Exercise session 5
Bayesian estimation, array signal processing, and Kalman filtering

ELEC-E5440 Statistical Signal Processing

Robin Rajamäki

December 2, 2020

1 General homework tips

1. MS and MAP estimators
   - We observe a (single) measurement of the form $Y = X + N$
   - What does the distribution of the measurement $f_Y(y)$ look like[1]?
   - Try visualizing the likelihood $f_{Y|X=x}(y)$ for different values of x
   - Both $f_Y(y)$ and $f_{Y|X=x}(y)$ have very simple expressions
   - If you're stuck, try the above steps assuming X is deterministic. For example, try $X = c$, for some constant $c \in \mathbb{R}$
   - Remember that a probability density function should integrate to 1
2. MS and MAP estimators
   - Straightforward, albeit slightly more cumbersome computations
3. Direction-of-arrival estimation using real-world data
   - Consider the conceptual difference between LS and TLS
   - Why is TLS-ESPRIT reasonable in our case?
   - Compare at least the DoA estimates of MUSIC and TSL-ESPRIT
   - Can you say something about the sensitivity of either method to the (model/algorithm specific) assumptions?
4. Target tracking using Kalman filter
   - See link in previous exercise handout

[1] Nice to know: In general, the distribution of Y is the convolution of the marginal distributions of X and N, for arbitrarily distributed but independent X and N. However, you do not need to evaluate the convolution integral in this exercise (but of course you may).

2 More examples of Bayesian estimation

See Example 11.2 on p. 351 in Kay's book [1] for an intuitive understanding of the interplay between the prior, likelihood, and posterior. For convenience, we solve a similar problem, namely Exercise 11.4 in [1, p. 370], below.

11.4 The data $x[n] = A + w[n]$ for $n = 0, 1, \ldots, N-1$ are observed. The unknown parameter A is assumed to have the prior PDF

$$p(A) = \begin{cases} \lambda \exp(-\lambda A) & A \geqslant 0 \\ 0 & A < 0 \end{cases}$$

where $\lambda > 0$, and $w[n]$ is WGN with variance σ^2 and is independent of A. Find the MAP estimator of A.

Solution: Since the noise is Gaussian, its PDF is

$$p(w[n]) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\Big(-\frac{w^2[n]}{2\sigma^2} \Big).$$

Consequently, the conditional probability of $x[n]$ given A is also a Gaussian process with mean A, i.e.,

$$p(x[n]|A) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x[n]-A)^2}{2\sigma^2} \right)$$

For simplicity, collect the observations into vector $\boldsymbol{x} = [x[0], x[1], \ldots, x[n-1]]^{\mathrm{T}}$. The likelihood of the i.i.d. observations is therefore the conditional PDF

$$p(\boldsymbol{x}|A) = \prod_{n=0}^{N-1} p(x[n]|A) = \frac{1}{\sqrt{2\pi\sigma^2}^N} \exp\left(-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n]-A)^2 \right).$$

The MAP estimator of A maximizes the posterior, that is,

$$\hat{A}_{\mathrm{MAP}} = \arg\max_A p(A|\boldsymbol{x}) = \arg\max_A p(\boldsymbol{x}|A)p(A) = \arg\max_A \big(\log p(\boldsymbol{x}|A) + \log p(A) \big).$$

For the given $p(A)$ and $p(\boldsymbol{x}|A)$, the MAP becomes (ignoring irrelevant constants)

$$\hat{A}_{\mathrm{MAP}} = \arg\max_A \underbrace{\left(-\frac{1}{2\sigma^2} \sum_{n=0}^{N-1} (x[n]-A)^2 - \lambda A \right)}_{f(A)}.$$

Setting the derivative of $f(A)$ w.r.t. A equal to zero yields

$$\frac{df(A)}{dA} = \frac{1}{\sigma^2} \sum_{n=0}^{N-1} (x[n]-A) - \lambda = 0 \iff A = \bar{x} - \frac{\lambda\sigma^2}{N},$$

where $\bar{x} = \frac{1}{N}\sum_{n=0}^{N-1} x[n]$ denotes the sample mean. Indeed, the above argument maximizes $f(A)$, since

$$\frac{d^2 f(A)}{dA^2} = -\frac{N}{\sigma^2} - \lambda < 0.$$

However, we also need to take into account[2] that $A \geq 0$, since it may occur that $\bar{x} < \lambda\sigma^2/N$. Consequently, the MAP estimator is

$$\hat{A}_{\text{MAP}} = \max\left(0, \bar{x} - \frac{\lambda\sigma^2}{N}\right).$$

Fig. 1 illustrates the MAP estimate, together with the prior and posterior PDFs, as well as the likelihood function, for the case $N = 1$ and $x = 2.0558$. We see that the prior biases the MAP estimate towards smaller values of A compared to the mode of the likelihood function, i.e., the maximum likelihood estimate (MLE). Note that the MLE is the sample mean, which equals the measurement x in this single observation case.

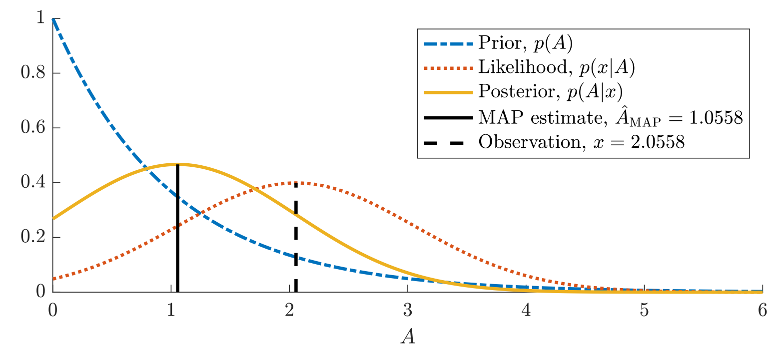

Figure 1: MAP estimate together with prior, likelihood, and posterior PDFs.

References

[1] S. M. Kay, *Fundamentals of statistical signal processing: Estimation theory*. Prentice Hall PTR, 1993.

[2] Since we previously solved an *unconstrained* optimization problem, for simplicity.