

DEEP LEARNING

Lesson 1: Introduction to Deep Learning

- Define Deep Learning
- Neural Networks
- Deep Learning Applications

Lesson 2: Perceptron

- What is a Perceptron
- Logic Gates with Perceptrons
- Activation Functions
- Sigmoid
- ReLU
- Softmax
- Hyperbolic Functions

Lesson 3: How to train ANNs

- Introduction
- Perceptron Learning Rule
- Gradient Descent Rule
- Minimize Cost Function
- Tuning Learning Rate
- Stochastic vs Batch Gradient Descent

Lesson 4: Multi-layer ANN

- Intro to MLP
- Forward propagation
- Minimize Cost Function
- Backpropagation
- Convergence in a neural net
- Overfitting and Capacity
- Hyperparameters in an ANN

Lesson 5: Introduction to TensorFlow

- Intro to TensorFlow
- Computational Graph
- Key highlights
- Creating a Graph
- Regression example
- Gradient Descent
- Saving and Restoring Models
- Tf.layers API
- Keras-based networks
- TensorBoard

Lesson 6: Training Deep Neural Nets

- Vanishing/Exploding Gradients
- Xavier Initialization

- Leaky ReLUs and ELUs
- Batch Normalization
- Transfer Learning
- Unsupervised Pre-training
- Optimizers
- Regularization
- Dropout

Lesson 7: Convolutional Neural Networks

- Intro to CNNs
- Convolution Operation
- Kernel filter
- Feature Maps
- Pooling
- CNN Architecture
- Implement CNN in TensorFlow

Lesson 8: Recurrent Neural Networks

- Intro to RNNs
- Unfolded RNNs
- Basic RNN Cell
- Dynamic RNN
- Training RNNs
- Time-series predictions
- LSTM
- Word Embeddings
- Seq2Seq Models
- Implement RNN in TensorFlow

Lesson 7: Other forms of Deep Learning

- Autoencoders
- Reinforcement Learning (RL)
- Generative Adversarial Networks (GANs)