



MACHINE LEARNING

Available from **Feb 2020**

Instructor: Dr. Nayyar Zaidi	Time/Place: TBA
	Duration: 1-3 Days
	Price: Ask for the quote
	Contact: info@datascience-works.com

Description:

Machine Learning has proven itself to be the technology of the future and has already resulted in countless applications. This comprehensive course provides you with a complete introduction to Machine Learning. The course is combination of both theory and practice, and is up-to-date with the latest research and trends in Machine Learning. For example, it covers topics such as Representation Learning (new emerging field with-in Machine Learning), Factorization Machines (very popular in online advertisement placement, large-scale learning, etc.), Feature Engineering (secret sauce behind all practical and effective algorithms), Deep Learning, etc. Other main topics include, fundamental problems such as classification, regression, prediction, anomaly detection, model selection, clustering, dimensionality reduction, recommender systems, etc.

Note, this workshop is not a training of (defacto standard) machine learning languages e.g., R, Python, etc. We expect basic understanding of either one of these languages from you. This workshop is training in underlying machine learning methodologies, fundamental concepts and algorithms that every data scientists should be familiar with. Of course, we illustrate concepts taught in the labs, delivered in Python.

Training Learning Outcome (TLO):

After the training, you are expected to:

- Have an excellent understanding of various topics in Machine Learning.
- Be able to comprehend, communicate, disseminate complex Machine Learning concepts.
- Have a basic to expert understanding of inner workings of various Machine Learning algorithms.
- Be able to scope a Machine Learning project.

Target Audience:

The training is for any aspiring or seasoned data scientists, and is perfect for:

- Computer Scientists and I.T Professionals,
- Engineers (Electrical, Mechanical, Industrial, etc),
- First year Ph.D. students in any field looking to break in Data Science,
- Post-doc fellows and Early Career Researchers in any field.

Duration:

The course is expected to be delivered in 2 days (9-5pm), but the duration can be adjusted based on audience experience and background.

Outline:

The following outline is tentative, and can be customized based on audience demand. The electives can be chosen and replaced from Table 2.

	Day 1	Day 2	Day 3
Session 1	Introduction	Unsupervised Machine Learning	Feature Engineering
Session 2	Supervised Machine Learning	Artificial Neural Networks	Recommender Systems
Session 3	Model Selection	Advanced ML Lab	Reinforcement Learning
Session 4	Supervised Machine Learning lab	Further Topics in Machine Learning	Deep Learning II

Table 1: 3 Days – Training Outline (Day 3 is optional).

Electives	Data Generation	Representation Learning	Audience Presentations
-----------	-----------------	-------------------------	------------------------

Table 2: Other Electives.

Let us delve deep into the details (outline of topics covered) of each session in the following. (E) denotes an elective session.

Introduction

- Machine learning, Artificial Intelligence, Statistics, Data Mining and More
- Machine learning applications
- Introduction to Data Science and Big Data
- Ingredients of Machine Learning – Data, Model and Process
- Training your first practical model

Supervised Machine Learning

- Regression
 - Linear Regression, Polynomial Regression
- Classification
 - Logistic Regression
 - LDA/QDA

- Naive Bayes, Decision Trees
 - Nearest Neighbour Methods
 - Generative vs. Discriminative Learning
 - Forecasting - Time Series Analysis
 - Seasonality and Trends
 - Moving Averages
 - Holt-Winters Method Family
 - ARIMA
-

Model Selection

- Bias and Variance Analysis
 - Achieving Low-variance
 - Regularization
 - Feature Selection
 - Achieving Low-bias
 - Feature Construction
 - Kernels and Kernel trick
 - Feature Engineering
 - Evaluating and Comparing Models
 - Cross-validation
 - Lift Charts, ROC, RPC, other metrics
 - Statistical Tests, Null-Hypothesis, Friedman Statistics, etc.
-

Feature Engineering (E)

- Ensemble-based Methods
 - Bagging, Stacking
 - Random Forest
 - Gradient Boosting
 - XGBoost
 - Extreme Non-linear Models
 - Generalized Linear Models
 - Factorization Machines
 - Support Vector Machines
 - Artificial Neural Networks
 - Need for Deep Learning
-

Artificial Neural Networks

- Foundations of Deep Learning
 - Backpropagation
 - Gradient Computation
 - Optimization
 - Gradient Vanishing/Exploding
 - Model Architecture
 - Batch Normalization
 - Error Surfaces of ANN
 - Miscellaneous Issues
 - Why Deep Learning Now?
-

Unsupervised Machine Learning

- Clustering
 - K-means, DB-Scan, Hierarchical
 - Density Estimation
 - Bayesian Networks
 - EM Algorithms for Clustering and Gaussian Mixture Models
 - Curse of Dimensionality
 - Similarity Measurements
 - Exact vs. Approximate Measures
 - Local Sensitive Hashing
 - Data Pre-processing
 - Dimensionality Reduction
 - Data Standardization
 - Data Munging
 - Feature Hashing
 - Overview of anomaly detection
 - Association rules and discovery
 - APriori Algorithm
-

Representation Learning (E)

- Introduction
 - Principal Component Analysis (PCA)
 - Singular-value-Decomposition (SVD)
 - First look at Embeddings
 - Unsupervised ANN
 - Restricted Boltzmann Machines (RBM)
 - Greedy Layer-wise Training
 - Contrastive Divergence
 - Auto-Encoders (AE)
 - Encoder/Decoder Architecture
 - Stacked Auto-Encoders
 - Sparse Auto-Encoders
 - Embeddings
 - Word2Vec (Skipgram, CBOW), Node2Vec, etc.
 - Transfer Learning
 - Tensor Decomposition
-

Deep learning I (E)

- Introduction
 - Convolution
 - Feature Maps, Max Pooling
- CNN Architectures
 - LeNet-5, AlexNet, VGG-16, ResNet
 - Transfer Learning Revisited
- Object Detection
 - Harnessing the value of Covolution for Object Detection
- Face Learning
 - Face Verification
 - * Triplet Loss

- * Siamese Network
 - FaceNet and DeepFace
 - CNN for non-images
 - 1-d, 3-d and 4-d Convolutions
-

Deep learning II (E)

- Introduction
 - Various Architectures
 - Applications
 - RNN Embeddings
 - LSTM, GRU
 - Attention-based models
-

Data Generation (E)

- Introduction
 - Importance/Rejection Sampling
 - Gibbs Sampling and MCMC
 - Pixel CNN and Pixel RNN
 - Variational Auto-Encoders (VAE)
 - Generative Adversarial Networks (GAN)
 - Conditional GAN
 - Applications
 - Adversarial Learning
 - Defence Methods
 - Attack Methods
-

Reinforcement Learning (E)

- Introduction
 - Markov Decision Process and RL
 - Introduction to Dynamic Programming
 - Q-Learning
 - Deep Q-Learning
 - Double-deep Q-Learning
 - Policy-gradient Methods
 - REINFORCE
 - Variance-control Methods
 - Actor-Critic Models
-

Recommender Systems (E)

- Recommender Systems – An Introduction
 - Evaluating Recommendations
 - Data Structure of Recommender Systems
- Content-based Filtering
 - Neighbourhood-based
 - Similarity-based
 - Model-based
- Collaborative Filtering
 - Baseline Estimates

- Neighbourhood-based
- Model-based
- Singular Value Decomposition (SVD)
 - Matrix Factorization
 - Constrained Matrix Factorization (Non-negative Matrix Factorization)
 - SVD++
 - Collaborative Filtering meets Content-based Filtering
- Advertising on the Web
 - Ad Placement
 - Multi-arm Bandits
 - Contextual Multi-arm Bandits

Further Topics in Machine Learning

- Bayesian Machine Learning
- Text Mining
 - Representing Texts (TF-IDFs)
 - Name Entity Recognition
 - Topic Models (pLSA)
- Learning to Rank
- Causality
 - Randomized Clinical Trials
 - AB Testing from Software Engineering Perspective

Audience Presentations Analytics (E)

- 3-5 minutes presentation by member of the audience on their previous data science project
- Discussion on how the lessons learned from this training can help improve that project

About the Instructor:

Dr. Nayyar Zaidi is the lead Data Scientist at DataScienceWorks and a Senior Lecturer of Computer Science at Deakin University. He received the B.S. degree in computer science and engineering from the University of Engineering and Technology, Lahore, in 2005, and the Ph.D. degree in Artificial Intelligence from Monash University, Melbourne, Australia, in 2011. He worked as a Research Fellow, a Lecturer, and a Research Fellow, from 2011 to 2013, from 2013 to 2014, and from 2014 to 2017, respectively, at the Faculty of Information Technology, Monash University. From 2017 to 2019, he worked as Research Scientist at Credit AI (Trusting Social) Melbourne Lab. His research interests include feature engineering, data generation explainable models and ethical AI. He is also interested in practical (applied) data science, machine learning engineering, and data science training.

About this Document:

©DataScienceWorks Pty Ltd. – All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law. For permission requests, write to the publisher at <http://www.datascience-works.com>.