HW #2

(1.) Bishop 4.9

(2.) Bishop 4.10

(3.) Bishop 4.15

(4.) Linear models for classification

Generate $N_1 = 150$ datapoints from $\mathcal{N}(\vec{x} \mid \vec{\mu}_1, \Sigma)$ and $N_2 = 300$ datapoints from $\mathcal{N}(\vec{x} \mid \vec{\mu}_2, \Sigma)$, where

$$\begin{cases} \vec{\mu}_1 = \overline{2,0} \\ \vec{\mu}_2 = \overline{-2,0} \end{cases} \qquad \Sigma = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}, \quad \text{where } \sigma = 3$$

Assume that the class-conditional densities are given by

$$\begin{cases} p(\vec{x} \mid C_1) = \mathcal{N}(\vec{x} \mid \vec{\mu}_1, \Sigma), \\ p(\vec{x} \mid C_2) = \mathcal{N}(\vec{x} \mid \vec{\mu}_2, \Sigma) \end{cases}$$

$\vec{x}$ ← 2D observation vector

(a) Using $\vec{\mu}_1, \vec{\mu}_2, \Sigma, N_1, N_2$ from above, find the exact decision boundary (DB) at which $p(C_1 \mid \vec{x}) = p(C_2 \mid \vec{x}) = 0.5$.

- Plot $p(\vec{x}|C_1)$, $p(\vec{x}|C_2)$ and the exact DB as a heatmap or a contour map.
- plot $p(C_1|\vec{x})$, $p(C_2|\vec{x})$ and the exact DB as a heatmap or a contour map.

(b) Estimate $p(C_1|\mathcal{D})$, $p(C_2|\mathcal{D})$, $\vec{\mu}_1$, $\vec{\mu}_2$, Σ by ML, find the DB using these estimated values. Add the ML DB to the plots in (a).

(c) Find the DB by logistic regression and add it to the plots in (a).

- Plot $p(C_1|\vec{x})$ & $p(C_2|\vec{x})$ ← as a heatmap or a contour plot, along with the logistic regression DB, and compare these plots qualitatively with the exact posterior probabilities from part (a).

Discuss how close the 3 DBs are to each other.