Insight into Semi-Convergence of Iterative Regularization Methods

Per Christian Hansen

Technical University of Denmark

This work was partially funded by the Villum Foundation via the project

CUQI

With thanks to Jasper Everink, Ken Hayami, and Michiel Hochstenbach

Lin. Alg. Appl., doi 10.1016/j.laa.2025.07.036

Intro to Semi-Convergence

The term "semi-convergence" was coined by Frank Natterer (1986) who writes about an iterative method applied to a noisy inverse problem:

> *"even if it provides a satisfactory solution after a certain number of steps, it deteriorates if the iteration goes on."*

- ▷ Initially the iterates approach the desired exact solution.
- ▷ Eventually the iterates converge to a very noisy and undesired solution.

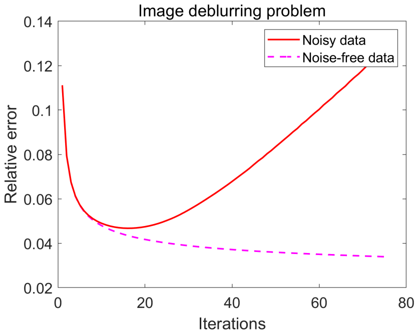

How to Study Semi-Convergence

Split the **reconstruction error** into 2 parts.

- **Iteration error**
  associated with noise-free data;
- **Noise error**
  associated with the data noise.

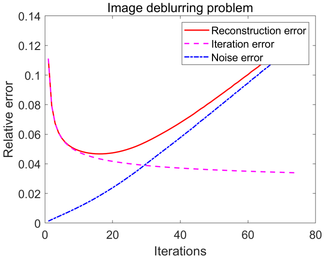

Good understanding of the iteration error.

Re. the noise error:

- we can derive an *upper bound*[1] (sometimes pessimistic);
- no lower bound, to verify that it actually grows.

We take a **statistical approach** and demonstrate that

- the noise error is very likely to increase with the number of iterations;
- hence semi-convergence is very likely to happen.

[1] T. Elfving, *Noise propagation in linear stationary iterations*, Numer. Algo., 2025.

Notation

$$A \in \mathbb{R}^{m\times n} \qquad \text{with} \quad m \geq n \quad \text{and} \quad \text{rank}(A) = n \ .$$

$$b = \bar{b} + e \ , \qquad \bar{b} = A\,\bar{x} \ , \qquad e \sim \mathcal{N}(0, \eta^2 I) \ .$$

Given the SVD

$$A = U\,\Sigma\,V^\top = \sum_{i=1}^{n} u_i\,\sigma_i\,v_i^T \ , \qquad U \in \mathbb{R}^{m\times n} \ , \quad \Sigma, V \in \mathbb{R}^{n\times n} \ ,$$

we can write the least squares solution as

$$x_{\text{LS}} = \sum_{i=1}^{n} \frac{u_i^T b}{\sigma_i}\, v_i = \bar{x} + \sum_{i=1}^{n} \frac{u_i^T e}{\sigma_i}\, v_i \ .$$

For inverse problems, where the σ_i decay towards zero, the second term (the "inverted noise") typically dominates over $\bar{x}$.

SVD Filtering to Suppress Noise

SVD filtering:

$$x_{\text{reg}} = \sum_{i=1}^{n} \phi_i \, \frac{u_i^\top b}{\sigma_i} \, v_i = \sum_{i=1}^{n} \phi_i \, \frac{u_i^\top \bar{b}}{\sigma_i} \, v_i + \sum_{i=1}^{n} \phi_i \, \frac{u_i^\top e}{\sigma_i} \, v_i \ .$$

The filter factors $\phi_1, \ldots, \phi_n$ filter or dampen the SVD components corresponding to small σ_i.

Examples: Tikhonov filters $\phi_i = \sigma_i^2/(\sigma_i^2 + \lambda^2)$, TSVD filters $\phi_i = 1$ or 0.

If $\Phi_{\text{reg}} = \text{diag}(\phi_1, \ldots, \phi_n)$ then

$$x_{\text{reg}} = A_{\text{reg}}^{\#} b \qquad \text{with} \qquad A_{\text{reg}}^{\#} = V \, \Phi_{\text{reg}} \, \Sigma^\dagger U^\top \ .$$

The *regularized inverse* $A_{\text{reg}}^{\#}$ allows us to study how information and noise propagate from the right-hand side to the regularized solution.

Split the Reconstruction Error (lots of notation)

$$x_{\mathrm{reg}} - \bar{x} = \underbrace{x_{\mathrm{reg}} - \bar{x}_{\mathrm{reg}}}_{\text{noise error}} + \underbrace{\bar{x}_{\mathrm{reg}} - \bar{x}}_{\text{reg. error}}$$

$$\bar{x} = \textstyle\sum_{i=1}^{n} \bar{\xi}_i \, v_i \ , \quad \bar{\xi}_i = v_i^{\top} \bar{x}$$

$$e = \textstyle\sum_{i=1}^{m} \varepsilon_i \, v_i \ , \quad \varepsilon_i = u_i^{\top} e$$

$$x_{\mathrm{reg}} = A_{\mathrm{reg}}^{\#} b = \sum_{i=1}^{n} \phi_i \, \frac{u_i^{\top} b}{\sigma_i} \, v_i = \sum_{i=1}^{n} \phi_i \, \bar{\xi}_i \, v_i + \sum_{i=1}^{r} \phi_i \, \frac{\varepsilon_i}{\sigma_i} \, v_i \ ,$$

$$\bar{x}_{\mathrm{reg}} = A_{\mathrm{reg}}^{\#} \bar{b} = \sum_{i=1}^{n} \phi_i \, \frac{u_i^{\top} \bar{b}}{\sigma_i} \, v_i = \sum_{i=1}^{r} \phi_i \, \bar{\xi}_i \, v_i \ .$$

The *regularization error* reveals the influence of the regularization:

$$\bar{x}_{\mathrm{reg}} - \bar{x} = A_{\mathrm{reg}}^{\#} \bar{b} - \bar{x} = \sum_{i=1}^{n} (\phi_i - 1) \, \bar{\xi}_i \, v_i \ .$$

The *noise error* reveals how the noise propagates:

$$x_{\mathrm{reg}} - \bar{x}_{\mathrm{reg}} = A_{\mathrm{reg}}^{\#} \, b - A_{\mathrm{reg}}^{\#} \, \bar{b} = A_{\mathrm{reg}}^{\#} \, e = \sum_{i=1}^{n} \phi_i \, \frac{\varepsilon_i}{\sigma_i} \, v_i \ .$$

Statistical Aspects of the Regularization Error

Using the previous relations, we get

$$\|x_{\text{reg}} - \bar{x}\|_2^2 = \sum_{i=1}^{n} (1-\phi)^2 \bar{\xi}_i^2 + \sum_{i=1}^{n} \frac{\phi_i^2}{\sigma_i^2}\, \varepsilon_i^2 + 2\sum_{i=1}^{n} (\phi_i - 1)\, \phi_i\, \frac{\bar{\xi}_i}{\sigma_i}\, \varepsilon_i$$

and the expected value is

$$\mathcal{E}(\|x_{\text{reg}} - \bar{x}\|_2^2) = \sum_{i=1}^{n} (1-\phi)^2 \bar{\xi}_i^2 + \sum_{i=1}^{n} \frac{\phi_i^2}{\sigma_i^2}\, \mathcal{E}(\varepsilon_i^2) + \xcancel{2\sum_{i=1}^{n} (\phi_i - 1)\, \phi_i\, \frac{\bar{\xi}_i}{\sigma_i}\, \mathcal{E}(\varepsilon_i)}$$

1$^{\text{st}}$ term: squared norm of the *regularization error* $\bar{x}_{\text{reg}} - \bar{x}$ caused by applying regularization to the noise-free data $\bar{b}$.

2$^{\text{nd}}$ term: expected value of the squared norm of the *noise error* $x_{\text{reg}} - \bar{x}_{\text{reg}}$.

3$^{\text{rd}}$ term: is zero because $\mathcal{E}(\varepsilon_i) = 0$.

Continuing From the Previous Relations

Expected value of reconstruction error:

$$\mathcal{E}(\|x_{\text{reg}} - \bar{x}\|_2^2) = \|\bar{x}_{\text{reg}} - \bar{x}\|_2^2 + \mathcal{E}(\|x_{\text{reg}} - \bar{x}_{\text{reg}}\|_2^2) \ ,$$

where

$$\|\bar{x}_{\text{reg}} - \bar{x}\|_2^2 = \sum_{i=1}^{n} (1 - \phi_i)^2 \, \bar{\xi}_i^2 \ ,$$

$$\mathcal{E}(\|x_{\text{reg}} - \bar{x}_{\text{reg}}\|_2^2) = \mathcal{E}(\|A_{\text{reg}}^{\#} e\|_2^2) = \eta^2 \sum_{i=1}^{n} \frac{\phi_i^2}{\sigma_i^2} \ .$$

Variance of ditto:

$$\mathcal{V}(\|x_{\text{reg}} - \bar{x}\|_2^2) = \mathcal{V}(\|x_{\text{reg}} - \bar{x}_{\text{reg}}\|_2^2)$$

$$= \sum \left(\frac{\phi_i^2}{\sigma_i^2} \right)^2 \mathcal{V}(\varepsilon_i^2) = \sum \frac{\phi_i^4}{\sigma_i^4} \, 2 \, \mathcal{V}(\varepsilon_i)^2 = 2 \, \eta^4 \sum \frac{\phi_i^4}{\sigma_i^4} \ .$$

TSVD Example with Test Problem gravity

Note:
log axis →

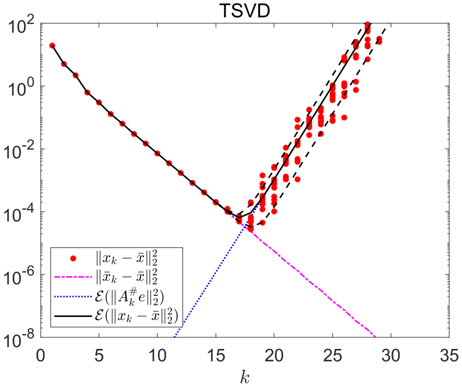

The dashed lines illustrate the standard deviation $\pm\mathcal{V}^{1/2}$.

The TSVD regularization error dominates for small k while the noise error dominates for large k where there is little filtering.

Test problems: see appendix.

Prelude to Regularizing Iterations

To set the notation, we write the noisy and the noise-free iterates as

$$x_k = A_k^{\#} b \,, \qquad \bar{x}_k = A_k^{\#} \bar{b} \,.$$

- The regularized inverse $A_k^{\#}$ is defined by the iterative method,
- $\bar{x}_k$ are the iterates that we compute if there were no noise.

We split the reconstruction error as:

$$x_k - \bar{x} = \underbrace{x_k - \bar{x}_k}_{\text{noise error}} + \underbrace{\bar{x}_k - \bar{x}}_{\text{it. error}}$$

and we refer to $\bar{x}_k - \bar{x}$ as the *iteration error*.

In classical convergence analysis we analyze the decay of $\|\bar{x}_k - \bar{x}\|_2$.

For regularizing iterations we study the growth of the noise error $x_k - \bar{x}_k = A_k^{\#} e$ as a function of k.

Illustration: Landweber's Method

$$x_{k,\mathrm{L}} = x_{k-1,\mathrm{L}} + \omega\, A^T (b - A\, x_{k-1,\mathrm{L}}) \; , \qquad k = 1, 2, \ldots$$

Example: `heat` with large and small noise levels.

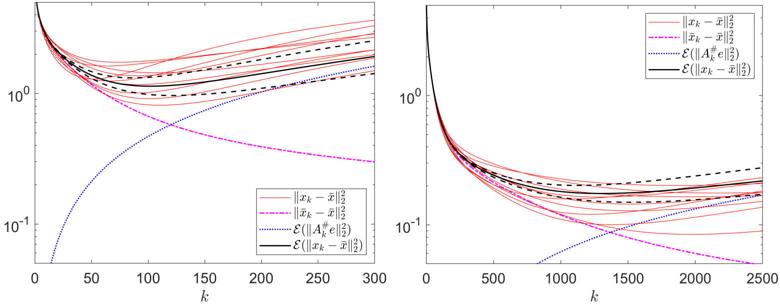

The expected value $\mathcal{E}(\cdot)$ — and the standard deviation $\mathcal{V}(\cdot)^{1/2}$ -- are computed with the expressions from slide 8.

CGLS Regularizing Iterations

For CGLS (and other Krylov subspace methods) the filter factors depend on both A and b. The kth CGLS iterates can be written as

$$x_k = \arg\min_x \|A\,x - b\|_2 \qquad \text{s.t.} \qquad x \in \mathcal{K}_k(A^\top A, A^\top \overset{\downarrow}{b})$$

with the Krylov subspace

$$\mathcal{K}_k(A^\top A, A^\top b) = \text{span}\{A^\top b, A^\top A\,A^\top b, \ldots, (A^\top A)^{k-1} A^\top b\} \ .$$

The CGLS filter factors are given by

$$\phi_i^{(k)} = 1 - \prod_{j=1}^{k} \frac{\theta_j^{(k)} - \sigma_i^2}{\theta_j^{(k)}} \ , \qquad i = 1, \ldots, n$$

where the Ritz values $\theta_j^{(k)}$ are the eigenvalues of the $k \times k$ symmetric tridiagonal matrix associated with applying CG to $A^\top A\,x = A^\top b$.

(Also, they are the squares of the singular values of the bidiagonal matrix generated by the Golub-Kahan algorithm underlying the LSQR method.)

The Important Role of the Ritz Values (ex. gravity)

Consider the residual norm (assuming that $A\,x = b$ is consistent)

$$\|r_k\|_2^2 = \sum_{i=1}^{n} \bigl(\mathcal{R}_k(\sigma_i^2)\ \beta_i\bigr)^2\ , \qquad \mathcal{R}_k(\theta) = \prod_{j=1}^{k} \frac{\theta_j^{(k)} - \theta}{\theta_j^{(k)}} = \text{Ritz polynomial.}$$

CGLS places the roots of $\mathcal{R}_k$ such that $\mathcal{R}_k(\sigma_i^2)$ is small where β_i^2 is large.

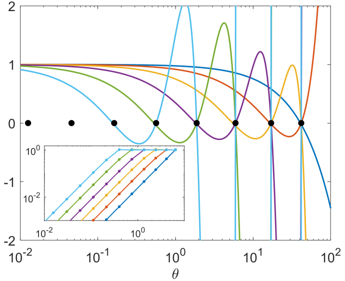

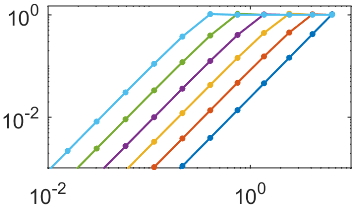

Can **not** guarantee monotonicity.

Plot shows Ritz pol. $\mathcal{R}_k(\theta)$ for $k = 1, \ldots, 6$ and the largest σ_i^2 (black dots).

In general, the filter factors satsify $\phi_i^{(k)} = 1 - \mathcal{R}_k(\sigma_i^2) \approx 1$ for $i = 1, \ldots, k$.

Two Different CGLS Iterates x_k and $\bar{x}_k$

Write the kth CGLS iterates $x_k \in \mathcal{K}_k$ as

$$x_k = A_k^{\#} b \ , \qquad A_k^{\#} = \big(I - \mathcal{R}_k(A^\top A)\big) A^\top \ .$$

The *noise-free CGLS iterates* $\bar{x}_k$ correspond to the noise-free data $\bar{b}$.
Hence, $\bar{x}_k$ lies in a *different* Krylov subspace associated with $\bar{b}$:

$$\overline{\mathcal{K}}_k(A^\top A, A^\top \overset{\downarrow}{\bar{b}}) = \mathrm{span}\{A^\top \bar{b}, A^\top A\, A^\top \bar{b}, \ldots, (A^\top A)^{k-1} A^\top \bar{b}\} \ .$$

Hence, we write $\bar{x}_k \in \overline{\mathcal{K}}_k$ as

$$\bar{x}_k = \bar{A}_k^{\#} \bar{b} \ , \qquad \bar{A}_k^{\#} = \big(I - \overline{\mathcal{R}}_k(A^\top A)\big) A^\top \ ,$$

where $\overline{\mathcal{R}}_k$ is the Ritz polynomial associated with $\overline{\mathcal{K}}_k$.

And Now: Split the CGLS Reconstruction Error

Now we can express the reconstruction error as follows:

$$\begin{aligned} x_k - \bar{x} &= A_k^{\#} b - \bar{x} = A_k^{\#}(\bar{b} + e) - \bar{x} + (\bar{x}_k - \bar{x}_k) \\ &= A_k^{\#}\bar{b} + A_k^{\#} e + -\bar{x} + \bar{x}_k - \bar{A}_k^{\#}\bar{b} \\ &= \underbrace{(A_k^{\#} - \bar{A}_k^{\#})\bar{b} + A_k^{\#} e}_{\text{noise error}} + \underbrace{\bar{x}_k - \bar{x}}_{\text{it. error}} \ . \end{aligned}$$

The iteration error involves the noise-free iterations lying in $\overline{\mathcal{K}}_k$.

Novel Insight. The noise error for CGLS consists of **two** components:

- the propagated noise $A_k^{\#} e$ and
- the *deviation* $(A_k^{\#} - \bar{A}_k^{\#})\bar{b}$ caused by the difference between the two Krylov subspaces $\mathcal{K}_k(A^\top A, A^\top b)$ and $\overline{\mathcal{K}}_k(A^\top A, A^\top \bar{b})$.

The latter component is unique to Krylov subspace methods, incl. CGLS.

Ex: gravity ($\eta = 10^{-4}$) and paralleltomo ($\eta = 0.1$)

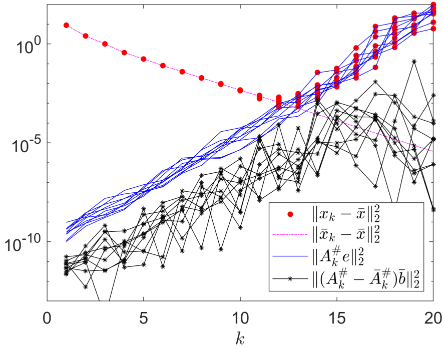

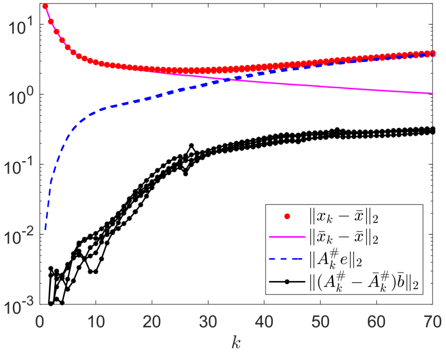

In these examples, the norm of the deviation $(A_k^{\#} - \bar{A}_k^{\#})\bar{b}$ is smaller than the norm of the propagated noise $A_k^{\#} e$.

This is not always true – see the paper.

Statistics of the Noise Error: Propagated Noise

For the propagated noise $A_k^{\#} e$, we have

$$\mathcal{E}(\|A_k^{\#} e\|_2^2) = \sum_{i=1}^{n} \mathcal{E}\left(\left(\phi_i^{(k)} \varepsilon_i\right)^2 \right) \frac{1}{\sigma_i^2} \ , \qquad \phi_i^{(k)} = 1 - \mathcal{R}_k(\sigma_i^2)\sigma_i^2$$

and $\phi_i^{(k)}$ are correlated with the noise via the Ritz polynomial $\mathcal{R}_k$ that depends on the "noisy" Krylov subspace $\mathcal{K}_k$.

Numerical experiments show that the correlation between ε_i and $\phi_i^{(k)}$ is very small, and hence we use the approximation

$$\mathcal{E}(\|A_k^{\#} e\|_2^2) \approx \sum_{i=1}^{n} \mathcal{E}\left(\left(\phi_i^{(k)}\right)^2 \right) \mathcal{E}\left(\varepsilon_i^2\right) \frac{1}{\sigma_i^2} = \eta^2 \sum_{i=1}^{n} \frac{\mathcal{E}\left(\left(\phi_i^{(k)}\right)^2 \right)}{\sigma_i^2} \ .$$

Cliffhanger: what to do about $\mathcal{E}\left(\left(\phi_i^{(k)}\right)^2 \right)$? $\rightarrow$ Next slide.

A Closer Look at the CGLS Filter Factors

Test problem `gravity`:

- violin plots of $\mathcal{E}\Big(\big(\phi_i^{(k)}\big)^2\Big)$, together with
- the noise-free $\big(\bar{\phi}_i^{(k)}\big)^2$
- — the sample mean of $\mathcal{E}\Big(\big(\phi_i^{(k)}\big)^2\Big)$.

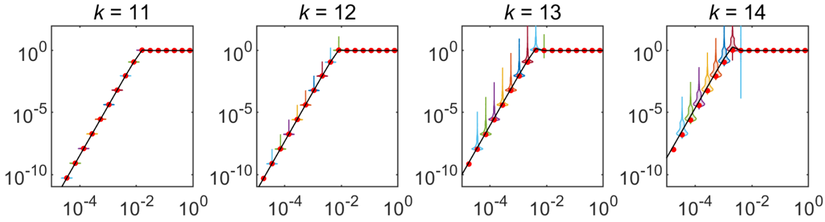

This motivates the approximation for the propagated noise:

$$\mathcal{E}(\|A_k^{\#} e\|_2^2) \approx \eta^2 \sum_{i=1}^{n} \frac{\big(\bar{\phi}_i^{(k)}\big)^2}{\sigma_i^2} \ .$$

Statistics of the Noise Error: The Deviation

For the deviation $(A_k^{\#} - \bar{A}_k^{\#})\bar{b}$, recall that the filter factors $\bar{\Phi}_k$ in $\bar{A}_k^{\#} = V\bar{\Phi}_k\Sigma^{-1}U^\top$ are different from those Φ_k of $A_k^{\#} = V\Phi_k\Sigma^{-1}U^\top$.

$$\mathcal{E}\left(\|(A_k^{\#} - \bar{A}_k^{\#})\bar{b}\|_2^2\right) = \sum_{i=1}^{n} \mathcal{E}\Big(\big(\phi_i^{(k)} - \bar{\phi}_i^{(k)}\big)^2\Big)\,\frac{\bar{\beta}_i^2}{\sigma_i^2} \,.$$

Approximation $\phi_i^{(k)} \approx \bar{\phi}_k^{(k)}$ is no good here → numer. exp. (`gravity`):

- violin plots of $\big(\phi_i^{(k)} - \bar{\phi}_i^{(k)}\big)^2\,\frac{\bar{\beta}_i^2}{\sigma_i^2}$

— the sample mean of its expected value.

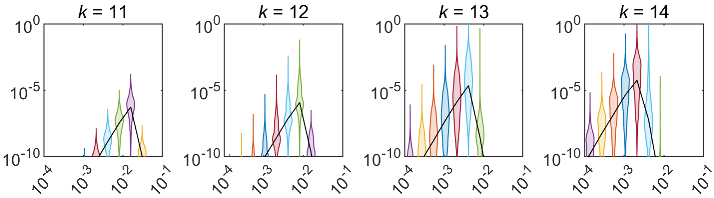

We can use the sample mean (it is quite small).

Putting It Together

Fat solid **brown** lines: expected values.

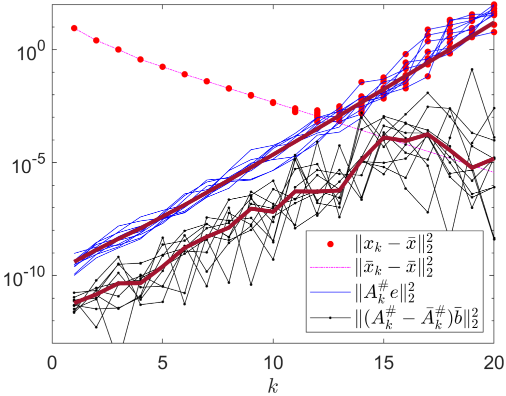

While the deviation $(A_k^{\#} - \bar{A}_k^{\#})\bar{b}$ is part of the CGLS noise error, we can often ignore it and consider only the iteration error and the propagated noise $A_k^{\#} e$. We cannot guarantee that $\|A_k^{\#} e\|_2^2$ is monotonic.

Conclusions

New insight from a (simple) statistical analysis:

- The noise error is very likely to grow for regularizing iterations.
- Semi-convergence is therefore very likely to occur.

For CGLS, we introduce a novel splitting of the noise error into a propagated noise term $A_k^{\#} e$ and deviation term $(A_k^{\#} - \bar{A}_k^{\#})\bar{b}$, the latter typically being small.

Next steps: GMRES and Kaczmarz.

Appendix: Our Test Problems

We use these test problems from Regularization Tools:

- `blur` – Gaussian image deblurring with an $N \times N$ image. The singular values have a very slow exponential decay; the cond. number is 31.
- `gravity` – 1D gravity surveying problem; the matrix is 128×128. The singular values decay approximately as $e^{-0.7i}$.
- `heat` – Inverse heat equation problem; the matrix is 128×128. The singular values decay exponentially from 0.3 to 10^{-6}.
- `phillips` – Test problem with no origin in applications; the matrix is 128×128. The singular values decay approximately as i^{-3}.

We also use the X-ray tomography test problem `paralleltomotomo` from from AIR Tools II with a 64×64 phantom, 64 detector pixels, and view angles $2.5^\circ, 5^\circ, \ldots, 180^\circ$. The leading singular values decay as $i^{-1/4}$ while the trailing ones decay faster; the condition number is 2 392.

Appendix: When the Singular Values Decay Slowly

For a slow decay of the singular values (e.g., for mildly ill-posed problems) CGLS does not necessarily make $\mathcal{R}_k(\sigma_i^2) \approx 0$ at the k largest β_i^2.

Instead, it will place the k roots such that the polynomial is small for many more than k pairs of (σ_i^2, β_i^2). This dampens the contributions to $\|r_k\|_2$ for many SVD components over a large interval.

Example `blur`: for $k = 6$ the first 80 filter factors are between 0.8 and 1.2 meaning that we capture about 80 components in the 6th iteration vector.