Qualifying Exam

Student Number and Name : ______________________________

NATIONAL UNIVERSITY OF SINGAPORE

Qualifying Exam Paper 4

(Semester 1 : AY2021/2022)

Time : 9 :00am – 12 :00pm

INSTRUCTIONS TO STUDENTS

1. Use A4 size paper and pen (blue or black ink) to write your answers. Write down your student number clearly on the top left of every page of the answers. Do not write your name.
2. Write on one side of the paper only. Write the question number and page number on the top right corner of each page (e.g. Q1 (i)).
3. This examination paper contains SIX questions and comprises THREE pages. Answer ALL questions.
4. This is an OPEN BOOK examination. You may use any calculator. However, you should lay out systematically the various steps in the calculations.
5. Join the Zoom conference and turn on the video setting at all time during the exam. Adjust your camera such that your face and upper body including your hands are captured on Zoom. You may go for a short toilet break (not more than 5 minutes) during the exam.
6. At the end of the exam, scan or take pictures of your work (make sure the images can be read clearly) together with the declaration form ; merge all your images into one pdf file (arrange them in the order : Declaration form, Q1 to Q6 in their page sequence) ; name the pdf file by Matric No_Moule Code (e.g. A123456R_MA4255).
7. Submit your answers via email to qianxiao@nus.edu.sg and mattxin@nus.edu.sg.

Question	1	2	3	4	5	6	Total
Marks	/15	/20	/15	/15	/15	/20	

Question 1 : Pokémon has $n(>1)$ different monster cards. You can get a random monster card every time you open a pack of chips. Let X be the number of pack of chips you need to collect all n different monster cards. What is the expectation of X ? What is the variance of X ? (Hint : let Y_t be the number of different monster cards you collected from t packs, you can show it is a Markov chain.)

Question 2 : Suppose we have i.i.d. samples $x_i \sim \mathcal{N}(0, \Sigma)$ and $y_i \sim x_i^T w^* + \xi_i, \xi_i \sim \mathcal{N}(0, 1)$. Let us use stochastic gradient descent to solve this problem. This yields a stochastic sequence :

$$w_{t+1} = w_t - \eta \nabla_w f(w, x_i, y_i), \quad w_0 = 0,$$

where $f(w, x_i, y_i) = \frac{1}{2}|y_i - w^T x_i|^2$ and η is the step size. Formulate an appropriate range of η so that w_t will become an ergodic process. Also find mean and covariance under the equilibrium distribution.

Question 3 : Consider a univariate process

$$dX_t = (X_t - X_t^3)dt + \sigma dW_t, \quad X_0 = 0,$$

where W_t is the standard Brownian motion.

(i) Show that X_t is an ergodic process.

(ii) Derive an upper bound for $\lim_{t\to\infty} \mathbb{E}X_t^4$. (The tighter the better)

Question 4 :
Let X be a finite set and 2^X the collection of subsets of X. Show that the following functions from $2^X \times 2^X \to \mathbb{R}$ are symmetric positive definite kernels

(i) $k(A, B) = \sum_{x\in X, y\in X} k_0(x, y)$ where $k_0 : X \times X \to \mathbb{R}$ is a symmetric positive definite kernel on X.

(ii) $k(A, B) = \exp(-\frac{1}{2}|A\Delta B|)$, where $A\Delta B = (A\setminus B)\cup(B\setminus A)$ is the symmetric difference, and $|A|$ is the number of elements in A.

Discuss two possible applications of any/both of these kernels in machine learning.

Question 5 :
Write down the definition of PAC learnability. Propose a PAC learning algorithm for the concept class of unions of two disjoint closed intervals, i.e.

$$\mathcal{C} = \{\mathbf{1}_{[a,b]\cup[c,d]} : a, b, c, d \in \mathbb{R}, a < b < c < d\}.$$

Prove a PAC-learning bound corresponding to your proposed algorithm. You should show your steps clearly.

Question 6 :
Consider a discrete-time Markov decision process with finite state space $\mathcal{S}$ and action space $\mathcal{A}$. We use the usual notation of $\{S_t, A_t, R_t\}$ to denote the state, action and reward at time t respectively. Let the transition probability kernel be

$$p(s', r|s, a) = \mathbb{P}[S_{t+1} = s', R_{t+1} = r|S_t = s, A_t = a].$$

(i) Define the value function v_π with respect to a policy π. You may assume that we consider a discount rate of $0 < \gamma < 1$ when computing the returns.

(ii) Write down the Bellman's optimality equation that the value function corresponding to an optimal policy should satisfy.

(iii) Show that there exists a unique solution to the Bellman's optimality equation.

(iv) Is an optimal policy always unique ? If so, prove this statement. If not, give a counter-example.

End of Paper