CS 3200: INTRODUCTION TO MACHINE LEARNING

Semester Hours:	3.0	*Contact Hours:* 3
Coordinator:	Shuteng Niu	
Text:	TBD	
Author(s):	TBD	
Year:	TBD	

SPECIFIC COURSE INFORMATION

Catalog Description:

The course provides an introduction to machine learning and implementation of simple machine learning algorithms. Topics include regression, classification, supervised learning, unsupervised learning, and ethics in Machine Learning. Practical applications are also considered using various machine learning libraries and tools. Prerequisites: C or Better in CS 2020 and a corequisite of MATH 3320 or equivalent.

Course type: **SELECTIVE ELECTIVE**

SPECIFIC COURSE GOALS

- I can design learning model for a simple application.
- I can explain the differences between supervised and unsupervised learning.
- I can implement a well-known supervised learning algorithm.
- I can implement a well-known unsupervised learning algorithm.
- I can identify an ethical problem in machine learning.

LIST OF TOPICS COVERED

- Introduction (~5%)
  - What is machine learning
  - Types of machine learning
  - Probability review
- Input to Output (~10%)
  - Functional learning, parametric and non-parametric functions

    - Bias and variance, Bias-variance trade-off
    - Overfitting and underfitting
- Linear Regression (~10%)
    - Linear fitting, least squares fit
    - Loss function
    - Polynomial and quadratic models
- Linear Discriminant Analysis (~15%)
    - Learning LDA models
    - Making prediction
    - Bayes Theorem
- Bayesian Classifiers – Naïve Bayes (~10%)
    - Bayesian classification
    - Naïve Bayes classifier
    - Examples
- K-Nearest Neighbors (~10%)
    - Instance-based classifiers
    - Nearest neighbor classifier
    - Distance metrics
    - Choosing K, scaling issues
- Neural Networks (~20%)
    - Artificial neurons, activation functions, learning networks
    - Perceptrons
    - Training, error functions
    - Backpropagation
    - Sigmoid unit, convergence, overfitting
    - Intro. to Convolutional Neural Networks
- Support Vector Machines (~10%)
    - Linearly separated classes
    - Computation of optimal hyperplane
    - Maximum margin
    - Non-linearity

    - Kernel trick
- Ethics (~10%)
    - Data
    - Algorithm
    - Results

RECOMMENDED REFERENCES

- Hands-on Machine Learning with Scikit-Learn & TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems by Aurelien Geron, O'Reilly Media, 2017
- Make Your Own Neural Network by Tariq Rashid, CreateSpace, 2016
- Deep Learning by Ian Goodfellow, Yoshua Bengio, and Aaron Courville, MIT Press, 2016