



SPECIALIZATION COURSE ON MACHINE LEARNING & DEEP LEARNING

COURSE DESIGN

High-quality videos, slides, hands-on examples, quizzes, automated assessments, case studies, and real-world projects.

COURSE MATERIAL

Lifetime access to cutting-edge self-paced learning content.

LAB

180 Days of [CloudxLab](#) access for hands-on practice.

SUPPORT

Email support to answer your queries and we've also launched [Discussions](#) - a Q&A site for Artificial Intelligence, Machine Learning, Deep Learning, Big Data & Data Science professionals.

CERTIFICATE

Earn certificate in Machine Learning and Deep Learning

LIVE SESSIONS

100+ hours of live online instructor-led training. Classes will be conducted every Saturday & Sunday between (7 AM - 10 AM Indian Standard Time) or (6:30 PM - 9:30 PM Pacific Time).

MACHINE LEARNING SPECIALIZATION - COURSE SYLLABUS

PYTHON FOR MACHINE LEARNING

1. Introduction to Linux
2. Introduction to Python
3. Hands-on using Jupyter on CloudxLab
4. Overview of Linear Algebra
5. Introduction to NumPy & Pandas

Project - Analyze your mailbox

COURSE ON MACHINE LEARNING

1. Introduction to Statistics
 - Statistical Inference
 - Types of Variables
 - Probability Distribution
 - Normality
 - Measures of Central Tendencies
 - Normal Distribution
2. Machine Learning Applications & Landscape
 - Introduction to Machine Learning,
 - Machine Learning Application
 - Introduction to AI
 - Different types of Machine Learning - Supervised, Unsupervised
 - Reinforcement
3. Building end-to-end Machine Learning Project
 - Machine Learning Projects Checklist,
 - Frame the problem and look at the big picture
 - Get the data
 - Explore the data to gain insights

- Prepare the data for Machine Learning algorithms
- Explore many different models and short-list the best ones
- Fine-tune model
- Present the solution
- Launch, monitor and maintain the system

4. Classifications

- Training a Binary classification,
- Performance Measures
- Confusion Matrix
- Precision and Recall
- Precision/Recall Tradeoff
- The ROC Curve
- Multiclass Classification
- Multilabel Classification
- Multi-output Classification

5. Training Models

- Linear Regression
- Gradient Descent
- Polynomial Regression
- Learning Curves
- Regularized Linear Models
- Logistic Regression

6. Support Vector Machines

- Linear SVM Classification
- Nonlinear SVM Classification
- SVM Regression

7. Decision Trees

- Training and Visualizing a Decision Tree
- Making Predictions
- Estimating Class Probabilities
- The CART Training Algorithm
- Gini Impurity or Entropy
- Regularization Hyperparameters
- Regression
- Instability

8. Ensemble Learning and Random Forests

- Voting Classifiers
- Bagging and Pasting
- Random Patches and Random Subspaces
- Random Forests
- Boosting
- Stacking

9. Dimensionality Reduction

- The Curse of Dimensionality
- Main Approaches for Dimensionality Reduction
- PCA
- Kernel PCA
- LLE
- Other Dimensionality Reduction Techniques

Projects

1. Predict the median housing prices in California
2. Classify handwritten digits in MNIST dataset
3. Noise removal from the images
4. Predict the class of flower in IRIS dataset
5. Predict which passengers survived in the Titanic shipwreck
6. Predict bikes rental demand
7. Build a spam classifier

COURSE ON DEEP LEARNING

1. Introduction to Deep Learning

- Deep Learning Applications,
- Artificial Neural Network,
- TensorFlow Demo,
- Deep Learning Frameworks

2. Up and Running with TensorFlow

- Installation,
- Creating Your First Graph and Running It in a Session,
- Managing Graphs,
- Lifecycle of a Node Value,
- Linear Regression with TensorFlow,
- Implementing Gradient Descent,
- Feeding Data to the Training Algorithm,
- Saving and Restoring Models,
- Visualizing the Graph and Training Curves Using TensorBoard,
- Name Scopes, Modularity,
- Sharing Variables

3. Introduction to Artificial Neural Networks

- From Biological to Artificial Neurons,
- Training an MLP with TensorFlow's High-Level API,
- Training a DNN Using Plain TensorFlow,
- Fine-Tuning Neural Network Hyperparameters

4. Training Deep Neural Nets

- Vanishing / Exploding Gradients Problems,
- Reusing Pretrained Layers,
- Faster Optimizers,
- Avoiding Overfitting Through Regularization,
- Practical Guidelines

5. Convolutional Neural Networks

- The Architecture of the Visual Cortex,
- Convolutional Layer,
- Pooling Layer,
- CNN Architectures

6. Recurrent Neural Networks

- Recurrent Neurons,
- Basic RNNs in TensorFlow,
- Training RNNs,
- Deep RNNs,
- LSTM Cell,
- GRU Cell,
- Natural Language Processing

7. Autoencoders

- Efficient Data Representations,
- Performing PCA with an Undercomplete Linear Autoencoder,
- Stacked Autoencoders,
- Unsupervised Pretraining Using Stacked Autoencoders,
- Denoising Autoencoders,
- Sparse Autoencoders,
- Variational Autoencoders

8. Reinforcement Learning

- Learning to Optimize Rewards,
- Policy Search,
- Introduction to OpenAI Gym,
- Neural Network Policies,
- Evaluating Actions: The Credit Assignment Problem,
- Policy Gradients,
- Markov Decision Processes,
- Temporal Difference Learning and Q-Learning,
- Learning to Play Ms. Pac-Man Using Deep Q-Learning

Projects

1. Build cats classifier using neural network
2. Classify large images using Inception v3
3. Classify clothes using TensorFlow
4. Predict the hourly rain gauge total

[*Click Here to Enroll Now!!*](#)

Please feel free to email your queries to reachus@cloudxlab.com

Regards,
The CloudxLab Team