CSE4820/CSE5819. Introduction to Machine Learning

The most up-to-date syllabus can be found in HuskyCT course page.

Instructor: Jinbo Bi
Time: TuTh 12:30-1:45PM
Place: MCHU (McHugh Hall) 306
Email: jinbo.bi@uconn.edu
Office Hours: M10-11AM, W11-12noon

Teaching Assistant
Graduate TA: Zhili Zhang (zhili.zhang@uconn.edu)
Office Hours: W3-4PM, Th2-3PM (ITEB221)
Undergraduate TA: Illy Hoang (illy.hoang@uconn.edu)
Office Hours: M4-5PM, Tu5-6PM (ITEB114)

This course offers an introduction to the basic concepts and techniques of machine learning, including supervised and unsupervised learning tasks and models, related optimization techniques, and methods for model validation and data preparation. Topics may include linear and logistic regression, regularization, support vector machines, neural networks, clustering, dimension reduction, as well as the gradient descent optimization algorithm. The materials covered in this course may vary slightly across different semesters.

The course consists of lectures, programming-based and non-programming homework assignments, and exams. Lectures will serve as the vehicle for the instructor to introduce concepts and knowledge to students. Exams are used to test if basic concepts have been mastered. Programming-based assignments will be used for students to get profound hands-on experience by programming or experimenting with certain machine learning algorithms.

This Fall2024 semester will be designed to test if ChatGPT can help pave individualized learning path for students.

Course Objectives:

- Master basic concepts of machine learning, such as supervised and unsupervised learning, overfitting, regularization, model validation, data-driven approach, and related optimization theory for model training
- Understand some of the widely used classic machine learning methods and model evaluation metrics, such as linear and logistic regression, k-means, receiver operating characteristic curve, dendrogram for clustering
- Get to know some of the modern deep learning methods, such as ReLU-based fully-connected networks, convolutional neural networks
- Hands-on experience with using some machine learning methods to solve practical problems

Prerequisites:

Math skills: Calculus (MATH 1131Q), Linear Algebra (MATH 2210Q), and Statistics (STAT 3025Q, or 3345Q or 3375Q, or MATH 3160)
Programming skills: fluent with Python

Textbook (Optional):

- *Hands-on Machine Learning with Scikit-Learn, Keras and TensorFlow* by Aurélien Géron, ISBN: 9781492032649
- *Pattern Recognition and Machine Learning* (Information Science and Statistics) by Christopher M. Bishop, ISBN-10: 0387310738
- *Machine Learning: A Probabilistic Perspective* by Kevin Murphy, ISBN-10: 0262018020
- *Introduction to Data Mining* by Pang-Ning Tan, Michael Steinbach, Vipin Kumar, ISBN-10: 0133128903
- *The Elements of Statistical Learning* by Trevor Hastie, Robert Tibshirani, Jerome Friedman, ISBN-10: 0387848576

The textbooks are not required. These textbooks are recommended reading, and some of course materials come from these books. Reading these books may help understand the course materials and expand the content discussed in lectures. Lectures will come with slide files, reading materials, and survey papers for students to study after lectures.

Grading:

For CSE4820,

1. HW assignments account for 60% of the final grade (change from 40% - 60%) We will remove the worst HW score, and use the average score of the remaining HWs to map to 60% of the final grade. Missed HW(s) will automatically receive a score of 0.
2. The three exams (two prelims and one final) will account for 40% of the final grade (change from 60% to 40%). We will remove the worst exam score, and use the average score of the other two exams to map to 40% of the final grade. Missed exams will automatically receive a score of 0. (For instance, if students have good scores for the two prelims, they may not come to the final.)

For CSE5819,

1. HW assignments account for 40% of their final grade (no change). We also remove the worst HW score as well to compute the average of HW scores. Missed HW(s) will automatically receive a score of 0.
2. The three exams (two prelims and one final) will account for 40% of the final grade (change from 60% to 40%). We will remove the worst exam score, and use the average score of the other two exams to map to 40% of the final grade. Missed exams will automatically receive a score of 0. (For instance, if students have good scores for the two prelims, they may not come to the final.)
3. A term project is required and accounts for 20% of the final grade. The term project can be performed in a group of maximal 5 students. By the end of the seven weeks into the semester, student groups need to submit their project ideas for evaluation of whether they are suitable for this course and whether they carry enough challenges. Student groups will give project presentations in the last week of lectures, and turn in a written report and all related codes/data for evaluation by the end of the final exam week.

For both CSE4820/CSE5819
Additional in-class quizzes (quizzes are not graded in terms of correctness but students need to take at least half of the quizzes to receive the credit) 5%.

Letter Grades will be given as **A** 100-93; **A-** 92-90; **B+** 89-87; **B** 86-83; **B-** 82-80; **C+** 79-77; **C** 76-73; **C-** 72-70; **D+** 69-67; **D** 66-63; **D-** 62-60; **F** 59 or less
We will round your final scores up for letter grade.

Each HW assignment will generally contains two parts: a ChatGPT interaction part; and a problem answering part. As long as the first part is turned in, students will get 20% of that HW grade. The second part will be graded with respect to the answer correctness, possessing 80% of the grade for that HW.

Each student can have a total of 5 free late calendar days to use for homework assignments. Each 24 hours or part thereof that an assignment is late are counted as one full late day. Once late days are exhausted, no late assignments will be accepted for any reason. Students are highly encouraged to reserve your late days for unavoidable emergencies or planned travel.

If you miss an exam due to sickness, you will need a doctor note that details the period during which you were medically incapable of taking the exam for a potential makeup. Notify the instructor immediately via email if you are going to miss an exam before the exam takes place unless medically impossible; or otherwise notify the instructor as soon as you return to class.

Tentative Weekly Schedule:

Week 1: Introduction
Week 2: Basic linear algebra and statistics
Week 3: Introduction to optimization
Week 4: Supervised learning – logistic regression (estimating probabilities from data)
Week 5: Supervised learning – linear/ridge regression (regularization principle)
Week 6: Review and prelim exam 1
Week 7: Support vector machines
Week 8: Perceptron, multi-layer perceptron
Week 9: Supervised learning algorithm and model evaluation
Week 10: Unsupervised learning – Clustering
Week 11: Review and prelim exam 2
Week 12: Deep learning – convolutional neural networks
Week 13: Deep learning – recurrent neural networks or transformers
Week 14: Additional DL topics and review for final

Course Policy:

1. For in-person lectures, computers are allowed in classroom for taking notes or any activity related to the current class meeting.
2. For online sessions, communications with the instructor and TA may happen in a specific timeslot or by appointments.
3. Exams and homework assignments are graded by a teaching assistant or an online grading system if one has been set up.
4. Prelim exams and the final are close-book, close-notes, and no computers are allowed. It is not allowed to search answers online or ask help from others. A calculator can be used.

HuskyCT:

A Husky CT site has been set up for the class. You can access it by logging in with your NetID and password. You must use HuskyCT for submitting assignments. The instructor uses the HuskyCT announcement to announce class materials, grades, problem clarifications, changes in class schedule, and other class announcements.

Academic Integrity:

You are expected to adhere to the highest standards of academic honesty. Use of published materials is allowed, but the sources should be explicitly stated in your solutions. Violations will be reviewed and sanctioned according to the University Policy on Academic Integrity.

“Academic integrity is the pursuit of scholarly activity free from fraud and deception and is an educational objective of this institution. Academic dishonesty includes, but is not limited to, cheating, plagiarizing, fabricating of information or citations, facilitating acts of academic dishonesty by others, having unauthorized possession of examinations, submitting work for another person or work previously used without informing the instructor, or tampering with the academic work of other students.”