



**University School of Automation and Robotics**  
**GURU GOBIND SINGH INDRAPRASTHA UNIVERSITY**  
 East Delhi Campus, Surajmal Vihar  
 Delhi - 110092

<b>Paper Code: ARI 314T</b>											L	T/P	Credits
<b>Subject: Deep Learning and Reinforcement Learning</b>											4	0	4
<b>Marking Scheme</b> Teachers Continuous Evaluation: As per university examination norms from time to time. End Term Theory Examination: As per university examination norms from time to time.													
<b>INSTRUCTIONS TO PAPER SETTERS: Maximum Marks: As per University Norms</b>													
1. There should be 9 questions in the end term examination question paper 2. Question No. 1 should be compulsory and cover the entire syllabus. This question should have objective or short answer type questions. It should be of 15 marks. 3. Apart from Question No. 1, rest of the paper shall consist of four units as per the syllabus. Every unit should have two questions. However, student may be asked to attempt only 1 question from each unit. Each question should be 15 marks. 4. The questions are to be framed keeping in view the learning outcomes of course/paper. The standard/ level of the questions to be asked should be at the level of the prescribed textbooks.													
<b>Course Outcomes [Bloom's Knowledge Level (KL)]:</b>													
<b>CO1:</b>	Ability of students to understand the basics concepts of Deep feed forward networks [K1, K2]												
<b>CO2:</b>	Ability of students to apply and analyze various deep learning applications with case studies [K3,K4]												
<b>CO3:</b>	Ability of students to understand the basics concepts of reinforcement learning and MDP [K1, K2]												
<b>CO4:</b>	Ability of students to understand deep and multi agent reinforcement learning and analyze various reinforcement learning applications and case studies [K1,K2,,K4]												
<b>Course Outcomes (CO) to Programme Outcomes (PO)</b>													
<b>CO/PO</b>	<b>PO1</b>	<b>PO2</b>	<b>PO3</b>	<b>PO4</b>	<b>PO5</b>	<b>PO6</b>	<b>PO7</b>	<b>PO8</b>	<b>PO9</b>	<b>PO 10</b>	<b>PO11</b>	<b>PO12</b>	
<b>CO1</b>	3	2	1	1	3	-	-	-	2	2	1	2	
<b>CO2</b>	3	3	3	3	3	-	-	-	3	3	4	2	
<b>CO3</b>	3	2	1	1	3	-	-	-	2	2	1	3	
<b>CO4</b>	3	3	3	3	3	-	-	-	3	3	3	3	
<b>Average</b>	3	2.5	2	2	3	-	-	-	2.5	2.5	2	2.5	
<b>Contents</b>												<b>Number of Lectures</b>	



<p><b>Unit I</b> <b>Deep Feedforward Networks:</b> Artificial Neural Networks, Artificial Neuron, Example: Learning XOR, Gradient-Based Learning, Hidden Units, Architecture Design, Back-Propagation and Other Differentiation Algorithms, Regularization for Deep Learning- Parameter Norm Penalties, Dataset Augmentation, Noise Robustness, Early Stopping, Dropout, Adversarial Training, Optimization for Training Deep Models- How Learning Differs from Pure Optimization? Challenges in Neural Network Optimization, Basic Algorithms- Stochastic Gradient Descent, momentum. Parameter Initialization Strategies, Algorithms with Adaptive Learning Rates, Optimization Strategies and Meta-Algorithms.</p>	[8]
<p><b>Unit II</b> <b>Convolutional Networks:</b> The Convolution Operation, Motivation, Pooling, Data Types, building block of CNN, Transfer Learning, Autoencoders- Under Complete, regularized, sparse Denoising, Generative Modeling with DL, Generative Adversarial Network Revisiting Gradient Descent, Momentum Optimizer, RMSProp, Adam <b>Deep Learning Applications and Case Studies:</b> Large Scale Deep Learning, Deep Learning in Computer Vision, Deep Learning in Speech Recognition, Deep Learning in Natural Language Processing, Deep Learning for Recommender Systems.</p>	[12]
<p><b>Unit III</b> <b>Introduction to Reinforcement Learning:</b> Introduction to Reinforcement Learning, The Reinforcement Learning Problem, Markov Decision Process (MDP)-Markov Process, Markov Reward Process, Markov Decision Process and Bellman Equations, Partially Observable MDPs. Exploration and Exploitation (Bandits), Multi-arm Bandits, Contextual Bandits and MDP Extensions.</p>	[10]
<p><b>Unit IV</b> <b>Planning by Dynamic Programming (DP):</b> Policy Evaluation, Value Iteration, Policy Iteration, DP Extensions and Convergence using Contraction Mapping <b>Policy Gradient Methods:</b> Finite-Difference, Monte-Carlo and Actor-Critic Methods <b>Reinforcement Learning applications and case Studies:</b> TD-Gammon, Samuel's Checkers Player, Watson's Daily-Double Wagering, Optimizing Memory Control, Human-Level Video game play, Mastering the game of Go, Personalized Web Services, Thermal Soaring.</p>	[10]
<p><b>Text Books:</b> [T1] Ian Goodfellow, Yoshua Benjio, Aaron Courville , (2016), Deep Learning, The MIT Press [T2] Richard S. Sutton and Andrew G. Barto; Reinforcement Learning: An Introduction; 2nd Edition, MIT Press, 2020.</p>	
<p><b>Reference Books:</b> [R1] Josh Patterson, Adam Gibson, (2017), Deep Learning: A Practitioner's Approach, O'Reilly [R2] Duda, R. O. &amp; Hart, P. E. (2006). Pattern Classification. John Wiley &amp; Sons. [R3] Csaba Szepesvári; Algorithms of Reinforcement Learning; Synthesis Lectures on Artificial Intelligence and Machine Learning, vol. 4, no. 1, 2010. [R4] Dimitri P. Bertsekas; Reinforcement Learning and Optimal Control; 1st Edition, Athena Scientific, 2019.</p>	