

Practical Machine Learning & Deep Learning for Developers

40 hours

Course Overview:

Machine Learning is in the technological forefront these days, offering innovative and exciting ways transforming data into knowledge and insights. These, in turn, open-up growth opportunities for companies across diverse industries, from healthcare to cyber-security, from e-commerce to finance - the potential is practically limitless.

The goal of this course is to propel developers from having just heard about Machine Learning, to making their first steps building Deep Learning models, in just a few days. The course covers some needed theoretical background, while focusing on putting the theory into practice using industry-standard tools and techniques.

Who Should Attend:

Developers who wish to take their first steps into the world of Data Science, Machine learning and Deep Learning, acquiring practical knowledge, tools and techniques.

Required Skills:

- Understanding of basic programming concepts
- Python programming skills

Course Contents:

- Python refresher
 - Environment set-up
 - Data types
 - Collections
 - Flow control
 - Functions
 - Object Oriented
 - Modules and packages
- Machine Learning Concepts
- The Machine Learning process / pipeline
- Data Acquisition and Analysis
 - Numpy
 - Multi-dimensional arrays
 - Slicing and “fancy indexing”
 - Broadcasting

- Functions
- Pandas & Seaborn
 - Series and Data-Frames
 - Selecting data
 - Operations and functions
 - Grouping and aggregation
- Data visualization using Matplotlib
 - Line and Bar plots
 - Histograms
 - Scatter plots
 - Sub-plots (multi-plot graphs)
 - Animated plots
 - Box plots
 - Pair plot
 - Heat maps
- Data Pre-processing
 - Cleanup
 - Dealing with outliers
 - Handling missing values and Imputation
 - Categorical feature encoding
 - Feature selection approaches
 - Feature scaling
 - Feature engineering
 - Splitting the dataset
- Machine Learning Algorithms
 - Major learning types - supervised, unsupervised, reinforcement
 - Introduction to Scikit-Learn
 - Supervised Learning Algorithms - theory and implementation
 - Linear Regression
 - Logistic Regression
 - Decision Trees
 - Random Forests
 - K-Nearest Neighbors
 - Support Vector Machines

- Gradient Boosting
- Unsupervised Learning Algorithms - theory and implementation
 - Clustering
 - K-Means
- Model Evaluation
 - Confusion Matrix
 - Metrics: Accuracy, Precision, Recall, F1 Score, ROC/AUC, R-squared
 - Which metric to optimize for?
 - Model Complexity
 - Overfitting and Underfitting
 - Complexity Graph
 - Managing model complexity
- Putting it all together - use case analysis from raw data to predictive model
- Deep Neural Networks using Keras
 - The building blocks of a neural network
 - Demystifying NN - behind the scenes walkthrough
 - Convolutional Neural Networks and computer vision
 - Gradient Descent - theory of operation
 - Building a Deep Neural Network Model
 - Introduction to Keras
 - Parameters and hyperparameters
 - Network structure and layer types
 - Picking an activation function
 - Picking a loss function
 - Setting the Learning Rate
 - Implementing a fully-connected NN
 - Implementing a Convolutional NN
 - Evaluation and tuning
 - Making predictions
 - Other NN types (RNN, Auto-Encoders, LSTM, ...)