CISC484: Intro to Machine Learning

Credits: 3

1. Instructor Information

Instructor: Dr. Xi Peng

Email: xipeng@udel.edu

Instructor information: https://sites.google.com/site/xipengcshomepage/

2. Prerequisites

- **Mathematics Background**:
  - Calculus; *(require)*
  - Linear Algebra; *(require)*
  - Statistics. *(recommend)*
- **Computer Sciences Background:**
  - Data Structure; *(require)*
  - Algorithm; *(recommend)*
  - Numerical Analysis. *(recommend)*
- **Programming background:**
  - Python. *(require)*

3. Course Calendar

4. Course Description

This course introduces the preliminary theory, models, and algorithms of machine learning. Topics covered include regression, classification, clustering, and deep learning. The students are required to accomplish a series of math (knowledge foundation) and mini-project (python programming) homework, as well as a midterm exam and a final project. The goal is two-fold: 1) understand fundamental machine learning concepts and their underlying mathematical background; 2) program machine learning models and algorithms to solve practical tasks and real-world problems.

Topics:

- **Machine Learning Models:**

    - Linear Regression
    - Logistic Regression
    - Support Vector Machine
    - Kernel Methods
    - Random Forest
    - Deep Neural Networks
    - Convolutional Neural Networks
    - Recurrent Neural Networks (tentative)
- **Machine Learning Concepts:**
    - Model Representation
    - Cost Function
    - Optimization
    - Regularization
- **Machine Learning Knowledge:**
    - Gradient Descent Algorithm
    - Decision Boundary
    - Large-margin Optimization
    - Kernel Method
    - Bias-Variance Tradeoff
- **Machine Learning Programming**
    - Python
    - Libraries: Numpy, Scipy, Scikit-learn, Matplotlib, ...

5. Resources

- **Course Slides (Required).**
- **Textbook** (for questions, references, and more):
    - "Machine Learning, A Probabilistic Perspective", K. Murphy (2012). *(recommend)*
    - "Pattern Recognition and Machine Learning", C. Bishop (2006). *(recommend)*
- **Online Resources:**
    - Machine Learning
        - Coursera-Machine Learning (Andrew Ng, Stanford)
        - Least Squares in Matrix Form
    - Python
        - A Visual Intro to Numpy and Data Representation

- Statistics
  - Probability Review (David Blei, Princeton) *(recommend)*
- Linear Algebra
  - Linear Algebra Tutorial (C.T. Abdallah, Penn) *(recommend)*
  - Linear Algebra Review and Reference (Zico Kolter and Chuong Do, Stanford)
  - Linear Algebra Lecture (Gilbert Strang, MIT)

6. Final Grade Breakdown

Course Component	Percentage of Total
Five math/programming homework (individual)	**50% (10% each)**
Five in-class quiz (Optional)	**10% (2% each)**
Mid-term exam (one-page cheating sheet)	**30%**
Final project (groups)	**15%**
Attendance	**5% (sign-in)**

7. Grading and Submission Policy

- **Homework (50%):**
  - All homework assignments are **individual** problems and must be done **individually**;
  - Homework will be released on Tuesdays, due in two weeks and at the same time release the next homework;
  - **100%** grade penalty if group work OR code sharing OR online copy is detected;
  - Late submission will be charged by **20%** penalty each late day and **3** days maximum;
  - Please submit the homework to **Canvas**;
  - Please **visit or contact TA first** for any issues regarding homework submission and grading.
- **In-class quiz (bonus 10%):**

- Five voluntary quiz to help you review and course materials;
- The in-class quiz consists of the course material covered by the exam day;
- You can choose to submit or not submit your answer.

- **Mid-term (30%):**
  - The 1-hours exam consists of the course material covered by the exam day;
  - Closed book except for **One A4 page double-side cheating sheet**.
- **Final project (15%):**
  - Group work (Totally 15 groups);
  - Students in the same group will receive the same score;
  - Proposal:
    - In-class presentation: ~5-page slides plus 7-min pitch (starting from the 4th class);
    - Group size: <= 2 students;
    - Approve or Revise;
  - Presentation (10%):
    - In-class presentation: ~15-page slides plus 15-min pitch;
    - Crowdsourcing grading;
  - Online Webpage (5%):
    - Create a webpage link to present the outcomes of the projects;
- **Attendance (5%):**
  - Attendance is **mandatory**;
  - Please wear masks;
  - >= 3 absences without excuses will get a 15% deduction of the final score.
- **Final grading curve:**
  - The score in each category is less important than the score relative to the class average;
  - There is no fixed curve. If everyone performs well then everyone can get top grades.