Days 1 & 2 – Introduction to Machine Learning

Module 1 – Machine Learning Fundamentals

- What is Machine Learning?
- Types of Machine Learning
- What is an Algorithm?
- Did the Machine Learn?
- Classification vs. Regression
- Why Decision Trees?

Module 2 – Exploring the Data

- Course Datasets
- Exploratory Data Analysis (EDA)
- Numeric EDA
- Data Profiling
- Visualizing Features
- Hands-On Lab #1

Module 3 – Classification Trees

- The Intuition
- Overfitting
- Gini Impurity
- Split Quality
- Categorical Data
- Numeric Data
- Hands-On Lab #2

Module 4 – Awesome Classification Trees

- Under/Overfitting
- The Bias-Variance Tradeoff
- Supervising the Data
- Model Tuning
- Classification Tree Pruning
- Measuring Awesomeness
- Splitting Data
- Hands-On Lab #3

Module 5 – Feature Engineering

- Intuition
- Data Leakage
- Decision Boundaries
- Engineering Numeric Features
- Engineering Categorical Features
- Engineering Date-Time Features
- Missing Data
- Hands-on Lab #4

Module 6 – Regression Trees

- Fundamentals
- Numeric Feature MSE
- Categorical Feature MSE
- Tuning
- Imputation
- Hands-On Lab #5

Module 7 – The Mighty Random Forest

- Bad, Tree! Bad!
- Ensembles
- Bagging
- Feature Randomization
- Hands-on Lab #6

Module 8 – Awesome Random Forests

- Feature Importance
- Tuning
- Model Testing

Module 9 – Workshop Wrap-Up

- Additional Resources
- Want to Kaggle?

Day 3 – Time Series Forecasting with Python

Module 1 – Forecasting Fundamentals

- What is Forecasting?
- What is a Time Series?
- Time Series Characteristics

Module 2 – Simple Forecasting Models

- Our First Dataset
- Naïve Forecasting Models
- Moving Average Forecasting Models
- Hands-On Lab #1

Module 3 – Evaluating Forecasting Models

- Is Your Model "Good"?
- Splitting Time Series Data
- Model Bias
- Model Mean Absolute Error (MAE)
- Validating Simple Forecasting Models
- Hands-On Lab #2

Module 4 – Lagged Features

- Time Series for Machine Learning
- What Are Lags?
- Creating Lagged Features
- Training an ML Forecasting Model
- Making Predictions
- Validating Predictions
- Hands-on Lab #3

Module 5 – Adding Features

- Endogenous vs. Exogenous Features
- Endogenous Features
- Exogenous Features
- Exogenous Features Are Everywhere
- Hands-on Lab #4

Module 6 – Feature Importance

- Are Your Features Good?
- Permutation Feature Importance

Module 7 – Differencing

- The Problem with Regression Forests
- Predicting Changes
- Differencing Example
- Using Differences
- Evaluating Differencing Models

Module 8 – Course Wrap-Up

- Continue Your Learning

Hands-on Lab #5