Confronto di vecchie e nuove tecniche per la determinazione di un parametro di regolarizzazione discreto

Lothar Reichel Giuseppe Rodriguez

Department of Mathematical Sciences, Kent State University,
Kent, OH 44242, USA.

Dipartimento di Matematica e Informatica, Università di Cagliari
viale Merello 92, 09123 Cagliari, Italy

Due Giorni di Algebra Lineare Numerica
Genova, 16–17 febbraio, 2012

We consider the least-squares problem

$$\min_{\mathbf{x}\in\mathbb{R}^n} \|A\mathbf{x} - \mathbf{b}\|, \qquad A \in \mathbb{R}^{m\times n}, \qquad m \geq n,$$

under the assumptions that

- the matrix A is severely ill-conditioned;
- the data vector $\mathbf{b}$ is error-contaminated;
- there is not an obvious way to define a numerical rank for A;
- the solution is *regular* $\rightarrow$ it *lives* in the span of the first singular vectors (discrete Picard condition).

A regularization method to approximate the noise-free solution of

$$\min_{\mathbf{x}\in\mathbb{R}^n} \|A\mathbf{x} - \mathbf{b}^\delta\|, \qquad \text{where } \|\mathbf{b} - \mathbf{b}^\delta\| \leq \delta\|\mathbf{b}\|,$$

is a pair

$$\{R_\lambda, \lambda(\delta, \mathbf{b}^\delta)\}$$

such that

$$\mathbf{x}_\lambda = R_\lambda(\mathbf{b}^\delta) \longrightarrow A^\dagger\mathbf{b}, \qquad \text{when } \delta \longrightarrow 0.$$

The parameter choice rule in a regularization method must depend on the noise level δ [Bakushinski 1984].

Often in real applications an accurate bound δ is unknown, so *error-free*, or heuristic methods, are used.

Small to medium-sized problems: TSVD

Let the singular value decomposition (SVD) of A be

$$A = U \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} V^T,$$

where U and V are are orthogonal, $\Sigma = \text{diag}(\sigma_i) \in \mathbb{R}^{n \times n}$, and

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

The best rank-k approximation of A w.r.to the spectral norm is

$$A_k = U \begin{bmatrix} \Sigma_k \\ O \end{bmatrix} V^T,$$

with

$$\Sigma_k = \text{diag}[\sigma_1, \sigma_2, \ldots, \sigma_k, \underbrace{0, \ldots, 0}_{n-k}].$$

Small to medium-sized problems: TSVD

Introduce the Moore-Penrose pseudoinverse of A_k

$$A_k^\dagger = V \begin{bmatrix} \Sigma_k^\dagger & O^T \end{bmatrix} U^T,$$

where

$$\Sigma_k^\dagger = \text{diag}[\sigma_1^{-1}, \sigma_2^{-1}, \ldots, \sigma_k^{-1}, \underbrace{0, \ldots, 0}_{n-k}] \in \mathbb{R}^{n \times n}.$$

The minimal-norm solution of the least-squares problem $\min_{\mathbf{x} \in \mathbb{R}^n} \|A_k \mathbf{x} - \mathbf{b}\|$ can be expressed as

$$\mathbf{x}_k = A_k^\dagger \mathbf{b} = \sum_{j=1}^{k} \frac{\mathbf{u}_j^T \mathbf{b}}{\sigma_j} \mathbf{v}_j,$$

This is the truncated SVD (TSVD) solution and $k = 1, \ldots, \ell$ is the regularization parameter.

Using a regularization matrix: TGSVD

Let us introduce a regularization matrix $L \in \mathbb{R}^{p \times n}$. We call *minimal L-norm solution* the vector $\mathbf{x}_L^\dagger$ which solves the problem

$$\min_{\mathbf{x} \in \mathcal{S}} \|L\mathbf{x}\|, \qquad \mathcal{S} = \{\mathbf{x} \in \mathbb{R}^n \mid A^T A\mathbf{x} = A^T \mathbf{b}\},$$

under the assumption

$$\mathcal{N}(A) \cap \mathcal{N}(L) = \{0\}.$$

By means of the *generalized singular value decomposition* (GSVD) of the matrix pair (A, L)

$$A = U\Sigma_A Z^{-1}, \qquad L = V\Sigma_L Z^{-1},$$

it is possible to define the truncated GSVD (TGSVD) solution $\mathbf{x}_k$, where $k = 0, 1, \ldots, p$ is the regularization parameter.

Discrepancy principle vs. heuristic rules

The discrepancy principle is an *a posteriori* rule, which selects

$$k(\delta, \mathbf{b}^\delta) = \min\{k \ : \ \|A\mathbf{x}_k - \mathbf{b}^\delta\| \leq \tau\delta\|\mathbf{b}\|\}.$$

The constant τ is usually larger than 1.

We used $\tau = 1.3$, knowledge of the noise level δ is needed.

Heuristic rules are based on what can be observed/measured

- residuals: $\|\mathbf{r}_k\| = \|A\mathbf{x}_k - \mathbf{b}\|$
- norms: $\|\mathbf{x}_k\|$, $\|L\mathbf{x}_k\|$
- ...

Residuals are not enough

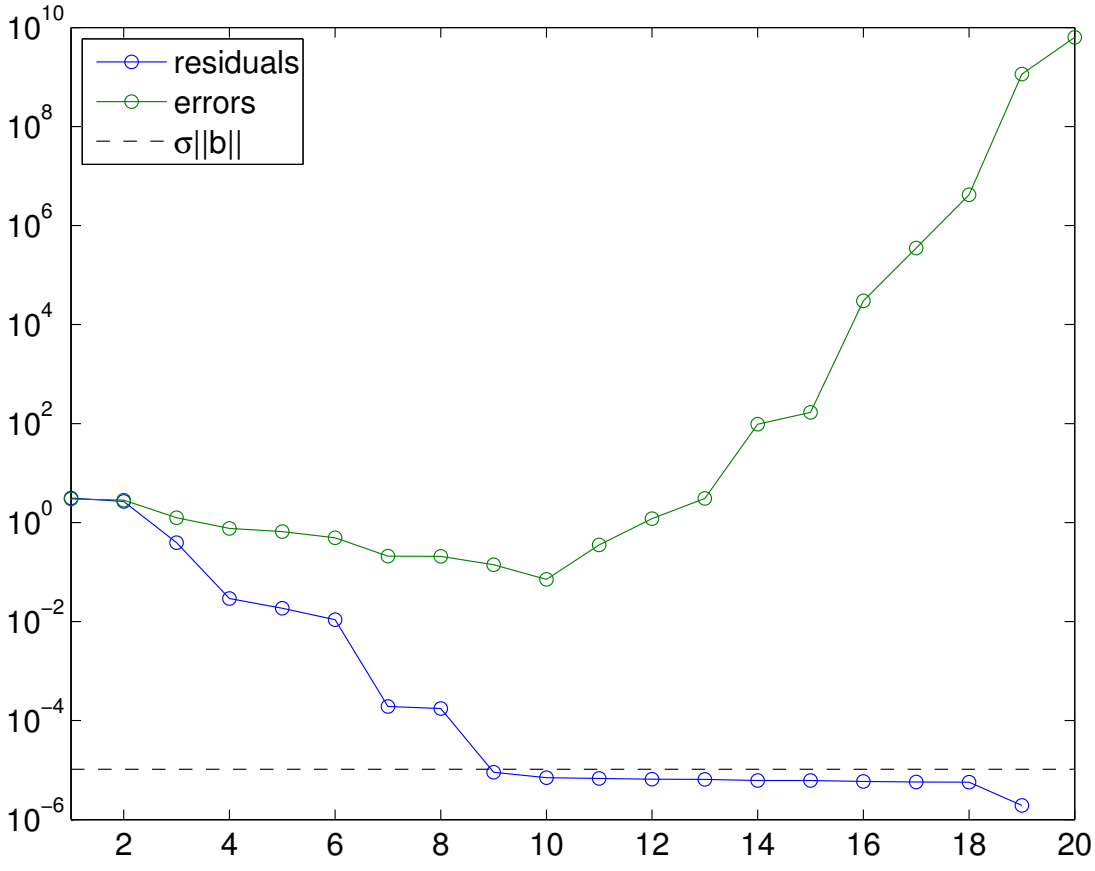

TSVD/*Shaw* example: 20×20, $\delta = 10^{-6}$

Some heuristic parameter choice rules

- L-curve [Hansen 1992], [Hansen, O'Leary 1993]
  - Corner algorithm [Hansen, Jensen, R 2006]
  - L-triangle [Castellanos, Gómez, Guerra 2002]
  - Residual L-curve [Reichel, Sadok 2008]
  - Condition L-curve [Calvetti, Lewis, Reichel 2002]
  - Regińska method [Regińska 1996]
  - restricted Regińska [Reichel, R 2012]
- Generalized Cross Validation (GCV) [Craven, Wahba 1979]
- Error Estimates [Brezinski, R, Seatzu 2008/9],
- Quasi-optimality criterion [Morozov 1984]
- Specific for LSQR
  - Quadrature [Hnětynková, Plešinger, Strakoš 2009]
  - Ratio [Reichel, R 2012]

The L-curve is defined as the graph that connects adjacent points in the sequence

$$\{\log \|A\mathbf{x}_j - \mathbf{b}\|, \log \|\mathbf{x}_j\|\}, \qquad j = 1, 2, \ldots, \ell.$$

The graph generally is L-shaped.

The L-curve criterion consist of selecting the value of the regularization parameter corresponding to the *corner* of the curve.

L-curve

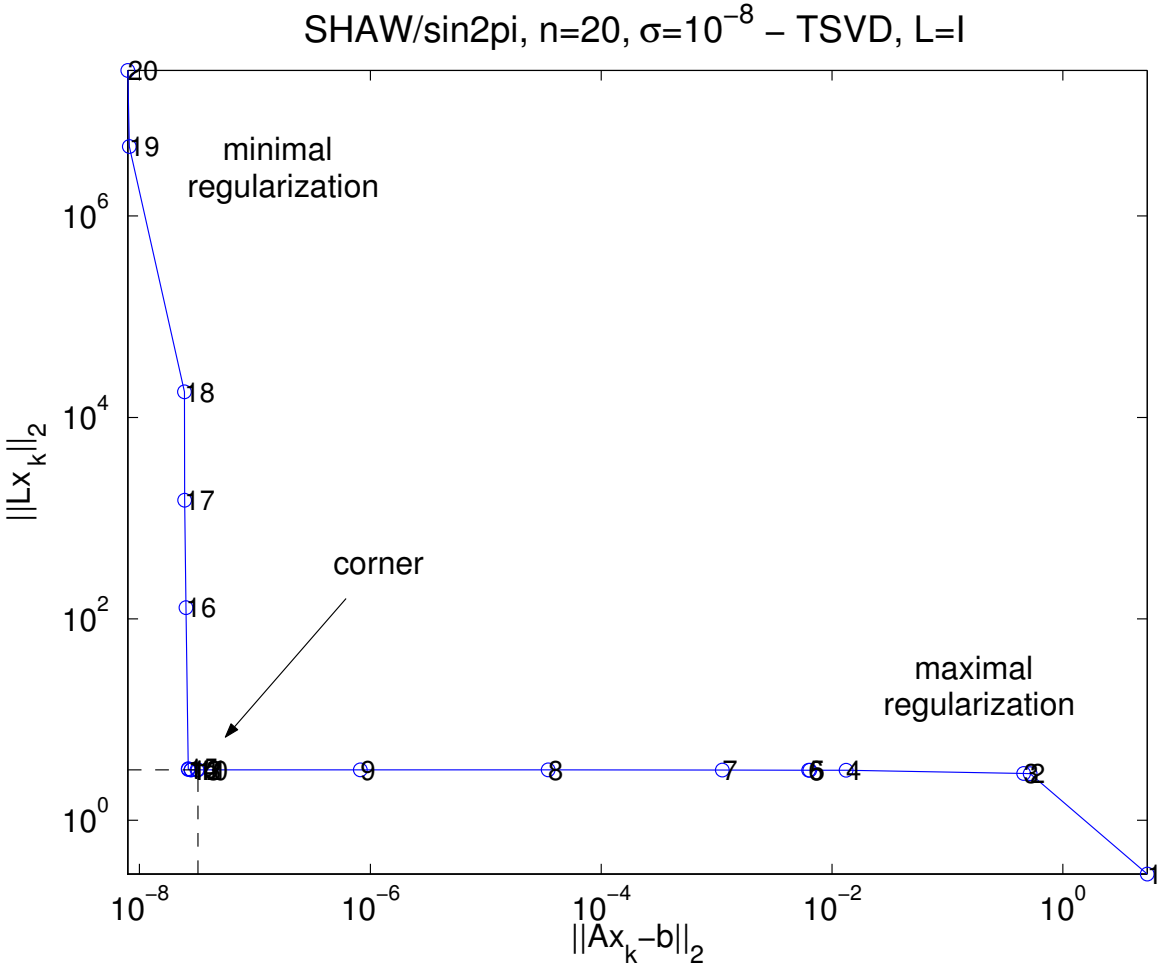

Shaw example: 40×20, $\delta = 10^{-8}$

Difficult situations: clusters of points

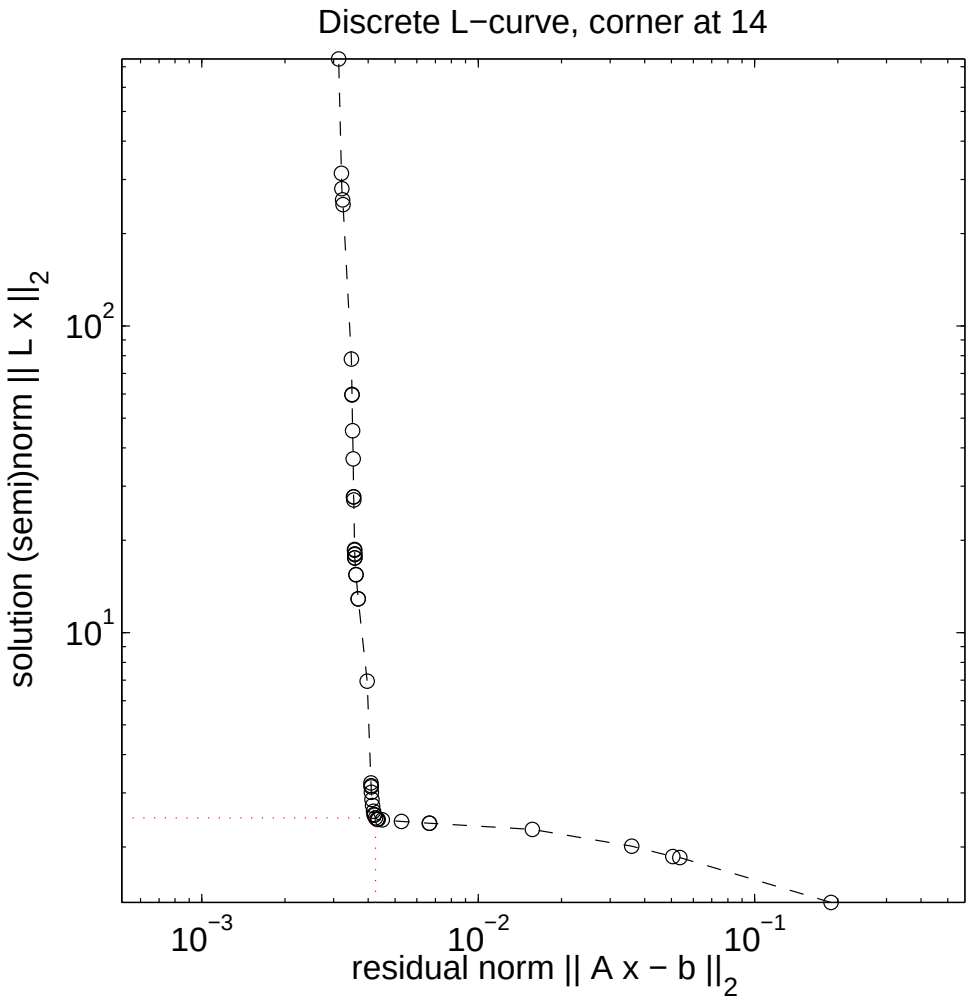

Heat(1) example: 100×50, $\delta = 10^{-2}$

Difficult situations: clusters of points

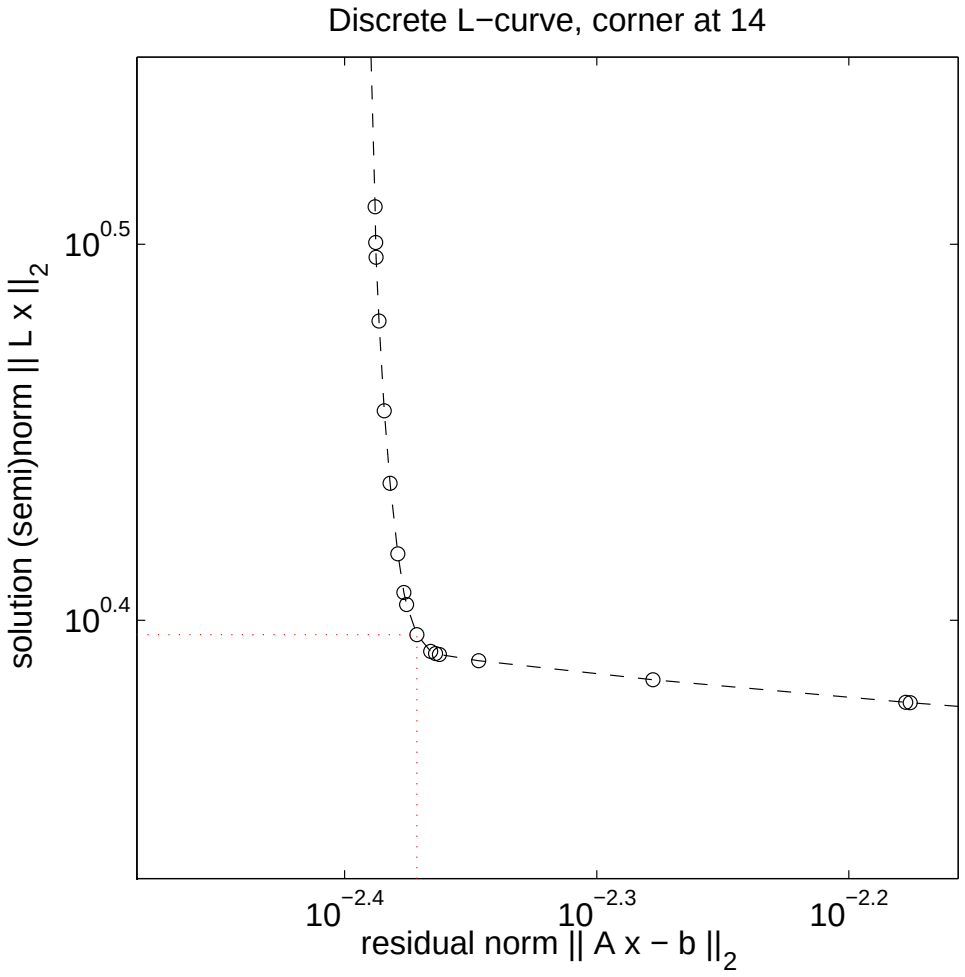

Heat(1) example: 100×50, $\delta = 10^{-2}$

Difficult situations: small var. in norms/residuals

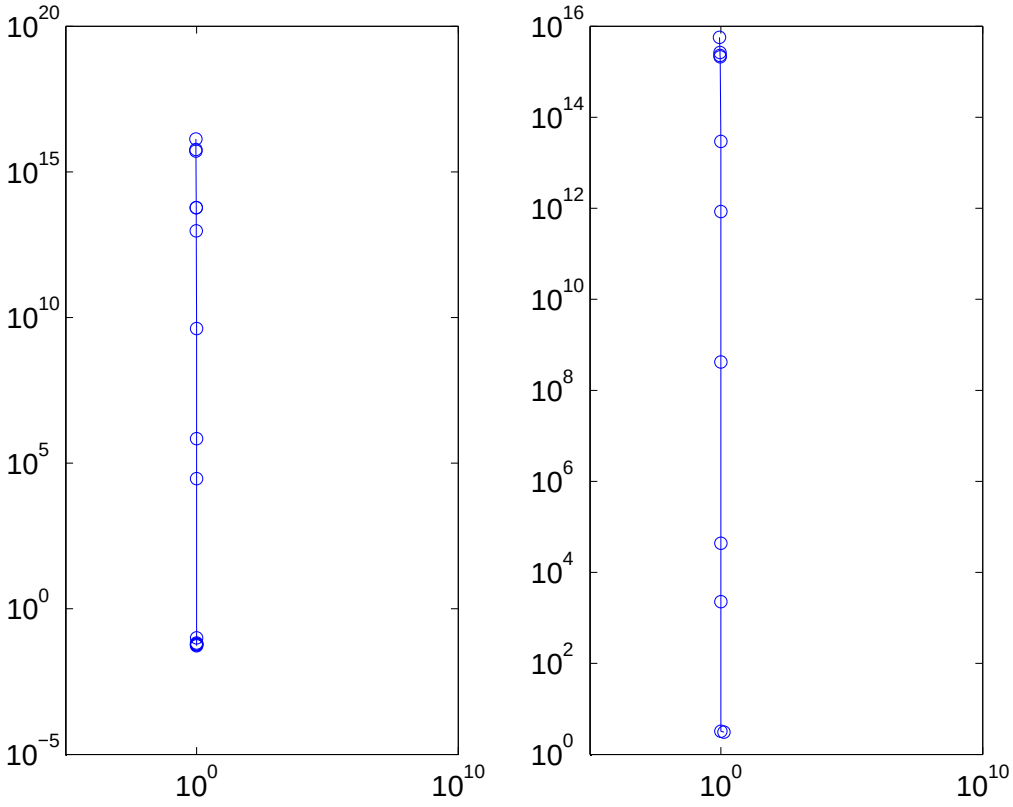

Shaw example: 40×20, $\delta = 10^{-8}$

Difficult situations: small var. in norms/residuals

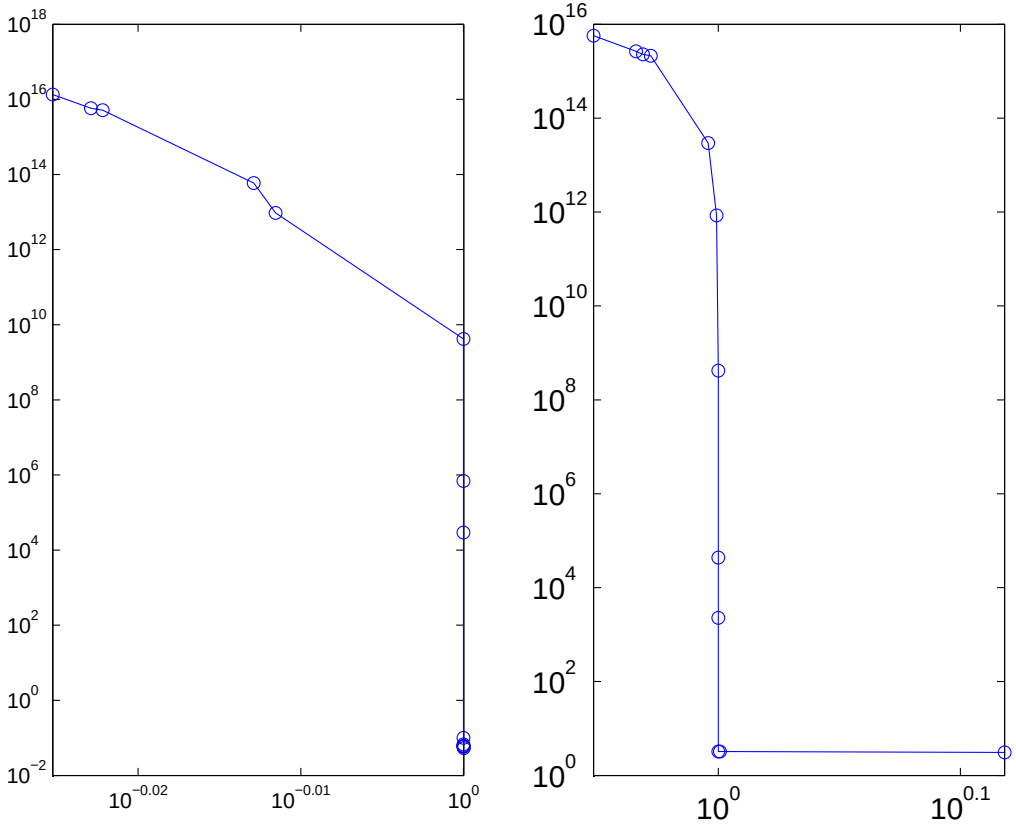

Shaw example: 40×20, $\delta = 10^{-8}$

left: $x \in \ker(L)$ right: $\min \|A\mathbf{x} - \mathbf{b}\| = 1$

The adaptive pruning (or *corner*) algorithm

The key idea is that if we remove the right amount of points from a discrete L-curve, then the corner can be easily found.

If too few points are removed we still maintain unwanted local features, and if too many points are removed the corner will be incorrectly located or may disappear.

[Hansen, Jensen, R 2006]

The adaptive pruning (or *corner*) algorithm

1. Stage one.
    1. We construct a sequence of *pruned* L-curves (selecting the largest line segments).
    2. For each curve, we select candidate corners by two algorithms:
        - the first is based on the local behaviour of the curve (angles);
        - the second is based on the global behaviour of the curve.
2. Stage two.
    1. We select a candidate from the list, so that
        - the curve is convex at that point;
        - the point is the last one before the residuals stagnate.

The *corner* algorithm at work

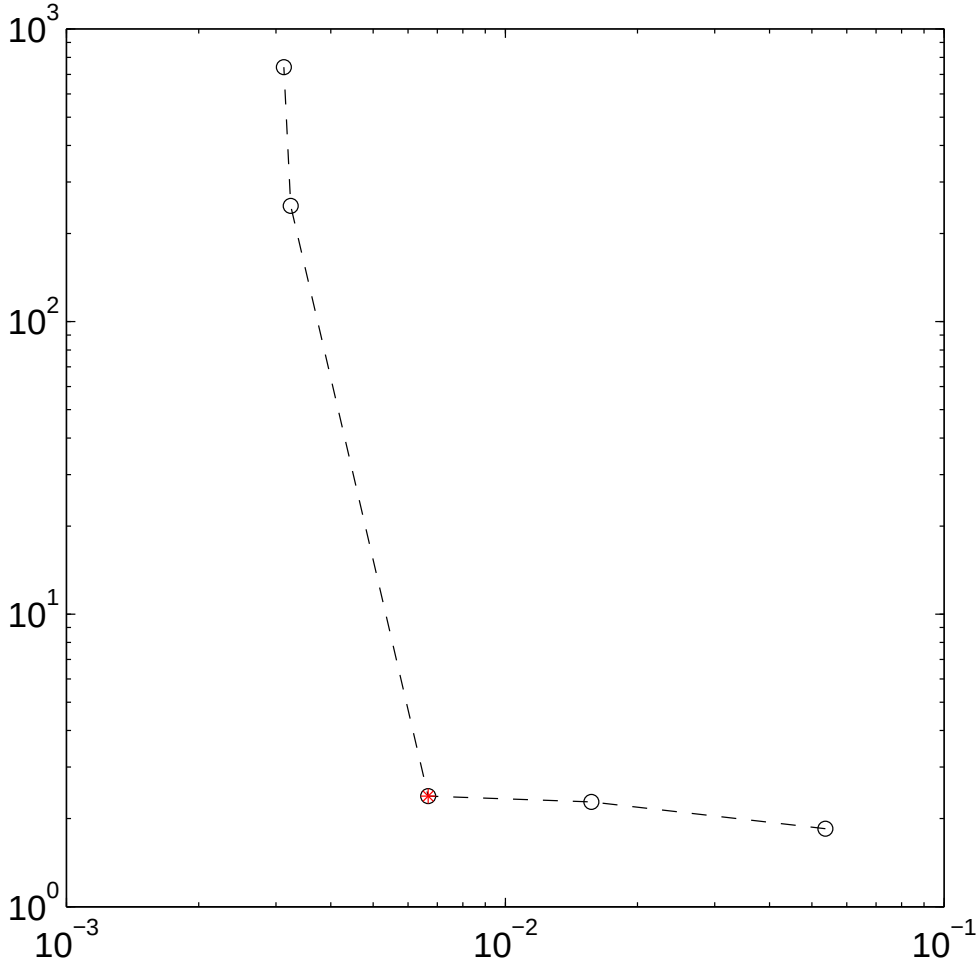

Heat(1) example: 100×50, $\delta = 10^{-2}$

The *corner* algorithm at work

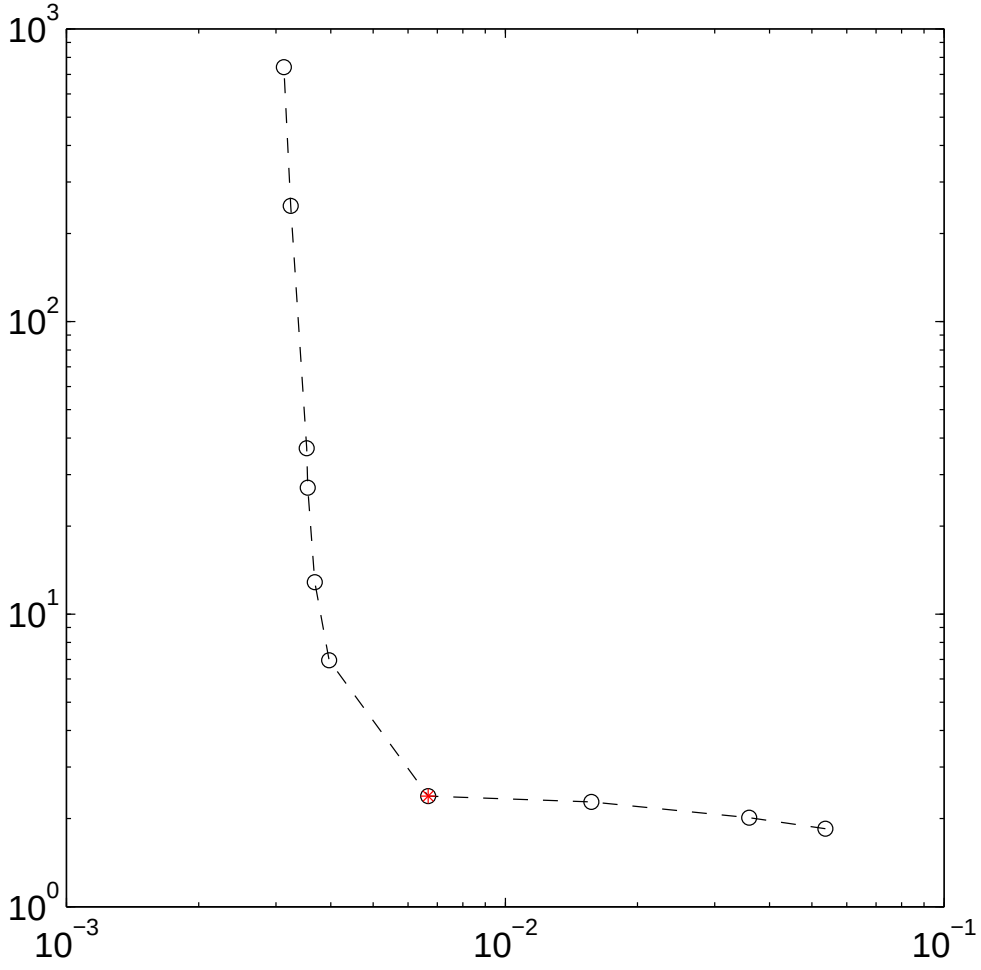

Heat(1) example: 100×50, $\delta = 10^{-2}$

The *corner* algorithm at work

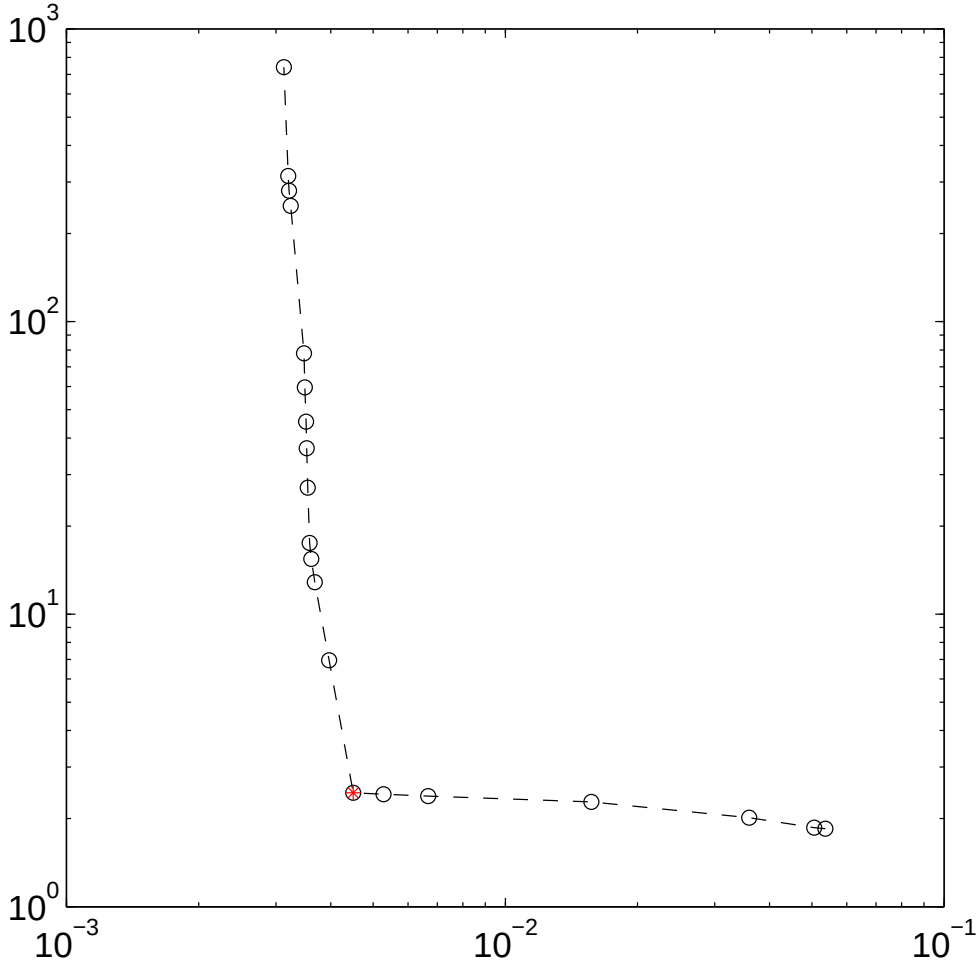

Heat(1) example: 100×50, $\delta = 10^{-2}$

The *corner* algorithm at work

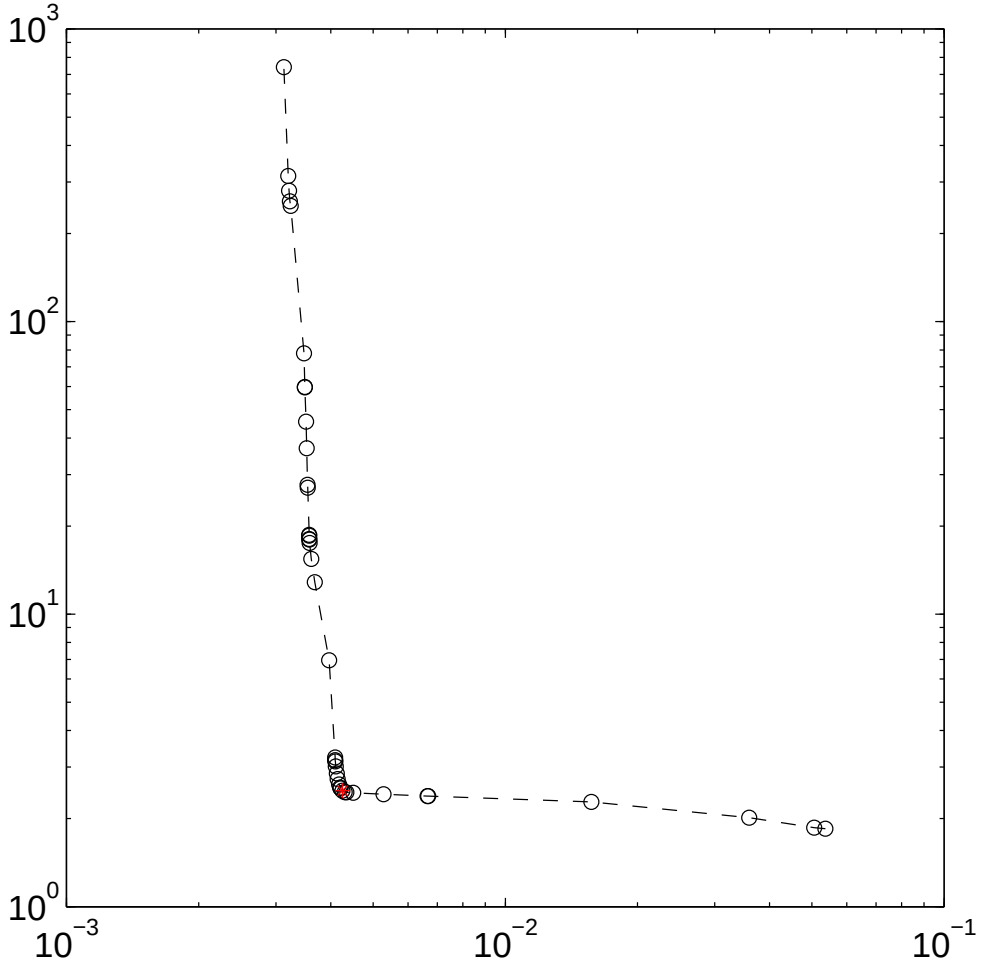

Heat(1) example: 100×50, $\delta = 10^{-2}$

The *corner* algorithm at work

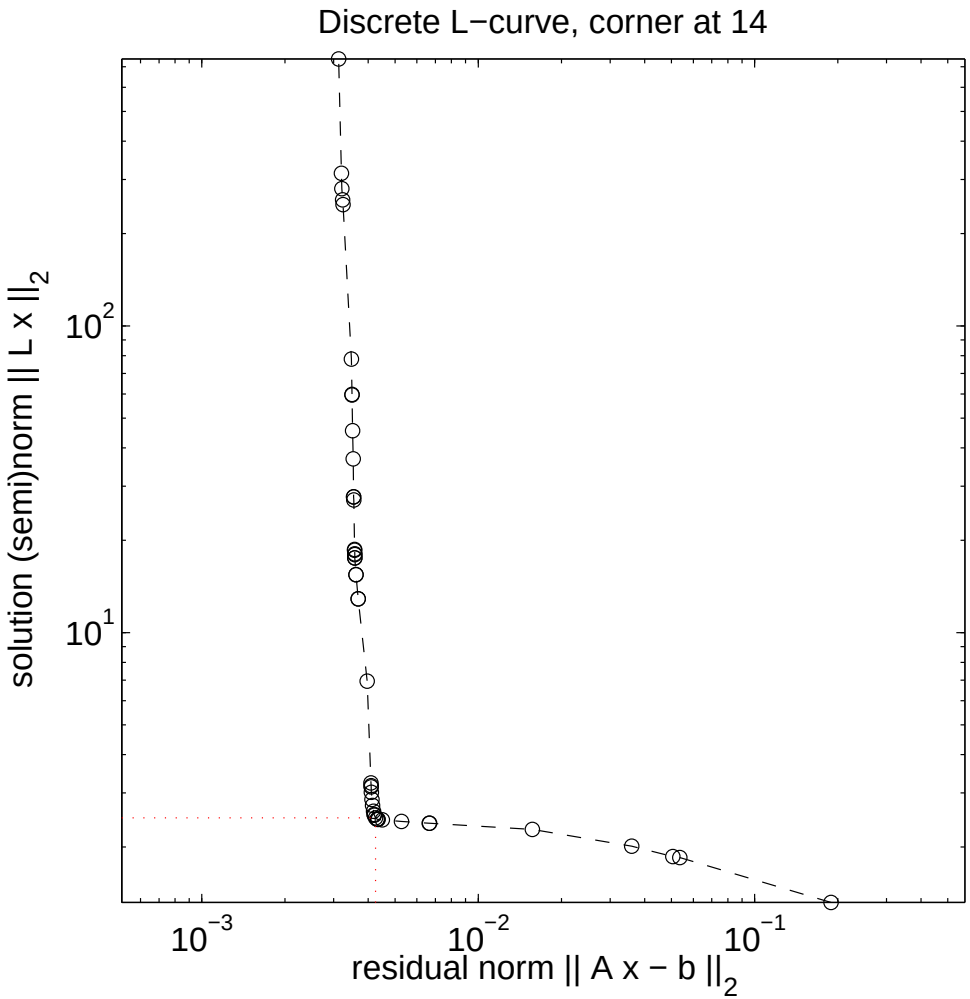

Heat(1) example: 100×50, $\delta = 10^{-2}$, $k_{\text{best}} = 12$

- Residual L-curve [Reichel, Sadok 2008]

$$\{\log \|\mathbf{r}_j\|, \log j\}, \quad j = 1, 2, \ldots, \ell.$$

- Condition L-curve [Calvetti, Lewis, Reichel 2002]

$$\{\log \|\mathbf{r}_j\|, \log \kappa(A_j)\}, \quad j = 1, 2, \ldots, \ell.$$

In both cases we select the TSVD truncation parameter by the *corner* algorithm.

In [Regińska 1996] an interesting analysis of the L-curve method, applied to Tikhonov regularization, is presented.

The choice rule proposed therein can be adapted to the situation when the regularization parameter is discrete.

In particular, it can be proved that minimizing the function

$$\phi_\nu(j) = \|\mathbf{r}_j\| \, \|\mathbf{x}_j\|^\nu, \quad j = 1, 2, \ldots, \ell - 1,$$

where $\nu > 0$ is a parameter, is *equivalent* to finding the corner of the L-curve. In fact

$$\phi_\nu(j) = \exp\left(\log \|\mathbf{r}_j\| + \nu \log \|\mathbf{x}_j\|\right) = \exp\left(x_j + \nu y_j\right)$$

is a norm of the vector from the origin to (x_j, y_j), on the L-curve.

The *restricted Regińska* (RR) method

The performance of the Regińska choice rule can be greatly enhanced by a preprocessing procedure:

- eliminate points $\{\|\mathbf{r}_j\|, \|\mathbf{x}_j\|\}$ with $\|\mathbf{x}_j\|$ very large;
- select the largest convex subset;
- apply Regińska's method.

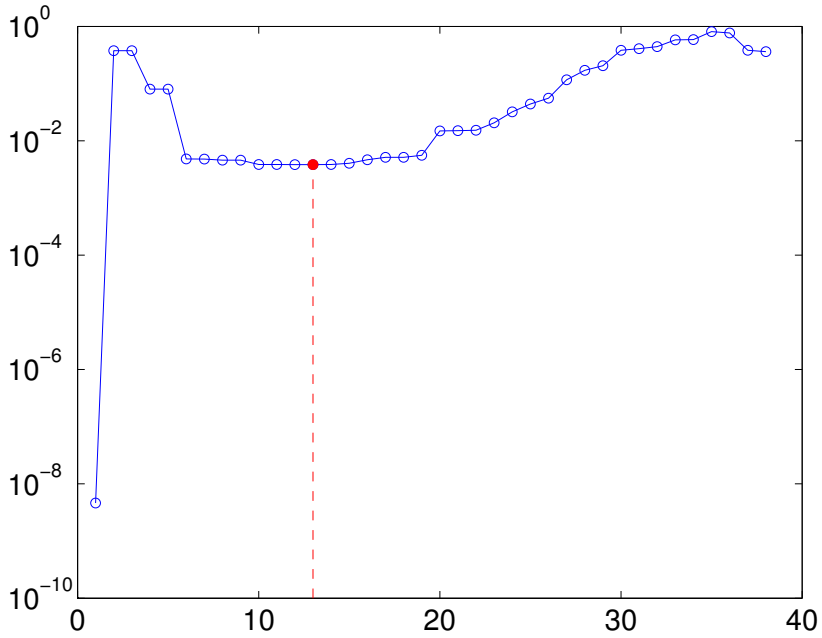

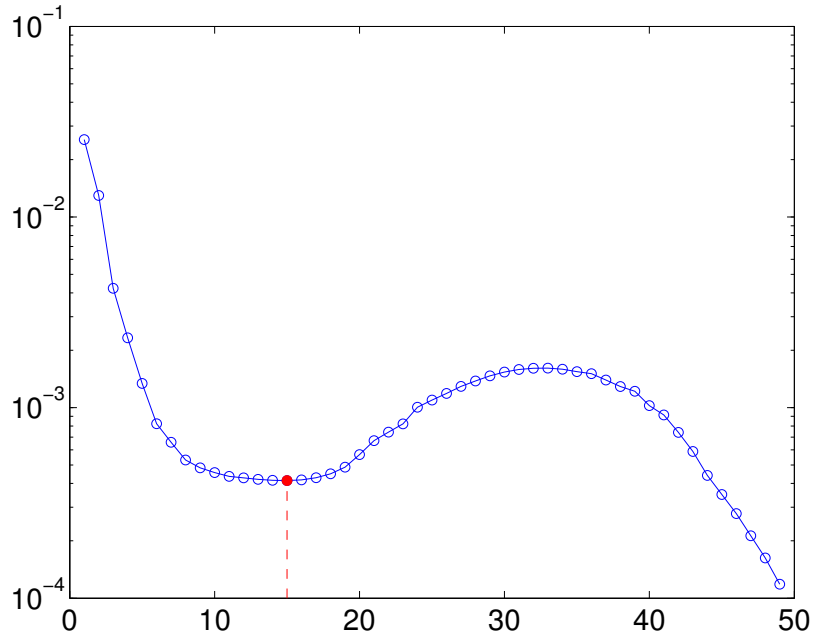

Error estimates based on extrapolation

Let $\mathbf{x}^\dagger$ be the normal solution to

$$\min_{\mathbf{x}\in\mathbb{R}^n} \|A\mathbf{x} - \mathbf{b}\|_2.$$

If $\mathbf{x}$ is an approximate solution and $\mathbf{r} = \mathbf{b} - A\mathbf{x}$, then

$$\|\mathbf{x}^\dagger - \mathbf{x}\|^2 \quad \simeq \quad \boxed{\eta_\nu^2 = d_0^{\nu-1} d_1^{5-2\nu} d_2^{\nu-3}} \qquad \nu \in \mathbb{R}$$

where

$$d_0 = \|\mathbf{r}\|^2, \quad d_1 = \|A^T\mathbf{r}\|^2, \quad d_2 = \|AA^T\mathbf{r}\|^2.$$

We select $k^* = \arg\min_k \eta_\nu(k)$.

Two examples: $\eta_2 = \frac{\|\mathbf{r}\|\cdot\|A^T\mathbf{r}\|}{\|AA^T\mathbf{r}\|}$, $\eta_3 = \frac{\|\mathbf{r}\|^2}{\|A^T\mathbf{r}\|}$ (Auchmuty).

[Brezinski, R, Seatzu 2008, 2009], [Reichel, R, Seatzu 2009].

Error estimates based on extrapolation

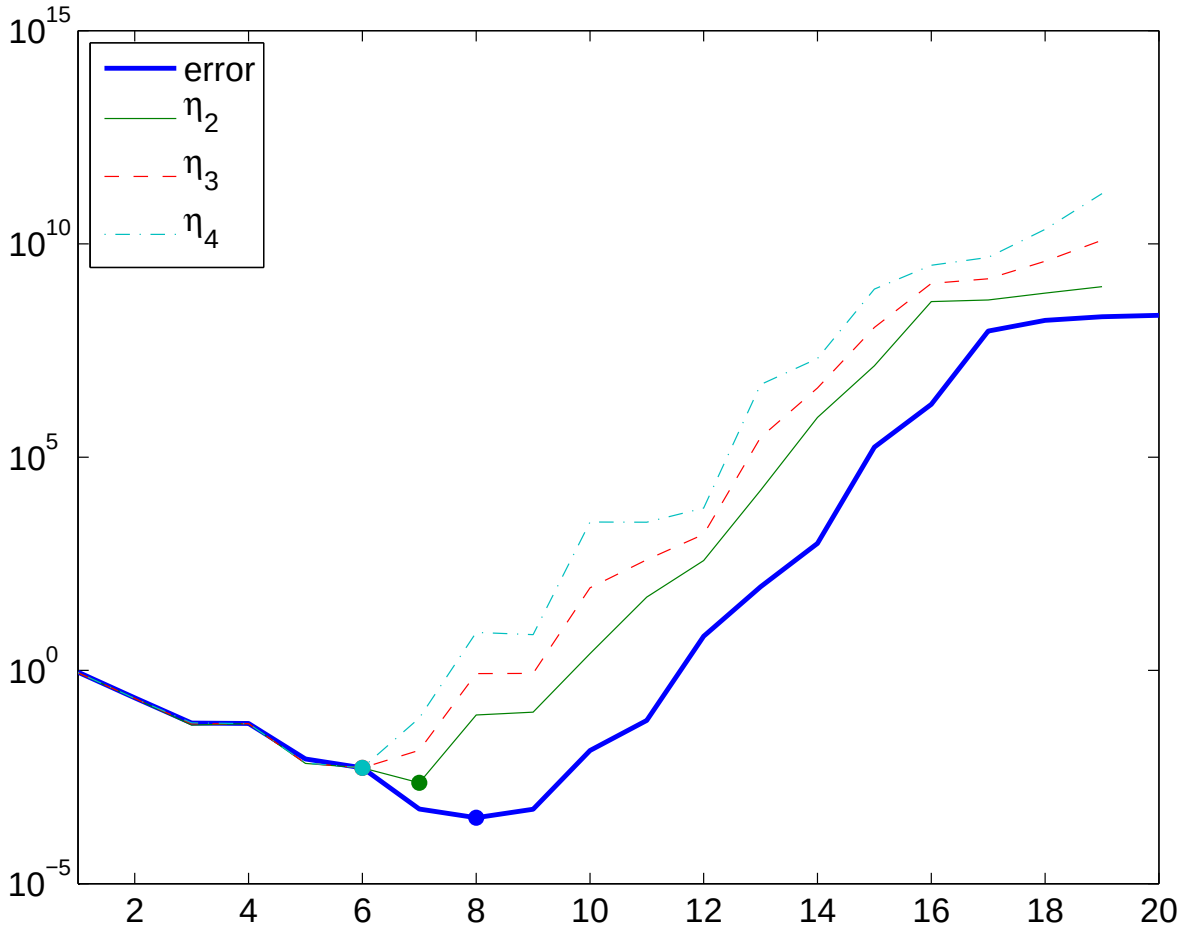

TSVD/*Shaw* example: 40×20, $\delta = 10^{-8}$

Quasi-optimality hybrid methods

The quasi-optimality criterion

$$\|\mathbf{x}_{k+1} - \mathbf{x}_k\| = \min_{1 \le j < \ell} \|\mathbf{x}_{j+1} - \mathbf{x}_j\|$$

was introduced in [Morozov 1984].

The function to be minimized typically has many local minima, which often lead to an unappropriate choice of the parameter.

We found it effective to use it in an hybrid approach:

- determine an initial index by a different method;
- apply the quasi-optimality criterion in a neighborhood of this estimate.

The iterates $\mathbf{x}_j \in \mathcal{K}_j(A^TA, A^T\mathbf{b})$ generated by LSQR are such that

$$\|A\mathbf{x}_j - \mathbf{b}\| = \min_{\mathbf{x}\in\mathcal{K}_j(A^TA,A^T\mathbf{b})} \|A\mathbf{x} - \mathbf{b}\|.$$

The method is implemented via Lanczos bidiagonalization of A:

$$AV_j = U_{j+1}\bar{C}_j, \qquad A^TU_j = V_jC_j^T,$$

where

- $U_{j+1}\mathbf{e}_1 = \mathbf{b}/\|\mathbf{b}\|$;
- the columns of U_{j+1} and V_j are orthonormal;
- $\bar{C}_j \in \mathbb{R}^{(j+1)\times j}$ is lower bidiagonal with leading submatrix C_j.

All the choice rules discussed can be inexpensively applied to LSQR to select the iterate $\mathbf{x}_j$ which best approximates $\mathbf{x}$.

Using a regularization matrix in LSQR

If we introduce a regularization matrix $L \in \mathbb{R}^{p \times n}$, the problem to be solved is the following

$$\min_{\mathbf{x} \in \mathcal{S}} \|L\mathbf{x}\|, \qquad \mathcal{S} = \{\mathbf{x} \in \mathbb{R}^n \mid A^T A\mathbf{x} = A^T \mathbf{b}\}.$$

In [Eldén 1982], the A-weighted generalized inverse $L_A^\dagger$ of L was introduced, with the property that the above problem reduces to

$$\min_{\mathbf{y} \in \mathbb{R}^n} \|AL_A^\dagger \mathbf{y} - \bar{\mathbf{b}}\|,$$

where $\bar{\mathbf{b}}$ is a suitable modification of $\mathbf{b}$.

We apply LSQR to the solution of this problem and use a choice rule to determine which iterate $\mathbf{y}_k$ to select. This iterate is then transformed to an approximate solution $\mathbf{x}_k$ of the original problem.

In [Hnětynková, Plešinger, Strakoš 2009], a new method to estimate the error $\mathbf{e}$ in $\mathbf{b}$ is presented, based on the connection between Lanczos bidiagonalization and Gauss quadrature originally introduced in [Golub, Meurant 1994].
Let

$$AA^T = W\Lambda W^T, \quad \Lambda = \mathrm{diag}[\lambda_1, \lambda_2, \ldots, \lambda_m] \in \mathbb{R}^{m\times m},$$

then the matrix function

$$\mathbf{b}^T f(AA^T)\mathbf{b} = \sum_{j=1}^{m} f(\lambda_j)\omega_j^2, \quad \omega_j = \mathbf{b}^T W\mathbf{e}_j,$$

can be viewed as a Stieltjes integral and approximated by a j-point Gauss quadrature formula. It can be shown that

$$G_j f := \|\mathbf{b}\|^2 \mathbf{e}_1^T f(C_j C_j^T)\mathbf{e}_1.$$

If we compute the SVD decomposition

$$C_j = W_j S_j \tilde{W}_j^T,$$

then

$$G_j f = \sum_{i=1}^{j} f\left((s_i^{(j)})^2\right)\left(\omega_i^{(j)}\right)^2,$$

where $S_j = \mathrm{diag}[s_1^{(j)}, s_2^{(j)}, \ldots, s_j^{(j)}]$ and $\omega_i^{(j)} = \|\mathbf{b}\|\mathbf{e}_1^T W_j \mathbf{e}_i$.

$(\omega_j^{(j)})^2$ is a decreasing function of j, which stagnates when j is large, where $(\omega_j^{(j)})^2 \approx \|\mathbf{e}\|^2$ [H., P., S. 2009]

They propose to stop the iteration when the stagnation first occur.

It is a modification of the Quadrature method. We consider

$$\rho_j = \frac{\omega_j^{(j)}}{s_j^{(j)}}, \qquad j = 1, 2, \ldots, i,$$

where $s_j^{(j)}$ is the smallest singular value of C_j.

We choose the iterate $\mathbf{x}_k$ such that

$$\rho_k = \min_{1 \le j \le i} \rho_j.$$

If a *plateau* is present around the minimum, we take the last point of the *plateau*.

LSQR: Quadrature vs. Ratio

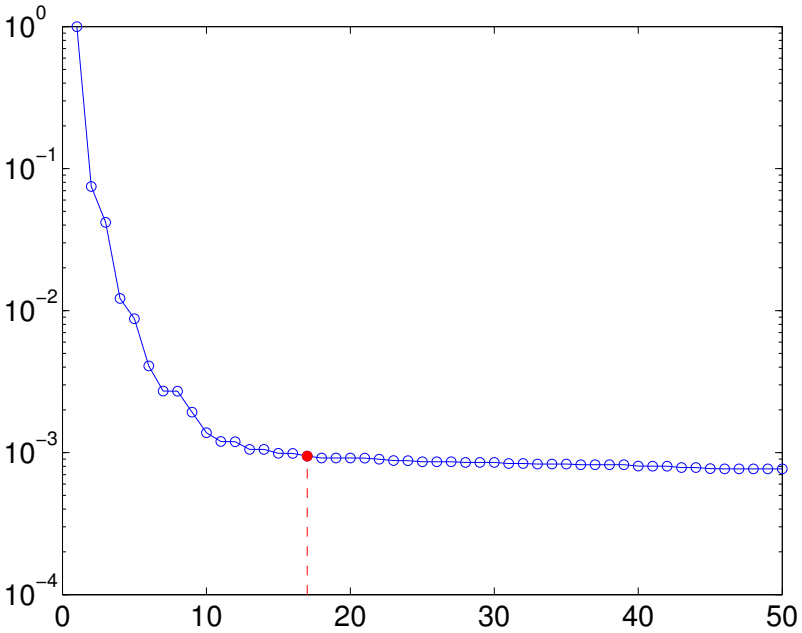

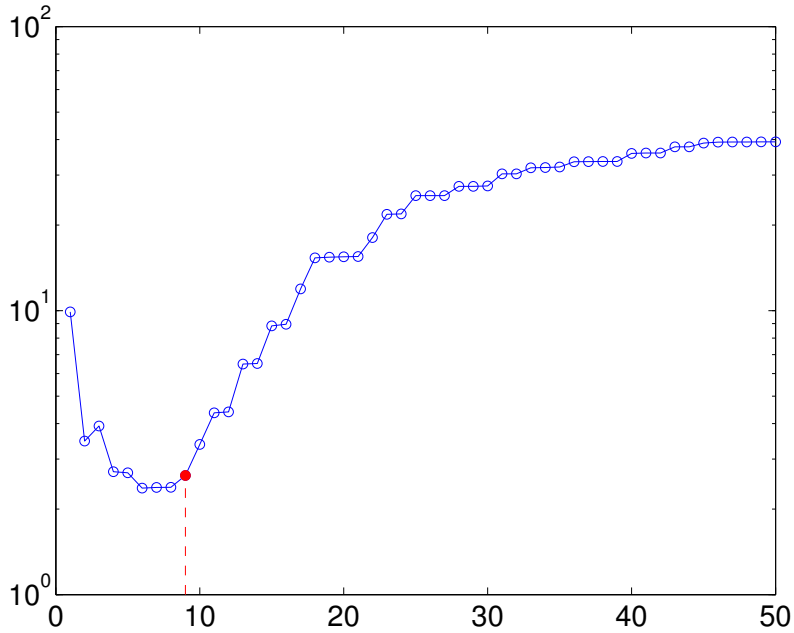

Method	Best	Restr. Regińska	Quasi-optimality	Quadrature	Ratio
k	11	15	49	17	9
error	0.29	0.44	8.42	0.73	0.32

Deriv2(2) example: 100×100, $\delta = 10^{-3}$

LSQR: the effect of reorthogonalization

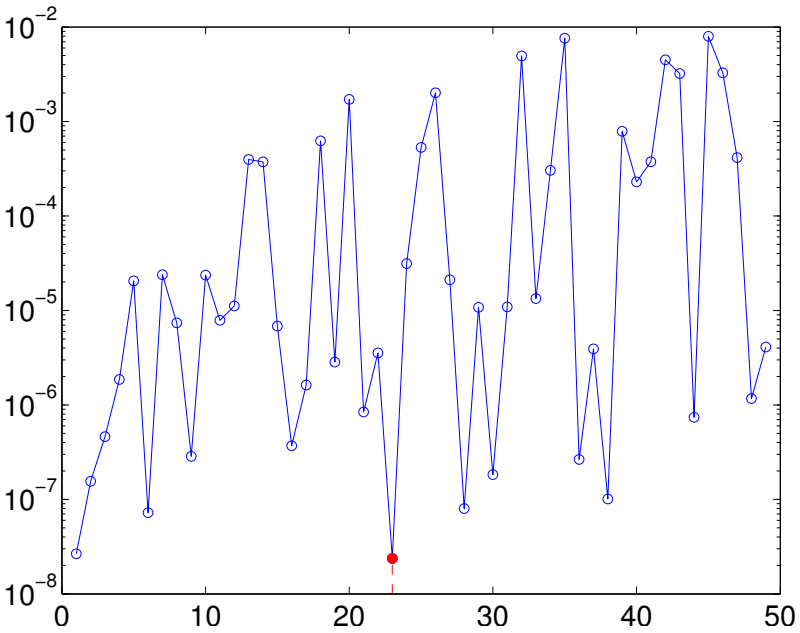

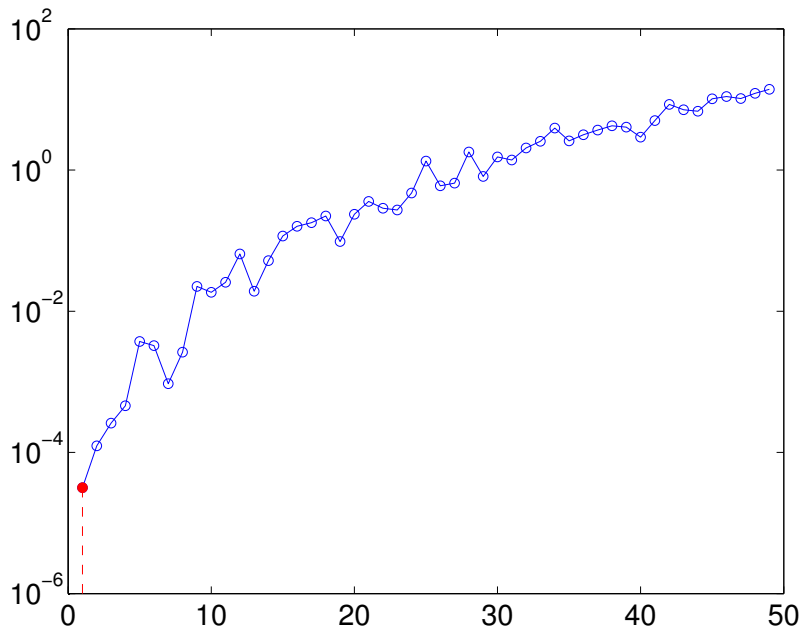

Method	Best	Restr. Regińska	Quad.	Ratio	Quasi-opt.	Q-opt./reorth.
k	1	1	15	1	23	1
error	0.021	0.021	9.21	0.021	27.21	0.021

Deriv2(2) example: 400×400, $\delta = 10^{-1}$, $H = D_2$

Numerical experiments: the *wild bunch*

We compared various methods on a set of 1200 linear systems:

- 10 test linear systems with prescribed solution
  - Matlab: Hilbert, Lotkin;
  - Regularization Tools: Baart, Deriv2(2), Foxgood, Gravity, Heat(1), Ilaplace(3), Phillips, Shaw.

Numerical experiments: the *wild bunch*

We compared various methods on a set of 1200 linear systems:

- 10 test linear systems with prescribed solution
- 3 (white) noise levels ($\delta = 10^{-3}, 10^{-2}, 10^{-1}$)
- 10 realizations of the noise ($\mathbf{b} = \hat{\mathbf{b}} + \delta \|\hat{\mathbf{b}}\| \cdot \mathbf{w}/\sqrt{n}$)
- the systems are either $(n \times n)$ or $(2n \times n)$, $n = 40, 100$ for TSVD, $n = 100, 400$ for LSQR.

Numerical experiments: the *wild bunch*

We compared various methods on a set of 1200 linear systems:

- 10 test linear systems with prescribed solution
- 3 (white) noise levels ($\delta = 10^{-3}, 10^{-2}, 10^{-1}$)
- 10 realizations of the noise ($\mathbf{b} = \hat{\mathbf{b}} + \delta \|\hat{\mathbf{b}}\| \cdot \mathbf{w}/\sqrt{n}$)
- the systems are either $(n \times n)$ or $(2n \times n)$, $n = 40, 100$ for TSVD, $n = 100, 400$ for LSQR.

Quality index:

$$Q(\mu) = \frac{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{est}}}\|}{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{opt}}}\|} \geq 1$$

failure: $Q(\mu) > 10$ *severe failure*: $Q(\mu) > 10^2$

Parameter choice rules for TSVD

Method	$n \times n$	$2n \times n$
L-corner	6%(0%)	4%(0%)
Res. L-curve	3%(1%)	1%(0%)
Cond. L-curve	2%(1%)	8%(6%)
Regińska	4%(0%)	2%(0%)
Restr. Regińska	3%(0%)	2%(0%)
Quasi-optimality	9%(1%)	7%(1%)
GCV	24%(20%)	8%(4%)
Extrapolation	5%(1%)	9%(0%)
Discrepancy	0%(0%)	0%(0%)

$$\frac{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{est}}}\|}{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{opt}}}\|} > 10, \qquad \frac{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{est}}}\|}{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{opt}}}\|} > 10^2$$

Parameter choice rules for LSQR

Method	$n \times n$	$2n \times n$
L-corner	4%(0%)	2%(0%)
L-triangle	4%(3%)	4%(4%)
Res. L-curve	13%(12%)	10%(10%)
Cond. L-curve	4%(4%)	13%(9%)
Regińska	3%(0%)	1%(0%)
Restr. Regińska	1%(0%)	1%(0%)
Quasi-optimality	2%(0%)	1%(0%)
Extrapolation	8%(0%)	12%(1%)
Quadrature	3%(1%)	2%(0%)
Ratio	0%(0%)	0%(0%)
Discrepancy	0%(0%)	0%(0%)

$$\frac{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{est}}}\|}{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{opt}}}\|} > 10, \qquad \frac{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{est}}}\|}{\|\mathbf{x} - \mathbf{x}_{\mu_{\text{opt}}}\|} > 10^2$$

Regularization matrix $L = D_2$

TGSVD	$n \times n$	$2n \times n$
Regińska	5%(0%)	6%(0%)
Restr. Regińska	2%(0%)	2%(0%)
Quasi-optimality	12%(0%)	6%(0%)
Extrapolation	8%(2%)	3%(0%)
Discrepancy	3%(0%)	2%(0%)

Std.form/LSQR	$n \times n$	$2n \times n$
Regińska	8%(0%)	7%(0%)
Restr. Regińska	1%(0%)	2%(0%)
Quasi-optimality	6%(0%)	4%(0%)
Ratio	11%(0%)	11%(1%)
Discrepancy	3%(0%)	3%(0%)

TSVD and Hybrid-Quasioptimality

Method	$n \times n$	$2n \times n$
Q-L-corner	6%(0%)	4%(0%)
Q-Res. L-curve	2%(1%)	1%(0%)
Q-Cond. L-curve	1%(1%)	5%(2%)
Q-Regińska	4%(0%)	2%(0%)
Q-Restr. Regińska	3%(0%)	2%(0%)
Q-GCV	20%(17%)	4%(1%)
Q-Extrapolation	1%(1%)	0%(0%)

Method	$n \times n$	$2n \times n$
L-corner	6%(0%)	4%(0%)
Res. L-curve	3%(1%)	1%(0%)
Cond. L-curve	2%(1%)	8%(6%)
Regińska	4%(0%)	2%(0%)
Restr. Regińska	3%(0%)	2%(0%)
GCV	24%(20%)	8%(4%)
Extrapolation	5%(1%)	9%(0%)

The End

Thanks!