Inverse Problems

Sommersemester 2023

Vesa Kaarnioja
vesa.kaarnioja@fu-berlin.de

FU Berlin, FB Mathematik und Informatik

Fourth lecture, May 8, 2023

Fredholm equation and its solvability

Separable Hilbert space

A Hilbert space is said to be *separable* if (and only if) there exists a *countable orthonormal basis* $\{\psi_j\}_{j=1}^{\infty}$ of H with respect to the inner product $\langle \cdot, \cdot \rangle_H$, that is,

$$\langle \psi_j, \psi_k \rangle_H = \delta_{j,k} \quad \text{and} \quad \left\| f - \sum_{j=1}^{\ell} \langle f, \psi_j \rangle_H \psi_j \right\|_H \xrightarrow{\ell \to \infty} 0 \quad \text{for all } f \in H.$$

This last condition is usually written as

$$f = \sum_{j=1}^{\infty} \langle f, \psi_j \rangle_H \psi_j.$$

Note that $\sum_{j=1}^{\ell} \langle f, \psi_j \rangle_H \psi_j$ is precisely the orthogonal projection onto the subspace spanned by $\psi_1, \ldots, \psi_\ell$.

Fredholm equation

Let us formalize the problem that we will concentrate on during the first part of the course.

Let H_1 and H_2 be separable real Hilbert spaces and let $A\colon H_1 \to H_2$ be a *compact* linear operator. We are interested in finding $x \in H_1$ such that

$$y = Ax,$$

where $y \in H_2$ is given. Recall that compact operators are the closure of finite-dimensional operators (loosely speaking: matrices) in the operator topology.

Examples:

- $H_1 = H_2 = L^2(a, b)$.
- $H_1 = \mathbb{R}^n$, $H_2 = \mathbb{R}^m$, and $A \in \mathbb{R}^{m \times n}$.

Singular value decomposition of a compact operator

Let us assume that H_1 and H_2 are separable real Hilbert spaces and let $A\colon H_1 \to H_2$ be a compact linear operator.

Then there exist (possibly countably infinite) orthonormal sets of vectors $\{v_n\} \subset H_1$ and $\{u_n\} \subset H_2$, and a sequence of positive numbers $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq \cdots \geq 0$ with $\lim_{n\to\infty} \lambda_n = 0$ in the countably infinite case such that

$$Ax = \sum_n \lambda_n \langle x, v_n \rangle u_n \quad \text{for all } x \in H_1. \tag{1}$$

In particular, since H_1 and H_2 are separable, we have

$$\overline{\mathrm{Ran}(A)} = \overline{\mathrm{span}\{u_n\}} \quad \text{and} \quad (\mathrm{Ker}(A))^{\perp} = \overline{\mathrm{span}\{v_n\}}.$$

The system (λ_n, v_n, u_n) is called a *singular system* of A, and (1) is a *singular value decomposition* (SVD) of A.

Singular value decomposition of matrices: $H_1 = \mathbb{R}^n$ and $H_2 = \mathbb{R}^m$

Let $H_1 = \mathbb{R}^n$ and $H_2 = \mathbb{R}^m$, meaning that

$$y = Ax$$

is a matrix equation with $A \in \mathbb{R}^{m \times n}$, $x \in \mathbb{R}^n$, and $y \in \mathbb{R}^m$.

Since this operator has finite rank ($\mathrm{rank}(A) := \dim \mathrm{Ran}(A) < \infty$), we have

$$Ax = \sum_{j=1}^{p} \lambda_j (x^{\mathrm{T}} v_j) u_j, \quad p := \mathrm{rank}(A) \leq \min\{n, m\},$$

where $\{v_j\}_{j=1}^p \subset \mathbb{R}^n$ and $\{u_j\}_{j=1}^p \subset \mathbb{R}^m$ are sets of orthonormal vectors and $\{\lambda_j\}_{j=1}^p$ are positive numbers such that $\lambda_j \geq \lambda_{j+1}$.

It is possible to complete the sequences of (orthonormal) singular vectors $\{v_j\}_{j=1}^{p} \subset \mathbb{R}^n$ and $\{u_j\}_{j=1}^{p} \subset \mathbb{R}^m$ with complementary orthonormal vectors $\{v_j\}_{j=p+1}^{n}$ and $\{u_j\}_{j=p+1}^{m}$ such that $\{v_j\}_{j=1}^{n}$ forms an orthonormal basis for $\mathbb{R}^n$ and $\{u_j\}_{j=1}^{m}$ forms an orthonormal basis for $\mathbb{R}^m$. This can be done, e.g., using the Gram–Schmidt process.

Define the matrices

$$V = [v_1, \ldots, v_n] \in \mathbb{R}^{n \times n},$$
$$U = [u_1, \ldots, u_m] \in \mathbb{R}^{m \times m}.$$

Due to the orthonormality of $\{v_j\}$ and $\{u_j\}$, the matrices V and U are orthogonal:

$$V^{\mathrm{T}} V = V V^{\mathrm{T}} = I \quad \text{and} \quad U^{\mathrm{T}} U = U U^{\mathrm{T}} = I.$$

Next, we define the matrix $\Lambda \in \mathbb{R}^{m \times n}$ as follows:

$$\Lambda = \left(\begin{array}{ccc|c} \lambda_1 & & & \\ & \ddots & & O_{m \times (n-m)} \\ & & \lambda_m & \end{array} \right) \quad \text{if } m < n,$$

$$\Lambda = \left(\begin{array}{ccc} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \\ \hline & O_{(m-n) \times n} & \end{array} \right) \quad \text{if } m > n,$$

and $\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_m)$ if $m = n$.

It is simple to check that

$$Ax = \sum_{j=1}^{p} \lambda_j u_j v_j^{\mathrm{T}} x = U \Lambda V^{\mathrm{T}} x \quad \text{for all } x \in \mathbb{R}^n,$$

which yields the matrix singular value decomposition (SVD)

$$A = U \Lambda V^{\mathrm{T}}.$$

In Python: `numpy.linalg.svd`. *In MATLAB:* `svd`.

Note that in the matrix SVD, the singular values $\{\lambda_j\}_{j=1}^{\min\{m,n\}}$ are *non-negative* and

$$\begin{aligned}
&\mathrm{Ran}(A) = \mathrm{span}\{u_j \mid 1 \le j \le p\},\\
&\mathrm{Ker}(A) = \mathrm{span}\{v_j \mid p+1 \le j \le n\},\\
&(\mathrm{Ran}(A))^{\perp} = \mathrm{span}\{u_j \mid p+1 \le j \le m\},\\
&(\mathrm{Ker}(A))^{\perp} = \mathrm{span}\{v_j \mid 1 \le j \le p\},
\end{aligned}$$

where $p = \mathrm{rank}(A) = \max_{1\le k\le \min\{m,n\}}\{k \mid \lambda_k > 0\}$.

Solvability of $y = Ax$

Let us assume that H_1 and H_2 are separable real Hilbert spaces and let $A\colon H_1 \to H_2$ be a compact linear operator. Let $P\colon H_2 \to \overline{\mathrm{Ran}(A)}$ be an orthogonal projection. This can be represented using the singular system of A as

$$Py = \sum_n \langle y, u_n \rangle u_n.$$

Theorem

Let $A\colon H_1 \to H_2$ be a compact operator with the singular system (λ_n, v_n, u_n). The equation $y = Ax$ has a solution iff

$$y = Py \quad \text{and} \quad \underbrace{\sum_n \frac{1}{\lambda_n^2} |\langle y, u_n \rangle|^2 < \infty}_{\text{“Picard criterion”}}.$$

In this case, the solution is of the form

$$x = x_0 + \sum_n \frac{1}{\lambda_n} \langle y, u_n \rangle v_n \quad \text{for arbitrary } x_0 \in \mathrm{Ker}(A).$$

Proof. "⇒" Suppose that $y = Ax$ has a solution $x \in H_1$. This implies that $y \in \mathrm{Ran}(A)$ (thus $y = Py$) and, moreover,

$$\langle y, u_j \rangle = \langle Ax, u_j \rangle = \langle x, A^* u_j \rangle = \lambda_j \langle x, v_j \rangle$$

$$\Rightarrow \sum_n \frac{1}{\lambda_n^2} |\langle y, u_n \rangle|^2 = \sum_n |\langle x, v_n \rangle|^2 \overset{\text{Bessel inequ.}}{\leq} \|x\|^2 < \infty.$$

"⇐" Next, suppose that $y = Py$ and the Picard criterion hold and define $x := x_0 + \sum_n \lambda_n^{-1} \langle y, u_n \rangle v_n$, where $x_0 \in \mathrm{Ker}(A)$ is arbitrary. We obtain

$$Ax = Ax_0 + \sum_n \frac{1}{\lambda_n} \langle y, u_n \rangle A v_n = \sum_n \langle y, u_n \rangle u_n = Py = y. \quad \square$$

Remark. In the above proof, it is helpful to note that if A has the SVD

$$Ax = \sum_n \lambda_n \langle x, v_n \rangle u_n,$$

then its adjoint A^* has the SVD

$$A^* y = \sum_n \lambda_n \langle y, u_n \rangle v_n.$$

Note that for any $x \in H_1$, we have

$$\|Ax - y\|^2 = \|Ax - Py\|^2 + \|(I - P)y\|^2 \geq \|(I - P)y\|^2.$$

Hence, if y has a nonzero component in the subspace orthogonal to the range of A (which may happen if y is contaminated by noise), the equation $Ax = y$ cannot be satisfied exactly. Thus, the best we can do is to solve the projected equation

$$Ax = PAx = Py.$$

However, there is in general no guarantee that the Picard criterion

$$\sum_n \frac{1}{\lambda_n^2} |\langle Py, u_n \rangle|^2 < \infty$$

is satisfied for a general $y \in H_2$ if $\operatorname{rank}(A) = \dim \operatorname{Ran}(A) = \infty$.

Truncated singular value decomposition (TSVD)

To recap: the best we can do is to solve the projected equation

$$Ax = Py.$$

However, the solution exists iff the very restrictive Picard criterion holds.

We begin by considering one of the simplest regularization techniques for linear inverse problems. By restricting the range of P onto a finite-dimensional subspace of the range of A, we obtain a well-defined approximation to the above problem.

Truncated singular value decomposition (TSVD)

Let us define a family of finite-dimensional orthogonal projections by

$$P_k : H_2 \to \mathrm{span}\{u_1, \ldots, u_k\}, \quad y \mapsto \sum_{n=1}^{k} \langle y, u_n \rangle u_n.$$

By the orthogonality of $\{u_n\}$,

$$P(P_k y) = \sum_{n} \langle P_k y, u_n \rangle u_n = \sum_{n=1}^{k} \langle y, u_n \rangle u_n = P_k y$$

and

$$\sum_{n} \frac{1}{\lambda_n^2} |\langle P_k y, u_n \rangle|^2 = \sum_{n=1}^{k} \frac{1}{\lambda_n^2} |\langle y, u_n \rangle|^2 < \infty.$$

Note that $k \leq \mathrm{rank}(A)$ if $\mathrm{rank}(A) < \infty$.

It follows that the problem

$$Ax = P_k y \tag{2}$$

is always solvable. Taking on both sides the inner product with u_n, we find that

$$\lambda_n \langle x, v_n \rangle = \begin{cases} \langle y, u_n \rangle, & 1 \leq n \leq k \\ 0, & n > k. \end{cases}$$

Hence the solutions to (2) are given by

$$x_k = x_0 + \sum_n \frac{1}{\lambda_n} \langle P_k y, u_n \rangle v_n = x_0 + \sum_{n=1}^{k} \frac{1}{\lambda_n} \langle y, u_n \rangle v_n \in H_1$$

for any $x_0 \in \mathrm{Ker}(A)$. Observe that since for increasing k,

$$\|Ax_k - Py\|^2 = \|(P - P_k)y\|^2 \xrightarrow{k \to \infty} 0,$$

the residual of the projected equation can be made arbitrarily small.

Finally, to remove the ambiguity of the sought solution due to the possible noninjectivity of A, we select $x_0 = 0$. This choice minimizes the norm of x_k since, by orthogonality,

$$\|x_k\|^2 = \|x_0\|^2 + \sum_{n=1}^{k} \frac{1}{\lambda_n^2} |\langle y, u_n \rangle|^2.$$

Definition

Let H_1 and H_2 be separable real Hilbert spaces and let $A\colon H_1 \to H_2$ be a compact linear operator with a singular system (λ_n, v_n, u_n). By the truncated SVD approximation (TSVD) of the problem $Ax = y$, we mean the problem of finding $x \in H_1$ such that

$$Ax = P_k y, \quad x \perp \mathrm{Ker}(A)$$

for some $k \geq 1$.

Theorem

The solution to the TSVD problem has a unique solution x_k, called the truncated SVD (TSVD) solution, given by

$$x_k = \sum_{n=1}^{k} \frac{1}{\lambda_n} \langle y, u_n \rangle v_n.$$

The TSVD solution satisfies

$$\|Ax_k - y\|^2 = \|(I - P)y\|^2 + \|(P - P_k)y\|^2 \xrightarrow{k\to\infty} \|(I - P)y\|^2.$$

Truncated SVD for a matrix $A \in \mathbb{R}^{m\times n}$

The truncated SVD solution, i.e., solution of

$$Ax = P_k y \quad \text{and} \quad x \perp \mathrm{Ker}(A), \quad 1 \le k \le p := \mathrm{rank}(A),$$

where $P_k : \mathbb{R}^m \to \mathrm{span}\{u_1, \ldots, u_k\}$ is an orthogonal projection, is given by

$$x_k = \sum_{j=1}^{k} \frac{1}{\lambda_j} \langle y, u_j \rangle v_j = \sum_{j=1}^{k} \frac{1}{\lambda_j} v_j (u_j^{\mathrm{T}} y) = V \Lambda_k^{\dagger} U^{\mathrm{T}} y,$$

where A has the SVD $A = U\Lambda V^{\mathrm{T}}$ and we define

$$\Lambda_k^{\dagger} = \begin{pmatrix} 1/\lambda_1 & 0 & \cdots & 0 & \cdots & 0 \\ 0 & 1/\lambda_2 & & & & \vdots \\ \vdots & & \ddots & & & \\ & & & 1/\lambda_k & & \\ & & & & 0 & \\ \vdots & & & & \ddots & \vdots \\ 0 & \cdots & & & \cdots & 0 \end{pmatrix} \in \mathbb{R}^{n\times m},$$

where $\lambda_1 \ge \cdots \ge \lambda_p > 0$ are the singular values of A (i.e., diagonal of Λ).

Moore–Penrose pseudoinverse of matrices

For the largest possible cut-off $k = p = \operatorname{rank}(A)$, the matrix

$$A^\dagger := A_p^\dagger = V\Lambda_p^\dagger U^{\mathrm{T}} =: V\Lambda^\dagger U^{\mathrm{T}}$$

is called the *Moore–Penrose pseudoinverse*. It follows from the above that $x^\dagger = A^\dagger y$ is the solution of the projected (matrix) equation

$$Ax = Py,$$

where $P\colon \mathbb{R}^m \to \operatorname{Ran}(A)$ is the orthogonal projection.

The solution $x^\dagger = A^\dagger y$ is called the *minimum norm solution* of the problem $y = Ax$ since

$$\|A^\dagger y\| = \min\{\|x\| : \|Ax - y\| = \|(I - P)y\|\},$$

where P is the projection onto the range of A. The minimum norm solution is the solution that minimizes the residual error and has the minimum norm.

Since the smallest singular value λ_p is extremely small in inverse problems, the use of the pseudoinverse is usually very sensitive to inaccuracies in the data y.

Spectral regularization using TSVD, i.e., discarding singular values below a certain threshold from the forward model, is a simple and popular technique used to render linear problems less ill-posed while improving the noise robustness of the numerical inversion procedure.

However, obtaining the singular values and vectors for large system matrices is usually very slow.

Numerical experiment: TSVD for X-ray tomography

Let us consider the inverse problem of recovering the attenuation coefficient (density) of an object given a set of X-ray measurements. Recall from last week that the mathematical model can be expressed as

$$y = Ax,$$

where $y \in \mathbb{R}^Q$ denotes the (noisy) measurements for Q X-rays, $A \in \mathbb{R}^{Q \times n^2}$ is the projection matrix subject to an $n \times n$ pixel discretization of the computational domain, and $x \in \mathbb{R}^{n^2}$ denotes the (piecewise constant) discretization of the unknown attenuation inside the object of interest.

The data y can be reshaped into an $n \times n$ array, which is a graphical representation of the X-ray measurements (sinogram). The unknown can be reshaped into an $n \times n$ image of the density of the imaged object.

Let us use TSVD to solve this inverse problem for *real-life measurement data*. We use the FIPS open dataset of carved cheese available at https://doi.org/10.5281/zenodo.1254210

The files `DataFull_128x15.mat` and `DataLimited_128x15.mat` contain sparse angle and limited angle tomography measurements, respectively. The data has been collected using 15 projections spanning either the full 360° circle in the first dataset, and 15 projections spanning a limited 90° angle of view in the second dataset. The computational domain is a 128×128 pixel grid in both cases. Each file contains a projection matrix `A` and a sinogram measurement matrix `m`.

By defining `y = m.reshape((m.size,1))` (in MATLAB: `y = m(:)`), the unknown `x` can be solved from the linear equation

$$\mathtt{y} = \mathtt{A}\mathtt{x}.$$

The reconstruction is the image `x.reshape((128,128))` (in MATLAB: `reshape(x,128,128)`).

See the files `tomo_tsvd.py` / `tomo_tsvd.m` on the course webpage!

TSVD for sparse angle tomography data

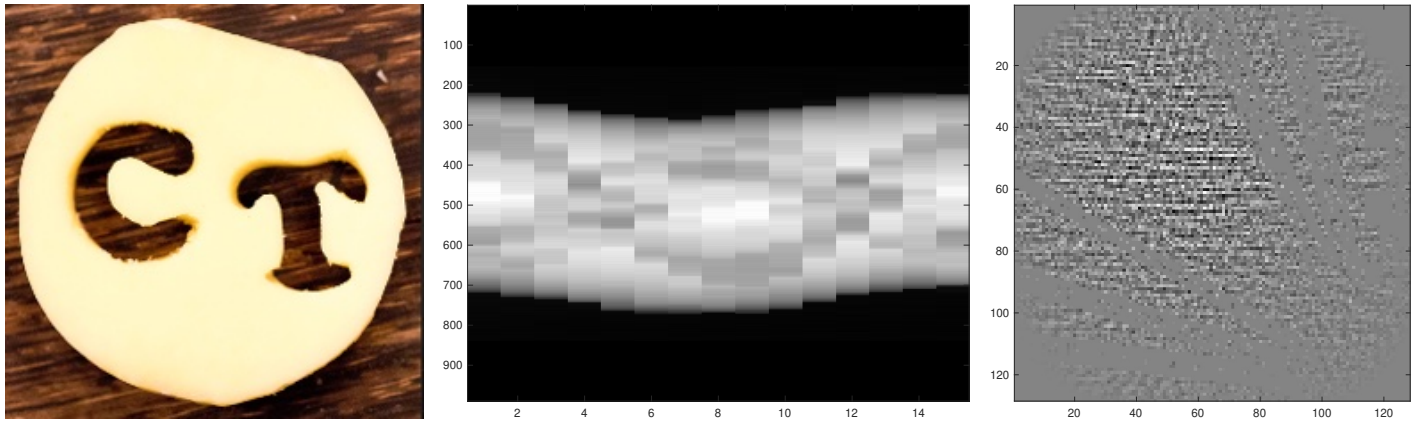

Left: the actual object. Middle: sinogram data for sparse angle tomography. Right: naïve reconstruction without any regularization.

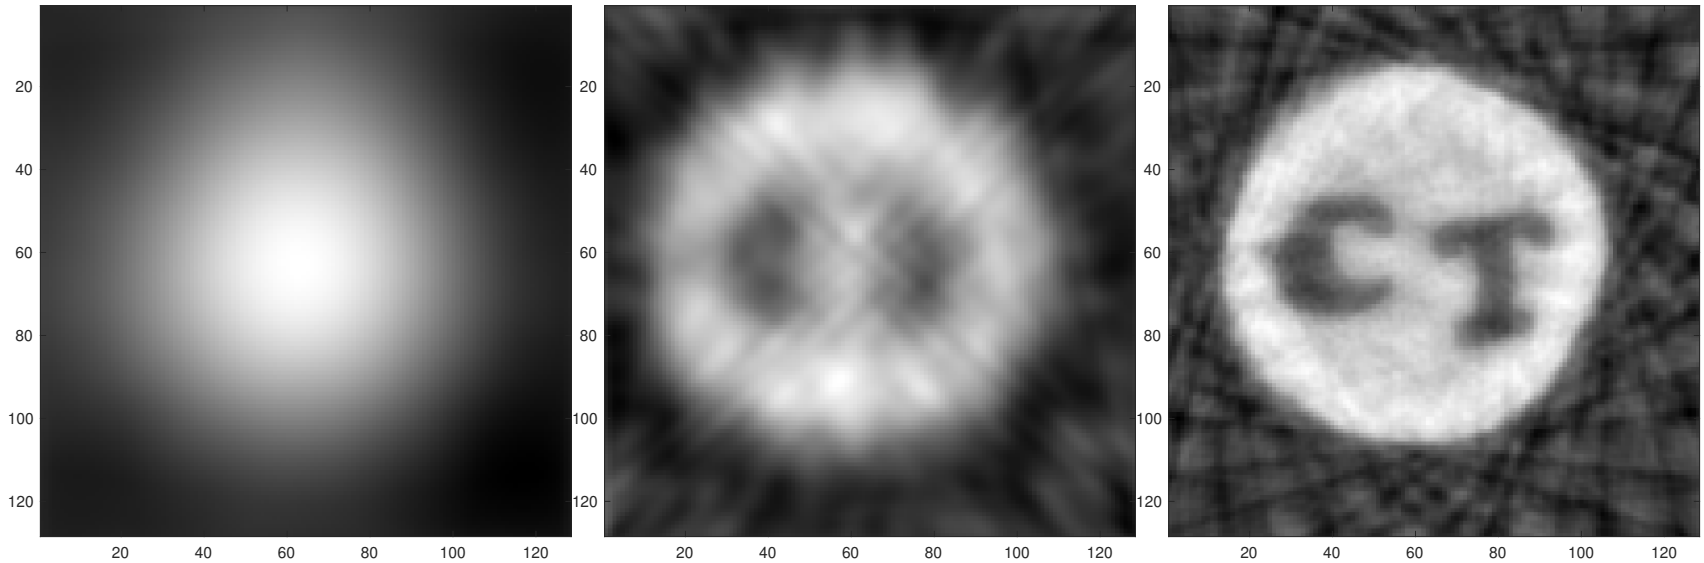

TSVD reconstructions with spectral cut-off $k = 10$ (left), $k = 100$ (middle), and $k = 1000$ (right).

TSVD for limited angle tomography data

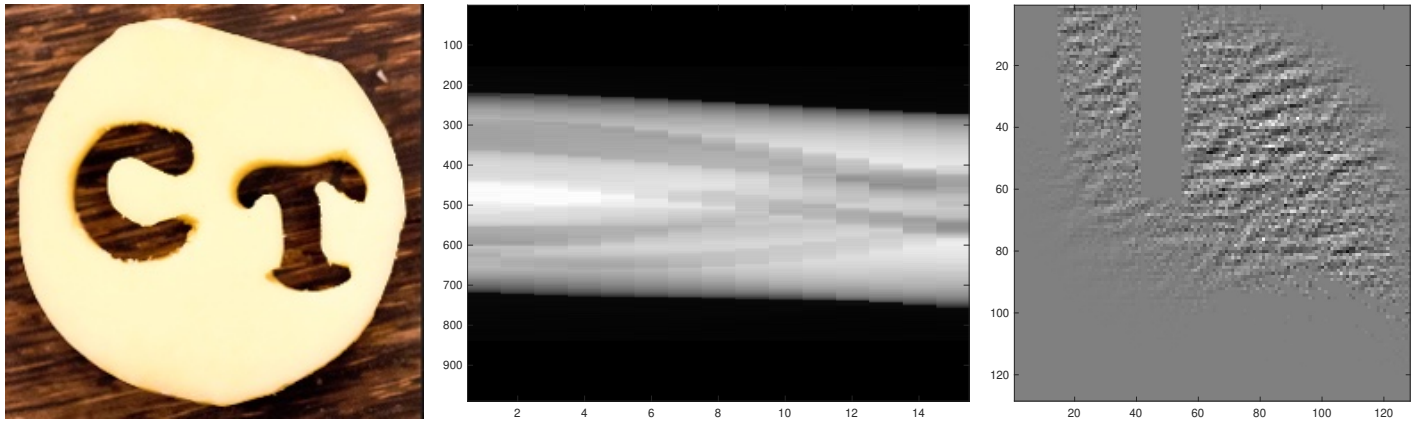

Left: the actual object. Middle: sinogram data for limited angle tomography. Right: naïve reconstruction without any regularization.

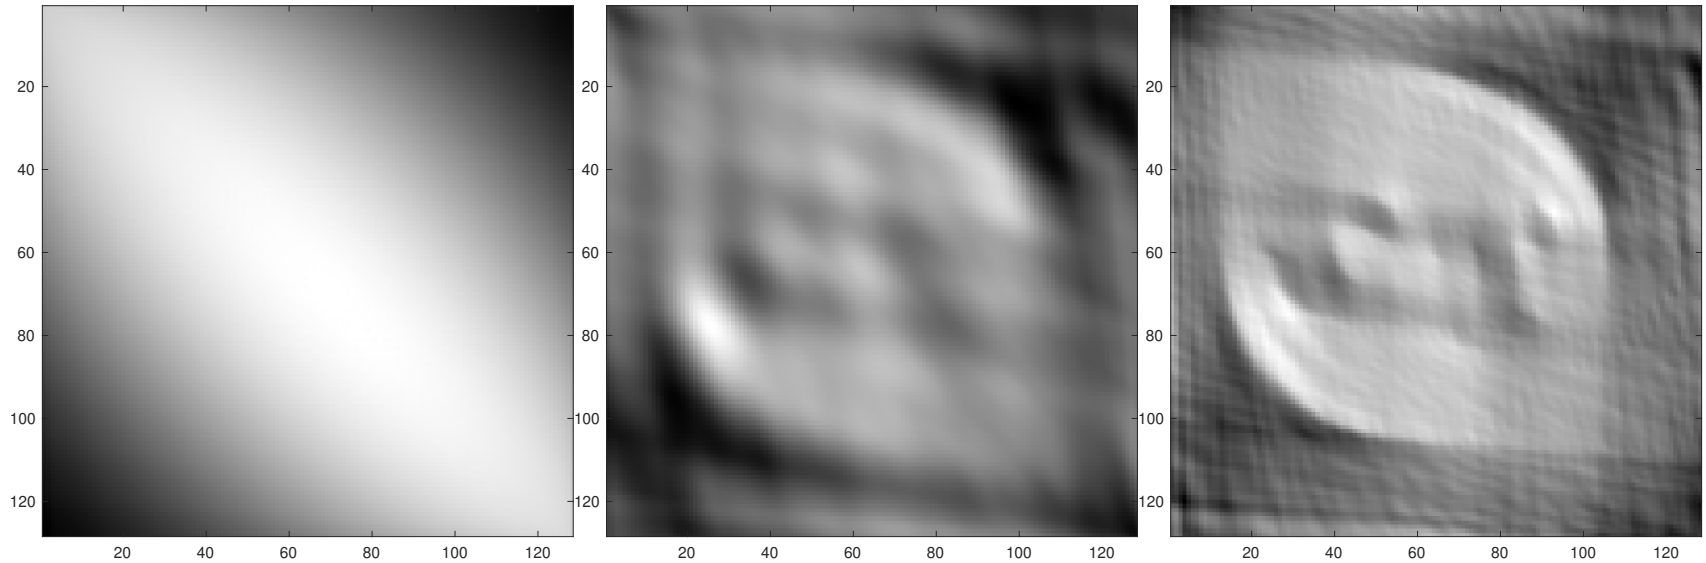

TSVD reconstructions with spectral cut-off $k = 10$ (left), $k = 100$ (middle), and $k = 1000$ (right).

In summary (matrix case $H_1 = \mathbb{R}^n$, $H_2 = \mathbb{R}^m$): Let the SVD of matrix $A \in \mathbb{R}^{m\times n}$ be given by

$$A = U\Lambda V^{\mathrm{T}},$$

where $\Lambda \in \mathbb{R}^{m\times n}$ has the non-negative singular values $\{\lambda_j\}_{j=1}^{\min\{m,n\}}$ on its diagonal and $V \in \mathbb{R}^{n\times n}$ and $U \in \mathbb{R}^{m\times m}$ are orthogonal matrices.[†]

The TSVD solution for $1 \le k \le p := \operatorname{rank}(A)$ is given by

$$x_k = V\Lambda_k^{\dagger} U^{\mathrm{T}} y,$$

where

$$\Lambda_k^{\dagger} = \begin{pmatrix} 1/\lambda_1 & 0 & \cdots & 0 & & \cdots & 0 \\ 0 & 1/\lambda_2 & & & & & \vdots \\ \vdots & & \ddots & & & & \\ & & & 1/\lambda_k & & & \\ & & & & 0 & & \\ \vdots & & & & & \ddots & \vdots \\ 0 & \cdots & & & & \cdots & 0 \end{pmatrix} \in \mathbb{R}^{n\times m}.$$

The matrix $A^{\dagger} = V\Lambda_p^{\dagger} U^{\mathrm{T}}$ is called the *Moore–Penrose pseudoinverse* of A.

[†]This means that the columns $\{v_j\}_{j=1}^{n}$ of V form an orthonormal basis for $\mathbb{R}^n$, and similarly the columns $\{u_j\}_{j=1}^{m}$ of U are an orhonormal basis of $\mathbb{R}^m$.

Morozov discrepancy principle

The implementation of TSVD raises the question: how to choose the spectral cut-off parameter k? If k is too small, the TSVD operator loses information about the forward operator. On the other hand, if k is chosen too large, then the forward operator becomes ill-conditioned and sensitive to measurement noise.

If the noise level of the data is known (or can be estimated), then one of the simplest criteria is to choose the spectral parameter as large as possible without fitting the solution to noise.

Morozov discrepancy principle

Let H_1 and H_2 be separable real Hilbert spaces and $A\colon H_1 \to H_2$ a compact linear operator.

How to choose the spectral cut-off index $k \geq 1$ in the TSVD problem

$$Ax = P_k y \quad \text{and} \quad x \perp \operatorname{Ker}(A)?$$

There is a rule of thumb called the *Morozov discrepancy principle:*

Suppose that the data $y \in H_2$ is a noisy approximation of noiseless "exact" data $y_0 \in H_2$. While y_0 is unknown to us, we may have an estimate on the noise level, e.g.,

$$\|y - y_0\| \approx \varepsilon > 0.$$

We choose the smallest $k \geq 1$ such that the residual satisfies

$$\|y - Ax_k\| \leq \varepsilon.$$

Intuitively, this means that we cannot expect the approximate solution to yield a smaller residual than the measurement error without fitting the solution to noise.

Q: *When does an index $k \geq 1$ satisfying $\|y - Ax_k\| \leq \varepsilon$ exist?*
A: When $\varepsilon > \|Py - y\|$ and $\operatorname{rank}(A) = \infty$, it follows from $\overline{\operatorname{Ran}(A)} = \operatorname{Ran}(P) \perp \operatorname{Ran}(I - P)$ that

$$\|Ax_k - y\|^2 = \|Ax_k - Py + Py - y\|^2 = \|Ax_k - Py\|^2 + \|(P - I)y\|^2$$
$$= \sum_{n=k+1}^{\infty} |\langle y, u_n \rangle|^2 + \|(P - I)y\|^2 \xrightarrow{k \to \infty} \|Py - y\|^2.$$

Due to the properties of the orthogonal projection, $\|Py - y\| = \inf_{z \in \operatorname{Ran}(A)} \|z - y\|$, so this is the best we can do. (Note however that there is no guarantee that prevents $\|x_k\|$ from blowing up as $k \to \infty$.)

On the other hand, if $p = \operatorname{rank}(A) < \infty$,

$$\|Ax_p - y\| = \|P_p y - y\| = \|Py - y\|.$$

One should usually avoid choosing the spectral cut-off to be this large in practice.

Numerical example: backward heat equation

Let us consider the backward heat equation:

$$\begin{cases} \partial_t u(x,t) = \partial_x^2 u(x,t) & \text{for } (x,t) \in (0,\pi) \times \mathbb{R}_+, \\ u(0,\cdot) = u(\pi,\cdot) = 0 & \text{on } \mathbb{R}_+, \\ u(\cdot,0) = f & \text{on } (0,\pi), \end{cases}$$

where $f : (0,\pi) \to \mathbb{R}$ is the initial heat distribution.

Forward problem: Given initial data $f : (0,\pi) \to \mathbb{R}$, determine the heat distribution $u(\cdot, T)$ at time $T > 0$.

Inverse problem: Reconstruct the initial state f based on noisy measurements of $u(\cdot, T)$ at time $T > 0$.

Let us consider a simple discretization of the PDE

$$\begin{cases} \partial_t u(x,t) = \partial_x^2 u(x,t) & \text{for } (x,t) \in (0,\pi) \times \mathbb{R}_+, \\ u(0,\cdot) = u(\pi,\cdot) = 0 & \text{on } \mathbb{R}_+, \\ u(\cdot,0) = f & \text{on } (0,\pi). \end{cases}$$

Let $x_j = jh$ for $j = 0, \ldots, 100$, where $h = \pi/100$ is the step size.

Zero Dirichlet boundary conditions imply that $u(x_0,t) = u(x_{100},t) = 0$.

The spatial second derivative can be discretized using the stencils

$$\partial_x^2 u(x_1,t) = \frac{-2u(x_1,t) + u(x_2,t)}{h^2} + \mathcal{O}(h^2),$$
$$\partial_x^2 u(x_j,t) = \frac{u(x_{j-1},t) - 2u(x_j,t) + u(x_{j+1},t)}{h^2} + \mathcal{O}(h^2) \quad \text{for } j = 2, \ldots, 98,$$
$$\partial_x^2 u(x_{99},t) = \frac{u(x_{98},t) - 2u(x_{99},t)}{h^2} + \mathcal{O}(h^2).$$

Denote $U(t) = (U_j(t))_{j=1}^{99} = (u(x_j,t))_{j=1}^{99}$ and $F = (f(x_j))_{j=1}^{99}$.

$$\frac{\partial}{\partial t}\begin{bmatrix} U_1(t) \\ U_2(t) \\ U_3(t) \\ \vdots \\ U_{98}(t) \\ U_{99}(t) \end{bmatrix} = \underbrace{\frac{1}{h^2}\begin{bmatrix} -2 & 1 & & & & \\ 1 & -2 & 1 & & & \\ & 1 & -2 & 1 & & \\ & & & \ddots & & \\ & & & 1 & -2 & 1 \\ & & & & 1 & -2 \end{bmatrix}}_{=:B} \begin{bmatrix} U_1(t) \\ U_2(t) \\ U_3(t) \\ \vdots \\ U_{98}(t) \\ U_{99}(t) \end{bmatrix}.$$

After spatial discretization, our PDE has been transformed into the initial value problem

$$\dot{U}(t) = BU(t), \quad U(0) = F.$$

At time $t = T > 0$, the discretized heat distribution $U := U(T)$ is given by

$$U = AF,$$

where $A = \mathrm{e}^{TB} \in \mathbb{R}^{99\times 99}$ and

$$\mathrm{e}^{M} := \sum_{k=0}^{\infty} \frac{1}{k!} M^k$$

is the *matrix exponential* (`scipy.linalg.expm` in Python / `expm` in MATLAB).

A note on simulating measurement data and inverse crimes

When simulating measurement data, one should take care not to use the same computational model for inversion as the one which was used to generate the measurements in the first place. This would lead to unreasonably good reconstructions, since this is akin to multiplying a matrix with its own inverse. This is known as an *inverse crime*. (Similar concerns also apply to non-linear problems.)

With real-life measurement data, we do not have worry about this phenomenon – measurements that come from nature are automatically independent of any computational model we end up using for practical inverse problems simulations.

A popular technique to avoid committing an inverse crime is using a higher resolution computational model to generate the measurements and interpolating the simulated data onto a coarser grid, where we plan to carry out the actual computational inversion. Another good option is to use an analytic solution, if one is readily available. We will use this technique with the heat equation.

The forward problem of the heat equation

$$\begin{cases} \partial_t u(x,t) = \partial_x^2 u(x,t) & \text{for } (x,t) \in (0,\pi) \times \mathbb{R}_+, \\ u(0,\cdot) = u(\pi,\cdot) = 0 & \text{on } \mathbb{R}_+, \\ u(\cdot,0) = f & \text{on } (0,\pi), \end{cases}$$

has the classical series solution

$$u(x,t) = \sum_{n=1}^{\infty} \hat{f}_n \, \mathrm{e}^{-n^2 t} \sin(nx),$$

where the coefficients $\hat{f}_n$ are the Fourier sine series coefficients of the initial heat distribution f satisfying

$$f(x) = \sum_{n=1}^{\infty} \hat{f}_n \sin(nx), \quad \hat{f}_n = \frac{2}{\pi} \int_0^{\pi} f(x) \sin(nx) \, \mathrm{d}x.$$

Let us fix the ground truth

$$f(x) = \begin{cases} 1 & \text{if } x \in [1,2], \\ 0 & \text{if } x \in (0,1) \cup (2,\pi). \end{cases}$$

It is easy to see that the Fourier sine coefficients are given by

$$\hat{f}_n = \frac{2}{n\pi}(\cos n - \cos 2n).$$

Let us plug these into the forward solution at time $t = T > 0$

$$u(x_j, T) = \sum_{n=1}^{\infty} \hat{f}_n \, \mathrm{e}^{-n^2 T} \sin(n x_j), \quad j = 1, \ldots, 99,$$

and add some simulated measurement noise!

We assume that the data $U(T) \in \mathbb{R}^{99}$ at time $T = 0.1$ is contaminated with mean-zero Gaussian noise with standard deviation 0.01, and that the discrepancy between the measured data and the underlying "exact" data equals the square root of the expected value of the squared norm of the noise vector, i.e.,

$$\varepsilon = \sqrt{99 \cdot 0.01^2} \approx 0.0995.$$

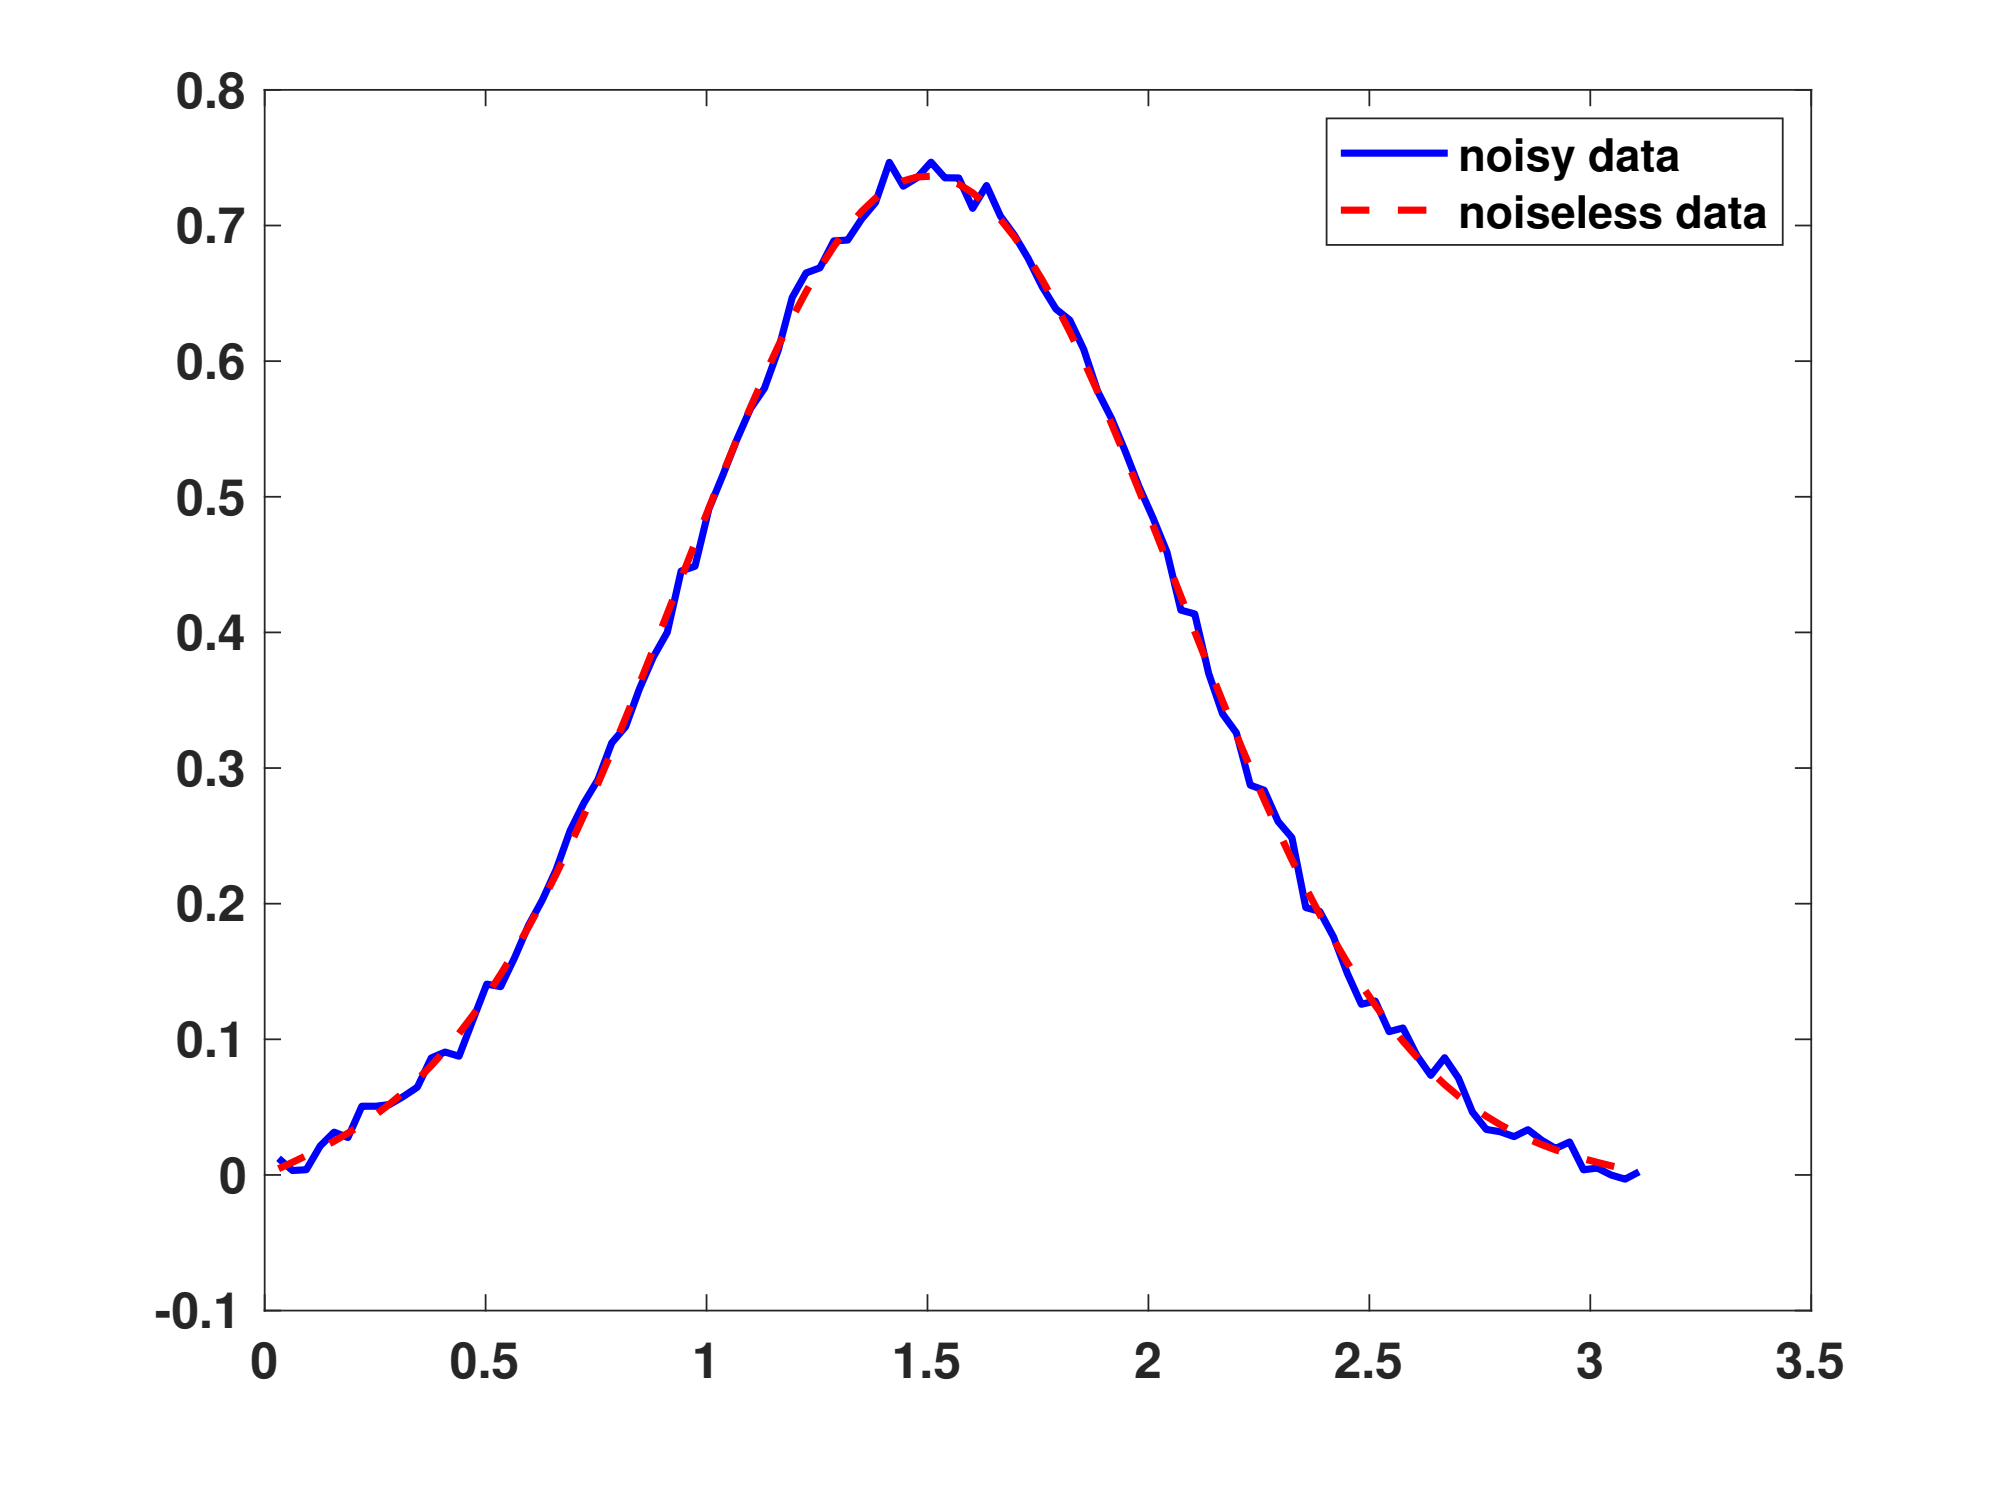
noisy data
noiseless data
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-0.1
0
0.5
1
1.5
2
2.5
3
3.5

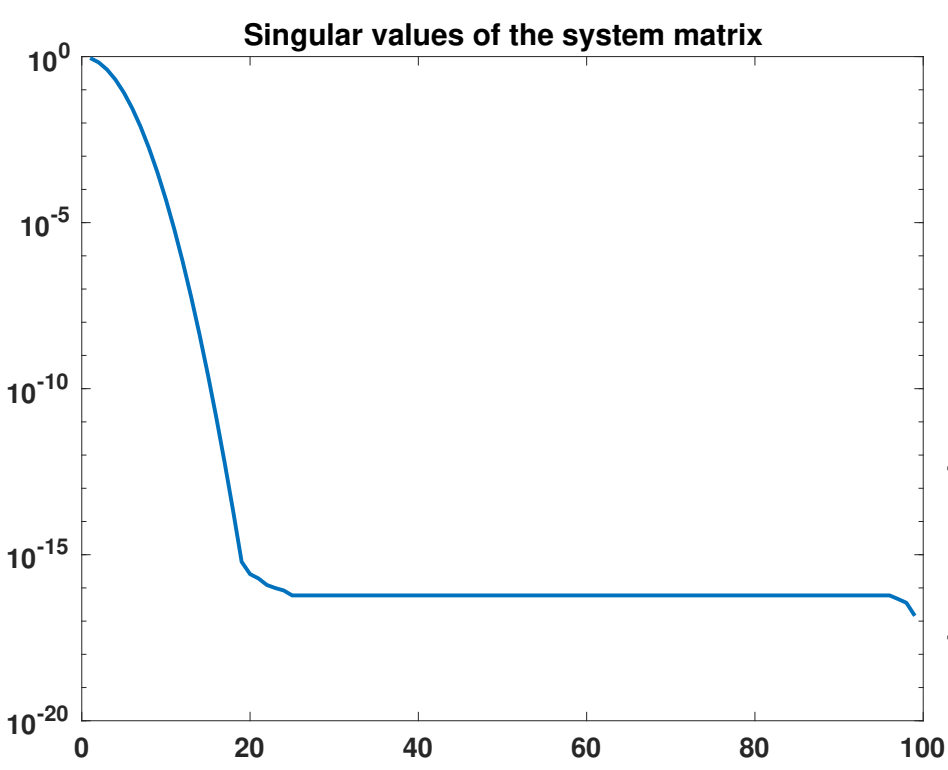
Singular values of the system matrix
10^0
10^-5
10^-10
10^-15
10^-20
0
20
40
60
80
100

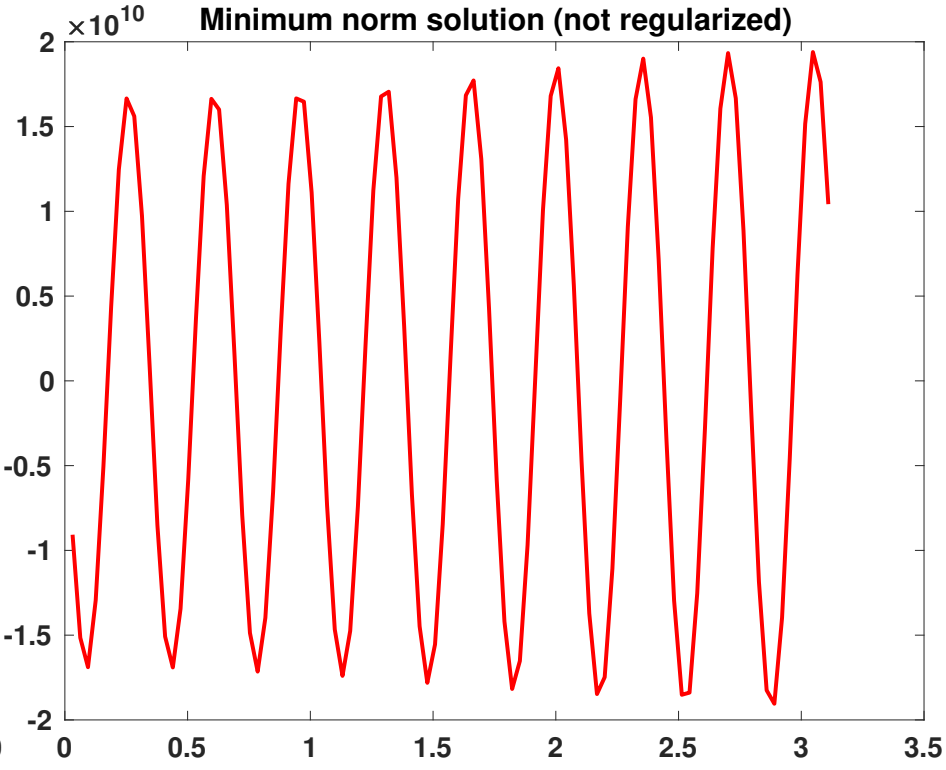
Minimum norm solution (not regularized)
×10^10
2
1.5
1
0.5
0
-0.5
-1
-1.5
-2
0
0.5
1
1.5
2
2.5
3
3.5

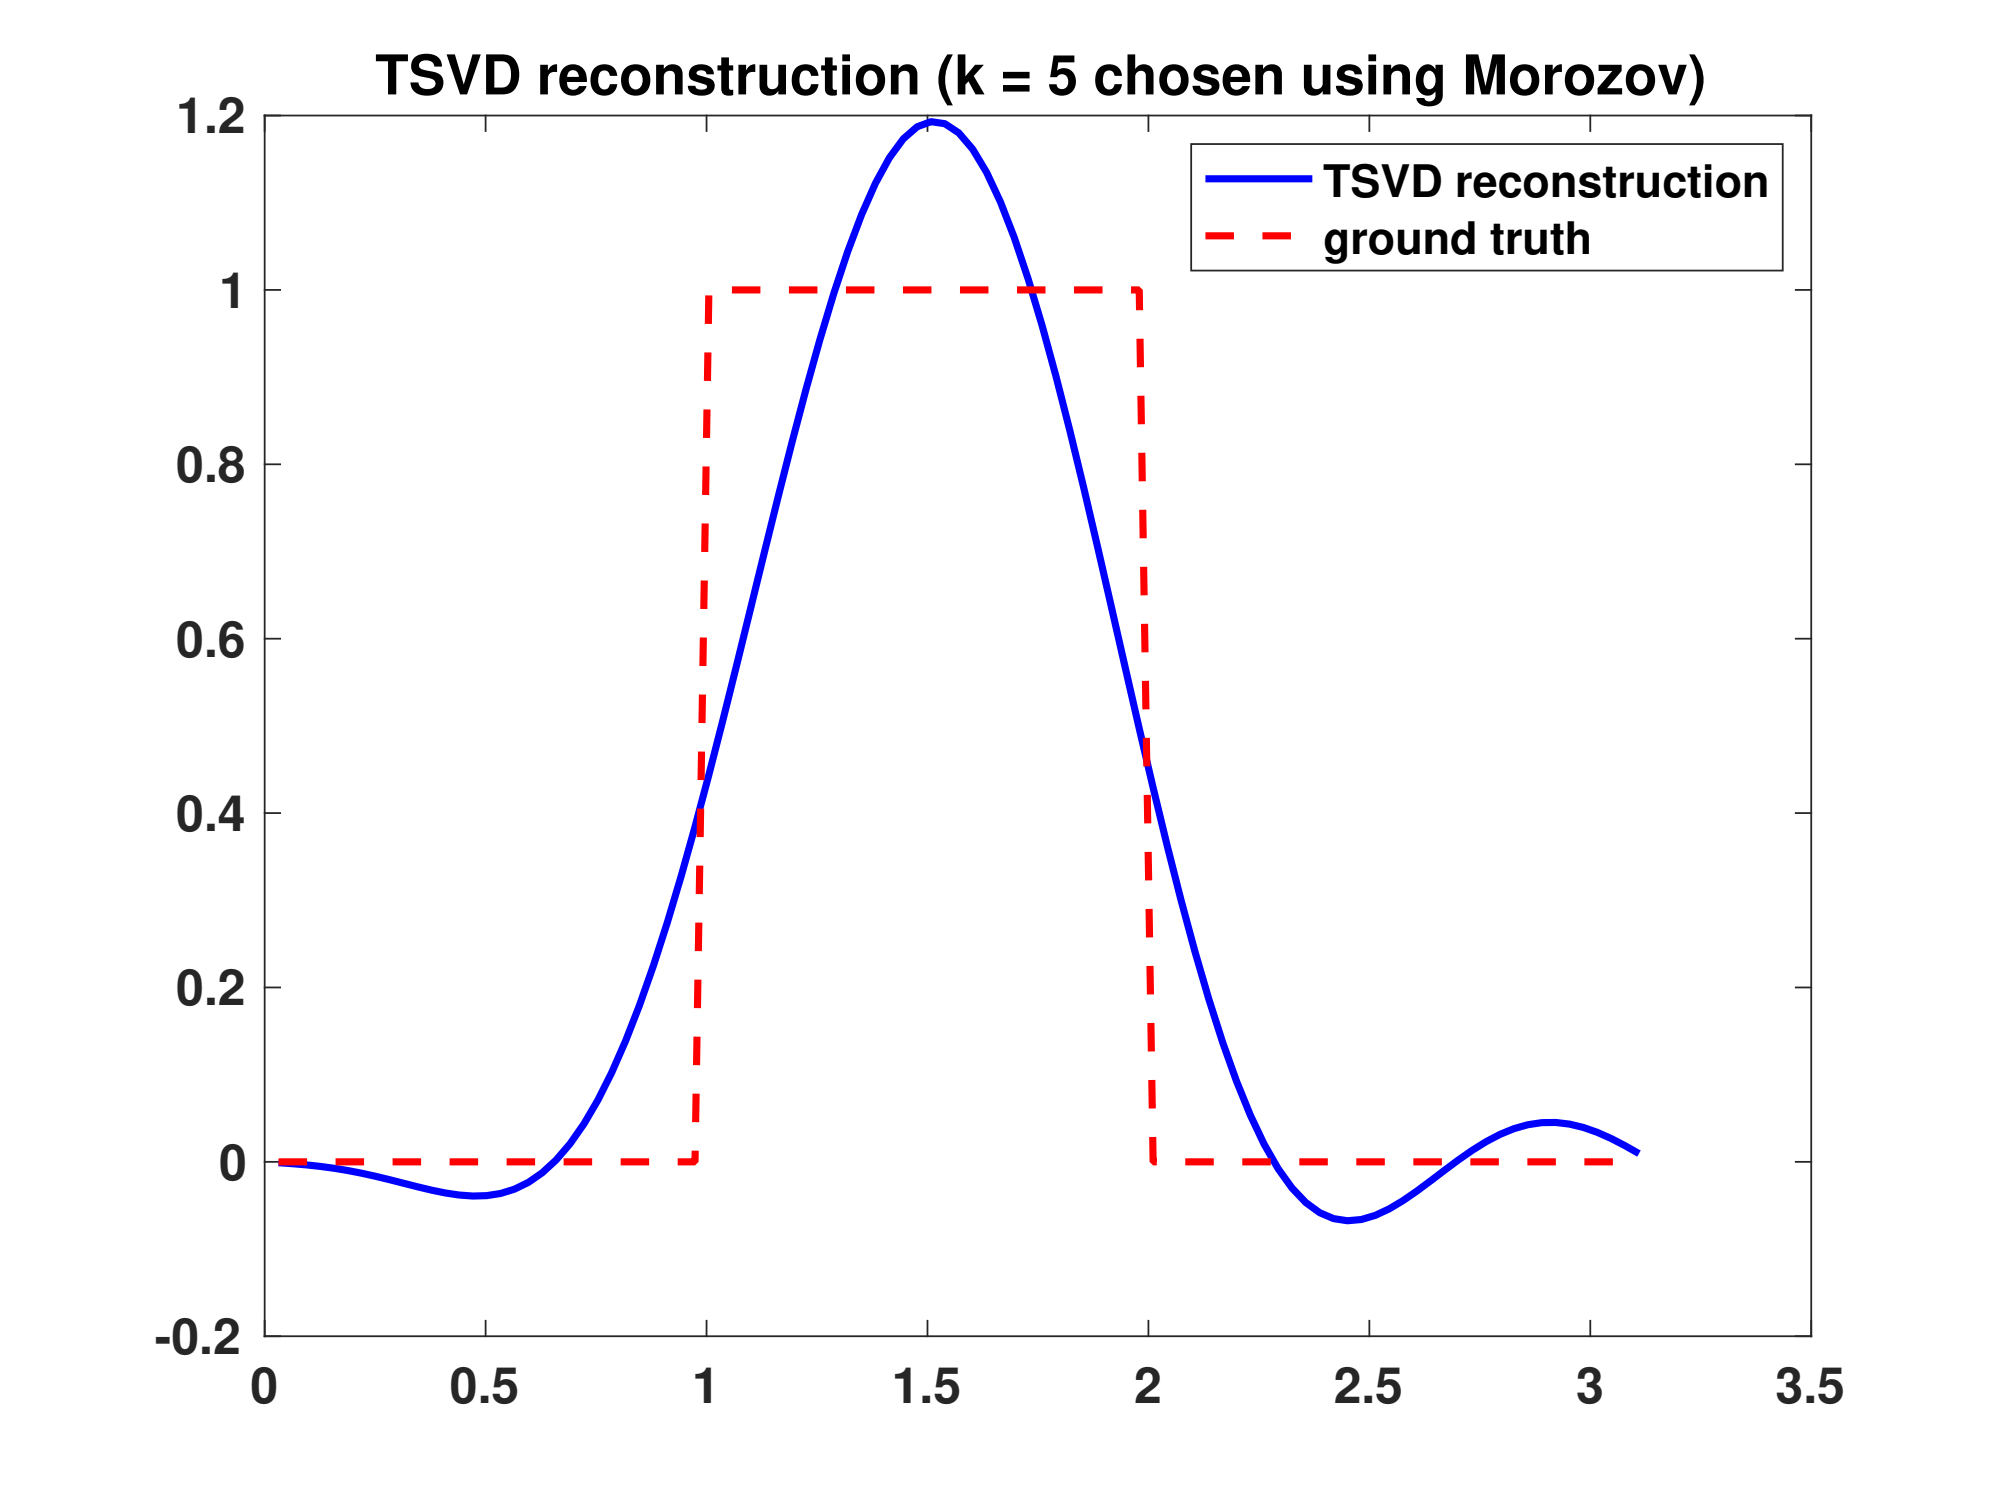
TSVD reconstruction (k = 5 chosen using Morozov)
TSVD reconstruction
ground truth
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
0.5
1
1.5
2
2.5
3
3.5

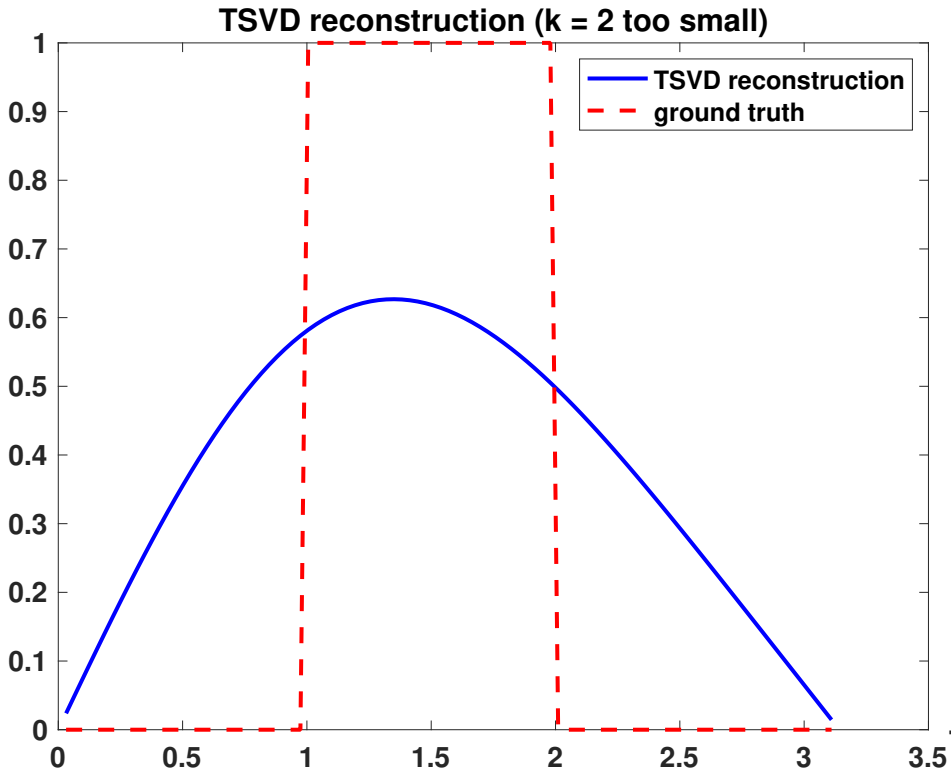

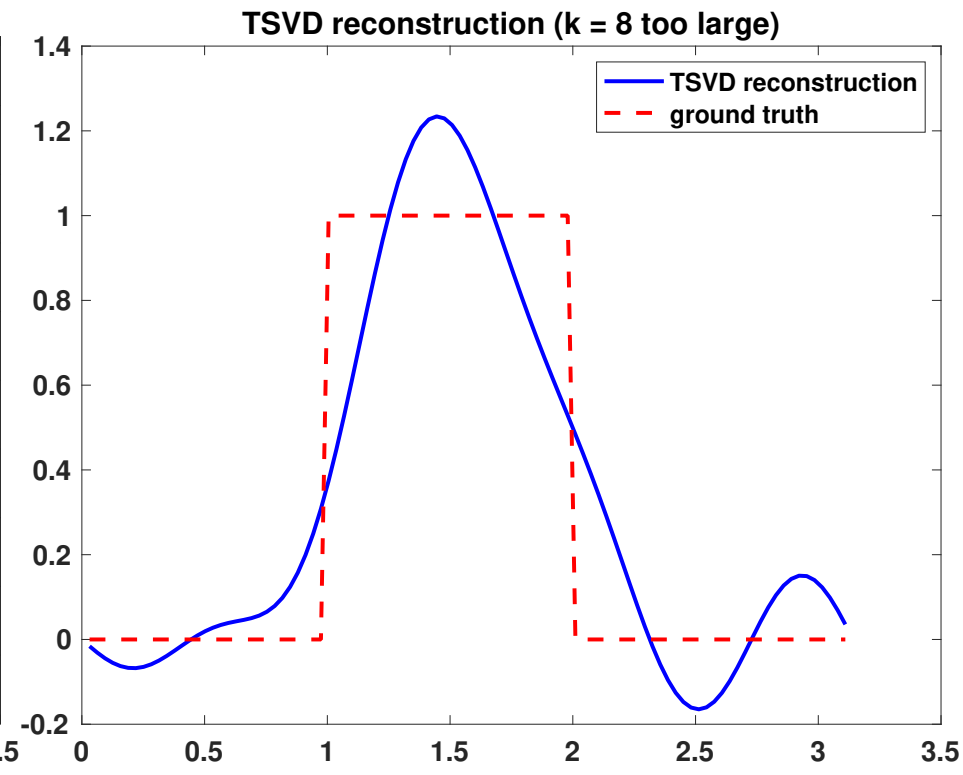

See the files `heateq_tsvd.py` / `heateq_tsvd.m` on the course webpage!