HW#3

1. Bishop 6.4

2. Bishop 6.11

3. Bishop 6.15

4. Bishop 6.18

5. Gaussian processes for regression

Generate $N = 20$ datapoints (equidistant or randomly chosen) from

$$f(x) = \sin(x) + \xi \quad , \quad x \in [0, 2\pi]$$

Here, $\xi = \mathcal{N}(\underbrace{0}_{\mu}, \underbrace{0.2^2}_{\sigma^2})$ is the noise distribution

Fit this data using the gaussian kernel:

$$k(x_n, x_m) = e^{-\frac{\theta}{2}(x_n - x_m)^2}$$

and the exponential kernel:

$$k(x_n, x_m) = e^{-\theta |x_n - x_m|}.$$

In both cases, find and report the optimal value of θ by maximizing $\log p(\vec{t}|\theta)$ wrt θ [report both θ_{opt} and $\log p(\vec{t}|\theta_{opt})$]

You may use a grid search or a gradient-based technique, but be careful with multiple local maxima.

In both cases, find the mean μ and σ of the predictive distribution with $\theta = \theta_{opt}$, and plot μ and $\mu \pm \sigma$ alongside $\sin(x)$ and the datapoints in the $[0, 2\pi]$ range.