Inverse Problems — Exercise 9

Sommersemester 2023

Return your written solutions either in person or by email
to vesa.kaarnioja@fu-berlin.de by Tuesday 4 July, 2023, 10:15

Please note that there are a total of 4 tasks in this exercise sheet.

1. Let $y \in \mathbb{R}^2$ and $x \in \mathbb{R}$ and

$$y = \begin{pmatrix} 2 \\ 1 \end{pmatrix} x + \eta, \quad \eta \sim \mathcal{N}(0, \gamma^2 I_2),$$

   where $I_2 \in \mathbb{R}^{2\times 2}$ is an identity matrix. Suppose that the prior distribution is given by $x \sim \mathcal{N}(0, 2)$, with x and η assumed independent. What is the posterior distribution if we observe $y = (1, 2)^{\mathrm{T}}$? What is the posterior variance? What happens to posterior distribution and variance under decreasing noise ($\gamma \downarrow 0$)?

2. Let us consider the high-dimensional integral

$$I_d := \int_{[0,1]^d} \cos\left(2\pi + \sum_{i=1}^{d} x_i\right) \mathrm{d}x_1 \cdots \mathrm{d}x_d.$$

   Estimate the value of this integral by implementing a Monte Carlo sampler in your favorite programming language.

   In this case, the exact value of this integral is $I_d = 2^d \cos\left(2\pi + \frac{d}{2}\right) \sin(\frac{1}{2})^d$ (you do not need to prove this). Compute the Monte Carlo integration error for sample sizes $n = 2^k$, $k = 0, 1, 2, \ldots, 20$. Try out several values for the dimension d, for example, $d = 10, 100, 1\,000$. What convergence rate do you observe for the error as a function of n? Does increasing the dimension d affect the convergence rate?
   MATLAB users: `rand(m,n)` produces an $m \times n$ array containing uniformly distributed random numbers between 0 and 1.
   Python users: the numpy library contains the function
   `numpy.random.uniform(low=0.0, high=1.0, size=(m,n))`
   which can be used to produce an $m \times n$ array containing uniformly distributed random numbers between 0 and 1.

3. Suppose we are given the posterior distribution

$$\pi^y(x) = \frac{1}{Z} g(x, y)\pi(x),$$

   where $x, y \in \mathbb{R}^2$, we have the prior density $\pi(x) = \frac{1}{2\pi} \exp(-\frac{1}{2}(x_1^2 + x_2^2))$, and

$$g(x, y) = \exp(-|x_1 - y_1^2| - |x_2^2 - y_2|).$$

   Here, $Z = \int_{\mathbb{R}^2} g(x, y)\pi(x)\,\mathrm{d}x$.

   Suppose we are given the observation $y = (3, 2)^{\mathrm{T}}$. Use importance sampling to estimate the posterior mean.

4. Consider the Bayesian inverse problem

$$y_j = F(x) + \eta_j, \quad j \in \{1, \dots, N\}, \tag{1}$$

where the unknown quantity $x \in \mathbb{R}^d$ is assumed to remain static, $F \colon \mathbb{R}^d \to \mathbb{R}^k$ is a function, we assume Gaussian observational noise $\eta_j \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \gamma^2 I)$, $\gamma > 0$, and we assume to have N independent observations $y_1, \dots, y_N \in \mathbb{R}^k$.

(a) What is the likelihood density $\mathbb{P}(y_1, \dots, y_N | x)$?

(b) Show that

$$\sum_{j=1}^{N} \|y_j - F(x)\|^2 = N\|F(x) - \overline{y}\|^2 + N\left(\frac{1}{N}\sum_{j=1}^{N} y_j^{\mathrm{T}} y_j - \overline{y}^{\mathrm{T}}\overline{y}\right),$$

where $\overline{y} = \frac{1}{N}\sum_{j=1}^{N} y_j$.

(c) Use parts (a) and (b) to deduce that the problem (1) is equivalent to the Bayesian inverse problem

$$\overline{y} = F(x) + \eta, \quad \eta \sim \mathcal{N}\left(0, \frac{\gamma^2}{N} I\right).$$

Interpretation: averaging a number of independent measurements of a static target results in variance reduction of the noise.