Error estimates in linear systems with an application to regularization

Claude BREZINSKI

University of Lille - France

- ➜ The norm of the error
- ➜ Extrapolation
- ➜ The formula for extrapolation
- ➜ Estimates of the norm of the error
- ➜ A generalization
- ➜ An application to regularization

The norm of the error

$$Ax = b$$

y an approximation of x and $e = x - y$.

We want to obtain an estimation of $\|\mathbf{e}\|$.

We set $r = b - Ay$. We have $Ae = r$ and the bounds

$$\frac{\|r\|}{\|A\|} \leq \|e\| \leq \|A^{-1}\| \cdot \|r\|.$$

These bounds require the knowledge of $\|A\|$ and $\|A^{-1}\|$, But $\|A^{-1}\|$ is difficult to compute, and the lower bound can be quite a bad estimate of $\|e\|$.

Estimates for the error in the conjugate gradient were given by Golub and Meurant.

EXTRAPOLATION

We will now obtain estimates of $\|e\|$ by an **extrapolation method**.

We have

$c_0 = (\mathbf{A^0 r}, \mathbf{A^0 r})$	$\mathbf{A^0}$	$\mathbf{A^0}$	$0 + 0 = 0$
$c_1 = (\mathbf{A^0 r}, \mathbf{A^1 r})$	$\mathbf{A^0}$	$\mathbf{A^1}$	$0 + 1 = 1$
$c_2 = (\mathbf{A^1 r}, \mathbf{A^1 r})$	$\mathbf{A^1}$	$\mathbf{A^1}$	$1 + 1 = 2$
$c_{-1} = (\mathbf{A^0 r}, \mathbf{A^{-1} r})$	$\mathbf{A^0}$	$\mathbf{A^{-1}}$	$0 + (-1) = -1$
$c_{-2} = (\mathbf{A^{-1} r}, \mathbf{A^{-1} r})$	$\mathbf{A^{-1}}$	$\mathbf{A^{-1}}$	$(-1) + (-1) = -2$

$$\mathbf{c_{-1}} = (\mathbf{A^0 r}, \mathbf{A^{-1} r}) = (\mathbf{e}, \mathbf{Ae})$$ A-norm of the error

$$\mathbf{c_{-2}} = (\mathbf{A^{-1} r}, \mathbf{A^{-1} r}) = (\mathbf{e}, \mathbf{e}) = \|\mathbf{e}\|^{\mathbf{2}}$$ norm of the error

We will **interpolate** the points $(\mathbf{0}, \mathbf{c_0}), (\mathbf{1}, \mathbf{c_1})$ and $(\mathbf{2}, \mathbf{c_2})$ by some function and then ....

extrapolate at the point $-\mathbf{2}$.

What is extrapolation ?

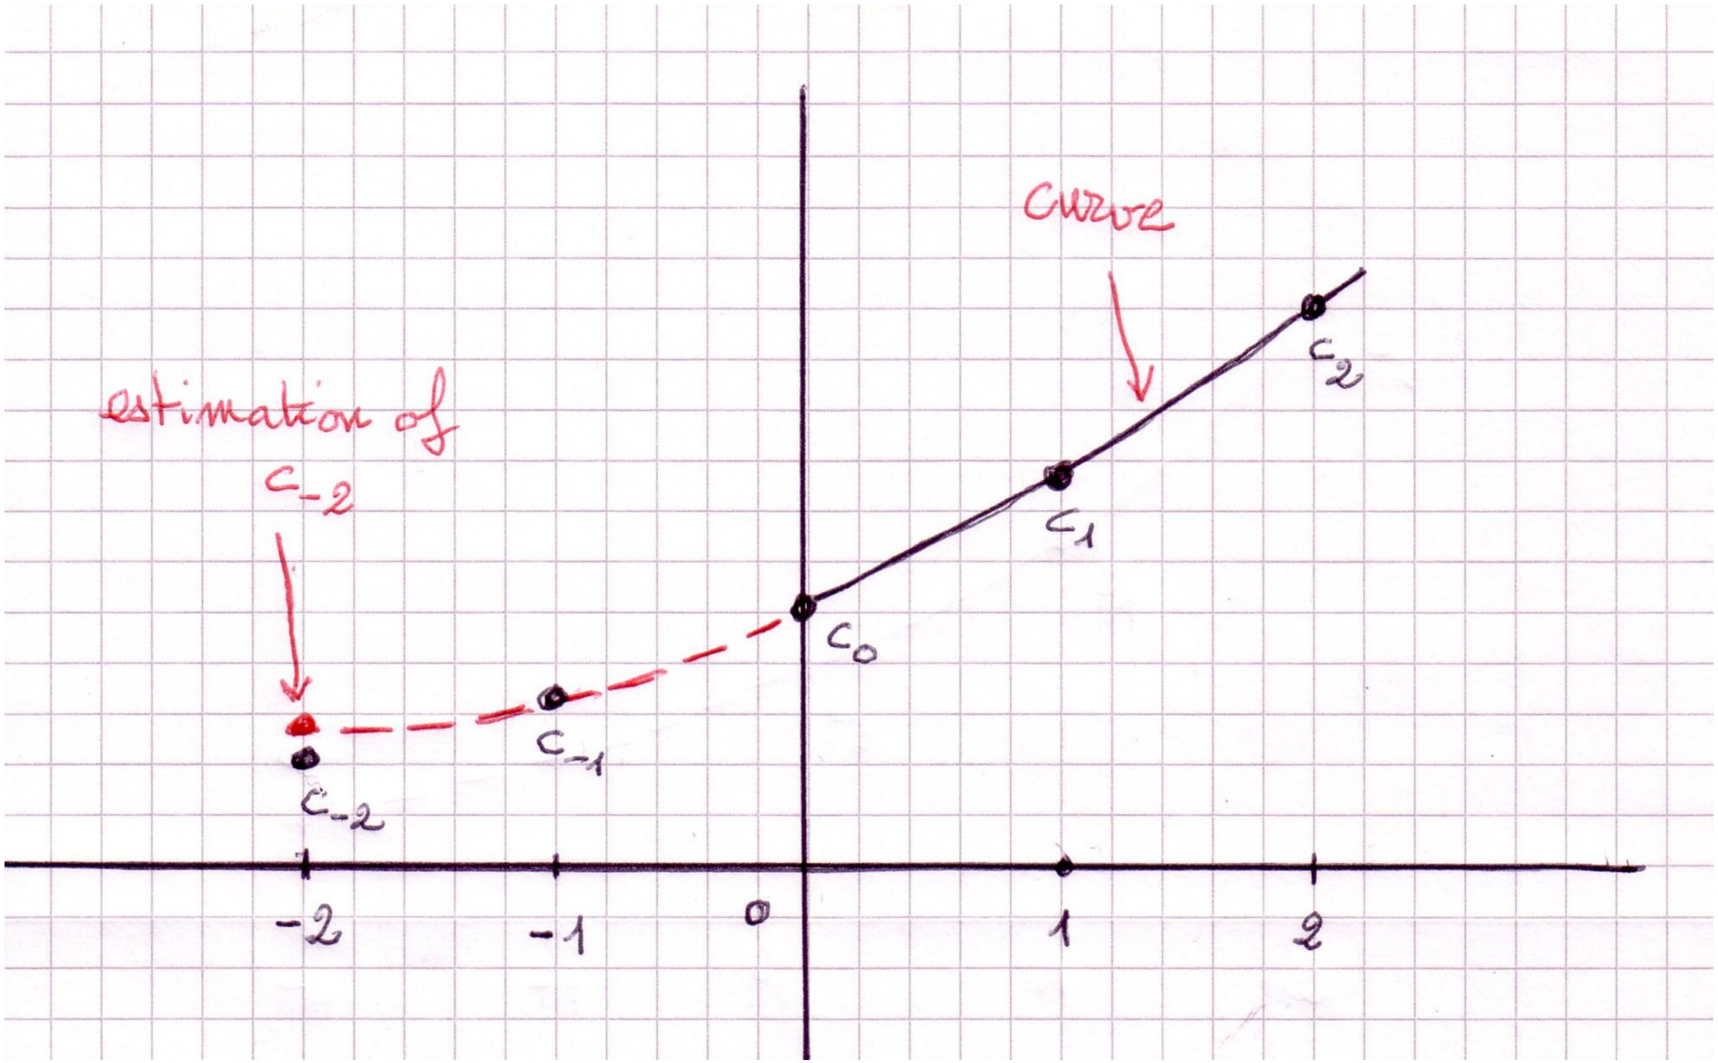

Which curve?

Answer: a curve which mimics the exact behavior of the $\mathbf{c_i}$

If the function mimics the behaviour of the c_i, then its value at -2 will be a good approximation of $\|e^2_{\|}$.

For choosing the interpolation function, we have to analyze the behaviour of c_0, c_1 and c_2.

So, let us now analyze this behavior.

We consider the **Singular Value Decomposition** (**SVD**) of the matrix A:

$$\begin{aligned} A &= U\Sigma V^T \\ U &= [u_1, \ldots, u_p] \\ V &= [v_1, \ldots, v_p] \end{aligned}$$

where
$UU^T = VV^T = I$,
$\Sigma =$diag$(\sigma_1, \ldots, \sigma_p)$ with $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p > 0$.

Let $\mathbf{v}$ be **any vector**. It holds

$$\begin{aligned} A\mathbf{v} &= \sum_{i=1}^{p} \sigma_i (v_i, \mathbf{v}) u_i \\ A^T\mathbf{v} &= \sum_{i=1}^{p} \sigma_i (u_i, \mathbf{v}) v_i \\ A^{-1}\mathbf{v} &= \sum_{i=1}^{p} \sigma_i^{-1} (u_i, \mathbf{v}) v_i. \end{aligned}$$

$$
\begin{aligned}
c_0 = (r,r) &= (U^T r, U^T r) = \sum_{i=1}^{p} \alpha_i^2, \quad \alpha_i = (u_i, r) \\
&= (V^T r, V^T r) = \sum_{i=1}^{p} \beta_i^2, \quad \beta_i = (v_i, r) \\
c_1 = (r, Ar) &= \sum_{i=1}^{p} \sigma_i \alpha_i \beta_i \\
c_2 = (Ar, Ar) &= \sum_{i=1}^{p} \sigma_i^2 \beta_i^2 \\
c_{-1} = (A^{-1} r, r) &= (e, Ae) = \sum_{i=1}^{p} \sigma_i^{-1} \alpha_i \beta_i \\
c_{-2} = (A^{-1} r, A^{-1} r) &= (e, e) = \sum_{i=1}^{p} \sigma_i^{-2} \alpha_i^2 .
\end{aligned}
$$

The formula for extrapolation

The function we will use for extrapolation has to mimic as closely as possible the behavior of the c_i.

So, we will keep only the **first term** in each of the preceding formulae, that is we will look for α, β and σ satisfying the **interpolation conditions**

$$
\begin{aligned}
c_0 &= \alpha^2 = \beta^2 \\
c_1 &= \sigma\alpha\beta \\
c_2 &= \sigma^2\beta^2.
\end{aligned}
$$

and then **extrapolate** for the values -1 and -2 of the index. Thus, c_{-1} and c_{-2} will be approximated by

$$\mathbf{c_{-1}} \simeq \sigma^{\mathbf{-1}}\alpha\beta \quad \text{and} \quad \mathbf{c_{-2}} = \|\mathbf{e}\|^{\mathbf{2}} \simeq \sigma^{\mathbf{-2}}\alpha^{\mathbf{2}}.$$

Estimates of the norm of the error

The preceding system has 3 unknowns and 4 equations which are not compatible. Thus, it has several solutions. For example, we get the following e_i^2 which are approximations of $\|e\|^2$

$$\begin{aligned} e_1^2 &= c_1^4/c_2^3 \\ e_2^2 &= c_0 c_1^2/c_2^2 \\ e_3^2 &= c_0^2/c_2 \\ e_4^2 &= c_0^3/c_1^2 \\ e_5^2 &= c_0^4 c_2/c_1^4. \end{aligned}$$

These estimates were numbered so that

$$e_1 \leq e_2 \leq e_3 \leq e_4 \leq e_5.$$

More estimates:

More estimates can be obtained by replacing c_2 in all formulae above by

$$\widetilde{c}_2 = (A^T r, A^T r) = \sum_{i=1}^{p} \sigma_i^2 \alpha_i^2,$$

and approximating it by $\sigma^2 \alpha^2$.

Similar results and properties are obtained.

They will be denoted by putting a $\sim$ over the letters.

It holds

$$\widetilde{e}_\nu^2 \leq \|e\|^2, \quad \forall \nu \leq 3.$$

The estimate $\widetilde{e}_3$ was given by Auchmuty in 1992.

NUMERICAL EXAMPLES:

$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$ is the **exact** solution of the linear system.

$\mathbf{y}$ is **any approximate** solution of it.

So, our estimates apply either to a **direct method** or to an **iterative method** for the solution of a system of linear equations.

They estimate **both** the rounding errors and the error of the method.

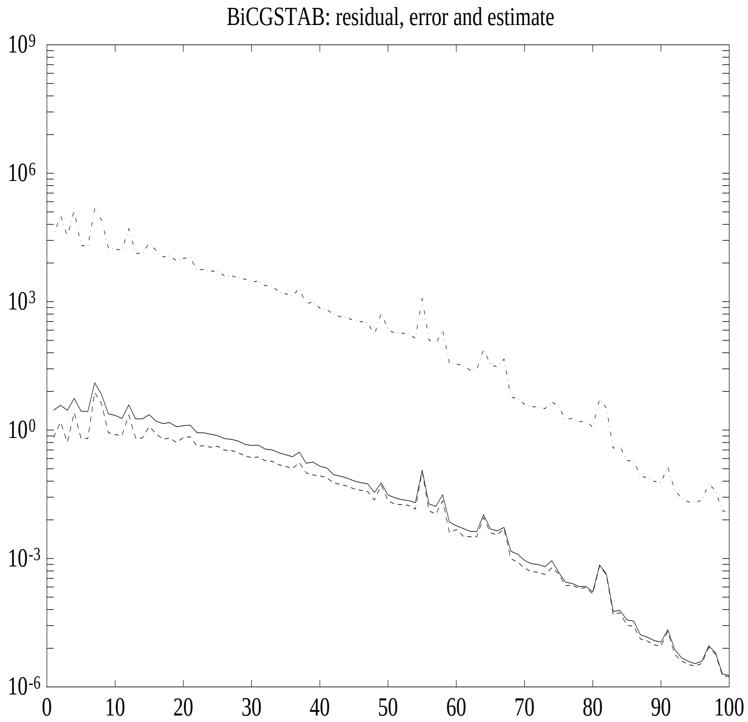

Figure 1: BiCGSTAB for I+50*CIRCUL(100); cond(A) =101

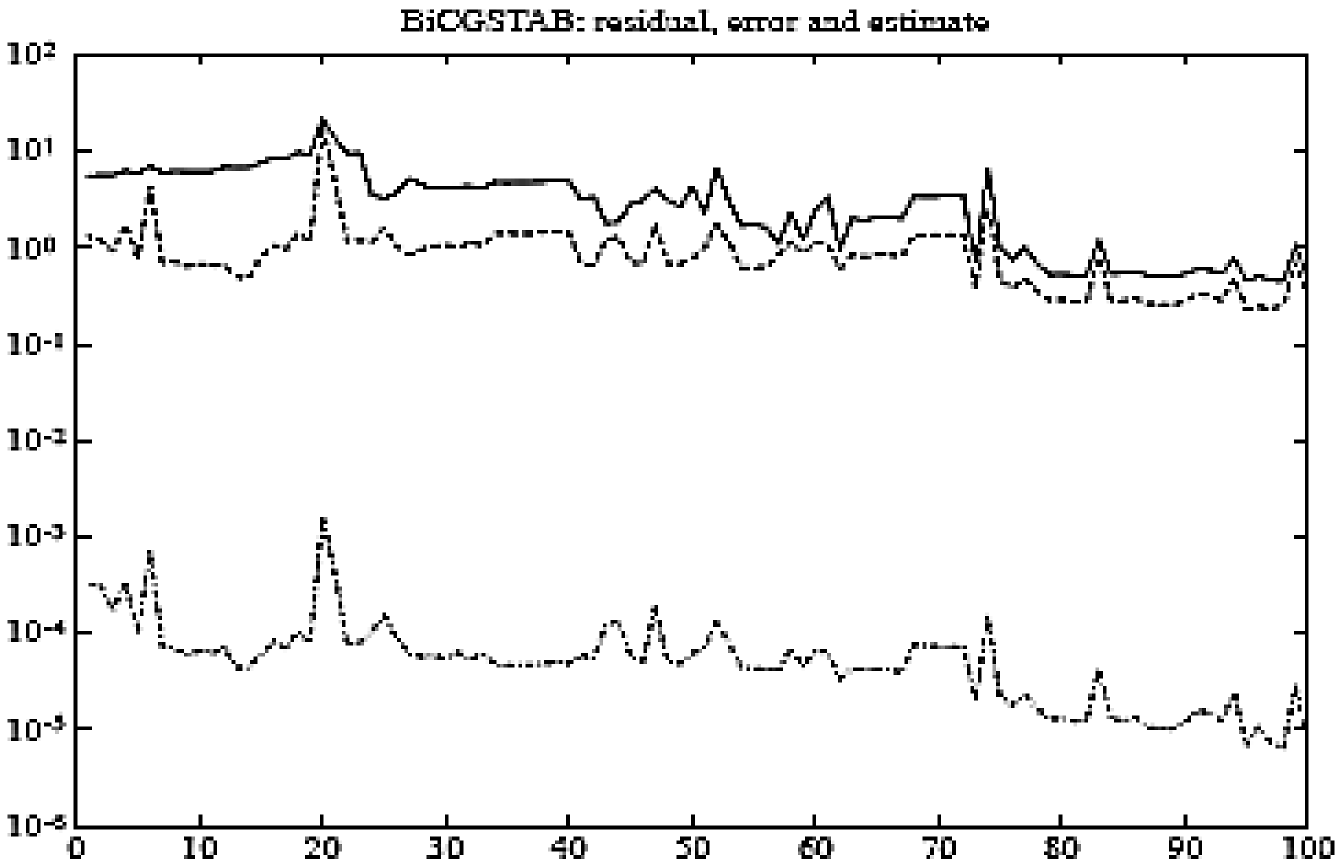

Figure 2: inv(I+50*circul(100)) $\kappa = 101.0408$,
$\|A\| = 4.0016 \cdot 10^{-4}$, $\|A^{-1}\| = 2.5250 \cdot 10^{5}$

A GENERALIZATION:

From here: joint work with G. Rodriguez and S. Seatzu (University of Cagliari, Italy).

The five estimates $e_1, \ldots, e_5$ can be gathered into only one formula

$$e_i^2 = c_0^{i-1}(c_1^2)^{3-i}c_2^{i-4}, \quad i = 1, \ldots, 5.$$

Moreover, this formula is not only valid for $i = 1, \ldots, 5$, but also for **any real number** ν, that is

$$\mathbf{e}_\nu^2 = \mathbf{c}_0^{\nu-1}(\mathbf{c}_1^2)^{3-\nu}\mathbf{c}_2^{\nu-4}, \quad \nu \in \mathbb{R}.$$

PROPERTIES:

We have

$$e_\nu^2 = \left(\frac{c_0 c_2}{c_1^2}\right)^\nu \cdot \left(\frac{c_1^6}{c_0 c_2^4}\right) = \rho^\nu e_0^2.$$

So, $\mathbf{e_\nu^2}$ is an **increasing** function of ν in $(-\infty, +\infty)$ since $\rho = (\mathbf{c_0 c_2})/\mathbf{c_1^2} \geq \mathbf{1}$.

Therefore, it exists $\nu_\mathbf{e}$ such that $\mathbf{e_{\nu_e}^2} = \|\mathbf{e}\|^\mathbf{2}$.

This ν_e is given by the formula

$$\nu_e = 2\ln(\|e\|/e_0)/\ln\rho.$$

AN APPLICATION TO REGULARIZATION

When a system is ill-conditioned, its solution cannot be computed accurately.

Tikhonov's regularization consists in computing the vector $\mathbf{x}_\lambda$ which **minimizes** the quadratic functional

$$J(\lambda, x) = \|Ax - b\|^2 + \lambda^2 \|Hx\|^2$$

over all vectors x, where λ is a parameter, and H a given $q \times p$ ($q \leq p$) matrix.

This vector x_λ is the solution of the system

$$(C + \lambda^2 E) x_\lambda = A^T b,$$

where $C = A^T A$ and $E = H^T H$.

If λ is close to zero, then x_λ is badly computed while, if λ is far away from zero, x_λ is well computed but the error $x - x_\lambda$ is quite large.

For decreasing values of λ, the norm of the error $\|x - x_\lambda\|$ first decreases, and then increases when λ approaches 0.

Thus **the error**, which is the sum of the theoretical error and the error due to the computer's arithmetic, **passes through a minimum** corresponding to the optimal choice of the regularization parameter λ.

Several methods have been proposed to obtain an effective choice of λ.

The **$\mathbf{L}$–curve** consists in plotting in $\log$–$\log$ scale the values of $\|Hx_\lambda\|$ versus $\|r_\lambda\|$. The resulting curve is typically L–shaped and the selected value of λ is the one corresponding to the corner of the L. However, there are many cases where the L–curve exhibits more than one corner, or no one at all.

The **Generalized Cross Validation** (GCV) searches for the minimum of a function of λ which is a statistical estimate of the norm of the residual. Occasionally, the value of the parameter furnished by this method may be inaccurate because the function is rather flat near the minimum.

But each of these methods can fail.

We are proposing **another test based on the preceding estimates of the norm of the error**.

Warning :
We don't pretend that this new procedure never fails!!!

There are two questions that have to be answered:

- **Is x_λ well computed?**

  For answering this question, we propose the following test.

  Remember that the vector x_λ is the solution of

  $$(C + \lambda^2 E)x_\lambda = A^T b,$$

  where $C = A^T A$ and $E = H^T H$.

  Set $r_\lambda = b - Ax_\lambda$.

  Since $A^T r_\lambda = \lambda^2 E x_\lambda$, it holds

  $$\frac{\lambda^2 \|E x_\lambda\|}{\|A^T r_\lambda\|} = 1.$$

  So, it could be checked if this ratio is close to 1 for all λ.

- **Is x_λ a good approximation of x?**

  For this question, the preceding estimates could be used.

  However, due to the ill-conditioning, $\widetilde{c}_2 = \|A^T r_\lambda\|^2$ is badly computed when λ approches zero.

  So, again, we will replace $A^T r_\lambda$ by $\lambda^2 E x_\lambda$ in $\|A^T r_\lambda\|$ and in $(r_\lambda, A r_\lambda) = (A^T r_\lambda, r_\lambda)$.

In order to find the best value of the parameter λ, we will now apply our estimates of the norm of the error to Tikhonov's regularization method.

Effecting the preceding substitutions, we finally obtain the error estimates

$$\widetilde{\mathbf{e}}_\nu^{\mathbf{2}} = \|\mathbf{r}_\lambda\|^{\mathbf{2}\nu-\mathbf{2}}(\mathbf{r}_\lambda, \mathbf{E}\mathbf{x}_\lambda)^{\mathbf{6}-\mathbf{2}\nu}\|\mathbf{E}\mathbf{x}_\lambda\|^{\mathbf{2}\nu-\mathbf{8}}\lambda^{-\mathbf{4}}.$$

Contrarily to the more general estimates which are always valid, this new formula is only valid for Tikhonov's regularization. So, it should lead to better numerical results. Testing the equality $\lambda^2\|Ex_\lambda\|/\|A^T r_\lambda\| = 1$ is also only valid for Tikhonov's regularization.

Let us remark that $(r_\lambda, Ex_\lambda) = (Hr_\lambda, Hx_\lambda)$ which avoids computing the matrix E and, in several cases, leads to a more stable procedure.

Example 1:

In this example we show how our estimates behave in a problem for which the L–curve method fails.

We consider the *Pascal* matrix of dimension 20 whose estimated condition number is $1.03 \cdot 10^{+21}$. The solution was chosen to be $x = (1, \dots, 1)^T$, the noise level on the right hand side was 10^{-8}, and the regularization matrix was the identity.

The thick line gives the Euclidean norm of the error. From the bottom to the top, the solid lines represent $\widetilde{e}_1, \widetilde{e}_3$ and $\widetilde{e}_5$ versus λ, while the dashed ones are $\widetilde{e}_2, \widetilde{e}_4$ and $\widetilde{e}_6$.

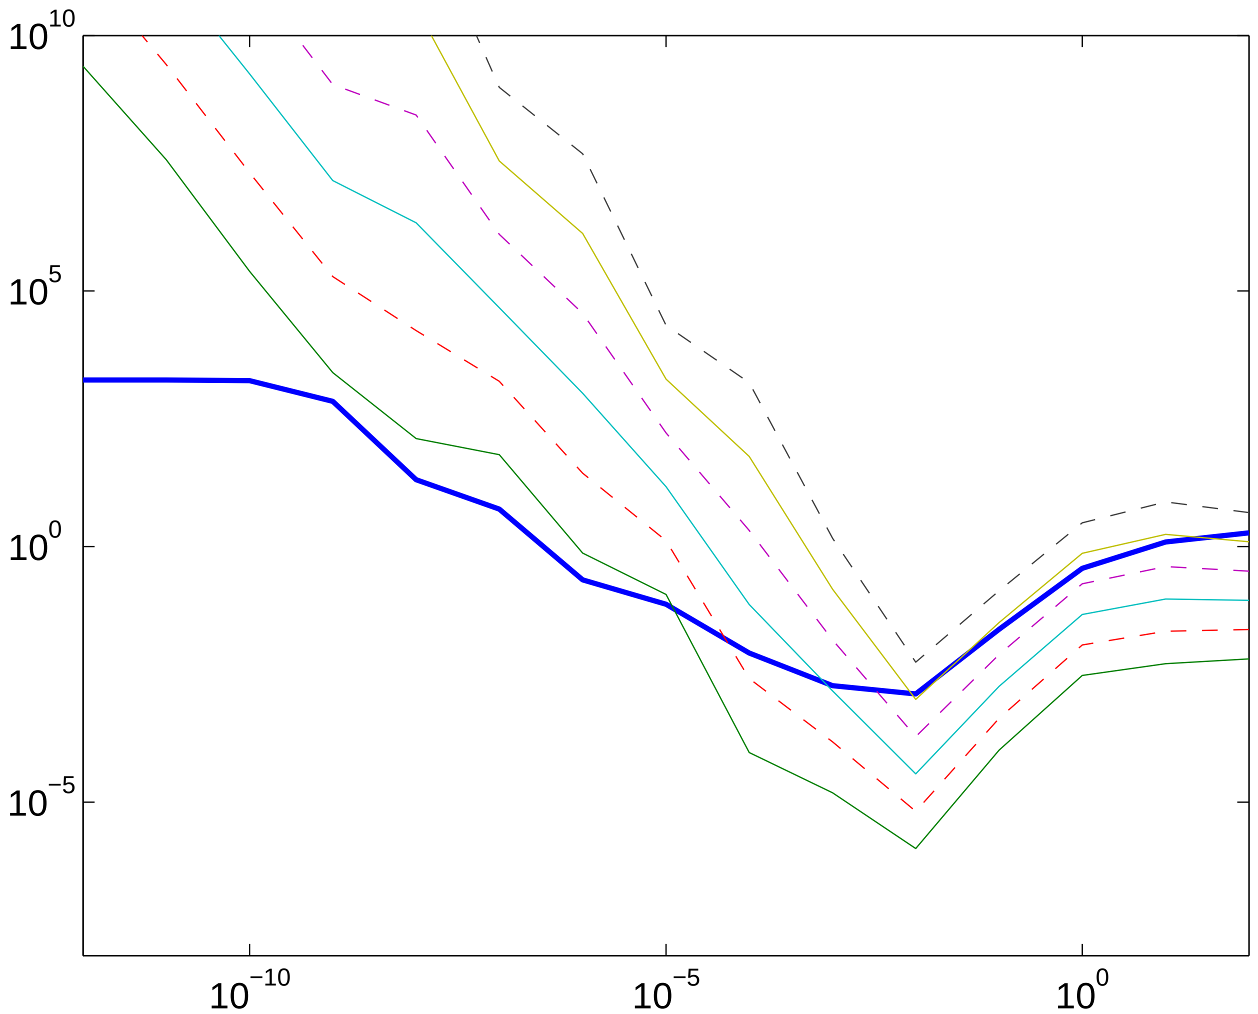

Figure 3: Error and estimates

- Using the SVD for computing x_λ, the minimal value for $\|x - x_\lambda\|$ is equal to $1.1 \cdot 10^{-3}$, and is reached for $\lambda = 6.1 \cdot 10^{-3}$.
- Our estimates furnish $\lambda = 7.8 \cdot 10^{-3}$, and the corresponding error is the optimal one, within the first two significant digits.
- The GCV provides $\lambda = 4.8 \cdot 10^{-4}$ with an error of $2.6 \cdot 10^{-3}$.
- The L–curve is displayed in the next Figure. It does not exhibit a recognizable corner (it is not even L-shaped), but the routine from Hansen's toolbox incorrectly locates a corner at $\lambda = 1.7 \cdot 10^{+9}$, with a corresponding error of 4.1.

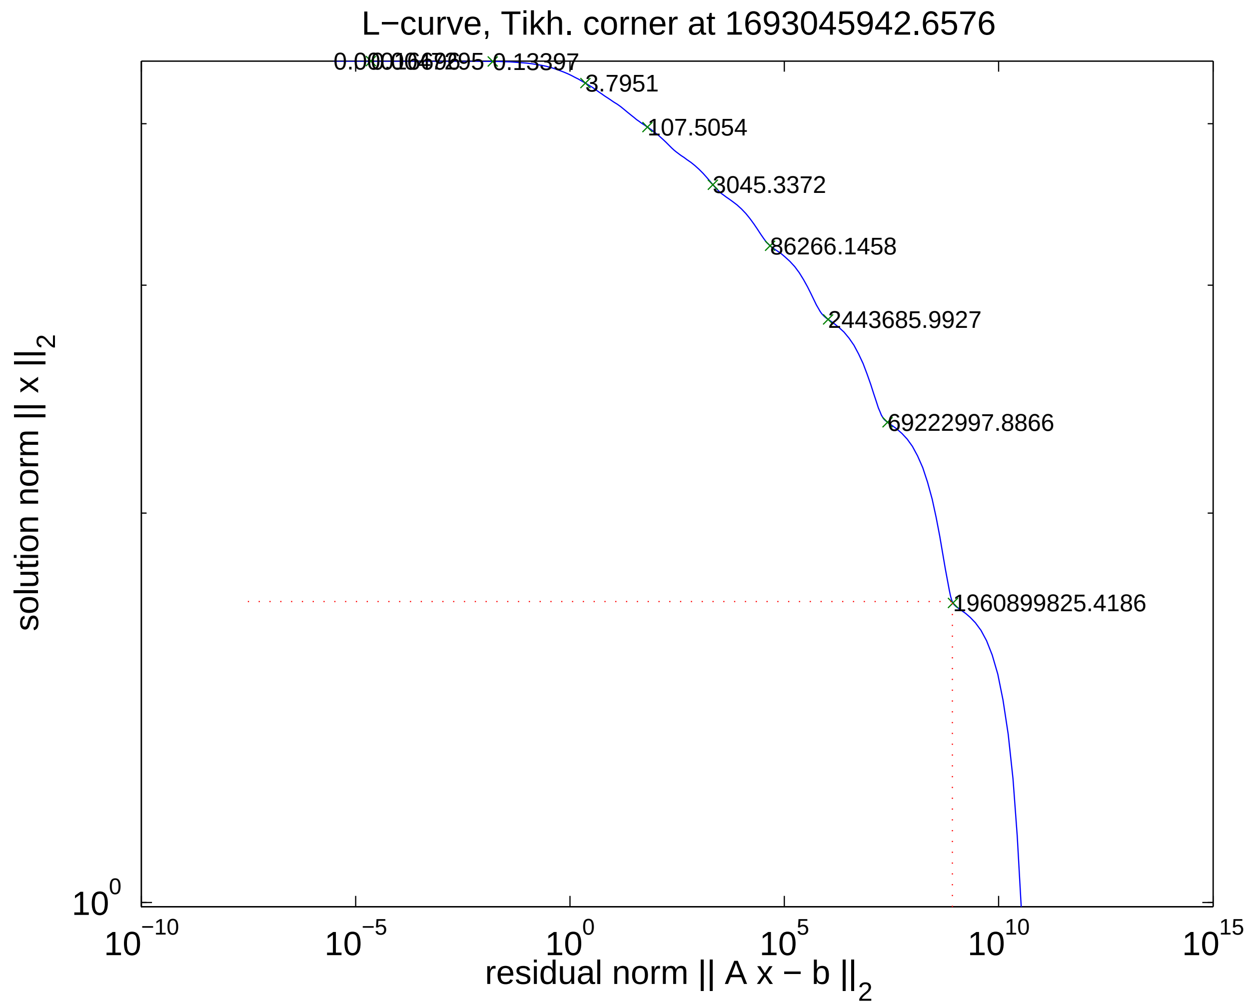

Figure 4: L–curve

EXAMPLE 2:

We consider the *Baart* matrix of dimension 20. The solution was chosen to be $x = (1, \ldots, 1)^T$, the noise level on the right hand side was 10^{-8}, and the regularization matrix was the discrete approximation of the second derivative.

For this example, we plot the ratio $\lambda^2 \|Ex_\lambda\| / \|A^T r_\lambda\|$ with respect to λ. This ratio must be equal to 1 for all λ. The vertical dashed line indicates the value of λ where $\|e_\lambda\|$ reaches its minimum. Thus, this ratio could also be used as a test for the correctness of the computation of x_λ, as mentioned above.

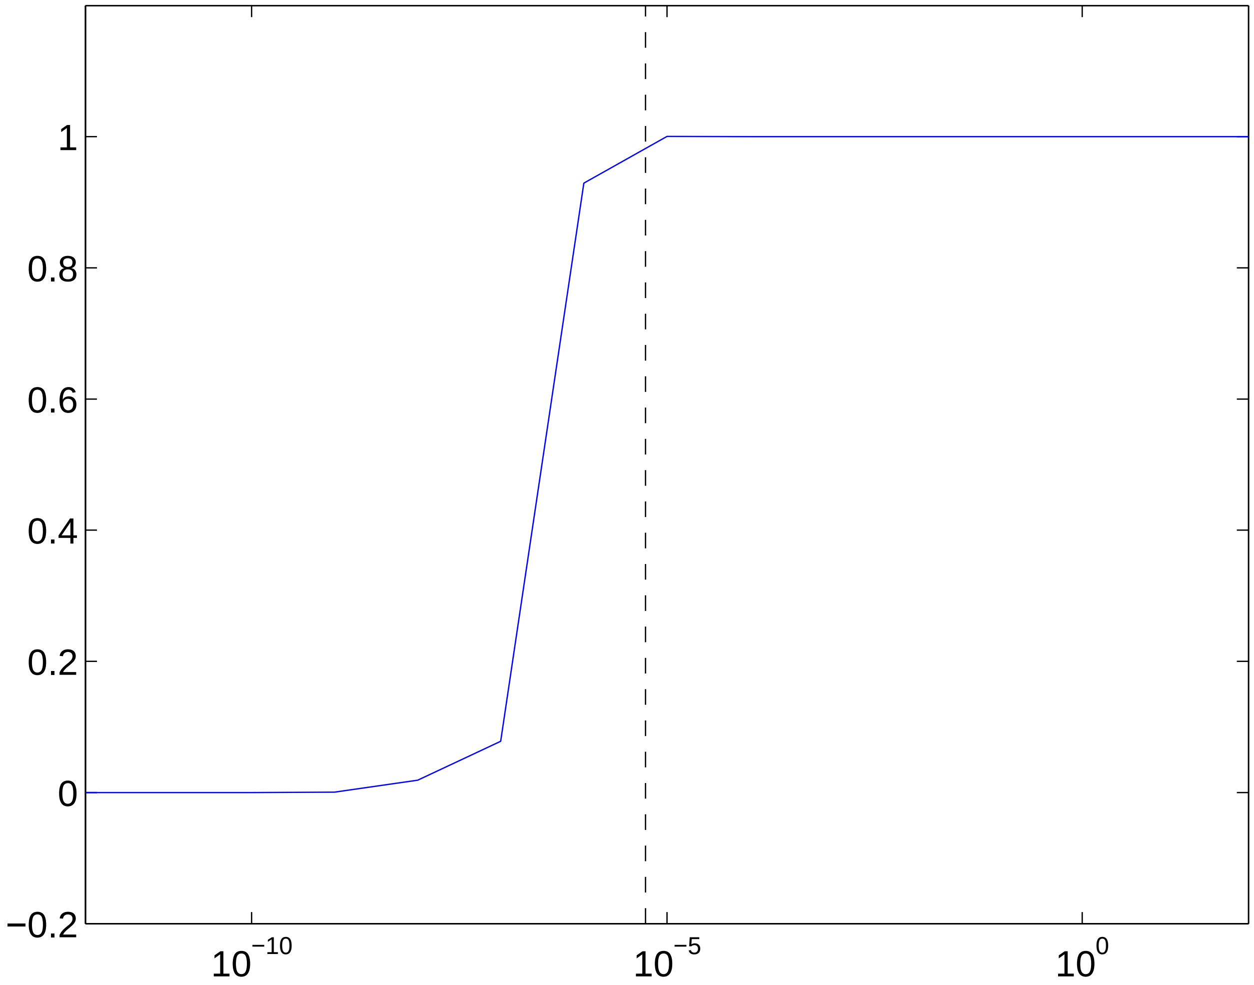

Figure 5: Ratio $\lambda^2\|Ex_\lambda\|/\|A^T r_\lambda\|$ for Baart matrix

EXAMPLE 3:

The conjugate gradient itself has a regularizing effect.

Let us now use our estimates for stopping the iterations of CG.

We take the *Gaussian* matrix of dimension 10000, with a unit solution, and a noise level of 10^{-4}.

Its asymptotic condition number, when the parameter σ is equal to 0.01, is $1.0 \cdot 10^{214}$.

Left: error (thick line), estimates $\widetilde{e}_\nu$ for $\nu = 1, \ldots, 5$, versus the iterations.

Right: (thick plain line), A-norm error, $(e, Ae)^{1/2}$ (thick dashed line), estimates $\widetilde{e}_3$ (thin plain line), $(\widehat{e}_3)^{1/2}$ (thin dashed line).

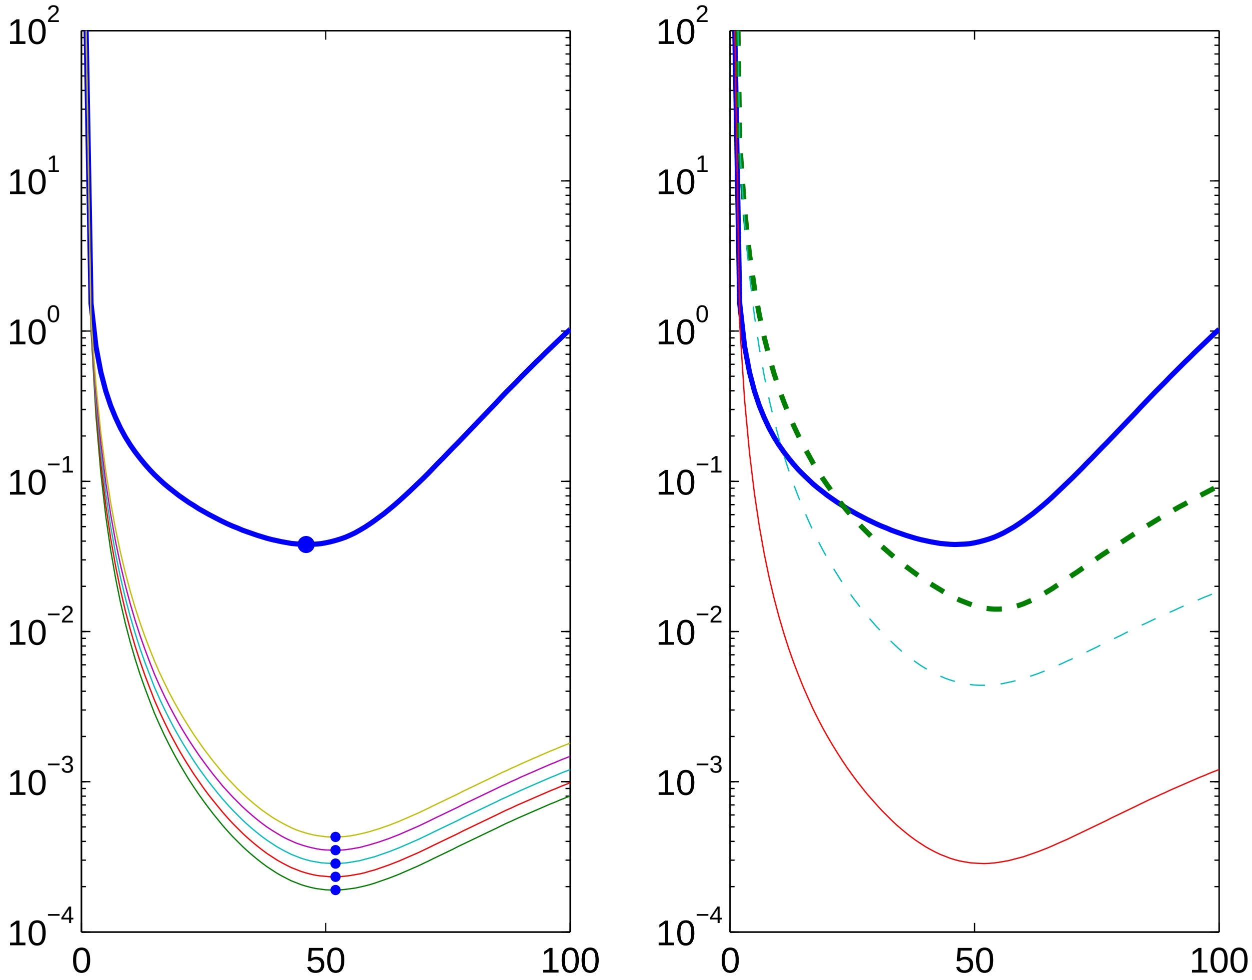

Figure 6: Regularizing CG: Gaussian matrix

Example 4:

Finally, let us solve an image deblurring problem by CG. The size of the image is 256×256, and so dimension of the system is $256^2 = 65536$.

We initially apply a Gaussian blur to a test image, displayed on the left of the Figure (see below), and contaminate it with a noise at level 10^{-4} and 10^{-2}.

The next Figure reports the graph of the error (thick lines) and of e_3 (thin lines) for the two noise levels.

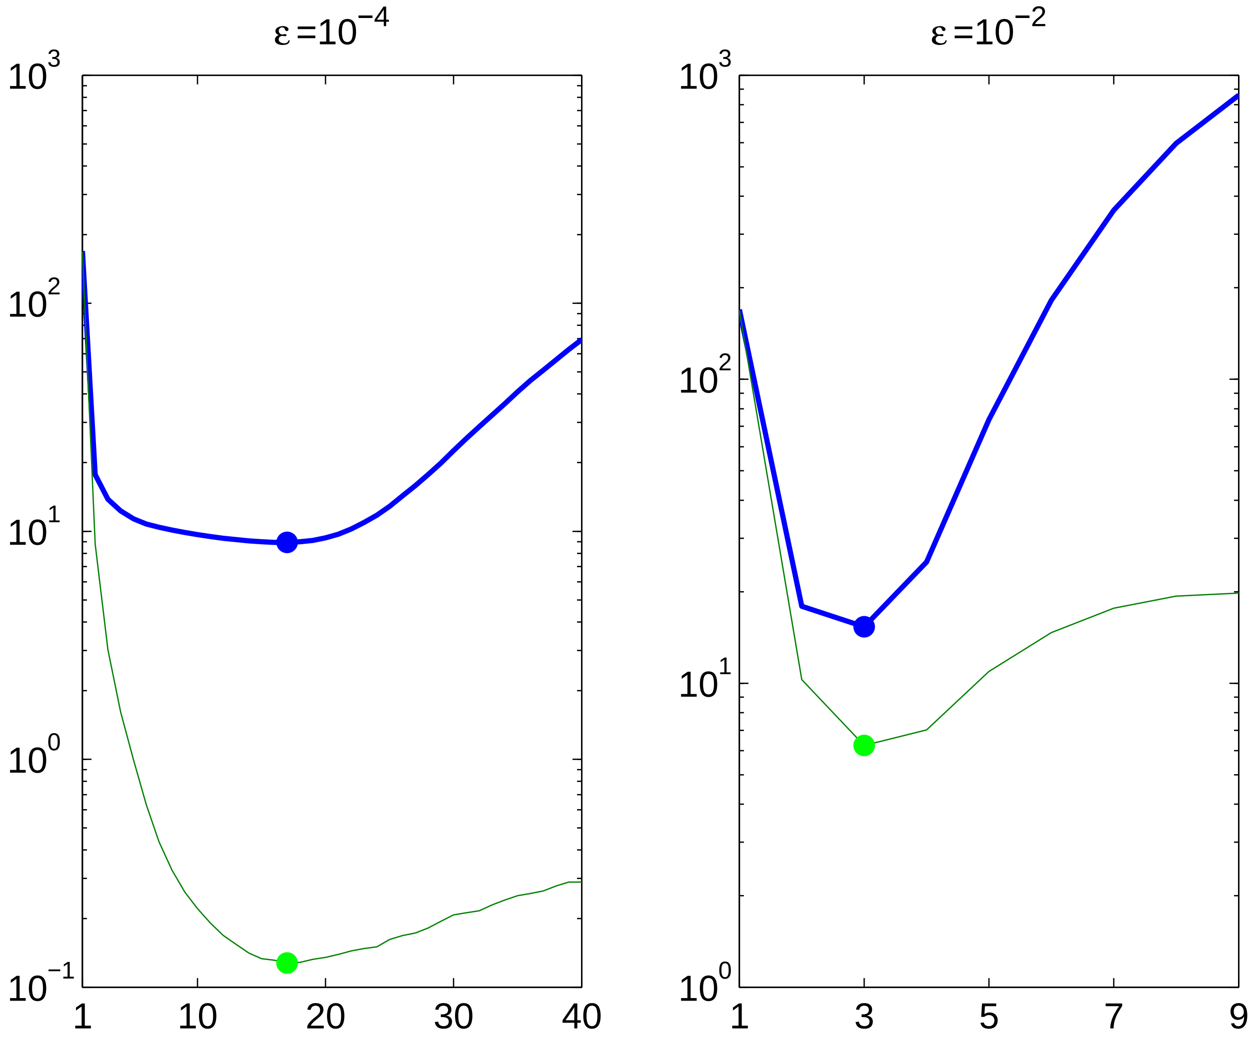

Figure 7: CG: deblurring problem

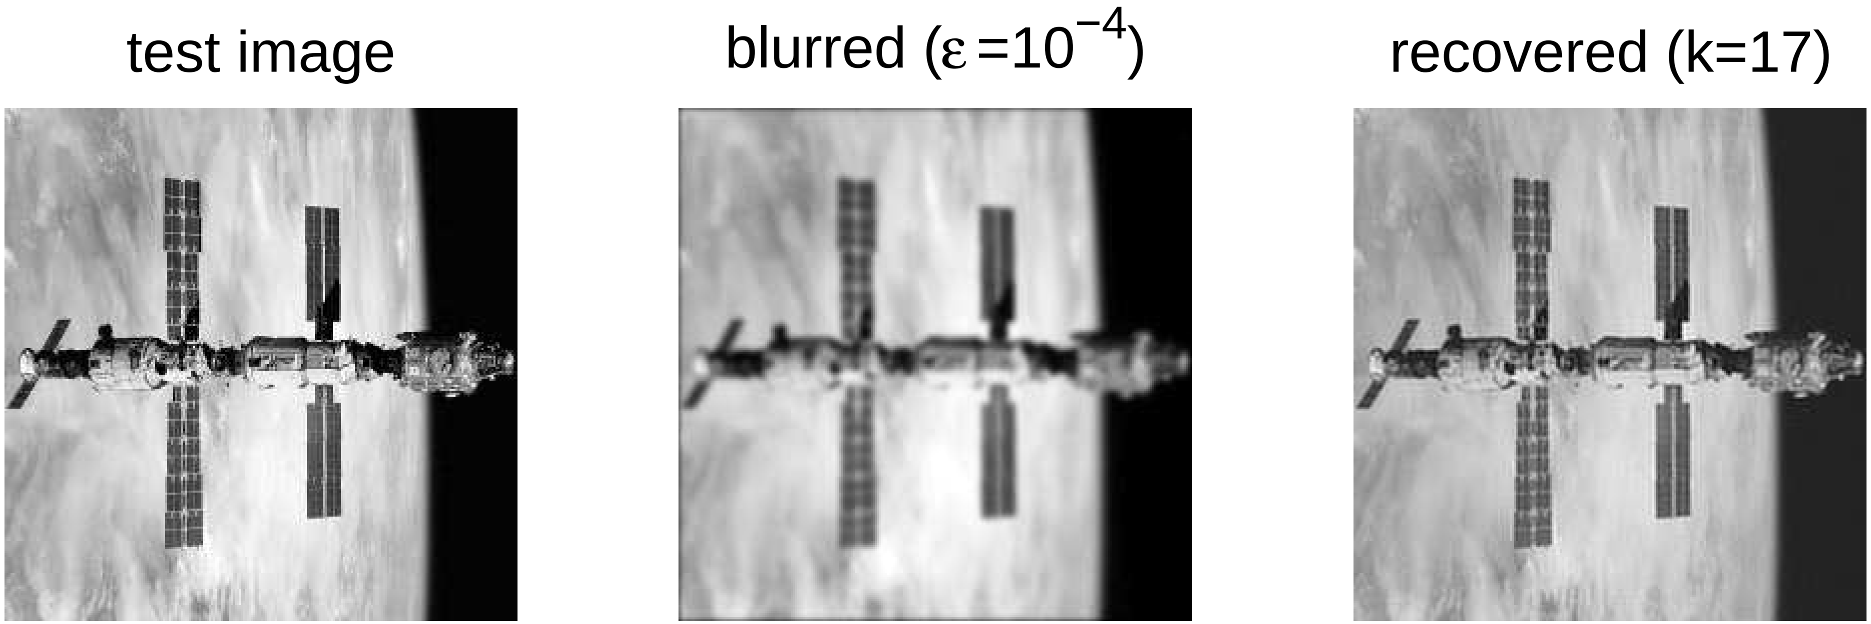

Figure 8: Images for the deblurring problem

RECENT WORKS

- ➜ **Least squares solution of rectangular systems**
  C.B., G. Rodriguez, S. Seatzu
- ➜ **Partial Lanczos bidiagonalization of the matrix**
  L. Reichel, G. Rodriguez, S. Seatzu

SOME REFERENCES

- G. Auchmuty, A posteriori error estimates for linear equations, Numer. Math., 61 (1992) 1–6
- C. Brezinski, Error estimates for the solution of linear systems, SIAM J. Sci. Comput., 21 (1999) 764–781.
- C. Brezinski, G. Rodriguez, S. Seatzu, Error estimates for linear systems with applications to regularization, Numer. Algorithms, 49 (2008), to appear.
- A. Galantai, A study of Auchmuty's error estimate, Comput. Math. with Appl., 42 (2001) 1093–1102.
- G.H. Golub, G. Meurant, *Matrices, Moments and Quadrature with Applications*, to appear.