CK0146: Exercise 03a (2017.1)

Exercise Q.1 (Bernoullis, betas, and binomials). Let the table

m	0	1	2	3	4	5
Occurances	6	10	14	13	6	1

summarise a sample of size 50 from a binomial distribution with $N = 5$ and unknown parameter μ

$$m \sim \text{Bin}(m|N, \mu) = \binom{N}{m} \mu^m (1-\mu)^{N-m}.$$

1) Plot the binomial distribution estimated by using frequency histograms.
2) Calculate the maximum likelihood estimate of μ and plot the resulting MLE distribution.
3) What is the maximum likelihood estimate of $P(m \geq 3)$?

A convenient choice for the prior distribution of μ is a beta distribution with parameters $a = b = 2$,

$$\mu \sim \text{Beta}(\mu|a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \mu^{a-1} (1-\mu)^{b-1}.$$

4) Plot the prior distribution of μ.
5) Write and plot the posterior distribution of μ.
6) What is the expected posterior of μ?

Exercise Q.2 (MLE and MAP of model parameters). We have the location $\mathbf{x}_i \in \mathcal{R}^2$ of N artificial noses. The noses are used to smell/detect the presence of hazardous substances. We want to use knowledge about the position of the noses ($\mathbf{x}_i$) and data on whether they detected an odour or not ($y_i \in \{0, 1\}$) to locate the emission's source. The data $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ are available here.

Let the unknown location of the source be $\mathbf{x} \in \mathcal{R}^2$. Each nose as a probability of detecting an odour that is dependent on its distance $\delta(\mathbf{x}_i, \mathbf{x})$ with the unknown source. Let the probability of the i-th nose to detect/miss (1/0) the odour be X_i, such that $Y_i \sim \text{Bern}(\mu_i)$ with $\mu_i = \exp(-\delta(\mathbf{x}_i, \mathbf{x}))$

$$p(Y_i|\mu_i) = \mu_i^{y_i}(1-\mu_i)^{1-y_i}, \quad y_i \in \{0, 1\}.$$

We encode our belief on the location of the source by modelling its position as a random variable X, such that $X \sim \mathcal{N}(\mathbf{g}, \sigma^2 \mathbf{I}_2)$, for some σ and some $\mathbf{g}$ between the noses who detected the odour[1].

Assuming that the X_i are independent, formally estimate $\mathbf{x}$ using

a) Maximum likelihood: $\mathbf{x}_{MLE} = \text{argmax}\ (p(\mathcal{D}|\mathbf{x}))$, $p(\mathcal{D}|\mathbf{x}) = \prod_{i=1}^N p(y_i|\mu_i)$
b) Maximum a posteriori: $\mathbf{x}_{MAP} = \text{argmax}\ (p(\mathbf{x}|\mathcal{D}))$, $p(\mathbf{x}|\mathcal{D}) \propto p(X|\mathbf{g}, \sigma^2 \mathbf{I}_2) \prod_{i=1}^N p(y_i|\mu_i)$

[1] You are free to defined your own strategy to choose $\mathbf{g} \in \mathcal{R}^2$ and σ.

Verify the solutions by plotting a discrete hypothesis space of $\mathbf{x}$, as depicted in the figure below.[2]

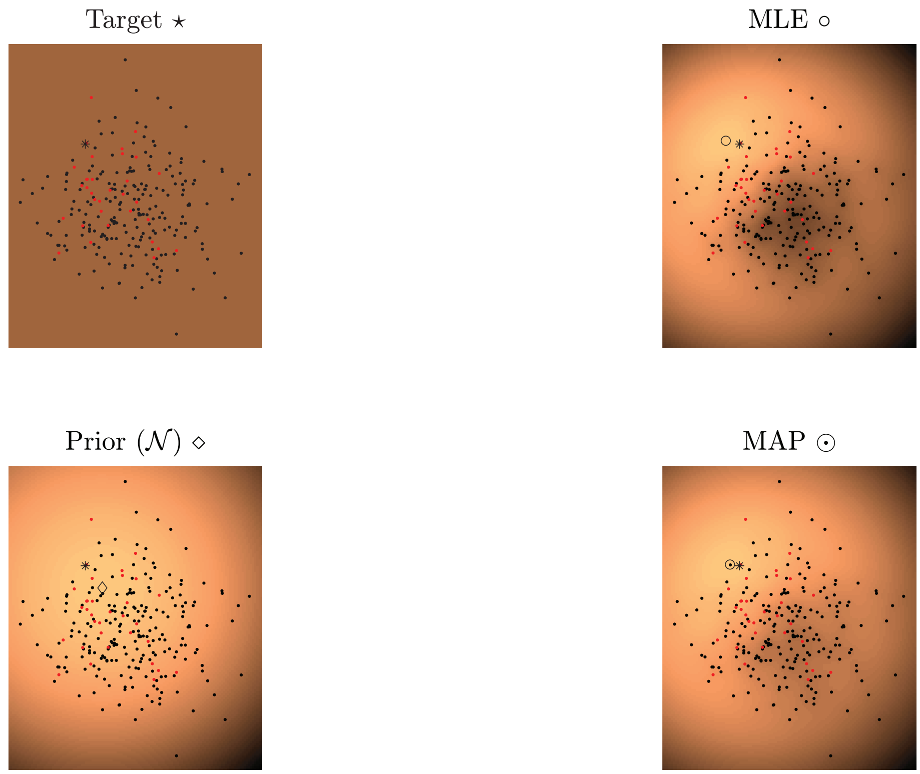

[2] Well, except that you do not know where the target $\star$ is actually located.