^T = V^T x, \\ b' &= [b'_1, b'_2, \ldots, b'_m]^T = U^T b. \end{aligned}$$

Then

$$\min_x \|Ax-b\|_2^2 = \min_y \|\Sigma y-b'\|_2^2 = \sum_{j=1}^{n} (\sigma_j y_j - b'_j)^2 + \sum_{j=n+1}^{m} (b'_j)^2.$$

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 y_j is undetermined and the least-squares solution not unique.

Often one is interested in the least-squares solution of minimal norm. Then undetermined elements 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

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

The matrix

$$A^\dagger = V\Sigma^\dagger U^T \in \mathbf{R}^{n\times m}$$

is known as the Moore–Penrose pseudoinverse of A.

The solution of the least-squares problem

$$\min_x \|Ax - b\|_2$$

of minimal Eucliden norm can be expressed as

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

Moreover,

$$A^\dagger A = I \in \mathbf{R}^{n\times n}, \qquad AA^\dagger = P_{\mathcal{R}(A)} \in \mathbf{R}^{m\times m}.$$

Note that

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

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,

$$\|A - A_k\|_2 = \min_{\text{rank}(B) \le k} \|A - B\|_2 = \sigma_{k+1},$$

i.e., A_k is the closest matrix of rank $\le k$ to A. Here $\|\cdot\|_2$ denotes the spectral norm.

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

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

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

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

might be huge and x a meaningless approximation of x_{exact}.

Recall that

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

is the 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 x_{exact} by

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

for some $k \leq \ell$.

Approximating x_{exact} by x_k is known at the truncated SVD (TSVD) method. How to choose k?

Example 1 cont'd: Singular values of A

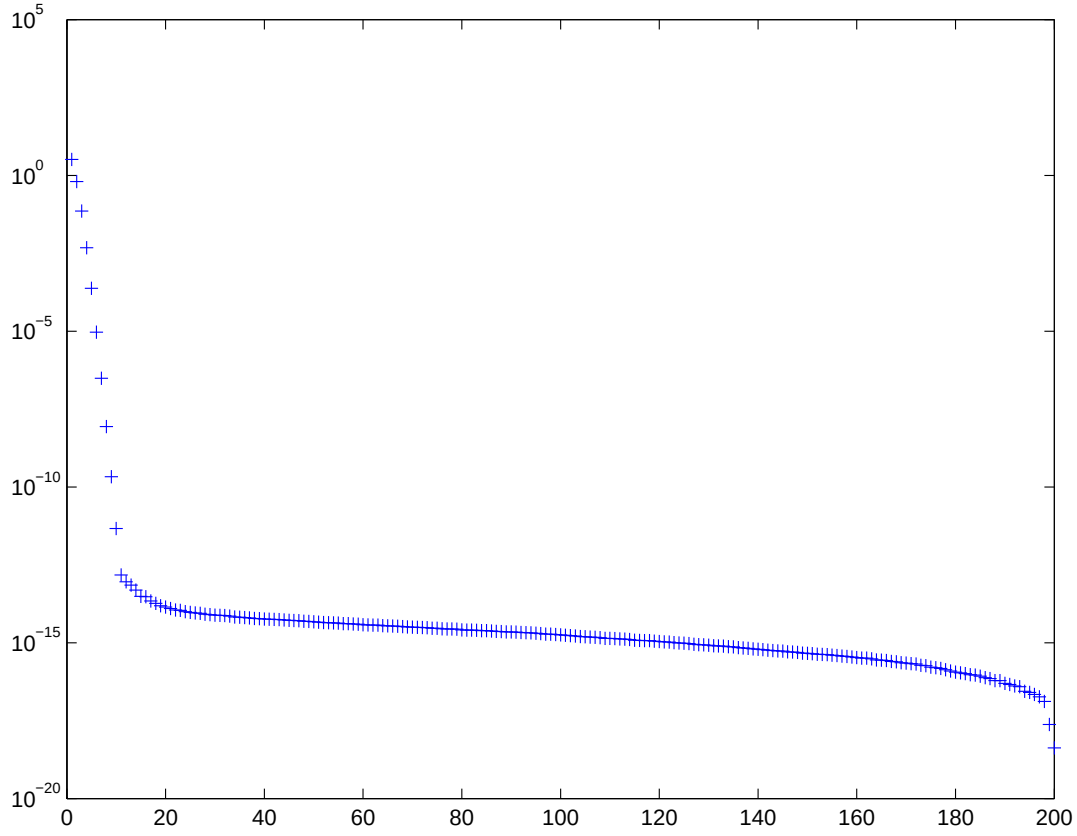

Example 1 cont'd: Right-hand side without noise

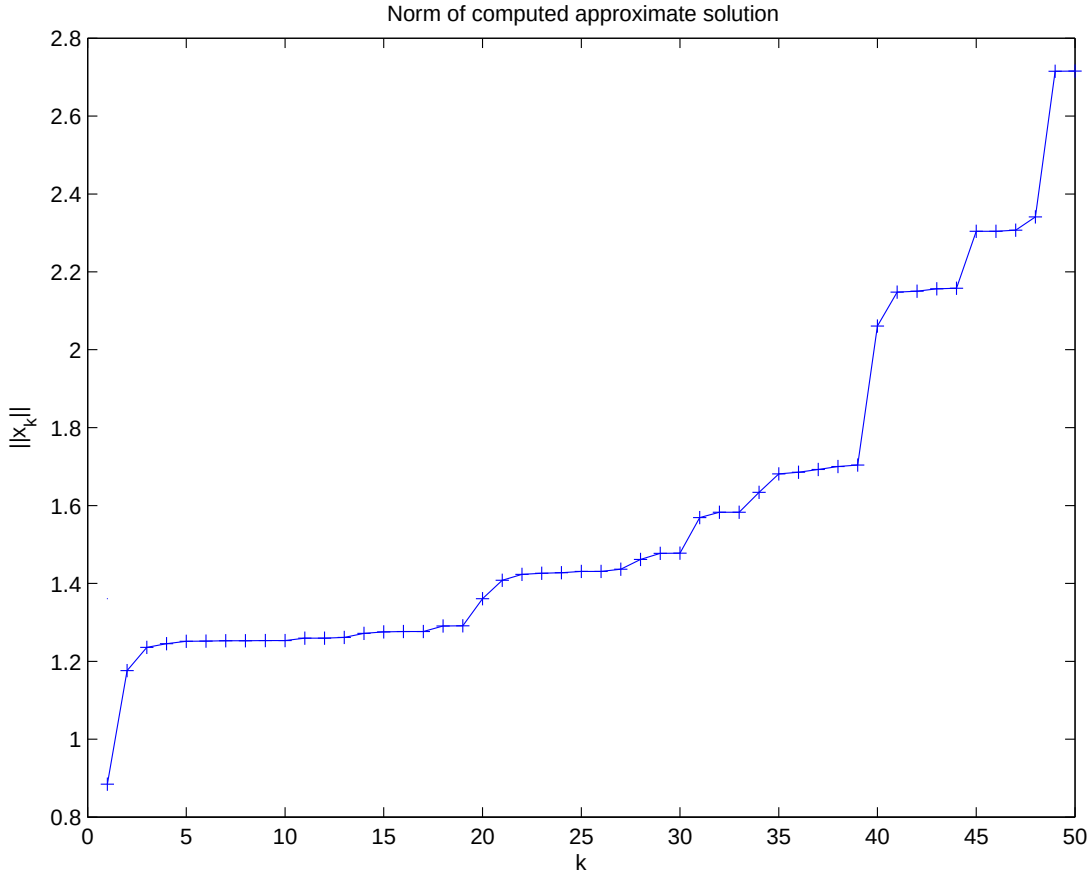

Example 1 cont'd: Right-hand side without noise

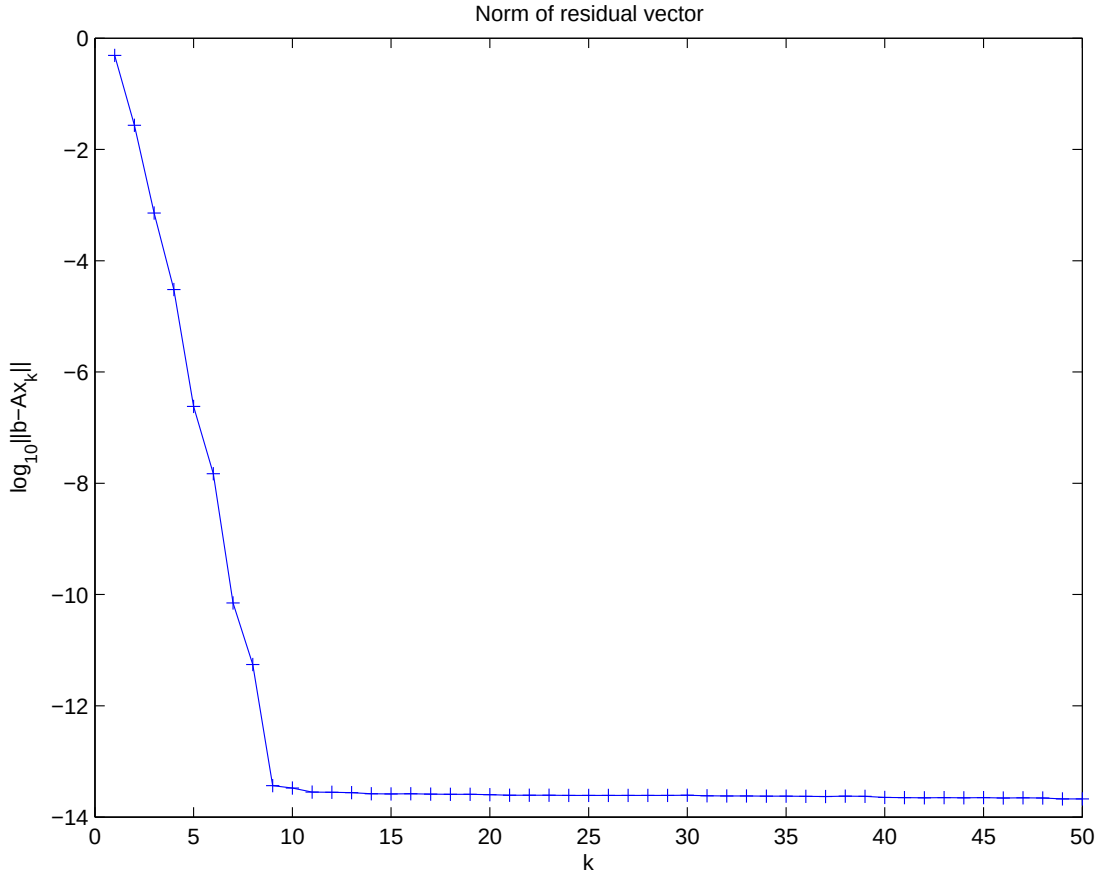

Example 1 cont'd: Right-hand side without noise: Exact and computed solutions

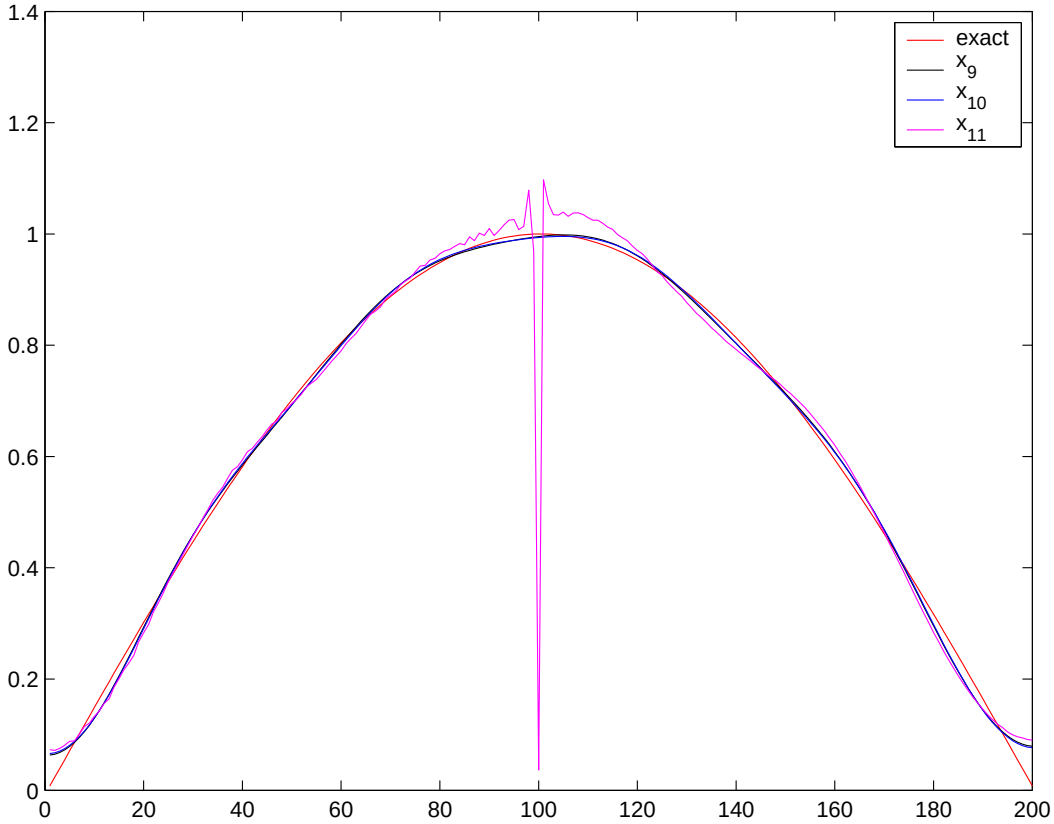

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

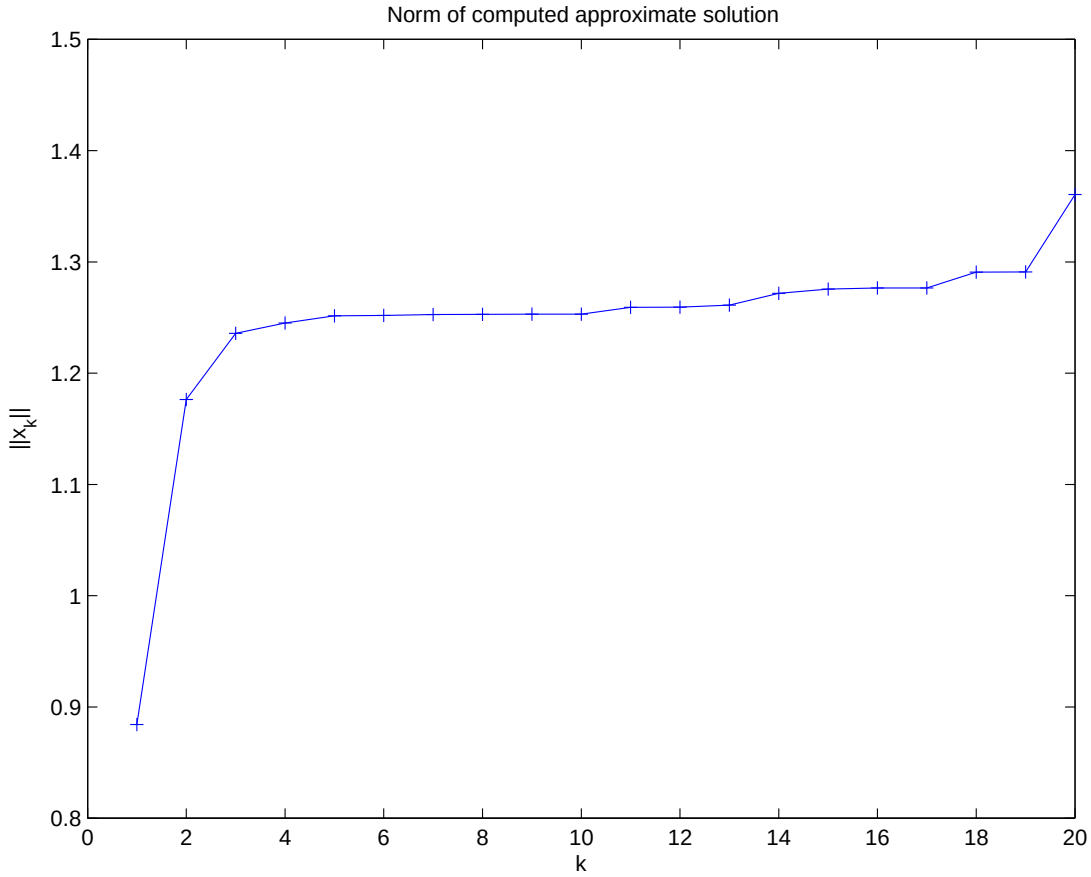

Example 1 cont'd: Right-hand side with relative noise 10^{-3}

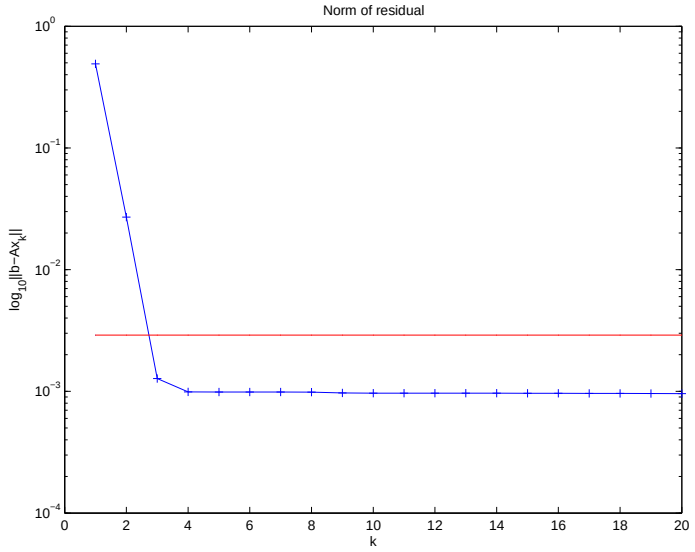

The discrepancy principle prescribes that the truncation index k be as large as possible so that, for some fixed $\eta > 1$,

$$\|Ax_k - b\|_2 \leq \eta \|b - b_{\text{exact}}\|_2.$$

Here $k = 3$.

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

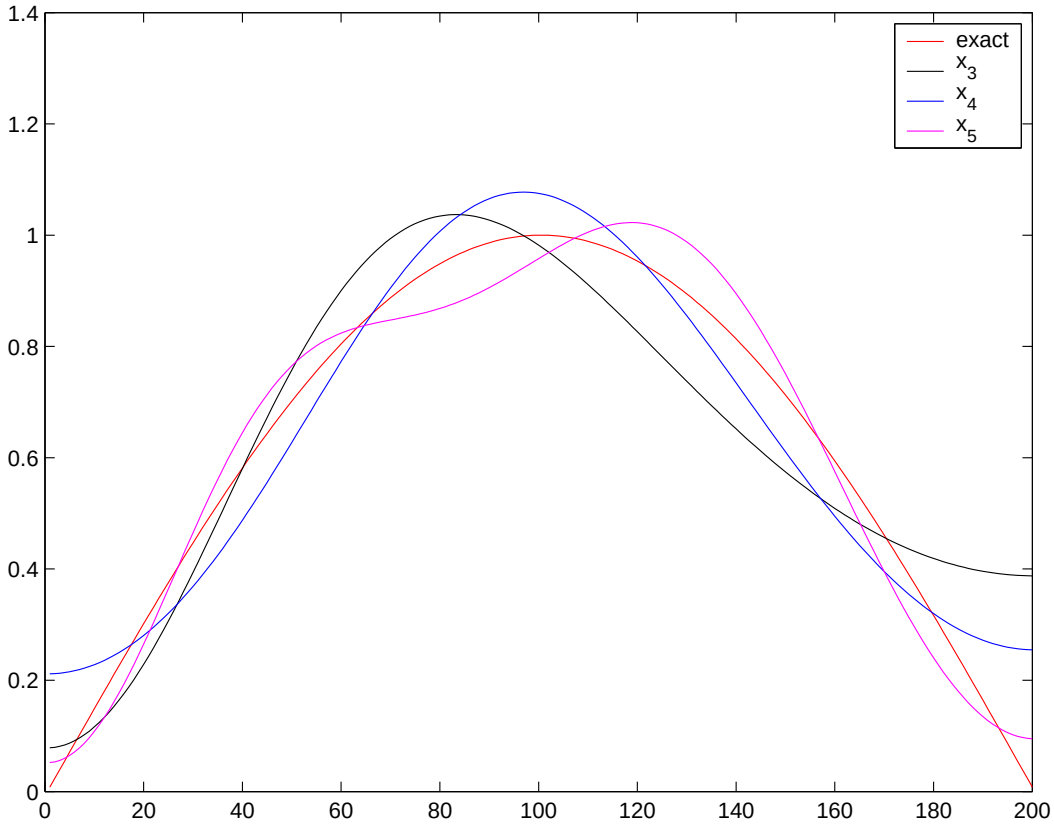

There are many ways to determine the truncation index k including:

- The discrepancy principle: Gives improved approximation of x_{exact} as $\|b - b_{\text{exact}}\|_2 \searrow 0$. Requires a bound for $\|b - b_{\text{exact}}\|_2$ and that $b_{\text{exact}} \in \mathcal{R}(A)$.
- The L-curve criterion: A method that often gives a suitable value of k when $\|b - b_{\text{exact}}\|_2$ is not too small.
- Cross validation and generalized cross validation: Statistically based methods.

- The quasi-optimality criterion: Is based on comparing consecutive norms

$$\|x_{j+1} - x_j\|_2, \qquad j = 1, 2, \ldots \ .$$

All methods, except for the discrepancy principle are referred to as "heuristic methods." They work well for certain problems, but may fail to determine a suitable truncation index for others.

Example 1 cont'd: The L-curve for right-hand side with relative noise 10^{-3}

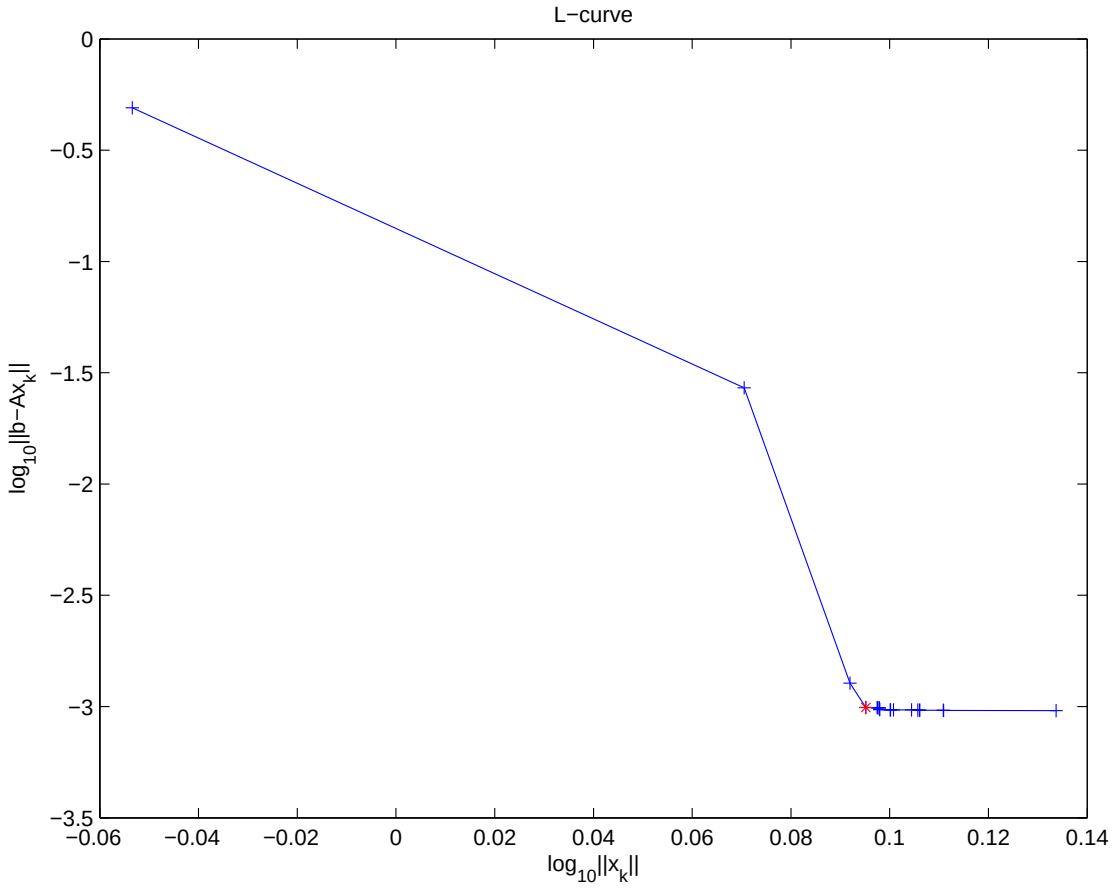

Krylov subspace iterative methods: Regularization by truncated iteration

The most popular iterative method is the conjugate gradient (CG) method applied to the normal equations

$$A^T A x = A^T b.$$

The most stable implementation is based on Golub–Kahan bidiagonalization applied to A. This is the basis for the LSQR algorithm by Paige and Saunders.

Let $A \in \mathbf{R}^{m\times n}$. Application of k bidiagonalization steps to A with initial vector b gives the decompositions

$$AV_k = U_{k+1}B_{k+1,k}, \quad A^TU_k = V_kB_{k,k}^T,$$

where the matrices

$$\begin{aligned} U_{k+1} &= [U_k, u_{k+1}] = [u_1, u_2, \ldots, u_{k+1}] \in \mathbf{R}^{m\times(k+1)}, \\ V_k &= [v_1, v_2, \ldots, v_k] \in \mathbf{R}^{n\times k}, \end{aligned}$$

have orthonormal columns with

$$\begin{aligned} \mathcal{R}(U_{k+1}) &= \mathbf{K}_{k+1}(AA^T, b) = \mathrm{span}\{b, (AA^T)b, \ldots, (AA^T)^k b\}, \\ \mathcal{R}(V_k) &= \mathbf{K}_k(A^TA, A^Tb) = \mathrm{span}\{A^Tb, (A^TA)A^Tb, \ldots, (A^TA)^{k-1}A^Tb\}. \end{aligned}$$

Moreover,

$$B_{k+1,k} = \begin{bmatrix} \beta_{1,1} & & & \\ \beta_{2,1} & \beta_{2,2} & & \\ & \ddots & \ddots & \\ & & \beta_{k,k-1} & \beta_{k,k} \\ & & & \beta_{k+1,k} \end{bmatrix} \in \mathbf{R}^{(k+1)\times k}$$

is lower bidiagonal, and $B_{k,k}$ is the leading $k \times k$ submatrix of $B_{k+1,k}$.

We assume that k is small enough so that the Golub–Kahan decomposition with the stated properties exists. Otherwise, breakdown occurs. This usually is beneficial.

The Golub–Kahan bidiagonalization algorithm.

1: **Input:** Matrix $A \in \mathbf{R}^{m \times n}$, initial vector $b \in \mathbf{R}^m$, number of steps k.

2: $v_0 = 0$, $\beta_{1,0} = \|b\|_2$, $u_1 = b/\beta_{1,0}$

3: **for** $j = 1$ **to** k

4: $\quad \widetilde{v} = A^T u_j - \beta_{j,j-1} v_{j-1}$, $\beta_{j,j} = \|\widetilde{v}\|_2$, $v_j = \widetilde{v}/\beta_{j,j}$

5: $\quad \widetilde{u} = A v_j - \beta_{j,j} u_j$, $\beta_{j+1,j} = \|\widetilde{u}\|_2$, $u_{j+1} = \widetilde{u}/\beta_{j+1,j}$

6: **end for**

7: **Output:** Golub–Kahan decompositions.

Main computational cost: k matrix-vector product evaluations with A and with A^T.

Application of Golub–Kahan bidiagonalization

$$\begin{aligned}
\min_{x \in \mathbf{K}_k(A^TA, A^Tb)} \|Ax - b\|_2 &= \min_{y \in \mathbf{R}^k} \|AV_k y - b\|_2 \\
&= \min_{y \in \mathbf{R}^k} \|U_{k+1}(B_{k+1,k} y - e_1 \|b\|_2)\|_2 \\
&= \min_{y \in \mathbf{R}^k} \|B_{k+1,k} y - e_1 \|b\|_2 \|_2 .
\end{aligned}$$

This shows that the iterative method is a minimal residual method.

The small least-squares problem on the right-hand side is solved for y_k by QR factorization of $B_{k+1,k}$. Then

$$x_k = V_k y_k \in \mathbf{K}_k(A^TA, A^Tb).$$

We would like that $x_k \approx x_{\text{exact}}$. How to choose k?

Implementation issues

Consider the QR factorization:

$$B_{k+1,k} = Q_{k+1} R_{k+1,k}.$$

Here $Q_{k+1} \in \mathbf{R}^{(k+1)\times(k+1)}$ is orthogonal and $R_{k+1,k} \in \mathbf{R}^{(k+1)\times k}$ is upper triangular and banded, with only 2 nonvanishing bands above the diagonal. The last row of $R_{k+1,k}$ vanishes.

The factorization can be computed in only $\mathcal{O}(k)$ flops with the aid of Givens rotations.

LSQR computes the solution $x_k \in \mathbf{K}_k(A^TA, A^Tb)$ of the least-squares problem without storing the the whole matrices V_k and U_{k+1}; only a few of the most recently generated columns are stored simultaneously.

Since $\mathbf{K}_{k-1}(A^TA, A^Tb) \subset \mathbf{K}_k(A^TA, A^Tb)$, the residual errors $r_k = b - Ax_k$ satisfy

$$\|r_k\|_2 \leq \|r_{k-1}\|_2, \qquad k = 1, 2, \ldots \ .$$

Generally, the inequality is strict.

Let

$$\delta = \|b - b_{\text{exact}}\|_2$$

and let $\eta > 1$ be independent of δ.

The discrepancy principle prescribes that the first iterate x_k that satisfies

$$\|Ax_k - b\|_2 \leq \eta\delta$$

be chosen as an approximation of x_{exact}. Denote the smallest largest k such that x_k satisfies the discrepancy principle by $k = k(\delta)$. Then $k(\delta)$ increases as δ decreases.

An iterative method is said to be a regularization method if

$$\lim_{\delta \searrow 0} \sup_{\|b - b_{\text{exact}}\|_2 \leq \delta} \|x_{k(\delta)} - x_{\text{exact}}\|_2 = 0.$$

A proof in Hilbert space that the iterative method described satisfies this condition is provided by Hanke and Nemirovskii.

Note that

- The above property is easy to show in finite dimensions for many iterative methods.
- The rate of convergence as $\delta \searrow 0$ of the computed iterates x_k towards x_{exact} may be slow.

Other iterative methods

Consider linear discrete ill-posed problems

$$Ax = b$$

with a symmetric matrix $A \in \mathbf{R}^{n\times n}$. Let $x_0 = 0$.

- When A is positive semidefinite, the conjugate gradient (CG) method can be used. The kth iterate, x_k, determined by CG satisfies $x_k \in \mathbf{K}_k(A, b)$ and

  $$(x_k - x_{\text{exact}})^T A(x_k - x_{\text{exact}}) = \min_{x\in\mathbf{K}_k(A,b)} (x - x_{\text{exact}})^T A(x - x_{\text{exact}}).$$

  Generally $k \ll n$. This method is not a regularization method. It gives poor approximations of x_{exact}.

- For symmetric indefinite problems SYMMLQ by Paige and Saunders can be applied. But SYMMLQ solves the same minimization problem as CG and also determines poor approximations of x_{exact}.
- The kth iterate, x_k, computed by the minimal residual (MINRES) method by Paige and Saunders satisfies $x_k \in \mathbf{K}_k(A, b)$ and

$$\|Ax_k - b\|_2 = \min_{x \in \mathbf{K}_k(A,b)} \|Ax - b\|_2.$$

  Regularization properties that are weaker than for LSQR can be shown. The error in b propagates somewhat faster into the iterates than for LSQR.

- The kth iterate, x_k, determined by the range restricted minimal residual (RRMINRES) method satisfies $x_k \in \mathbf{K}_k(A, Ab)$ and

$$\|Ax_k - b\|_2 = \min_{x \in \mathbf{K}_k(A,Ab)} \|Ax - b\|_2.$$

  This method gives more accurate approximations of x_{exact} than MINRES. The computed solution is orthogonal to $\mathcal{N}(A)$, because $x_k \in \mathcal{R}(A) = \mathcal{N}(A)^{\perp}$. Hanke showed it is a regularization method.

The storage requirement of CG, SYMMLQ, MINRES, and RRMINRES can be bounded independently of the number of iterations k.

The coding of the progressive form of RRMINRES is somewhat tricky. It is based on decompositions of the form

$$A\widehat{V}_k = V_{k+2}\widehat{H}_{k+2,k},$$

where

- $\widehat{V}_k \in \mathbf{R}^{n\times k}$ has orthonormal columns that span $\mathbf{K}_k(A, Ab) = \text{span}\{Ab, A^2b, \ldots, A^kb\}$.
- $V_{k+2} \in \mathbf{R}^{n\times(k+2)}$ has orthonormal columns that span $\mathbf{K}_{k+2}(A, b) = \text{span}\{b, Ab, \ldots, A^{k+1}b\}$ with first column $b/\|b\|_2$.
- $\widehat{H}_{k+2,k} \in \mathbf{R}^{(k+2)\times k}$ is lower Hessenberg with two nonvanishing subdiagonals.

The approximate solution determined at step k of RRMINRES satisfies $x_k \in \mathbf{K}_k(A, Ab)$ and

$$\begin{aligned} \|Ax_k - b\|_2 &= \min_{x \in \mathbf{K}_k(A,Ab)} \|Ax - b\|_2 \\ &= \min_{y \in \mathbf{R}^k} \|V_{k+2}(\widehat{H}_{k+2,k} y - e_1 \|b\|_2)\|_2 \\ &= \min_{y \in \mathbf{R}^k} \|\widehat{H}_{k+2,k} y - e_1 \|b\|_2 \|_2. \end{aligned}$$

The least-squares problem in the right-hand side easily can be solved by QR factorization of $\widehat{H}_{k+2,k}$. Only a few of the most recently generated columns of V_{k+2} and $\widehat{V}_k$ have to be stored simultaneously.

Consider linear discrete ill-posed problems

$$Ax = b$$

with a nonsymmetric square matrix $A \in \mathbf{R}^{n\times n}$. Let $x_0 = 0$.

- The kth iterate, x_k, determined by the generalized minimal residual (GMRES) method satisfies $x_k \in \mathbf{K}_k(A, b)$ and

$$\|Ax_k - b\|_2 = \min_{x\in\mathbf{K}_k(A,b)} \|Ax - b\|_2.$$

  This is a regularization method in the same sense as MINRES.

- The kth iterate, x_k, computed by the range restricted GMRES (RRGMRES) method satisfies $x_k \in \mathbf{K}_k(A, A^j b)$ and

$$\|Ax_k - b\|_2 = \min_{x \in \mathbf{K}_k(A, A^j b)} \|Ax - b\|_2.$$

  This is a regularization method under the same conditions as GMRES.

  When A is a low-pass filter, Ab contains less high-frequency error than b. This results in improved approximations of x_{exact}. Best results are achieved for $j = 1$ or $j = 2$.

The most common implementation of GMRES is based on partial Arnoldi decomposition of A. Application of k steps of the Arnoldi process to $A \in \mathbf{R}^{n\times n}$ with initial vector b gives the decomposition

$$AV_k = V_{k+1}H_{k+1,k},$$

where

$$V_{k+1} = [V_k, v_{k+1}] = [v_1, v_2, \ldots, v_{k+1}] \in \mathbf{R}^{n\times(k+1)}$$

has orthonormal columns with $V_{k+1}e_1 = b/\|b\|_2$, and $H_{k+1,k} \in \mathbf{R}^{(k+1)\times k}$ is upper Hessenberg.

The kth iterate determined by GMRES satisfies

$$\begin{aligned}
\|Ax_k - b\|_2 &= \min_{x \in \mathbf{K}_k(A,b)} \|Ax - b\|_2 \\
&= \min_{y \in \mathbf{R}^k} \|AV_k y - b\|_2 \\
&= \min_{y \in \mathbf{R}^k} \|V_{k+1}(H_{k+1,k} y - e_1 \|b\|_2)\|_2 \\
&= \min_{y \in \mathbf{R}^k} \|H_{k+1,k} y - e_1 \|b\|_2 \|_2 .
\end{aligned}$$

Solve the small least-squares problem on the right-hand side by using the QR factorization

$$H_{k+1,k} = Q_{k+1} R_{k+1,k},$$

where $Q_{k+1} \in \mathbf{R}^{(k+1)\times(k+1)}$ is orthogonal and $R_{k+1,k} \in \mathbf{R}^{(k+1)\times k}$ is upper triangular with vanishing last row.

Denote the solution by y_k. Then the kth iterate produced by GMRES is given by

$$x_k = V_k y_k.$$

The matrix V_k has to be available when computing x_k. Therefore, the storage requirment for GMRES grows linearly with k.

For many discrete ill-posed problems k can be chosen quite small.

The Arnoldi process

0. Input: Matrix $A \in \mathbf{R}^{n \times n}$, vector $b \in \mathbf{R}^n$, number of steps k;
1. Let $v_1 = b/\|b\|_2$;
2. `for` $j = 1, \ldots, k$ `do`
   - 2.1. $v = Av_j$;
   - 2.3. `for` $i = 1, \ldots, j$ `do`

     $$h_{i,j} = v^T v_i;\ v = v - h_{i,j} v_i;$$

   - 2.4. `end for`
   - 2.5. $h_{j+1,j} = \|v\|_2$;
   - 2.6. $v_{j+1} = v/h_{j+1,j}$;
3. `end for`

The RRGMRES method can be implemented by using a range restricted Arnoldi decomposition of A. Application of k steps of the range restricted Arnoldi process to $A \in \mathbf{R}^{n\times n}$ with initial vector b gives the decomposition

$$A\widehat{V}_k = V_{k+j}H_{k+j,k},$$

where

- $\widehat{V}_k \in \mathbf{R}^{n\times k}$ has orthormnal columns that span $\mathbf{K}_k(A, A^j b)$,
- $V_{k+j} \in \mathbf{R}^{n\times(k+j)}$ has orthonormal columns that span $\mathbf{K}_{k+j}(A, b)$ with the first column $b/\|b\|_2$.
- $H_{k+j,k} \in \mathbf{R}^{(k+j)\times k}$ is a generalized upper Hessenberg matrix with j nontrivial diagonal bands below the diagonal.

- The iterates x_k can be computed by using the range restricted Arnoldi decomposition:

$$\begin{aligned} \min_{x \in \mathbf{K}_k(A, A^j b)} \|Ax - b\|_2 &= \min_{y \in \mathbf{R}^k} \|A\widehat{V}_k y - b\|_2 \\ &= \min_{y \in \mathbf{R}^k} \|H_{k+j,k} y - e_1 \|b\|_2 \|_2. \end{aligned}$$

The solution y_k of the small least-squares problem on the right-hand side determines $x_k = \widehat{V}_k y_k$.

- When $j = 1$, the (standard) Arnoldi decomposition and (standard) GMRES are recovered.

- The storage equirement for RRGMRES grows linearly with the number of steps, k. It is about twice as large as for GMRES.

Example. Consider the integral equation

$$\int_{-\pi/2}^{\pi/2} \kappa(\tau,\sigma)x(\sigma)d\sigma = g(\tau), \qquad -\frac{\pi}{2} \le \tau \le \frac{\pi}{2},$$

where

$$\kappa(\sigma,\tau) = (\cos(\sigma)+\cos(\tau))\left(\frac{\sin(\xi)}{\xi}\right)^2, \quad \xi = \pi(\sin(\sigma)+\sin(\tau)).$$

Let $g(\tau)$ be a smooth function. Discretize by a Nyström method based on the trapezoidal rule with equidistant nodes. This gives a nonsymmetric matrix $A \in \mathbf{R}^{1000\times 1000}$ and $b_{\text{exact}} \in \mathbf{R}^{1000}$. Adding 1% Gaussian noise gives $b \in \mathbf{R}^{1000}$.

Best iterates generated by RRGMRES(j) for $j = 0, 1, \ldots, 4$, and smallest errors

iterative method	error for best iterate
RRGMRES(0)	$\|x_6^{(0)} - \hat{x}\| = 4.30$
RRGMRES(1)	$\|x_5^{(1)} - \hat{x}\| = 4.01$
RRGMRES(2)	$\|x_7^{(2)} - \hat{x}\| = 2.06$
RRGMRES(3)	$\|x_7^{(3)} - \hat{x}\| = 2.06$
RRGMRES(4)	$\|x_7^{(4)} - \hat{x}\| = 2.05$

Hybrid regularization methods

In the iterative methods described above, regularization is achieved by truncated iteration. This regularization technique is analogous to the TSVD method and performs well for many linear discrete ill-posed problems.

However, there are problems for which GMRES-type methods determine poor approximations of x_{exact}.

Reasons for this include:

- The solution subspace determined by a few steps of Arnoldi-type methods are poorly suited to approximate x_{exact}. This includes images that have been contaminated by significant motion blur.
- The solution subspace determined by a few steps of Arnoldi-type methods contains an accurate approximation of x_{exact}, but GMRES-type methods are unable to determine it.

Both types of difficulties will be illustrated in these lectures. They can be remedied as follows:

- When the solution subspace is poorly suited to represent x_{exact}, one often can use a preconditioner M to remedy this situation, i.e., one solves

$$AMy = b, \qquad x = My,$$

and determines an approximetion x_k in the solution subspace $M\mathbf{K}_k(AM, b)$.

- When solution subspace determined by the Arnoldi process contains an accurate approximation of x_{exact}, but GMRES fails to compute it, it may be beneficial regularize by the TSVD method applied to the Hessenberg matrix generated. This allows that more Arnoldi steps are carried out than when regularizing by truncated iteration.