ProTech *"Charting the Course ...*

protechtraining.com

... to Your Success!"

Deep Learning Theory: Hands-on Intro with TensorFlow and Keras

Course Summary

Description

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. 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.

Objectives

By the end of this course, students will be able to:

- Articulate why and how modern deep-learning approaches work, what sorts of problems they are suited to solving, and common issues and challenges
- Code, train, and evaluate your own neural network models in addition to using established patterns and pre-trained models.

Topics

- Predictive Modeling and Function Approximation
- Limited and General Approximators
- TensorFlow: Intro and Setup
- Arithmetic and Optimization with TensorFlow
- Backpropagation
- Intro to Keras
- Using Keras and scikit-learn to build and train a shallow linear model
- Activation Functions and Feature Learning
- Derivatives, TensorFlow, and Autodifferentiation
- Initializing Weights
- Nonlinear shallow networks with Keras
- Basic Deep Network Topologies
- Network Structure and Dataset Size
- MNIST Modeling
- Tuning: Overfitting, Regularization, Dropout
- Deeper MNIST Model
- Convolution and Convolutional Networks
- MNIST ConvNet
- LeNet and derivatives
- Sequences, Feed-forward networks, and Markov Chains
- Recurrent Neural Nets and Backpropagation Through Time
- Building a RNN
- Overview: Generative Networks, Reinforcement Learning
- Adding Deep Learning to Application or Service Architecture

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.

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. Coding is done in Python.

Duration

Two days