CAP 5625 Computational Foundations of Artificial Intelligence

Credits: 3 credits

Textbook, title, author, and year: The Elements of Statistical Learning: Data Mining, Inference, and Prediction, by Trevor Hastie, Robert Tibshirani, and Jerome Friedman. Springer, 2009, 2nd Ed, 2009, ISBN-13: 978-0-3878-4857-0.

**** Free ebook** from author website https://web.stanford.edu/~hastie/ElemStatLearn/

Reference materials: *An Introduction to Statistical Learning: with Applications in R*, by Gareth James, Daniela Witten, Trevor. Hastie, and Robert Tibshirani. Springer, 2017, ISBN-13: 978-1-4614-7137-0.

**** Free ebook** from author website http://faculty.marshall.usc.edu/gareth-james/ISL/

Pattern Recognition and Machine Learning, by Christopher M. Bishop. Springer, 2006, ISBN-13: 978-0-3873-1073-2.

**** Free ebook** from author website https://www.microsoft.com/en-us/research/people/cmbishop/

Specific course information

Catalog description: This course covers the mathematical and programming foundations of artificial intelligence (AI) and machine learning (ML) using contemporary programming languages and tools. As a result, students develop familiarity with mathematical methods (and associated notation, software packages and libraries) that are widely used in AI and ML projects and literature

Prerequisites: Graduate standing or permission of instructor

Specific goals for the course: By the end of the course, students will be able to:
Understand the mathematical foundations of machine learning. Demonstrate proficiency in solving machine learning problems.
Identify and apply statistical and computational models to machine learning problems.
Analyze the performance of particular machine learning models, and justify their use and limitations.

Brief list of topics to be covered:

- Linear regression
- Linear and non-linear regression and model selection
- Feature selection and regularization
- Advanced regularization techniques
- Principal components analysis and regression
- Discriminant analysis
- Logistic regression
- Support vector machines
- Neural networks
- Random forests and boostin
- Unsupervised learning
- Student presentations on special topics in machine learning