Knowledge Quiz

CSE 599B-24SP: Reinforcement Learning
Prof. Abhishek Gupta
TAs: Patrick Yin and Zoey Chen

1 Optimization

Question: Derive the dual of the following primal Linear Program. Hint: construct the Lagrangian

$$\begin{aligned} \min_{x} \quad & c^T x \\ \text{s.t.} \quad & Ax \leq b \\ & x \geq 0 \end{aligned} \tag{1}$$

2 Probability

Question: Prove the Kullback-Leibler divergence is always non-negative

$$D_{KL}(P \,||\, Q) \geq 0 \tag{2}$$

with $D_{KL}(P \,||\, Q) = 0$, if and only if $P = Q$. Hint: use Jensen's inequality and the concavity of the log function, which says

$$\log\left(\mathbb{E}_{p(x)}\left[f(x)\right]\right) \geq \mathbb{E}_{p(x)}\left(\log\left[f(x)\right]\right) \tag{3}$$

3 ML and Python

Background: Regression is a fundamental tool in machine learning. The goal is to minimize the discrepancy between the predicted values and the ground truth (actual values) across the dataset, assuming a linear model.

Question:

1. Find the weight parameter W that minimizes the squared l^2 norm of the difference between WX and Y. Express this in terms of matrices X and Y,

$$\min_{W} \|\mathbf{WX} - \mathbf{Y}\|_2^2 . \tag{4}$$

2. Use your results above to implement the function: `linear_regression_normal_equation (X, y)`

```python
import numpy as np
from sklearn.linear_model import make_regression
# Create data set.
X, y = make_regression(n_samples=100, n_features=1,
                       n_informative=1, noise=10, random_state=10)

def linear_regression_normal_equation(X, y):
    W=None
    # Your implementation here
    return W

W = linear_regression_normal_equation(X, y)
```

Listing 1: solving linear regression

Please fill in the evaluation form on the website once you finish the quiz.