



Machine Learning with Python

Overview

Machine learning (ML) is a category of algorithm that allows software applications to become more accurate in predicting outcomes without being explicitly programmed. The basic premise of machine learning is to build algorithms that can receive input data and use statistical analysis to predict an output while updating outputs as new data becomes available

Extremely Hands-On...

Incredibly Practical...

Unbelievably Real!

Upon completing this course you will know:

- Know about the various machine learning algorithms.
- Develop and run regression and classification models on small datasets with good accuracies.
- Know about various libraries in Python for machine learning, like, Numpy, Pandas, Sci-kit learn etc.

Targeted Audience

Students having prior knowledge of basic programming and interested to choose their career as Machine learning engineers and/or data scientists.

Course Agenda

Unit 1. Introduction to Machine Learning

What is Machine learning?
Why is there so much HYPE for it?
Machine learning + data
Types of Machine learning
Machine learning + Python

Unit 2. Python Crash course

Basics of Python
Basic Data Types and Objects
Conditioning in python
Looping and breaks
Class definition on python

Unit 3. Supervised learning- Linear Regression with one variable

Model Representation
Cost Function
Gradient Descent
Gradient descent for Linear Regression
Learning rate and its significance
Implementation in Python modules - Numpy, Pandas, Sci-kit learn.

Unit 4. Introducing Python Modules

Numpy

Working with Numpy.
Fast analysis and data handling with Pandas.
Exercise 1 : Average gold,silver and bronze medal problem

Pandas

Working with pandas data structures.
Working with pandas visualization.
Reading and Writing files with pandas

Matplotlib and Seaborn visualization

Working with matplotlib : creating figures and adding multiple axes.
Working with seaborn :add-on regression, distribution and matrix plots.

Unit 5. Supervised learning - Linear regression with Multiple variables

Multiple Features.

Gradient descent for Multiple Linear Regression

Feature Scaling

Features and polynomial regression

Normal equation

When to choose what - Gradient descent vs Normal Equation

Unit 6. Supervised learning - Logistic Regression

Classification

Hypothesis Representation

Decision Boundary

Simplified Cost Function and Gradient Descent

Advanced optimisation

Multiclass classification : One-vs-all

Implementation in Python modules - Numpy, Pandas, Sci-kit learn.

Unit 7. Regularization techniques

Notion of Overfitting

Cost function

Regularized Linear regression

Regularized Logistic regression

Implementation in Python module - sci-kit learn

Projects included

1. Boston Housing prediction task.
2. Titanic Survival prediction task.

Disclaimer

All the assignments and discussion links will be provided after the lecture of current topic.