CS 7200 : MACHINE LEARNING

Semester Hours:	3.0	*Contact Hours:* 3
Coordinator	Qing Tian	
Text	1. Pattern Recognition and Machine Learning 2. Deep Learning (Adaptive Computation and Machine Learning Series)	
Authors:	1. BISHOP 2. GOODFELLOW, BENGIO, AND COURVILLE	
Year	1. 2011 2. 2016	

SPECIFIC COURSE INFORMATION

Catalog Description:

The course provides foundations of machine learning, mathematical derivation and implementation of the algorithms and their applications. Topics include supervised learning, learning theory, graphical model, reinforcement learning, Bayesian techniques, and deep learning. In addition, practical applications are considered using the machine learning algorithms. The course also requires an open-ended research program. Prerequisites: CS 5200 and STAT 5020, or permission of instructor. Credit cannot be received for both DATA 7200 and CS 7200.

Course type: **ELECTIVE**

SPECIFIC COURSE GOALS

- I am able to explain the mathematical foundations of machine learning.
- I am able to explain the differences between supervised and unsupervised learning.
- I am able to design learning model for a real-world application.
- I am able to develop and implement well-known supervised learning algorithms.
- I am able to utilize modern frameworks and tools for deep learning.

LIST OF TOPICS COVERED

- Introduction (~5%)
    - Overview
    - Probability review, loss function, maximum likelihood
    - Linear regression, gradient descent, Newton method
- Classification (~10%)
    - K-nearest neighbors, Naïve Bayes
    - Linear & Gaussian discriminant analysis
    - General linear models
- Kernel Methods (~10%)
    - Kernel density estimation, kernel regression
    - Support vector machines
    - Convex optimization*
- Regularization (~10%)
    - L2 regularization
    - L1 regularization, sparsity and feature selection
    - Bias-variance tradeoff, overfitting
    - Developing basic machine learning algorithms
- Neural Networks (~10%)
    - Perceptron, multilayer perceptron
    - Back-propagation
- Learning Theory (~10%)
    - Sample complexity
    - Probably Approximately Correct (PAC) learning
    - Error bounds
    - Vapnik-Chervonenkis (VC) dimension*
- Graphical Model (~10%)
    - Bayesian networks
        - Representation, inference, maximum likelihood estimation, hidden Markov models
    - Structure learning
    - Bayesian inference and learning*
- Unsupervised Learning (~10%)
    - Clustering: K-means
    - PCA
    - Gaussian mixtures*
- Reinforcement Learning (~10%)
    - Markov Decision Processes (MDP)
    - Value function approximation
    - Dynamic programming*
- Deep Networks & Learning (~15%)

- Deep feedforward networks
- Regularization for Deep Learning
- Optimization for training deep models
- Convolutional networks
- Deep generative models*
- Practical methodology, applications

(*Optional, if time allows.)

RECOMMENDED REFERENCES

- Machine Learning: A Probabilistic Perspective, by Kevin Murphy, MIT Press, 2012
- Elements of Statistical Learning, by Hastie, Tibshirani, Friedman, Springer, 2010
- Bayesian Reasoning and Machine Learning, by David Barber, Cambridge University Press, 2012
- Reinforcement Learning: An Introduction, by Sutton and Barto, MIT Press, 1998