HW # 1

(1.) Bishop 3.4

(2.) Bishop 3.7

(3.) ML and Bayesian curve fitting

Consider $y(x) = a_0 + a_1 x^2$ with

$a_0 = 1$, $a_1 = 2$ and $x \in [-3, 3]$

generate $N = 100$ datapoints by:

(i) randomly sampling x in the $[-3,3]$ range using a uniform distribution

(ii) computing $y(x)$

(iii) computing $t = y(x) + \xi$, where ξ is a random variable sampled from

$\mathcal{N}(\xi | 0, \underbrace{0.01}_{\sigma^2})$. [i.e. $\sigma = 0.1$] **Try sigma=0.5 also**

Consider a linear model of the form

$y(x, \vec{w}) = w_0 + w_1 x + w_2 x^2$.

(a) Find the ML weights and plot $y(x, \vec{w}_{ML})$ alongside $y(x)$ [report $\vec{w}_{ML}$ as well]

(b) Find β_{ML} using $\vec{w}_{ML}$.

Use β_{ML} and $\alpha = 1.0$ to compute the predictive distribution in the $x \in [-3, 3]$ range. Plot the mean of the predictive distribution alongside with $\pm \sigma_N(x)$ curves [cf. Fig. 3.8] and $y(x)$, the "true" curve.

Draw $K=10$ samples from the posterior distribution for $\vec{w}$ and plot the corresponding $y(x, \vec{w})$ curves, alongside with $y(x)$ [cf. Fig. 3.9].