CISC684: Intro to Machine Learning

Semester: Fall 2023 | Credits: 3 | Section: 010
Class Hours: **Mon & Wed 5:20 PM - 6:40 PM** | In-person

1. Instructor Information

Instructor: Dr. Xi Peng
Email: xipeng@udel.edu
Office hours by schedule
Office: FinTech 416C

TA: Kien Nguyen
E-mail address: kxnguyen@udel.edu
Office hours: Friday 4PM - 6PM
Office: Smith 102A

2. Prerequisites

- **Mathematics Background**:
    - Calculus; *(require)*
    - Linear Algebra; *(require)*
    - Statistics. *(require)*
- **Computer Sciences Background:**
    - Data Structure & Algorithm; *(require)*
    - Intro to Machine Learning or AI; *(require)*
    - Other machine learning related courses. *(recommend)*
- **Programming background:**
    - Python. *(require)*

3. Course Description

This course introduces the preliminary theory, models, and algorithms of machine learning and deep 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 an exam and a final project. The goal is twofold: 1) understand fundamental machine learning concepts and their underlying mathematical background; 2) program machine learning and deep learning models and algorithms to solve practical tasks and real-world problems. More specifically, topics include linear/logistic regression, support vector machine, DNN, CNN, RNN, GAN, etc.

Topics (tentative):

- **Conventional machine learning topics (6 weeks):**
  - Regression Model: linear regression, polynomial regression
  - Classification Model: logistic regression, Bayesian classification, support vector machine
  - Machine Learning Foundation: regularization, gradient descent, maximum likelihood estimation
- **Deep learning topics (4 weeks):**
  - Deep Neural Network
  - Convolutional Neural Network
  - Recurrent Neural Network
  - Deep Learning Theory
- **Advanced deep learning topics (4 weeks):**
  - Generative Adversarial Network (GAN) and Theory
  - Meta Learning and Deep Transfer Learning
  - Graph Neural Network (GNN)
  - Large Language Model (LLM)
  - Explainable Machine Learning (XML)
- **Machine learning programming**
  - PyTorch, Numpy, Scipy, Scikit-learn, Matplotlib, ...

4. Resources

- **Course slides:**
  - All the slides will be uploaded before/after the lecture.
  - This is the main learning resource.
  - All the textbooks are recommended but not required.
- **Textbook**:
  - "Deep Learning," I. Goodfellow (2015). *(recommend)*
  - "Machine Learning, A Probabilistic Perspective," K. Murphy (2012). *(recommend)*
  - "Pattern Recognition and Machine Learning," C. Bishop (2006). *(recommend)*
- **Online Resources**
  - 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)
    - Python
        - A Visual Intro to Numpy and Data Representation
    - Machine Learning
        - Coursera-Machine Learning (Andrew Ng, Stanford)
        - Least Squares in Matrix Form

5. Final Grade Breakdown

Course Component	Percentage of Total
Five programming homework (individual) (10% each)	**50%**
Final project (group) • Proposal (5-min) • Presentation (15-min) (10%) • Report (4-page) (10%)	**20%**
Exam	**30%**

6. Grading and Submission Policy

- **Homework (50%):**
    - All homework assignments are **individual** problems and must be done **individually**;
    - PDF report to include all results;
    - **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**;
- **Final project (20%):**

    - Group of at least 3 at most 4;
    - Topics include but are not limited to:
        - Sentiment Analysis
        - Recommender System
        - Forecasting (i.e., stock, weather, sales)
        - Object Detection
        - Style Transfer (GAN)
    - Proposal:
        - Sign up your group's slot in the Course Schedule Google Sheet
        - In-class presentation: **5-page slides plus 5-min pitch**
        - Approve or Revise
    - Presentation (10%):
        - In-class presentation: **15-page slides plus 15-min presentation**
    - Report (10%):
        - **4-page PDF** minimum
        - Crowdsourcing grading
- **Exam (30%):**
    - In-person.
    - Last day of class (Dec 13th).
- **Attendance:**
    - Attendance is **required** with a sign-in sheet occasionally.
    - At most 3 absences without excuse. Send emails to instructor and TA before absences.
- **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.