



DEEP 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:

This comprehensive training provides you with a complete introduction to Deep Learning. Deep Learning has revolutionized analytics in just over five years. The field itself is changing very quickly, with interesting developments every day. This workshop is aimed to give you a complete introduction to Deep Learning through which you can develop a good understand of various Deep Learning concepts. It also provides a good understanding of problem-domains that can be solved by Deep Learning. Practical component will involve building Deep Learning solutions with TensorFlow.

Main topics include, Artificial Neural Networks (ANN), Deep ANN, AutoEncoders, Deep Belief Networks, Convolutional Neural Networks, Recurrent Neural Networks, LSTM, Attention-based methods, Generative Adversarial Networks (GANs), Variational Auto-encoders, Deep Reinforcement Learning, etc.

Training Learning Outcome (TLO):

After the training, you are expected to:

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

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.

	Day 1	Day 2	Day 3
Session 1	Introduction	Representation Learning	Attention-based Models
Session 2	Artificial Neural Networks	Deep Learning II	Generative Neural Networks
Session 3	Deep Learning I	Deep Reinforcement Learning	Deep Learning with GPUs
Session 4	DL Lab	DL Lab II	DL Lab III

Table 1: 2 Days – Training Outline (Day 3 Optional).

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

Introduction

- Machine learning, Artificial Intelligence, Statistics, Data Mining and More
- Linear/Logistic Regression
- Optimization
 - Gradient Descent, Stochastic Gradient Descent, etc.
- Model Selection
 - Regularization
 - Feature Engineering

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?

Deep Learning I

- 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

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

Reinforcement Learning

- 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
-

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
-

Attention-based Models (E)

- Why Attention?
- Transformers

- Different Architectures
 - Applications
 - Text Summarization
 - Building Chatbots
-

Deep Learning with GPUs (E)

- Introduction to CUDA
 - GPUs Set-up
 - Running Deep Learning Models on GPUs
-

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>.