HW #3

1. Ex. 5.2

2. Ex. 6.2

3. Logistic regression

generate $N_1 = 250$ datapoints from $\mathcal{N}(\underbrace{\vec{x}}_{\text{2-D vector}} | \vec{\mu}_1, \Sigma)$ (class 1) and $N_2 = 250$ datapoints from $\mathcal{N}(\vec{x} | \vec{\mu}_2, \Sigma)$ (class 2), where

$$\begin{cases} \vec{\mu}_1 = \begin{pmatrix} 2 \\ 0 \end{pmatrix}, \\ \vec{\mu}_2 = \begin{pmatrix} -2 \\ 0 \end{pmatrix} \end{cases} \qquad \Sigma = \begin{pmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}, \quad \sigma = 2.5.$$

Divide the data into training and test sets as you see fit.

(a) Using MAP logistic regression, plot training & test classification accuracy as a function of λ (use log-scale). What is the best value of $\lambda \Rightarrow \lambda^*$?

(b) Using λ^* from part (a), draw the logistic regression DB on a 2D contour or heatmap plot of $\underbrace{p(c=1|\vec{x})}_{\text{"}\mathcal{N}(\vec{x}|\vec{\mu}_1, \Sigma)}$ & $\underbrace{p(c=2|\vec{x})}_{\text{"}\mathcal{N}(\vec{x}|\vec{\mu}_2, \Sigma)}$.

Briefly discuss the location and the orientation of the DB.