

TRAINING OFFERING



DEEP LEARNING WITH TENSORFLOW AND KERAS

SUBJECT MATTER EXPERT

OVERVIEW

This class is designed to cover key theory and background elements of deep learning, along with hands-on activities using both TensorFlow and Keras – two of the most popular frameworks for working with neural networks. In order to gain an intuitive understanding of deep learning approaches together with practice in building and training neural nets, this class alternates theory modules and hands-on labs.

PREREQUISITES

The class communicates the mathematical aspects of deep learning in a clear, straightforward way, and does not require a background in vector calculus, although some background in calculus, linear algebra, and statistics is helpful. All code examples and labs are done with Python, so previous experience with Python is recommended.

TARGET AUDIENCE



This class is ideal for engineers or data scientists who want to gain an understanding of neural net models and modern techniques, and start to apply them to real-world problems.

FORMAT

25% Lecture/Discussion

75% Hands-On Labs

OBJECTIVES

- - Predictive Modeling and Function Approximation
 - - Limited and General Approximators
- - Intro to Neural Networks and Keras
 - - Perceptrons and feed-forward networks
 - - Using Keras and scikit-learn to build and train linear models

- - Activation Functions and Feature Learning
- - Nonlinear shallow networks with Keras
- - Basic Deep Network Topologies
- - TensorFlow: Intro and Setup
 - - Arithmetic and Optimization with TensorFlow
- - Training
 - - Initializing Weights
 - - Network Structure and Dataset Size
 - - Tuning: Overfitting, Regularization, Dropout
- - Convolution and Convolutional Networks
 - - Simple examples with ConvNets
 - LeNet Derivatives
- Recurrent Nets
 - - Sequences, Feed-forward networks, and Markov Chains
 - - Recurrent Neural Nets and Backpropagation Through Time
 - - Building a RNN
 - - LSTM, GRU
- - Generative Networks
 - - Goals and common key approaches
 - - Variational Autoencoder
 - - GAN (Generative Adversarial Networks)
- - Reinforcement Learning
 - - Policy gradient overview
 - - Deep Q Learning in detail
- - Operations
 - - Adding Deep Learning to Application or Service Architecture
 - - Challenges with Deep Learning
 - - Resources and future directions

HANDS-ON LABS

- -Building Baseline Models
- Training a Feedforward Network on Business Data
- Using Improved Activation Functions
- Solving Problems with Tensorflow
- Vectorized Solutions with Tensorflow (optional)
- Tuning a Convolutional Network for Image Recognition
- Exploring a Recurrent Models
- Learning about the World with Generative Models

~

- Training and Challenging a Reinforcement Learning Agent