Introduction to Machine Learning

Summary

This course will provide an overview of the fundamentals of machine learning. Students will learn about the type of problems that can be solved, the building blocks and the fundamentals of building models in machine learning. A number of key algorithms will be explored. By the end of this course, students will leave with practical knowledge in a number of supervised learning algorithms along with an understanding of key concepts like under and overfitting, regularization, and cross validation. Students will be able to identify the type of problem they're trying to solve, choose the right algorithm, tune parameters, and validate a model.

The curriculum is structured around 12 weeks of lectures and exercises. Each week requires ~3 hours to complete. The exercises are implemented in Python, so familiarity with the language is encouraged though not required (you can learn along the way)

Week1: Introduction to Machine Learning and Tools

Week 1 introduces the basic data science toolset – Jupyter notebooks for interactive coding; Numpy, Scipy and Pandas for numerical computation; Matplotlib and Seaborn for data visualization and Scikit Learn for machine learning libraries. You'll be using these tools to work through the exercises each week.

Week 2: Introduction to Supervised Learning and K Nearest Neighbors

Week 2 introduces the basic concepts and vocabulary of machine learning. You'll learn about Supervised Learning and how it can be applied to regression and classification problems. This lesson explores the K-Nearest Neighbor (KNN) algorithm for classification.

Week 3: Train Test Splits Validation Linear Regression

Week 3 covers some of the core model generalization principles. You will walk away knowing the difference between over-fitting and under-fitting a model, bias-variance tradeoff, finding the optimal training and test data set splits, cross-validation and model complexity vs error. This lesson also introduces the Linear Regression model for supervised learning.

Week 4: Regularization and Gradient Descent

Week 4 builds on concepts taught in previous weeks. You will learn about cost functions, regularization, feature selection, hyper-parameters and more complex statistical optimization algorithms like Gradient Descent and its application to linear regression.

Week 5: Logistic Regression and Classification Error Metrics

What is Logistic Regression and how is it different from Linear Regression? What are the metrics for classification error and in what scenarios can each be used?

Week 6: Naïve Bayes

This week takes you back to the basics of probability theory and its application to the Naïve Bayes classifier. What are the different types of Naïve Bayes classifiers and how do you train a model using this algorithm?

Week 7: SVM and Kernels

Advanced supervised learning algorithms – This week covers Support Vector Machines (SVMs), a popular algorithm used for classification problems. Through examples, you will learn their similarity to logistic regression, calculating the cost function of SVMs, regularization in SVMs and some tricks to obtain non-linear with SVMs.

Week 8: Decision Trees

Advanced supervised learning algorithms – What are decision trees and how can they be used for classification problems? How can you identify the best split? What are the factors for splitting? What are the strengths and weaknesses of decision tress? Also learn about regression trees that help with classifying continuous values.

Week 9: Bagging

Advanced supervised learning algorithms – In week 8, you learnt how decision trees tend to overfit the data. How can this variance be reduced? This week, you will learn about the concepts of bootstrapping and aggregating (commonly known as “Bagging”) to reduce variance. You will also learn about the Random Forest algorithm that further reduces the correlation seen in bagging models.

Week 10: Boosting and Stacking

Advanced supervised learning algorithms – This week, you will learn about Boosting algorithm that helps reduce variance and bias.

Week 11: Introduction to Unsupervised Learning and Clustering Methods

So far, the course focus has heavily been on supervised learning algorithms. This week, we will learn about unsupervised learning and how they can be applied to clustering and dimensionality reduction problems.

Week 12: Dimensionality Reduction and Advanced Topics

Dimensionality refers to the number of features in the dataset. Theoretically, more features should mean better models, but this is not true in practice. Too many features could result in spurious correlations, more noise and slower performance. This week, you will learn algorithms like Principal Component Analysis (PCA), Multi-Dimensional Scaling (MDS) that can be used to achieve a reduction in dimensionality.

Introduction to Deep learning

Summary

This course provides an introduction to deep learning. Deep learning has gained significant attention in the industry by achieving state of the art results in computer vision and natural language processing. Students will learn the fundamentals of deep learning and modern techniques to build state of the art models.

By the end of this course, students will have a firm understanding of the techniques, terminology, and mathematics of deep learning. They’ll know fundamental neural network architectures, feed-forward networks, convolutional networks, and recurrent networks, and how to appropriately build and train these models. An understanding of various deep learning applications will be gained. Finally, students will be able to use pre-trained models for best results.

Week 1: ML101 Review
Recap of the Machine Learning course. Students that are experts in machine learning can skip to the next week's class.

Week 2: Introduction to Neural Networks
The inspiration for neural networks comes from biology. This week, you will learn the basic nomenclature in deep learning – what is a neuron (and it's similarity to a biological neuron), the architecture of a feed-forward neural network, activation functions and weights.

Week 3: Backpropagation
In week 2, you learnt how a neural network computes the output given an input in a single forward pass. How can this network be used to train a model? This week, you will learn how you can calculate the loss and adjust weights using a technique called backpropagation. Different types of activation functions will also be introduced.

Week 4: Training Neural Nets Keras
This week, you will learn about techniques to improve training speed and accuracy. What are the pros and cons of using Gradient Descent, Stochastic Gradient Descent and mini-batches? With the foundational knowledge on neural networks covered in week 2 – week 4, you will learn how to build a basic neural network using Keras with Tensorflow as the backend.

Week 5: Regularization Dropout
How can you prevent overfitting (regularization) in a neural network? In week 5, you will learn about penalized cost function, dropout, early stopping, momentum and some optimizers like AdaGrad and RMSProp that help with regularizing a neural network.

Week 6: Introduction to Convolutional Neural Networks
In week 6, you will learn about Convolutional Neural Networks (CNN) and contrast with the Fully-Connected neural networks we have seen so far. You will also learn how to build a CNN by choosing the grid size, padding, stride, depth and pooling.

Week 7: Transfer Learning
Using the LeNet-5 topology, you will apply all of the CNN concepts learnt in week 6 to the MNIST (Modified National Institute of Standards and Technology) dataset for handwritten digits. With a trained neural network, how can the primitive features learned in the first few layers be generalized across image classification tasks? Learn how Transfer Learning can help with this.

Week 8: Network Architectures
Deep learning literature talks about many image classification topologies like AlexNet, VGG-16/VGG-19, Inception and ResNet. In this week, you will learn how these topologies are designed and the usage scenarios for each.

Week 9: CNN Advanced Techniques
One practical obstacle to building image classifiers is obtaining labeled training data. In this week, we explore how we can make the most of the available labeled data using data augmentation. You will then implement data augmentation using Keras.

Week 10: Text Word Vectors

So far, we have used images as inputs to neural networks. Image values are essentially numbers (grayscale or RGB). But, how do we work with text? How can we build a neural network to work with pieces of text of variable length? How do we convert words into numerical values? In this lecture, you will learn about Recurrent Neural Networks and their application to Natural Language Processing (NLP).

Week 11: Recurrent Neural Networks

In week 11, you will learn more advanced topics of developing a RNN. You will learn how the concept of recurrence can be used to solve the issue with variable sequence and ordering of words. Take out your notebook and pencil and work through the math of RNNs.

Week 12: LSTM RNNs

Standard RNNs have poor memory. In NLP, it is important to have a structure that can carry forward some of the signal over many steps. In this week, you will learn about the Long Short Term Memory (LSTM) that addresses this problem.