Computational inverse problems

Nuutti Hyvönen and Juha-Pekka Puska

nuutti.hyvonen@aalto.fi, juha-pekka.puska@aalto.fi

Sixth lecture, March 17, 2021.

Summary of the previous lecture

A Minimization problem: Let $A \in \mathbb{R}^{n\times n}$ be symmetric and positive definite. Instead of solving the original equation $Ax = y$ directly, we consider minimizing the functional

$$\phi(x) = (x_*-x)^{\mathrm{T}}A(x_*-x) = e^{\mathrm{T}}Ae = (y-Ax)^{\mathrm{T}}A^{-1}(y-Ax) = r^{\mathrm{T}}A^{-1}r,$$

where $x_* = A^{-1}y$ is the actual solution, and e and r are called the error and the residual corresponding to the approximate solution x. The unique minimizer of this functional is the solution of the original problem, i.e., x_*.

A sequence of minimizers: Given an initial guess x_0 and a set of non-zero search directions $\{s_j\}_{j=0}^{k} \subset \mathbb{R}^n$, we define the approximate solution x_{j+1}, $j = 1, \ldots, k$, recursively as the minimizer of the functional ϕ on the line

$$\mathcal{S}_j \;=\; \{x \in \mathbb{R}^n \mid x = x_j + \alpha s_j,\ \alpha \in \mathbb{R}\}.$$

This can be done through the iteration

$$x_{j+1} \;=\; x_j + \alpha_j s_j, \qquad \text{with} \;\; \alpha_j \;=\; \frac{s_j^{\mathrm{T}} r_j}{s_j^{\mathrm{T}} A s_j}, \qquad j = 0, \ldots, k,$$

where $r_j = y - Ax_j$ is the residual corresponding to x_j.

A-conjugate search directions: The non-zero vectors $\{s_j\}_{j=0}^k$ are called A-conjugate if

$$\langle s_i, s_j \rangle_A \;=\; s_i^{\mathrm{T}} A s_j \;=\; 0 \qquad \text{for } i \neq j.$$

If the search directions are chosen this cleverly, the iterate x_{k+1} is the minimizer of ϕ over the whole hyperplane

$$\mathcal{S}_k \;=\; \{x \in \mathbb{R}^n \mid x = x_0 + S_k h,\; h \in \mathbb{R}^{k+1}\},$$

i.e., over all vectors of the form $x = x_0 + \sum_{j=0}^k h_j s_j$, where $h_0, \ldots, h_k$ are real numbers. This minimizer can be given explicitly as

$$x_{k+1} \;=\; x_0 + S_k h_*, \qquad h_* \;=\; (S_k^{\mathrm{T}} A S_k)^{-1} S_k^{\mathrm{T}} r_0,$$

where $S_k = [s_0, \ldots, s_k] \in \mathbb{R}^{n \times (k+1)}$. In particular, x_n is the global minimizer, i.e., $x_n = x_*$.

Conjugate gradient method

To sum up, we have arrived at the following algorithm

Choose x_0.

Set $k = 0$, $r_0 = y - Ax_0$, $s_0 = r_0$;

Repeat until the chosen stopping rule is satisfied:

$\alpha_k = (s_k^{\mathrm{T}} r_k)/(s_k^{\mathrm{T}} A s_k)$;

$x_{k+1} = x_k + \alpha_k s_k$;

$r_{k+1} = r_k - \alpha_k A s_k$; % Note: $r_{k+1} = y - Ax_k - \alpha_k A s_k$

$\beta_k = -(s_k^{\mathrm{T}} A r_{k+1})/(s_k^{\mathrm{T}} A s_k)$;

$s_{k+1} = r_{k+1} + \beta_k s_k$;

$k \leftarrow k + 1$;

end

However, the algorithm is usually presented in a slightly different form. Assuming that the iteration has not yet converged at the iterate x_k, we deduce the following formulae:

Since $r_k \perp s_{k-1}$,

$$s_k^{\mathrm{T}} r_k \;=\; (r_k + \beta_{k-1} s_{k-1})^{\mathrm{T}} r_k \;=\; \|r_k\|^2,$$

resulting in

$$\alpha_k = \frac{\|r_k\|^2}{s_k^{\mathrm{T}} A s_k}.$$

In particular, since $r_{k+1} \perp \operatorname{span}\{s_0, \ldots, s_k\} = \mathcal{K}_{k+1} \ni r_k$, this means that

$$\|r_{k+1}\|^2 \;=\; r_{k+1}^{\mathrm{T}}(r_k - \alpha_k A s_k) \;=\; -\frac{\|r_k\|^2}{s_k^{\mathrm{T}} A s_k}\, r_{k+1}^{\mathrm{T}} A s_k \;=\; \beta_k \|r_k\|^2.$$

Solving for β_k and plugging the obtained formulae for α_k and β_k into the preliminary conjugate gradient algorithm leads to the standard form of the method:

Choose x_0.

Set $k = 0$, $r_0 = y - Ax_0$, $s_0 = r_0$;

Repeat until the chosen stopping rule is satisfied:

$$\alpha_k = \|r_k\|^2 / (s_k^{\mathrm{T}} A s_k);$$
$$x_{k+1} = x_k + \alpha_k s_k;$$
$$r_{k+1} = r_k - \alpha_k A s_k;$$
$$\beta_k = \|r_{k+1}\|^2 / \|r_k\|^2;$$
$$s_{k+1} = r_{k+1} + \beta_k s_k;$$
$$k \leftarrow k + 1;$$

end

NB: *There is an error in the update formula for x_{k+1} in the textbook.*

Conjugate gradient method for inverse problems

According to the above construction, if you apply the conjugate gradient method to the equation

$$Ax = y,$$

where $A \in \mathbb{R}^{n \times n}$ is symmetric and positive definite, you obtain the exact solution — up to rounding errors — in at most n iteration steps, i.e., $x_n = x_* = A^{-1}y$. However, such extensive iterating is not usually necessary: The algorithm typically converges satisfactorily much quicker; see, e.g., 2. exercise of the 3. session, where a (pessimistic) convergence rate is provided.

When dealing with ill-posed problems, one should be even more careful and terminate the iterations well before convergence, in order to avoid fitting the solution to noise. One should, actually, be extremely cautious because the conjugate gradient method often converges very fast.

Let us be a bit more precise and consider a general ill-posed matrix equation

$$Ax = y,$$

where $A \in \mathbb{R}^{m \times n}$ and $y \in \mathbb{R}^m$ are given.

In some cases, one may have $m = n$ and, in addition, some prior information stating that A is — at least in theory — positive (semi-)definite. In such situation, one can apply the conjugate gradient algorithm directly on this original equation.

In the general case, one may still consider the normal equation

$$A^{\mathrm{T}}Ax = A^{\mathrm{T}}y,$$

which corresponds, in essence, to solving the original equation in the least squares sense.

Now, the system matrix $A^{\mathrm{T}}A = (A^{\mathrm{T}}A)^{\mathrm{T}} \in \mathbb{R}^{n\times n}$ is symmetric and positive semi-definite:

$$u^{\mathrm{T}}A^{\mathrm{T}}Au \;=\; \|Au\|^2 \;>\; 0 \qquad \text{for all } \; u \in (\mathbb{R}^n \setminus \mathrm{Ker}(A)).$$

Hence, the conditions of the conjugate gradient algorithm are almost satisfied, and one may look for the solution of the inverse problem by using the conjugate gradient algorithm with A replaced by $A^{\mathrm{T}}A$ and y by $A^{\mathrm{T}}y$. (When implementing the algorithm in Matlab, bear in mind that matrix-matrix products are typically far more expensive than matrix-vector products.)

As a stopping condition, one may try, e.g., the Morozov principle for the original equation: Terminate the iteration when

$$\|y - Ax_k\| \;\leq\; \epsilon$$

for some $\epsilon > 0$, which measures the amount of noise in y in some sense.

An example: Heat distribution in a rod (revisited)

Let us once again consider the discretized inverse heat conduction problem in an insulated rod. We simulate the data in the exactly same way as above and add the same amount of noise.

The system matrix $A = e^{TB}$, $T = 0.1$, is symmetric since B is symmetric. Moreover, the infinite-dimensional version of A, i.e., E_T, is positive definite, and thus it is not far-fetched to assume that A is, at least, close to being positive semi-definite. (A symmetric matrix is positive definite if and only if all of its eigenvalues are positive; according to Matlab the eigenvalues of A are either positive or extremely close to zero.) Hence, it seems reasonable to try applying the conjugate gradient method directly to the original equation.

If we use the same value $\epsilon = \sqrt{99 \cdot 0.001^2} = 9.95 \cdot 10^{-3}$ for the Morozov discrepancy principle as in the previous examples, the conjugate gradient method becomes unstable before the stopping rule is satisfied. However, for the value $1.2 \cdot \sqrt{99 \cdot 0.001^2}$ the stopping rule is satisfied after seven iterations.

In the following, we visualize the evolution of the conjugate gradient iteration, show the norm of the residual $\|y - Ax_k\|$ as a function of k, and plot the solution corresponding to the (fine-tuned) discrepancy principle.

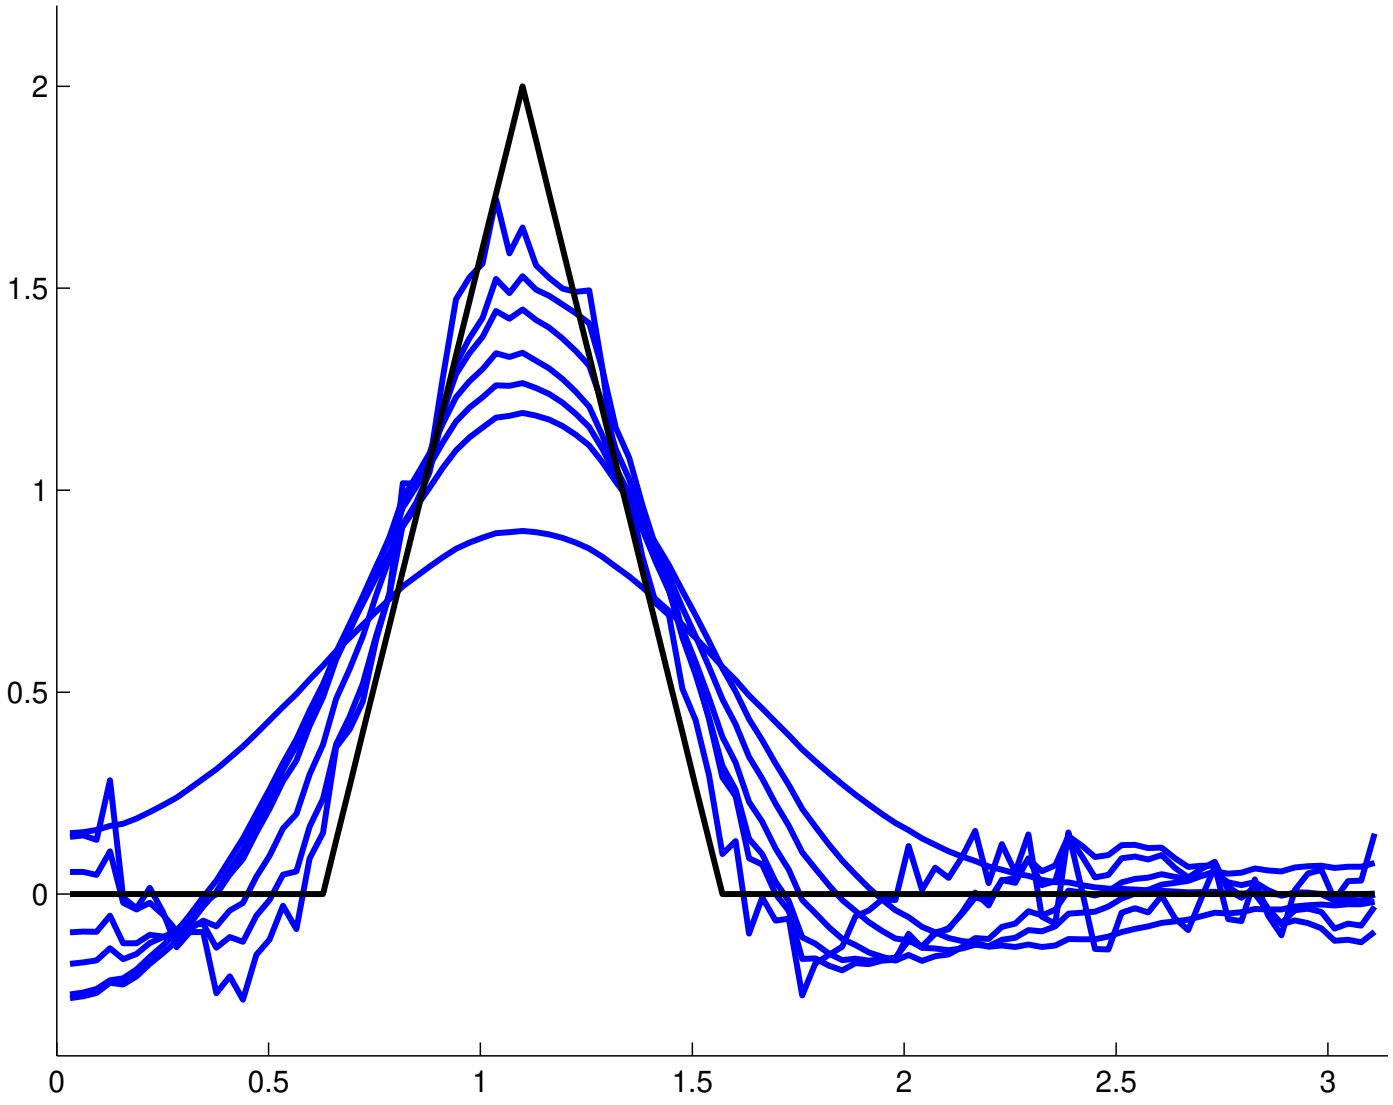
2
1.5
1
0.5
0
0
0.5
1
1.5
2
2.5
3

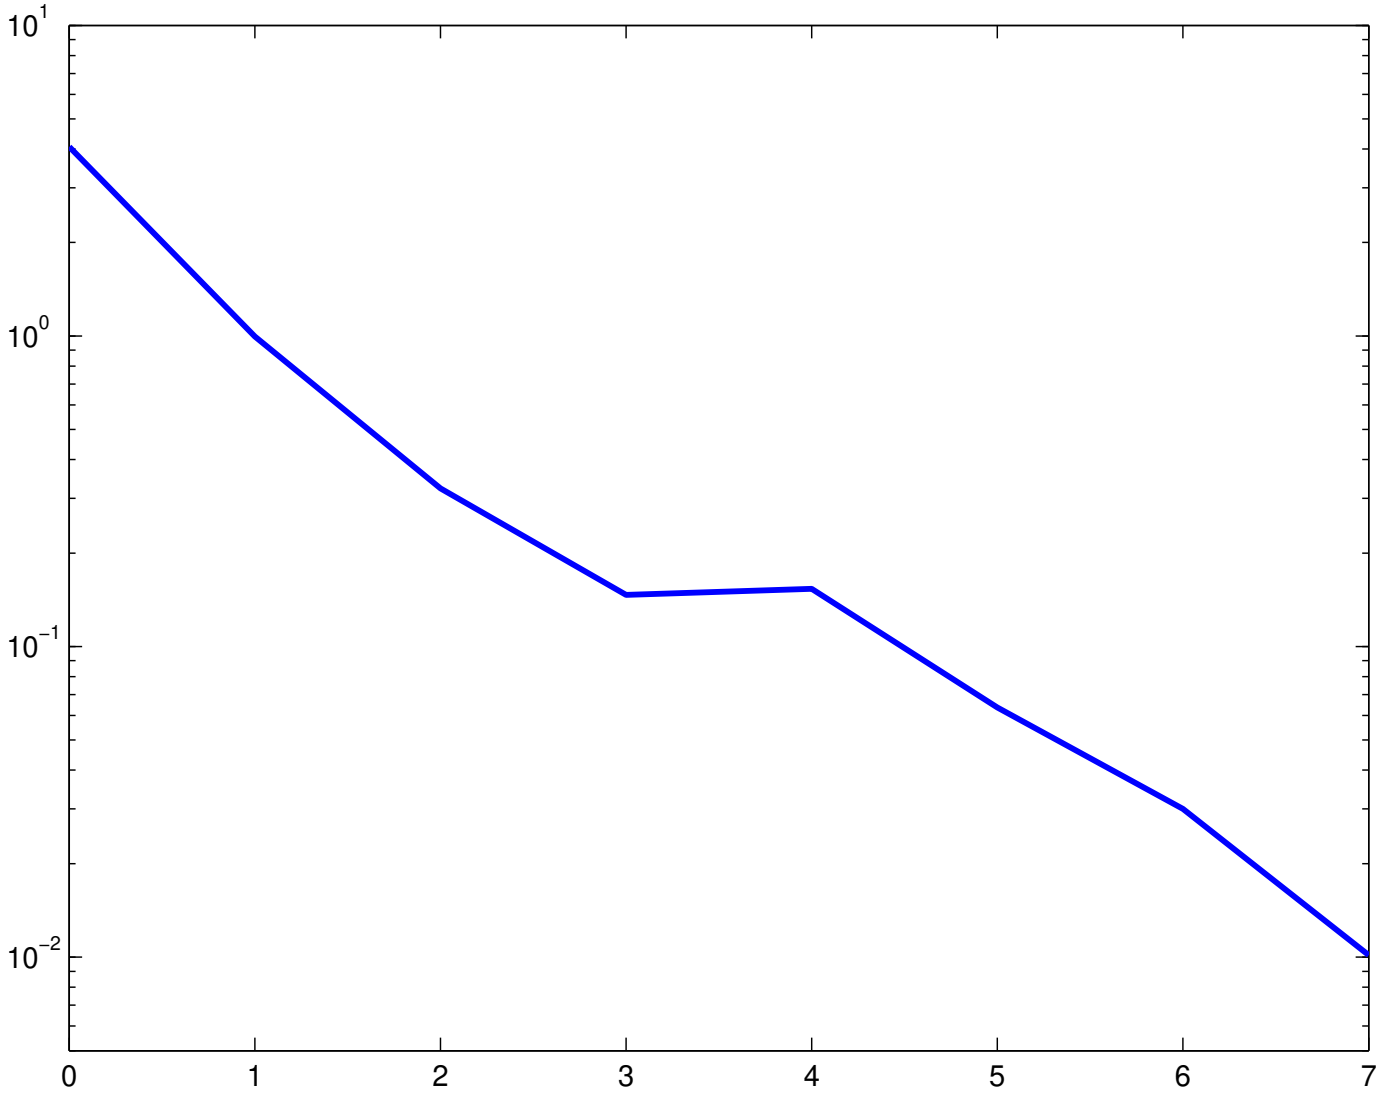
10^{1}
10^{0}
10^{-1}
10^{-2}
0
1
2
3
4
5
6
7

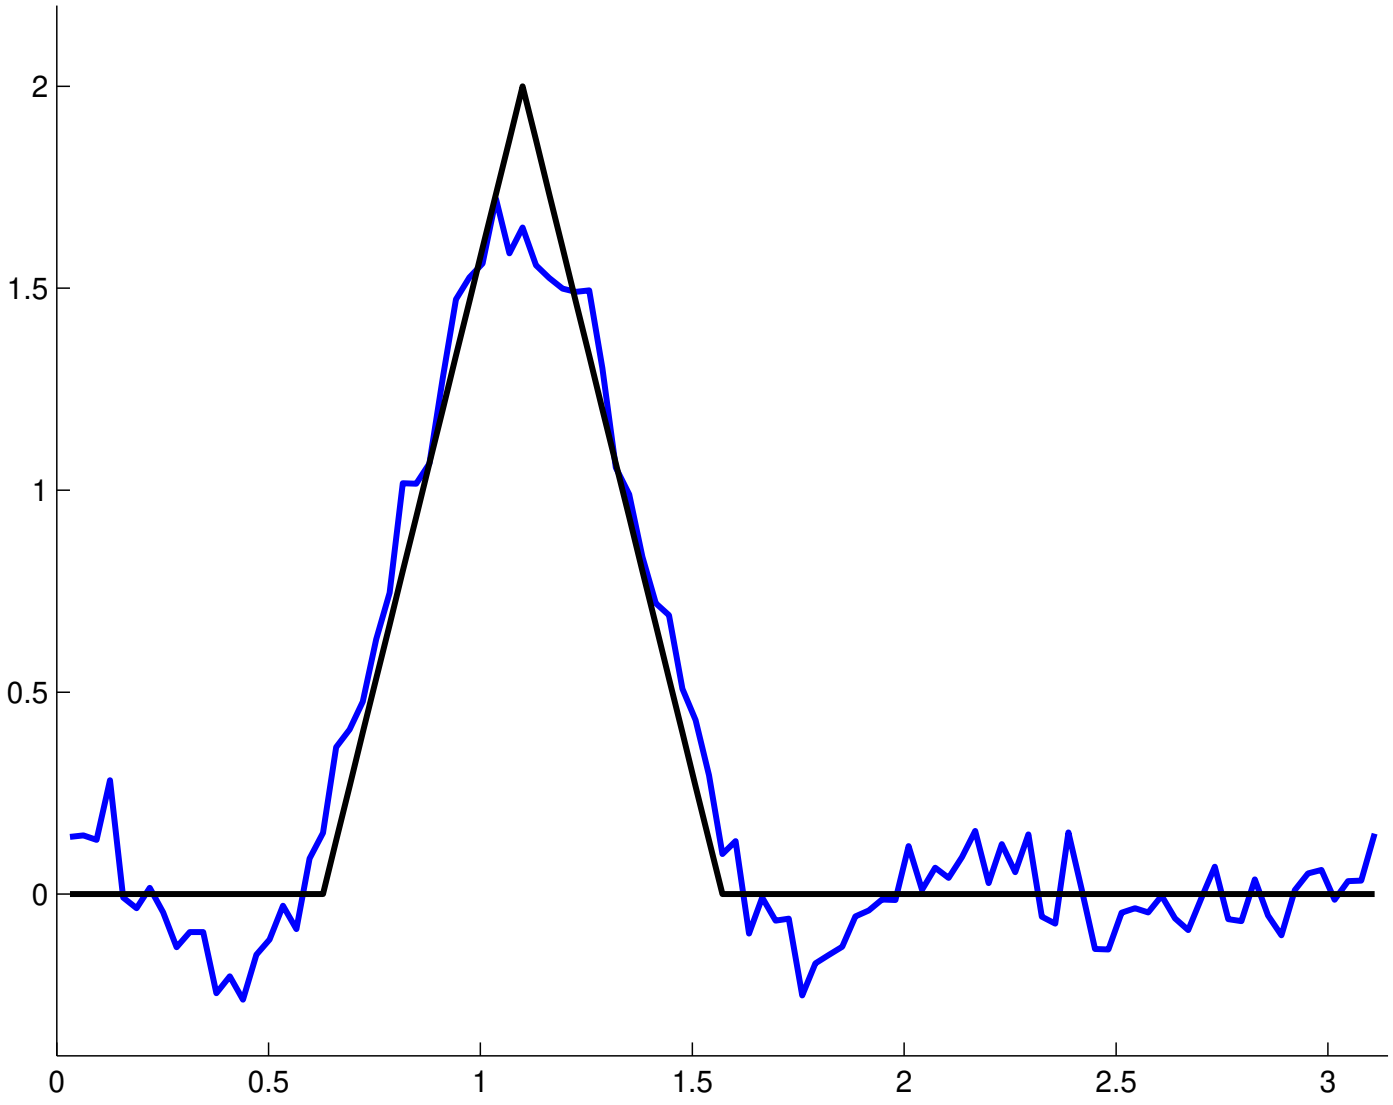
2
1.5
1
0.5
0
0
0.5
1
1.5
2
2.5
3

Next, we consider the exactly same problem, but this time apply the conjugate gradient method to the normal equation. As a stopping rule we use the Morozov discrepancy principle for the original equation, i.e., we stop the iteration when

$$\|y - Ax_k\| \leq \epsilon,$$

where we use the 'standard' $\epsilon = \sqrt{99 \cdot 0.001^2} = 9.95 \cdot 10^{-3}$.

For some reason, the use of the normal equation makes the algorithm more stable: the discrepancy principle for this 'original' ϵ is satisfied after seven iterations and the solution looks nicer than when applying the algorithm directly to the original equation. (Bear in mind, however, that considering the normal equation makes the algorithm slower since more matrix-matrix or matrix-vector products need to be computed.)

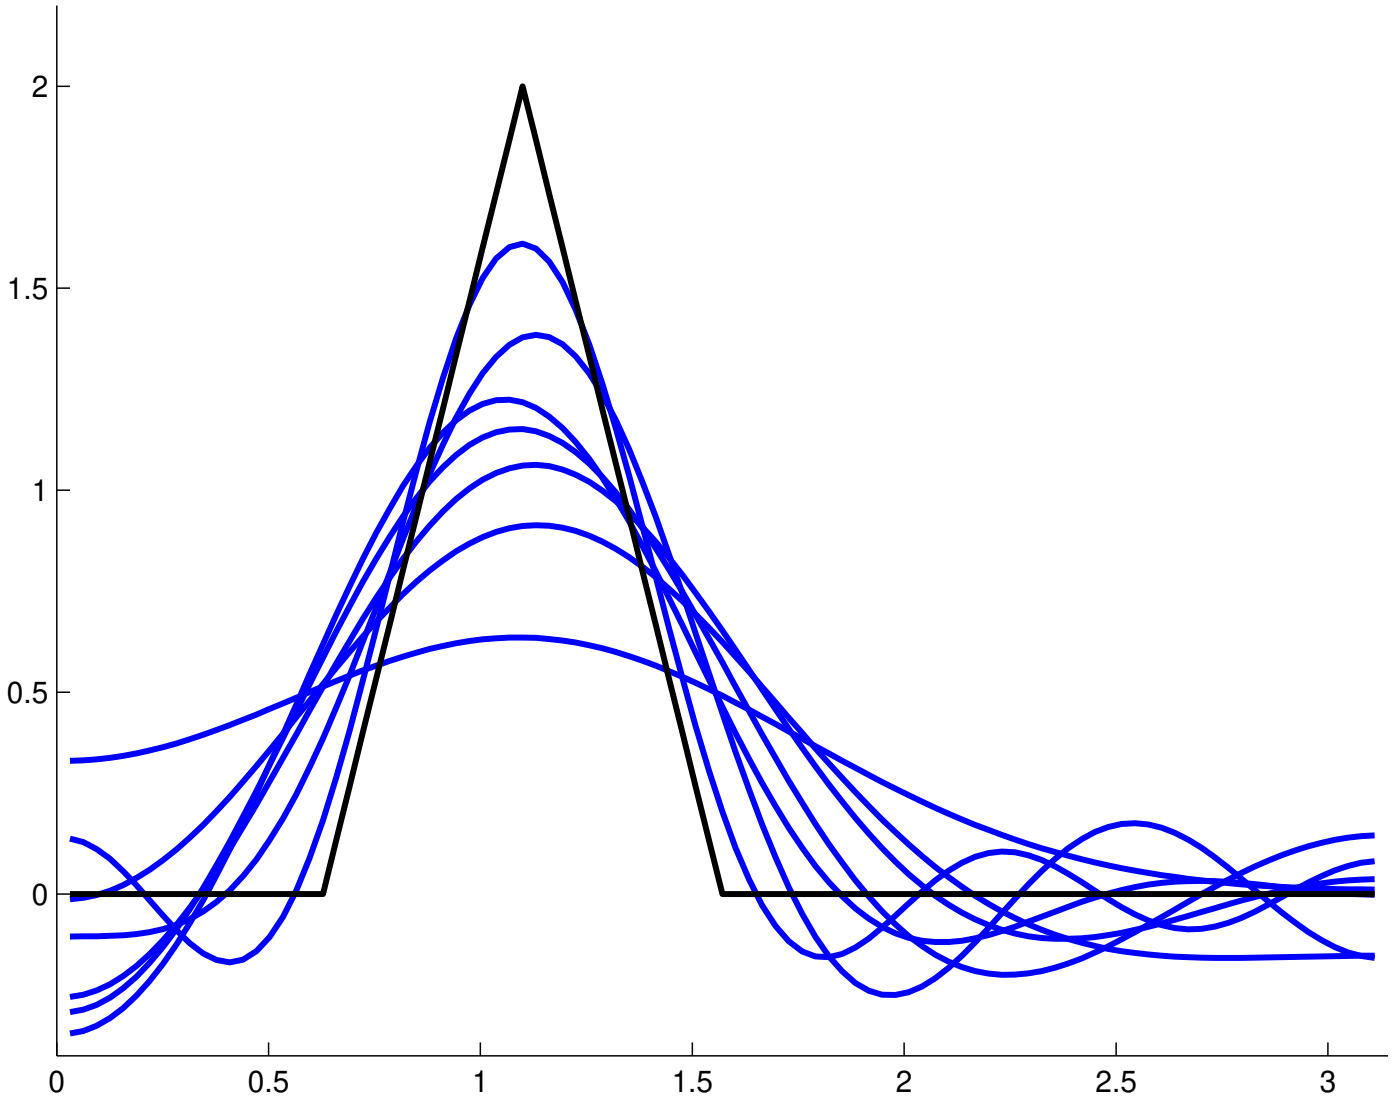
2
1.5
1
0.5
0
0
0.5
1
1.5
2
2.5
3

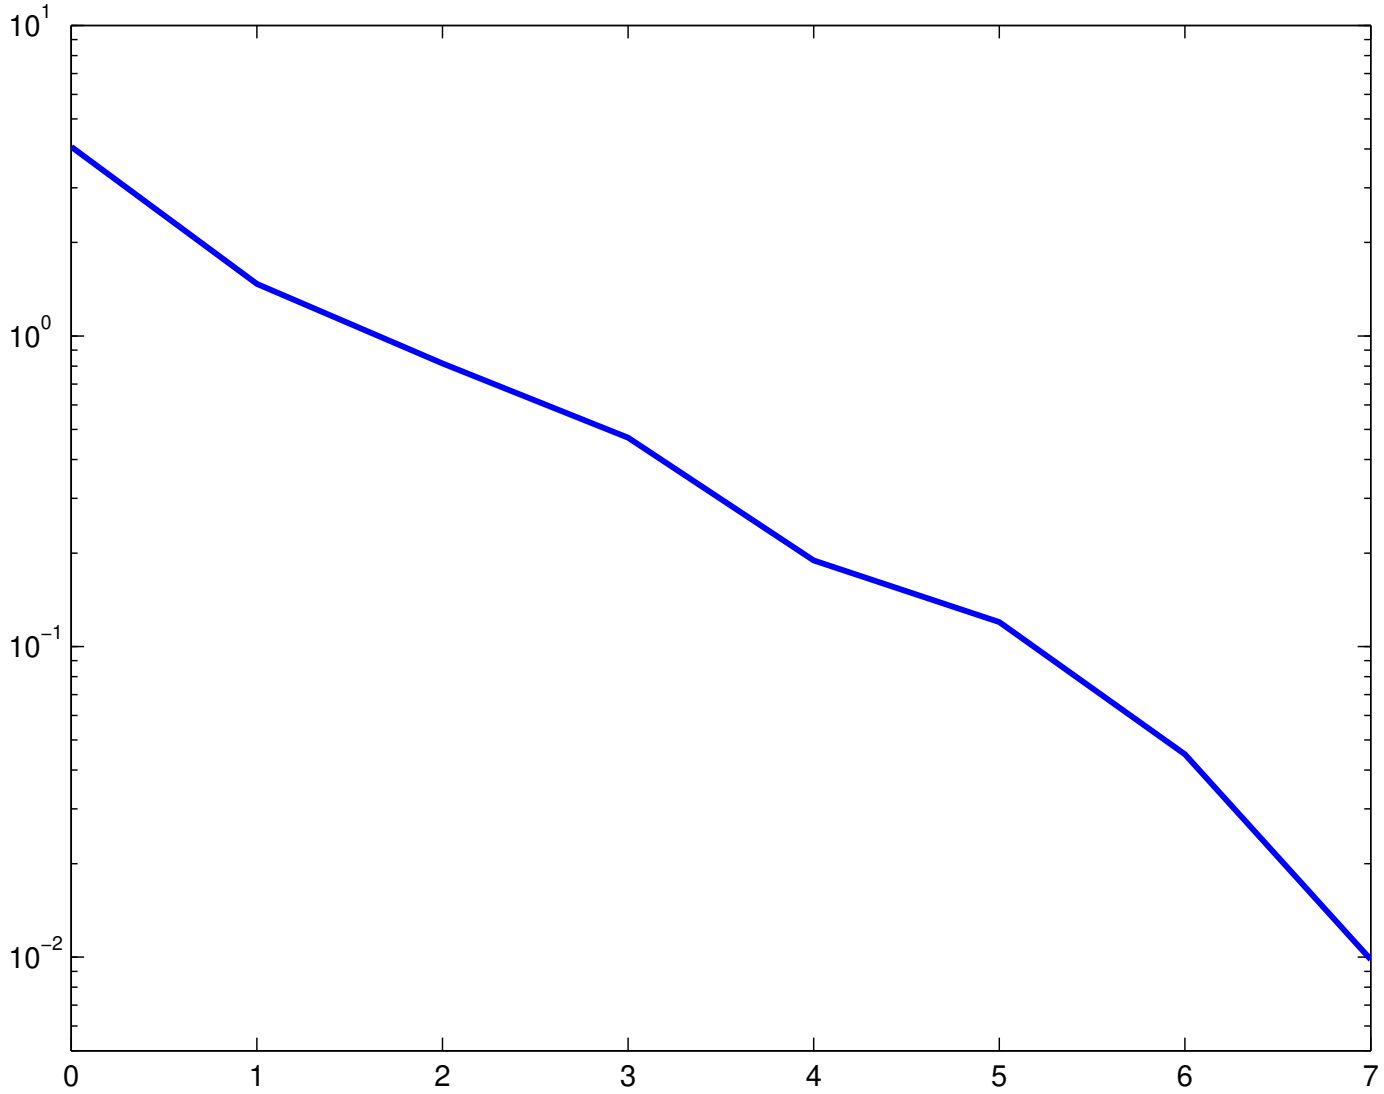
10^{1}
10^{0}
10^{-1}
10^{-2}
0
1
2
3
4
5
6
7

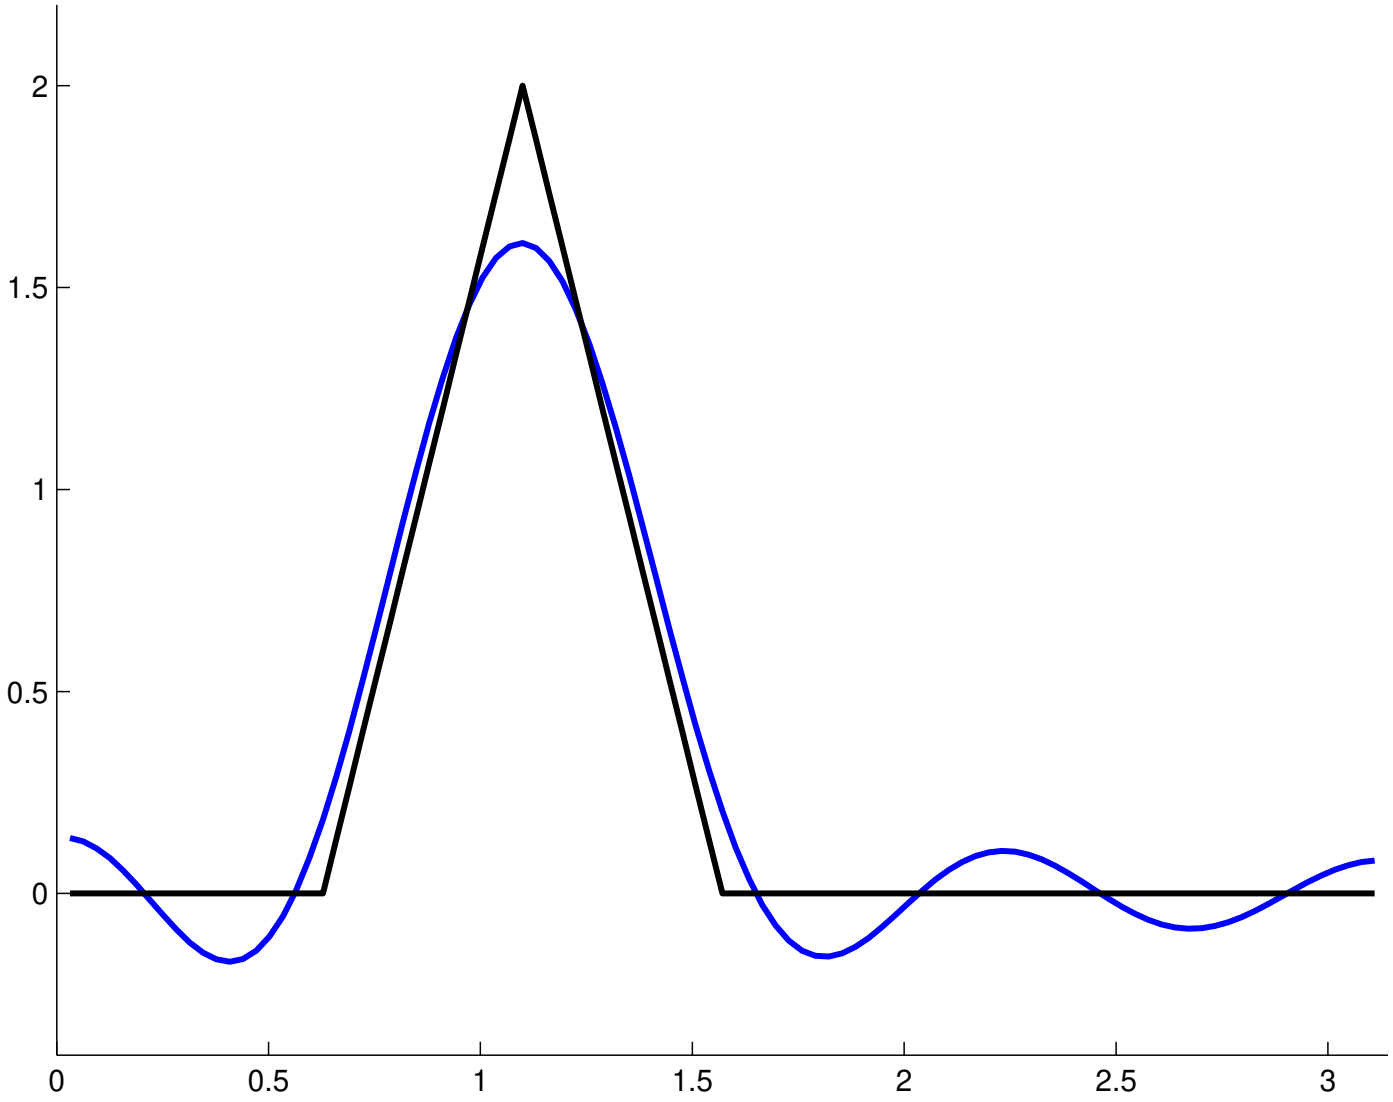
2
1.5
1
0.5
0
0
0.5
1
1.5
2
2.5
3

3. Statistical inversion

Computational inverse problems, part II

The second part of the course concentrates on the Bayesian approach to inverse problems.

The lectures are still mainly based on the books:

- "J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*, Springer, 2005" (parts of Chapter 3),
- "D. Calvetti and E. Somersalo, *Introduction to Bayesian Scientific Computing. Ten Lectures on Subjective Computing*, Springer, 2007".

Statistical inversion

In the statistical approach to inverse problems, the leading idea is to recast the inverse problem in the form of statistical *quest for information*.

- Quantities are either directly observable or unobservable.
- Some of the unobservable quantities are of primary interest, others may be considered to be of secondary interest.
- Quantities depend on each other through models.
- The objective of statistical inversion is to extract information on the unknown quantities of interest based on all available knowledge about the measurements, models coupling the parameters, and information that is available prior to the measurement.

The statistical approach is based on the following principles:

1. All variables are modelled as random variables.
2. The randomness describes our degree of (or lack of) information on their realizations.
3. The information concerning the values of the random variables is coded in probability distributions.
4. The solution of the inverse problem is the *posterior* probability distribution of the quantities of interest (given the measurement).

A classical regularization method typically produces a single estimate, using often a more or less ad hoc removal of the ill-posedness of the problem.

In the statistical framework, the solution is a probability distribution that contains all information on the possible values of the variable of interest. This distribution can be used to obtain different estimates and to evaluate their reliability, e.g., single estimates and credibility intervals. The statistical approach removes the ill-posedness by considering a well-posed extension of the inverse problem in the space of probability densities. When constructing the well-posed extension, the prior beliefs are more explicitly stated than in traditional regularization.

Subjective probability

Example: Tossing a coin.

Assume that the odds of getting heads or tails are equal, i.e.,

$$P(\text{heads}) = P(\text{tails}) = \frac{1}{2}.$$

Such an assumption is generally accepted and can be verified empirically (empirical probability). This example reflects the *frequentist* view, where probability can be seen as the relative frequency of occurrence in a set of repeated experiments.

In connection to Bayesian approach, one sometimes talks about *subjective* probabilities. The inference process commonly incorporates subjective components that reflect the beliefs of, e.g., the person doing the inference (e.g., in the form of prior beliefs about the behaviour of the unknown).

Examples:

What is the probability of rain tomorrow?

What is the probability that Finland will win a gold medal in the next Olympic games?

On random variables and probability densities

Probabilities and events (very informal)

Let Ω contain all possible outcomes, and consider a subset $E \subset \Omega$. For the probability $P(E)$ of an event E, we require

$$0 \leq P(E) \leq 1.$$

Furthermore, it is assumed that

$$P(\Omega) = 1 \qquad \text{and} \qquad P(\emptyset) = 0.$$

Additivity: If $A \cap B = \emptyset$ for $A, B \subset \Omega$, then

$$P(A \cup B) = P(A) + P(B).$$

Two events A and B are called *independent*, if

$$P(A \cap B) = P(A)P(B).$$

The *conditional probability* of A given B is the probability that A happens provided that B happens,

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}.$$

If A and B are mutually independent,

$$P(A \mid B) = P(A), \quad P(B \mid A) = P(B).$$

Real valued random variables (still informal)

We denote random variables by capital letters and their realizations with lower case letters. Let $X : \Omega \to \mathbb{R}$ be a real valued random variable and denote its *probability density* by $\pi(x) = \pi_X(x) \geq 0$.

The probability of the event $x \in B$, $B \subset \mathbb{R}$ is obtained through integration

$$P\{X(\omega) \in B\} = P(X^{-1}(B)) = \int_B \pi(x)dx.$$

In particular,

$$P\{X(\omega) \in \mathbb{R}\} = P(\Omega) = \int_{-\infty}^{\infty} \pi(x)dx = 1.$$

The *expectation* is the center of mass of the probability density

$$\mathrm{E}(X) = \int_{\mathbb{R}} x\pi(x)dx =: \bar{x}.$$

The *variance* is the expectation of the squared deviation from the expectation

$$\mathrm{var}(X) = \sigma_X^2 = \mathrm{E}\{(X - \bar{x})^2\} = \int_{\mathbb{R}} (x - \bar{x})^2 \pi(x)dx.$$

The *joint probability density* $\pi(x, y) = \pi_{X,Y}(x, y)$ of two random variables X and Y is defined via

$$P\{X \in A,\, Y \in B\} = \iint_{A\times B} \pi(x, y)dxdy.$$

The random variables X and Y are *independent* if

$$\pi(x, y) = \pi(x)\pi(y).$$

The *covariance* of X and Y is

$$\mathrm{cov}(X,Y) = \mathrm{E}\{(X-\bar{x})(Y-\bar{y})\}.$$

Note that

$$\mathrm{cov}(X,Y) = \mathrm{E}\{XY\} - \mathrm{E}\{X\}\mathrm{E}\{Y\}.$$

The *correlation coefficient* of X and Y is

$$\mathrm{corr}(X,Y) = \frac{\mathrm{cov}(X,Y)}{\sigma_X \sigma_Y}, \quad \sigma_X = \sqrt{\mathrm{var}(X)},\ \sigma_Y = \sqrt{\mathrm{var}(Y)},$$

or, equivalently, with the help of normalized random variables,

$$\mathrm{corr}(X,Y) = \mathrm{E}\{\tilde{X}\tilde{Y}\}, \quad \tilde{X} = \frac{X-\bar{x}}{\sigma_X}, \quad \tilde{Y} = \frac{Y-\bar{y}}{\sigma_Y}.$$

Random variables are *uncorrelated* if their covariance (or correlation coefficient) vanishes,

$$\mathrm{cov}(X,Y) = 0.$$

If X and Y are independent, they are uncorrelated, since

$$\mathrm{E}\{(X-\bar{x})(Y-\bar{y})\} = \mathrm{E}\{X-\bar{x}\}\mathrm{E}\{Y-\bar{y}\} = 0.$$

On the other hand, uncorrelated random variables are not necessarily independent.

Given two random variables X and Y with joint probability density $\pi(x,y)$, the *marginal density* of X when Y may take any value, is

$$\pi(x) = \int_{\mathbb{R}} \pi(x,y)dy.$$

Analogously, the marginal density of Y is

$$\pi(y) = \int_{\mathbb{R}} \pi(x,y)dx.$$

The *conditional probability density* of X given Y is the probability density of X assuming that $Y = y$:

$$\pi(x \,|\, y) = \frac{\pi(x, y)}{\pi(y)} \qquad \text{if } \pi(y) \neq 0.$$

Note that by the symmetry of the roles of X and Y, we have

$$\pi(x, y) = \pi(x \,|\, y)\pi(y) = \pi(y \,|\, x)\pi(x),$$

which leads to an important identity

$$\pi(x \,|\, y) = \frac{\pi(y \,|\, x)\pi(x)}{\pi(y)},$$

known as the *Bayes* formula.

The *conditional expectation* or the *conditional mean* is the expectation of X given that $Y = y$:

$$\mathrm{E}\{X \mid y\} = \int_{\mathbb{R}} x\pi(x \mid y)dx.$$

The expectation of X can be computed also via its conditional expectation:

$$\begin{aligned}
\mathrm{E}\{X\} &= \int x\pi(x)dx = \int x \left(\int \pi(x, y)dy \right) dx \\
&= \int x \left(\int \pi(x \mid y)\pi(y)dy \right) dx \\
&= \int \left(\int x\pi(x \mid y)dx \right) \pi(y)dy \\
&= \int \mathrm{E}\{X \mid y\}\pi(y)dy.
\end{aligned}$$

Multivariate random variables

A multivariate random variable is a random variable

$$X = \begin{bmatrix} X_1 \\ \vdots \\ X_n \end{bmatrix},$$

where each component X_i is a real scalar valued random variable.

The probability density of X is the joint probability density $\pi_X(x) = \pi(x) = \pi(x_1, \ldots, x_n)$ of its components.

The corresponding expectation is

$$\bar{x} = \int_{\mathbb{R}^n} x\pi(x)dx \ \in \mathbb{R}^n,$$

or, componentwise,

$$\bar{x}_i = \int_{\mathbb{R}^n} x_i \pi(x) dx \ \in \mathbb{R}, \quad 1 \leq i \leq n.$$

The *covariance matrix* is defined as

$$\mathrm{cov}(X) = \int_{\mathbb{R}^n} (x - \bar{x})(x - \bar{x})^{\mathrm{T}} \pi(x) dx \ \in \mathbb{R}^{n \times n},$$

or, componentwise,

$$\mathrm{cov}(X)_{ij} = \int_{\mathbb{R}^n} (x_i - \bar{x}_i)(x_j - \bar{x}_j)^{\mathrm{T}} \pi(x) dx \ \in \mathbb{R}, \quad 1 \leq i, j \leq n.$$

The covariance matrix is symmetric and positive semi-definite.

The symmetry is implicit in the definition of the covariance matrix, whereas the positive semi-definiteness follows by writing for $v \in \mathbb{R}^n$ that

$$\begin{aligned} v^{\mathrm{T}}\mathrm{cov}(X)v &= \int_{\mathbb{R}^n} [v^{\mathrm{T}}(x-\bar{x})][(x-\bar{x})^{\mathrm{T}}v]\pi(x)dx \\ &= \int_{\mathbb{R}^n} (v^{\mathrm{T}}(x-\bar{x}))^2\pi(x)dx \geq 0. \end{aligned}$$

Note that the above expression measures the variance of X in the direction v.

The diagonal entries of the covariance matrix are the variances of the individual components of X. Indeed, let us denote by $x_i' \in \mathbb{R}^{n-1}$ the vector x with the ith component deleted, i.e.,

$$x_i' = [x_1, x_2, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n]^{\mathrm{T}}.$$

Then, we have

$$\begin{aligned}
\mathrm{cov}(X)_{ii} &= \int_{\mathbb{R}^n} (x_i - \bar{x}_i)^2 \pi(x) dx \\
&= \int_{\mathbb{R}} (x_i - \bar{x}_i)^2 \left(\int_{\mathbb{R}^{n-1}} \pi(x_i, x_i') dx_i' \right) dx_i \\
&= \int_{\mathbb{R}} (x_i - \bar{x}_i)^2 \pi(x_i) dx_i \\
&= \mathrm{var}(X_i).
\end{aligned}$$

The marginal and conditional probabilities for multivariate random variables are defined by the same formulas as for the univariate random variables.