

RAJIV GANDHI PROUDYOGIKI VISHWAVIDYALAYA, BHOPAL
New Scheme Based On AICTE Flexible Curricula
Computer Science and Engineering, VII-Semester
Departmental Elective – CS702 (B) Deep & Reinforcement Learning

Pre-Requisite: Machine Learning

Course Outcomes:

After completing the course student should be able to:

1. Describe in-depth about theories, models and algorithms in machine learning.
2. Compare and contrast different learning algorithms with parameters.
3. Examine the nature of a problem at hand and find the appropriate learning algorithms and its parameters that can solve it efficiently enough.
4. Design and implement of deep and reinforcement learning approaches for solving real-life problems.

Course Contents:

Unit 1: History of Deep Learning, McCulloch Pitts Neuron, Thresholding Logic, Activation functions, Gradient Descent (GD), Momentum Based GD, Nesterov Accelerated GD, Stochastic GD, AdaGrad, RMSProp, Adam, Eigenvalue Decomposition. Recurrent Neural Networks, Backpropagation through time (BPTT), Vanishing and Exploding Gradients, Truncated BPTT, GRU, LSTMs, Encoder Decoder Models, Attention Mechanism, Attention overimages.

Unit 2: Autoencoders and relation to PCA, Regularization in autoencoders, Denoising autoencoders, Sparse autoencoders, Contractive autoencoders, Regularization: Bias Variance Tradeoff, L2 regularization, Early stopping, Dataset augmentation, Parameter sharing and tying, Injecting noise at input, Ensemble methods, Dropout, Batch Normalization, Instance Normalization, Group Normalization.

Unit 3: Greedy Layerwise Pre-training, Better activation functions, Better weight initialization methods, Learning Vectorial Representations Of Words, Convolutional Neural Networks, LeNet, AlexNet, ZF-Net, VGGNet, GoogLeNet, ResNet, Visualizing Convolutional Neural Networks, Guided Backpropagation, Deep Dream, Deep Art, Recent Trends in Deep Learning Architectures.

Unit 4: Introduction to reinforcement learning(RL), Bandit algorithms – UCB, PAC, Median Elimination, Policy Gradient, Full RL & MDPs, Bellman Optimality, Dynamic Programming - Value iteration, Policy iteration, and Q-learning & Temporal Difference Methods, Temporal-Difference Learning, Eligibility Traces, Function Approximation, Least Squares Methods

Unit 5: Fitted Q, Deep Q-Learning , Advanced Q-learning algorithms , Learning policies by imitating optimal controllers , DQN & Policy Gradient, Policy Gradient Algorithms for Full RL, Hierarchical RL, POMDPs, Actor-Critic Method, Inverse reinforcement learning, Maximum Entropy Deep Inverse Reinforcement Learning, Generative Adversarial Imitation Learning, Recent Trends in RL Architectures.

Text Books:

1. Deep Learning, An MIT Press book, Ian Goodfellow and Yoshua Bengio and Aaron Courville
2. Pattern Classification- Richard O. Duda, Peter E. Hart, David G. Stork, John Wiley & Sons Inc.
3. Reinforcement Learning: An Introduction, Sutton and Barto, 2nd Edition.
4. Reinforcement Learning: State-of-the-Art, Marco Wiering and Martijn van Otterlo, Eds