Homework 0

6501-003 Reinforcement Learning (Spring 2025)

This problem set is for self-assessment only—no submission required. It tests your foundation in high-school math, calculus, probability, and linear algebra.

Please try to work through these problems independently first. While you may look up concepts and formulas when stuck, avoid searching for complete solutions. If you struggle with more than 1 problem even after reviewing the concepts, please either consult with me or strengthen the prerequisites before taking the course.

P.S. These are not random questions, but the exact arguments we will use in the course.

1. Prove that for any $n \in \mathbb{N}$, $2\sqrt{n+1} - 2 < 1 + \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{3}} + \cdots \frac{1}{\sqrt{n}} < 2\sqrt{n}$.

2. Find $a, b > 0$ that minimizes $\frac{1}{a^2} + \frac{a}{b} + b$. Justify your answer.

3. Let $x_1, \ldots, x_n \in \mathbb{R}^d$ and $y_1, \ldots, y_n \in \mathbb{R}$. Let $\lambda > 0$ and define

$$F(\theta) = \sum_{i=1}^{n} (x_i^\top \theta - y_i)^2 + \lambda \|\theta\|_2^2.$$

   Prove that

$$\theta = \left(\lambda I + \sum_{i=1}^{n} x_i x_i^\top\right)^{-1} \left(\sum_{i=1}^{n} x_i y_i\right)$$

   minimizes $F(\theta)$.

4. Let $z_1, \ldots, z_n$ be independent variables following the standard normal distribution $\mathcal{N}(0,1)$, and let $x_1, \ldots, x_n \in \mathbb{R}^d$ be fixed vectors. Define $X = \sum_{i=1}^{n} z_i x_i$. Prove that the covariance matrix of X, i.e., $\mathbb{E}[XX^\top]$, is equal to $\sum_{i=1}^{n} x_i x_i^\top$.

5. Let $A, B \in \mathbb{R}^{d\times d}$ be symmetric positive definite matrices such that $A \succ B \succ 0$. Prove that $B^{-1} \succ A^{-1} \succ 0$. ($A \succ B$ means that for any $x \in \mathbb{R}^d \backslash \{0\}$, $x^\top (A - B)x > 0$.)

6. Let $y \in \Delta_d$ and $\theta \in \mathbb{R}^d$, where Δ_d is the d-dimensional probability space defined as $\{y \in \mathbb{R}^d : \sum_{i=1}^{d} y_i = 1 \text{ and } y_i \geq 0 \text{ for all } i\}$. Let $F(x) = x^\top \theta + \mathrm{KL}(x, y)$. Prove that the minimizer of $F(x)$ over $x \in \Delta_d$ satisfies

$$x_i = \frac{y_i e^{-\theta_i}}{\sum_{j=1}^{d} y_j e^{-\theta_j}}.$$

   Recall: the KL divergence between two distributions $x, y \in \Delta_d$ is defined as $\mathrm{KL}(x, y) = \sum_{i=1}^{d} x_i \log \frac{x_i}{y_i}$.

7. Let $\pi_\theta(x)$ be a function of x parameterized by θ. Prove that

$$\nabla_\theta^2 \log \pi_\theta(x) = \frac{\nabla_\theta^2 \pi_\theta(x)}{\pi_\theta(x)} - (\nabla_\theta \log \pi_\theta(x))(\nabla_\theta \log \pi_\theta(x))^\top.$$