

Data Science II

Description:

This class will provide a solid introduction to the field of deep learning and students will learn to use them to build powerful predictive modeling training deep neural networks and apply them to real-world problems. This course covers a variety of topics including SVM, Deep Neural Networks, Convolutional Neural Networks, Recurrent Neural Networks, and Reinforcement Learning.

Course Objectives:

At the end of this course, students should be able to (1) understand several advanced topics in Machine Learning, (2) learn to build and train a deep learning model, (3) applying deep learning techniques to different problems.

Prerequisite:

Students are expected to have pre-existing knowledge of probability, linear algebra, statistics, machine learning, and basic programming.

Course schedule:

The approximate number of lectures is tentative. Class time will also be used to cover programming sessions, conduct exam questions, and review sessions. Chapter sections are listed below.

1. ML Review
2. Decision Tree
3. Ensemble Methods
4. Support Vector Machine
5. Unsupervised Learning
6. Neural Network
7. Deep Learning
8. Convolutional Neural Networks
9. Recurrent Neural Networks
10. Reinforcement learning

Course Materials:

The course material will be available periodically.

There is no official textbook for the class but a number of the supporting readings will come from: “The Elements of Statistical Learning: Data Mining, Inference, and Prediction”, Trevor Hastie, Robert Tibshirani, Jerome Friedman, 2001.

“Machine Learning”, Tom Mitchell. McGraw-Hill, 1997.

“Deep Learning”, Ian Goodfellow, Yoshua Bengio, and Aaron Courville.

“Reinforcement Learning: An Introduction”, Sutton and Barto, 2nd Edition.