K. J. Somaiya Institute of Technology, Sion, Mumbai-22
(Autonomous College Affiliated to University of Mumbai)

May-June 2024
M. Tech Program: Artificial Intelligence Scheme I/II/IIB/III: II
Regular Examination: FY Semester: II
Course Code: PCEC201 and Course Name: Deep and Reinforcement Learning
Date of Exam: 15-05-2024 Duration: 10:30 am to 01:00 pm Max. Marks: 60

Instructions:
(1) All questions are compulsory.
(2) Draw neat diagrams wherever applicable.
(3) Assume suitable data, if necessary.

		Max. Marks	CO	BTL
Qu-1	Solve any **SIX** questions out of **EIGHT.**	12		
i)	What are the limitations of deep learning?	2	CO1	**1**
ii)	Explain the concept of gradient descent and how it is combined with backpropagation for training neural networks.	2	CO2	**2**
iii)	What is Recursive Neural Networks?	2	CO3	**1**
iv)	Define the agent-environment framework in reinforcement learning.	2	CO4	**1**
v)	Suppose every-visit Monte Carlo (MC) was used instead of first-visit MC on the blackjack task. Would you expect the results to be very different? Why or why not?	2	CO5	**1**
vi)	Explain in short how deep learning is used in sentiment analysis.	2	CO6	**1**
vii)	Explain linearly separable and linearly non-separable patterns with example.	2	CO1	**2**
viii)	List the advantages of Temporal-Difference (TD) prediction methods.	2	CO5	**1**
Qu-2	Solve any **FOUR** questions out of **SIX**.	16		
i)	Illustrate/How are ANNs able to handle noise in data?	4	CO1	**3**
ii)	What are the different activation functions used in artificial neuron modeling, and how do they vary in terms of their suitability for different applications?	4	CO2	**3**
iii)	Illustrate the operation of pooling layer in CNN with simple example.	4	CO3	**2**
iv)	Describe the Upper Confidence Bound (UCB) algorithm and explain how it selects actions based on both the estimated action values and uncertainty.	4	CO4	**2**
v)	Describe the basic workings of the ε-greedy algorithm in the context of bandit problems. How does the algorithm balance exploration and exploitation, and what role does the ε parameter play in its operation?	4	CO5	**2**
vi)	Elaborate emerging research areas in Reinforcement learning and provide an example application that illustrates the importance of these areas.	4	CO6	**2**

K. J. Somaiya Institute of Technology, Sion, Mumbai-22
(Autonomous College Affiliated to University of Mumbai)

May-June 2024
M. Tech Program: Artificial Intelligence Scheme I/II/IIB/III: II
Regular Examination: FY Semester: II
Course Code: PCEC201 and Course Name: Deep and Reinforcement Learning
Date of Exam: 15-05-2024 Duration: 10:30 am to 01:00 pm Max. Marks: 60

Qu-3	Solve any **TWO** questions out of **THREE**.	16		
i)	Recommend possible methods for pre-training the input and output layers in the machine translation approach with sequence-to-sequence learning.	8	CO3	**3**
ii)	A single neuron network using *f(net)=sgn(net)* has been trained using the pairs *(xi, di)* as given below: $X_1 = [1 \ -2 \ 3 \ -1]^t$ $d_1 = -1$ $X_2 = [0 \ -1 \ 2 \ -1]^t$ $d_2 = 1$ $X_3 = [-2 \ 0 \ -3 \ -1]^t$ $d_3 = -1$ Final weights obtained using Perceptron rule are $W_4=[3 \ 2 \ 6 \ 1]^t$. knowing that correction has been performed in each step. C=1 determine weights W_3, W_2, W_1 by backtracking the training.	8	CO2	**6**
iii)	Is there any difference between LSTM and RNN? Justify your answer.	8	CO3	**3**
Qu-4	Solve any **TWO** questions out of **THREE**.	16		
i)	A mobile robot has the job of collecting empty soda cans in an office environment. It has sensors for detecting cans, and an arm and gripper that can pick them up and place them in an onboard bin; it runs on a rechargeable battery. The robot's control system has components for interpreting sensory information, for navigating, and for controlling the arm and gripper. High-level decisions about how to search for cans are made by a reinforcement learning agent based on the current charge level of the battery. Assume that only two charge levels can be distinguished, comprising a small state set S = {high, low}. In each state, the agent can decide whether to: a. Actively search for a can for a certain period of time, b. Remain stationary and wait for someone to bring it a can, or c. Head back to its home base to recharge its battery. Represent the dynamics of a Mobile Robot as a transition graph and transition table considering finite MDP	8	CO4	**2**
ii)	What is Sarsa? Explain how the choice of exploration strategy, such as epsilon-greedy, impacts the convergence and performance of the Sarsa algorithm in learning optimal policies.	8	CO5	**2**
iii)	List the various case studies of Deep Learning Applications and explain any one in detail.	8	CO6	**2**

****