

Objectives:

- To introduce the idea of artificial neural networks and their architecture.
- To introduce techniques used for training artificial neural networks.
- To enable design of an artificial neural network for classification
- To enable design and deployment of deep learning models for machine learning problems

Keywords:

Neural Networks, Convolutional Neural Networks, Autoencoder, Long Short-Term Memory.

Contents:

Artificial Neural Networks: The Neuron-Expressing Linear Perceptrons as Neurons-Feed-Forward Neural Networks- Linear Neurons and Their Limitations- Sigmoid, Tanh, and ReLU Neurons- Softmax Output Layers. Training Feed-Forward Neural Networks.-Gradient Descent-Delta Rule and Learning Rates- Gradient Descent with Sigmoidal Neurons- The Backpropagation Algorithm-Stochastic and Minibatch Gradient Descent- Test Sets, Validation Sets, and Overfitting- Preventing Overfitting in Deep Neural Networks. Local Minima in the Error Surfaces of Deep Networks- Model Identifiability- Spurious Local Minima in Deep Networks- Flat Regions in the Error Surface - Momentum-Based Optimization- Learning Rate Adaptation.

Convolutional Neural Networks: Architecture -Accelerating Training with Batch Normalization- Visualizing Learning in Convolutional Networks. Embedding and Representation Learning: Autoencoder Architecture- Denoising- Sparsity in Autoencoders. Models for Sequence Analysis: Recurrent Neural Networks- Vanishing Gradients- Long Short-Term Memory (LSTM) Units-Augmenting Recurrent Networks with Attention.

Deep Reinforcement Learning: Markov Decision Processes (MDP), Explore Versus Exploit, Policy Versus Value Learning, Pole-Cart example with Policy Gradients. Q-Learning and Deep Q-Networks, Improving and Moving Beyond DQN. Case studies on domain specific applications.

Outcomes:

CO 1: Able to understand the mathematics behind functioning of artificial neural networks

CO 2: Able to analyze the given dataset for designing a neural network based solution.

CO 3: Able to carry out design and implementation of deep learning models for signal/image processing applications.

CO – PO Mapping:

CO/PO	PO 1	PO 2	PO 3	PO 4	PO 5
CO 1	3	-	3	-	3
CO 2	3	-	3	-	3
CO 3	-	3	3	3	3

TEXT BOOKS/REFERENCES:

1. Nikhil Buduma, *Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms*, O'Reilly, 2017.
2. Shaoshan Liu, Liyun Li, Jie Tang, Shuang Wu and Jean-Luc Gaudiot. "Creating Autonomous Vehicle Systems", Morgan & Claypool Publishers,2018.
3. Ian Goodfellow, Yoshua Bengio and Aaron Courville, *Deep Learning*, MIT Press, 2016.
4. Aurélien Géron, *Hands-On Machine Learning with Scikit-Learn and TensorFlow*, O'Reilly, 2017.
5. Nikhil Ketkar, *Deep Learning with Python: A Hands-on Introduction*, Apress, 2017.