ML2025 Fall Homework Assignment 5 Handwritten

Yu-Cheng Lin
b11901152@ntu.edu.tw

Yi-Chen Lee
b12901024@ntu.edu.tw

November 2025

Tools You Need to Know/Learn

- Markov decision process, Q value function and value funciton
- Stationary and Deterministic Policy
- SVM
- Adaboost
- EM algorithm

You are expected to know (or learn through this homework) the definitions and usages of the above concepts, and are encouraged to discuss with TAs during TA hour if you encounter difficulties.

Homework Policy

- The homework is graded out of 100 points, with up to 15 additional bonus points available (only awarded for completely correct solutions).
- The official submission deadline will follow the schedule announced on NTU COOL.
- Homework may be handwritten or typed (e.g. using LaTeX), but must be submitted in **PDF format**.
- If you discuss the homework with classmates, you should state their student IDs in your submission.
- Plagiarism is strictly prohibited. Serious violations will be dealt with according to NTU regulations.

Problem 1 (Trace Optimization)(20 pts)

1. (10 pts) Let $\Sigma \in R^{m\times m}$ be a symmetric positive semi-definite matrix, $\mu \in R^m$. Please construct a set of points $x_1, ..., x_n \in R^m$ such that

$$\frac{1}{n}\sum_{i=1}^{n}(x_i-\mu)(x_i-\mu)^T = \Sigma, \quad \frac{1}{n}\sum_{i=1}^{n} x_i = \mu$$

   Hint: n is not given by the problem. WLOG, you may assume $\mu = 0 \in \mathbb{R}^d$.

2. (10 pts) Let $1 \le k \le m$, solve the following optimization problem and justify with proof:

$$\begin{array}{ll} \text{minimize} & Trace(\Phi^T \Sigma \Phi) \\ \text{subject to} & \Phi^T \Phi = I_k \\ \text{variables} & \Phi \in R^{m\times k} \end{array}$$

   In other words, you need to find Φ and verify that your Φ minimize the trace.

Problem 2 (Gradient Boosting)(20 pts)

Consider the binary classification problem, where we are given training data set $\{(x_i, y_i)\}_{i=1}^N$ with $x_i \in \mathbb{R}^d$ and $y_i \in \{1, -1\}$. Let $F = \{f \mid f : \mathbb{R}^d \to \{1, -1\}\}$ be the collection of classifiers. Given number of epochs $T \in \mathbb{N}$. Suppose that we want to find the function

$$g(x) = \sum_{t=1}^{N} \alpha_t f_t(x)$$

where $f_t \in F$ and $\alpha_t \in \mathbb{R}$ for all $t = 1, \ldots, T$, by which the aggregated classifier is given by

$$h(x) = \begin{cases} 1, & \text{if } g(x) > 0 \\ -1, & \text{if } g(x) \leq 0. \end{cases}$$

Please apply gradient boosting to show how the functions f_t and the coefficients α_t are computed with an aim to minimize the following loss function

$$L(g) = \sum_{i=1}^{N} \log\left(1 + e^{-y_i g(x_i)}\right).$$

Problem 3 (EM algorithm for mixture of exponential model)(20 pts)

Given N samples $x_1, \ldots, x_N \in [0, \infty)$, we would like to cluster them into K clusters. Assume the samples are generated according to Exponential mixture models

$$X \sim \sum_{j=1}^{K} \pi_j Exp(\tau_j)$$

where $\pi_1 + \ldots + \pi_K = 1$, and $Exp(\tau)$ denotes the exponential distribution with probability density function

$$f_\tau(x) = \begin{cases} (1/\tau)e^{-x/\tau} & , x \geq 0 \\ 0 & , x < 0. \end{cases}$$

We would like to apply Expectation Maximization algorithm to find the maximum likelihood estimation of parameters $\theta = \{(\pi_k, \tau_k)\}_{k=1}^{K}$.

(a) (10 pts) Please write down the E-step and M-step and show that the parameters are updated from $\theta^{(t)} = \left\{(\pi_k^{(t)}, \tau_k^{(t)})\right\}_{k=1}^{K}$ to $\theta^{(t+1)} = \left\{(\pi_k^{(t+1)}, \tau_k^{(t+1)})\right\}_{k=1}^{K}$ in the following form:

$$\tau_k^{(t+1)} = \frac{\sum_{i=1}^{N} \delta_{ik}^{(t)} x_i}{\sum_{i=1}^{N} \delta_{ik}^{(t)}}, \quad \pi_k^{(t+1)} = \frac{1}{N} \sum_{i=1}^{N} \delta_{ik}^{(t)}$$

(b) (10 pts) What is the closed form expression of $\delta_{ik}^{(t)}$?

Problem 4 (Sparse SVM)(20 pts)

Given training data of N input-output pairs $\mathscr{D} = ((x_i, y_i))_{i=1}^N$, where $x_i \in \mathcal{X}$ and $y_i \in \{\pm 1\}$. One can give two types of arguments in favor of the SVM algorithm: one based on the sparsity of the support vectors, another based on the notion of margin. Suppose instead of maximizing the margin, we choose instead to maximize sparsity by minimizing the p-norm of the vector $\alpha = (\alpha_1, ..., \alpha_N)$ that defines the weight vector $\mathbf{w}$, for some $p \geq 1$. In this question we consider the case $p = 2$, which leads to the following optimization problem:

$$\begin{array}{ll} \text{minimize} & f(\alpha, b, \boldsymbol{\xi}) = \frac{1}{2}\sum_{i=1}^N \alpha_i^2 + \sum_{i=1}^N C_i \xi_i \\ \text{subject to} & y_i\left(\sum_{j=1}^N \alpha_j y_j \mathbf{x}_i \cdot \mathbf{x}_j + b\right) \geq 1 - \xi_i,\ i \in \{1, ..., N\} \\ \text{variables} & b \in \mathbb{R}, \alpha_i \geq 0, \xi_i \geq 0,\ i \in \{1, ..., N\} \end{array}$$

which can be rewritten in the following primal problem:

$$\begin{array}{ll} \text{minimize} & f(\alpha, b, \boldsymbol{\xi}) = \frac{1}{2}\sum_{i=1}^N \alpha_i^2 + \sum_{i=1}^N C_i \xi_i \\ \text{subject to} & \left.\begin{array}{l} g_{1,i}(\alpha, b, \boldsymbol{\xi}) = 1 - \xi_i - y_i\left(\sum_{j=1}^N \alpha_j y_j \mathbf{x}_i \cdot \mathbf{x}_j + b\right) \leq 0 \\ g_{2,i}(\alpha, b, \boldsymbol{\xi}) = -\alpha_i \leq 0 \\ g_{3,i}(\alpha, b, \boldsymbol{\xi}) = -\xi_i \leq 0 \end{array}\right\} i \in \{1, ..., N\} \\ \text{variables} & \alpha = (\alpha_1, ..., \alpha_N) \in \mathbb{R}^N, b \in \mathbb{R}, \boldsymbol{\xi} = (\xi_1, ..., \xi_N) \in \mathbb{R}^N \end{array} \tag{1}$$

as well as its Lagrangian dual problem:

$$\begin{array}{ll} \text{maximize} & \theta(\boldsymbol{\omega}, \beta, \boldsymbol{\gamma}) = \inf_{\alpha \in \mathbb{R}^N, b \in \mathbb{R}, \boldsymbol{\xi} \in \mathbb{R}^N} L(\alpha, b, \boldsymbol{\xi}, \boldsymbol{\omega}, \beta, \boldsymbol{\gamma}) \\ \text{subject to} & \omega_i \geq 0, \beta_i \geq 0, \gamma_i \geq 0,\ i \in 1, N \\ \text{variables} & \boldsymbol{\omega} = (\omega_1, ..., \omega_N) \in \mathbb{R}^N, \beta = (\beta_1, ..., \beta_N) \in \mathbb{R}^N, \boldsymbol{\gamma} = (\gamma_1, ..., \gamma_N) \in \mathbb{R}^N \end{array} \tag{2}$$

1. (2 pts) Write down the Lagrangian function $L(\alpha, b, \boldsymbol{\xi}, \boldsymbol{\omega}, \beta, \boldsymbol{\gamma})$ in explicit form of $\alpha, b, \boldsymbol{\xi}, \boldsymbol{\omega}, \beta, \boldsymbol{\gamma}$.
2. (2 pts) Show that the duality gap between (1) and (2) is zero.
3. (2 pts) Derive $\theta(\boldsymbol{\omega}, \beta, \boldsymbol{\gamma})$ in explicit form of dual variables $\boldsymbol{\omega}, \beta, \boldsymbol{\gamma}$.
4. (4 pts) Show that the dual problem can be simplified as

$$\begin{array}{ll} \text{maximize} & \sum_{i=1}^N \omega_i - \frac{1}{2}\sum_{i=1}^N \left(\sum_{j=1}^N \omega_j y_j y_i \mathbf{x}_j \cdot \mathbf{x}_i\right)_+^2 \\ \text{subject to} & \sum_{i=1}^N \omega_i y_i = 0 \\ \text{variables} & 0 \leq \omega_i \leq C_i,\ i = 1, \ldots, N \end{array} \tag{3}$$

5. (10 pts) Suppose $(\bar{\alpha}, \bar{b}, \bar{\boldsymbol{\xi}})$ and $(\bar{\boldsymbol{\omega}}, \bar{\beta}, \bar{\boldsymbol{\gamma}})$ are the optimal solutions to problems (1) and (2) respectively. Denote $\bar{\mathbf{w}} = \sum_{j=1}^N \bar{\alpha}_j y_j \mathbf{x}_j$.

   (a) (2 pts) Prove that

$$\bar{\alpha}_i = \max\left(\sum_{j=1}^N \bar{\omega}_j y_j y_i \mathbf{x}_j \cdot \mathbf{x}_i, 0\right) \quad \forall i = 1, \ldots, N \tag{4}$$

   (b) (3 pts) Prove that

$$\bar{b} = \arg\min_{b \in \mathbb{R}} \sum_{i=1}^N C_i \max\left(1 - y_i\left(\bar{\mathbf{w}} \cdot \mathbf{x}_i + b\right), 0\right), \tag{5}$$

   (c) (2 pts) Prove that $\bar{\xi}_i = \max\left(1 - y_i\left(\bar{\mathbf{w}} \cdot \mathbf{x}_i + \bar{b}\right), 0\right)$ for all $i = 1, \ldots, N$.

   (d) (3 pts) Prove that

$$\left.\begin{array}{ll} \bar{\omega}_i = C_i, & \text{if } y_i\left(\bar{\mathbf{w}} \cdot \mathbf{x}_i + \bar{b}\right) < 1 \\ \bar{\omega}_i = 0, & \text{if } y_i\left(\bar{\mathbf{w}} \cdot \mathbf{x}_i + \bar{b}\right) > 1 \\ 0 \leq \bar{\omega}_i \leq C_i, & \text{if } y_i\left(\bar{\mathbf{w}} \cdot \mathbf{x}_i + \bar{b}\right) = 1 \end{array}\right\} \forall i = 1, \ldots, N$$

Problem 5 (Policy Improvement Theorem for Stochastic Policies)(20 pts)

Consider a discounted Markov Decision Process (MDP) with discount factor $0 \le \gamma < 1$. A stochastic policy $\pi(a \mid s)$ induces the Bellman operator $\mathbb{B} \in \ell^\infty(S) \to \ell^\infty(S)$.

$$(\mathbb{B}^\pi V)(s) \quad = \quad \sum_a \pi(a \mid s) \Big[r(s,a) + \gamma \sum_{s'} \mathbb{P}(s' \mid s,a)\, V(s') \Big].$$

For any policy π, let V^π and Q^π denote its value function and Q-function.

(a) (4 pts) Derive the Bellman equations for V^π and Q^π under a stochastic policy:

$$V^\pi(s) = \sum_a \pi(a \mid s) \Big[r(s,a) + \gamma \sum_{s'} \mathbb{P}(s'|s,a) V^\pi(s') \Big],$$

$$Q^\pi(s,a) = r(s,a) + \gamma \sum_{s'} \mathbb{P}(s'|s,a) V^\pi(s').$$

Show that V^π is a fixed point of the Bellman operator:

$$V^\pi = \mathbb{B}^\pi V^\pi.$$

(b) (6 pts) Prove that $\mathbb{B}^\pi$ is a γ-contraction in the ℓ_∞ norm:

$$\|\mathbb{B}^\pi V - \mathbb{B}^\pi W\|_\infty \ \le \ \gamma \|V - W\|_\infty, \qquad \forall\, V, W : \mathcal{S} \to \mathbb{R}.$$

(c) (4 pts) Show that $\mathbb{B}^\pi$ has a unique fixed point, and derive its fixed point.

(d) (4 pts) Let π' be another stochastic policy. Assume that for every state s,

$$\sum_a \pi'(a \mid s)\, Q^\pi(s,a) \ \ge \ V^\pi(s).$$

Show that the new policy π' is no worse than π:

$$V^{\pi'}(s) \ \ge \ V^\pi(s), \qquad \forall\, s \in \mathcal{S}.$$

(Hint: show that $\mathbb{B}^{\pi'} V^\pi \ge V^\pi$ and show that if $V \le V'$ then $\mathbb{B}^{\pi'} V \le \mathbb{B}^{\pi'} V$)

(e) (2 pts) Define the stochastic greedy policy π_{greedy} by assigning nonzero probability only to actions

$$\arg\max_a Q^\pi(s,a).$$

Show that π_{greedy} satisfies the condition in part (d), and conclude that

$$V^{\pi_{\text{greedy}}}(s) \ \ge \ V^\pi(s), \qquad \forall\, s.$$

(Notice that $S = \arg\max_a Q^\pi(s,a) \subset \mathcal{A}$ may not be a singleton.)