CS 167: Machine Learning
Fall 2024

Instructor: Md Alimoor Reza

Assistant Professor of Computer Science
Department of Mathematics and Computer Science
Drake University

Logistics: The lectures and most course content will be hosted at `https://drake.blackboard.com`. Coursework will include attending weekly lecture sessions, and completing individual (and collaborative) activities: programming assignments, in-class activities, quizzes, participation in polls, and two projects.

Class Meeting Time and Place: We will meet in person twice weekly for lecture sessions. Section-specific information is as follows:

- Location: Meredith Hall # 235
- Time: Monday/Wednesday 9:30 am - 10:45 am (CRN#10790)
- Time: Tuesday/Thursday 11:00 am - 12:15 pm (CRN#8363)

Office hours: The instructor will hold weekly office hours, either in person or online via Zoom. Please see Blackboard for up-to-date times and URLs. If you cannot make office hours, contact the instructor to make a separate appointment.

- Office hour#1: Tuesday: 12:00 pm - 2:30 pm
- Office hour#2: Thursday 12:00 pm - 2:30 pm
- Location: Collier-Scripps#323

Course Communication: The course content will be posted on the course Blackboard page. Any announcements will be sent using Blackboard's announcement feature (as well as email). It is a personal goal of mine to respond to student emails within 24 hours of receipt.

Course overview: This course introduces approaches to developing computer programs that learn from data. Both foundational and contemporary machine learning algorithms will be covered in the context of a variety of data and problem types. Specific topics will vary but may include:

- *k-nearest-neighbors*
- *decision trees*
- *random forests*
- *support vector machines*
- *artificial neural networks*
- *convolutional neural networks*
- *recurrent neural networks*
- *transformers*
- *other relevant advanced machine-learning techniques*

Students will develop their own implementations of the algorithms as well as utilizing modern machine learning software and programming libraries.

Learning outcome: After successfully taking this course, you will be able to:

- Apply a variety of modeling techniques to classification, regression, and unsupervised learning problems using data in different formats (such as typical structured data, text, and images).
- Create software that utilizes machine-learning programming libraries in order to conduct machine-learning-based data analysis.
- Develop and conduct machine-learning-based data analysis experiments, and they will be able to interpret and explain the results.

- Feel comfortable with using industry-standard tools such as Google Colab, GitHub, etc
- Understand fundamentals of machine learning
- Gain an understanding of the advantages and disadvantages of different learning paradigms so that students can choose appropriate solutions given a problem description
- Receive hands-on experience with commonly used algorithms and software tools within machine learning

Schedule, readings, and resources will be available via Blackboard, `https://drake.blackboard.com`.

Grading and requirements:

- *Notebook assignments (30%):* Take home assignments (submit on CodePost).
- *In-class activities (15%):* Simple coding activities or paper-based tasks (submit on Blackboard).
- *Quizzes (30%):* 3 quizzes, 10% each.
- *Projects (20%):* 2 projects, each worth 10% of final grade.
- *Participation (05%):* physical presence, participation in polls (not based on correctness).

Grading scale: The tentative grading scale for this course would be as follows (subject to change upon Instructor's discretion):

A (93%-100%)	A- (90%-92.9%)	B+ (87%-89.9%)
B (84%-86.9%)	B- (80%-83.9%)	C+ (77%-79.9%)
C (74%-76.9%)	C- (70%-73.9%)	D (60%-69.9%)
F (0%-59.9%)		

Notebooks Assignments (30%) Throughout the course, we will be learning to apply machine learning principles using Python machine learning tools. Machine learning code is often developed in and communicated using an interactive integrated development environment called Jupyter Notebooks which support a natural interleaving of code, output/results, and mark-up documentation–what you're seeing right now is actually a Jupyter notebook. You will regularly submit notebook files (files with the extension .ipynb) to demonstrate your proficiency with the Python tools we are using. Given the long computation times of the programs you write, I will not usually be executing your code, so it is critical that the results from your executions are preserved in the notebook. You can expect to submit 6 notebooks throughout the course.

In-class Activities (15%) You'll also regularly submit notebook files with simple coding or paper-based activities. These should be easier since they follow the lecture and are usually done during or right after class. I won't run your code, so please make sure to save the results in the notebook. Expect to submit about 10-12 in-class activities throughout the course.

Quizzes (30%) There will be 3 quizzes that will be administered via Blackboard. They will not be timed, and you will have a few days to complete them. Quizzes should be completed individually and will be due before class on the following Monday. There is no time limit on these quizzes, but they must be submitted before we convene for the next class. As in the real world, you will be allowed to use external resources like the class notes and the internet. You will be required to cite any sources that you used while completing these quizzes other than the class notes.

Projects (20%): The two projects in this course will emphasize the design, execution, and interpretation of machine learning experiments. The grading emphasis will be on how well you explain your data and experiment as well as your written interpretation. For these, you will submit Jupyter Notebooks with more extensive writing in the mark-up cells than for your regular notebook assignments.

Attendance/Participation (05%): This class is highly interactive, meaning that active participation is both expected and the norm. You will receive credit for your participation, and it will be counted towards your final grade. I will keep track of your involvement using a signature sheet. Throughout the course, I will pose questions using polling software and conduct in-class Q&A sessions to better understand how the class is grasping the content. These responses will not be evaluated for correctness but rather for completion.

I respect your privacy. If you encounter challenges (physical health, mental health, or life in general) that interfere with your ability to participate in the course or complete your work, I will not require any kind of documentation. You also do not need to explain; you can simply inform me that you are experiencing problems and we will work together to figure out a plan that will enable you to complete the course if you want to. For example, if for some reason, you are unable to make the in-person class session, please email me and I will provide you with the Zoom link for the day so you can attend class virtually.

If you are unable to participate in the course for a prolonged period, we will discuss whether an incomplete is the best option.

Textbooks and materials: We won't use a specific textbook for this course. Occasionally, I'll suggest reading chapters from *Probabilistic Machine Learning: An Introduction - Kevin P. Murphy*, which will be available electronically. These readings are optional but will enhance the in-class activities and conceptual understanding.

Academic Integrity Policy: *We take academic integrity very seriously.* You are required to abide by the Drake University policy on academic integrity, as described in the Statement on Academic Dishonesty: Cheating and Plagiarism (`https://www.drake.edu/studentlife/handbook-resources/handbook/academic/`). It is your responsibility to understand these policies. Students agree that by taking this course, papers and source code submitted to us may be subject to textual similarity review, for example, by Turnitin.com. These submissions may be included as source documents in reference databases solely for the purpose of detecting plagiarism of such papers or codes.

Academic Accommodations: Drake University is committed to providing equitable access to learning opportunities for all students. The Disability Services office (107 Old Main) collaborates with students who have disabilities to provide and/or arrange reasonable accommodations. If you have, or think you may have, a disability (e.g., mental health, attentional, learning, autism spectrum disorders, chronic health, traumatic brain injury and concussions, vision, hearing, mobility, or speech impairments), please contact: **Michelle Laughlin**, Student Disability Services Coordinator (x1835), *michelle.laughlin@drake.edu* to arrange a confidential discussion regarding equitable access and reasonable accommodations. The process for communicating academic accommodations is as follows:

- When students request accommodations, they must provide documentation to Access & Success.
- First-year students, students new to requesting accommodations, and transfer students are required to make an appointment with Access & Success. Access & Success work together with the student to determine reasonable accommodations for each of the student's classes.
- Returning students can fill out an accommodation request form, found on the Access & Success website, www.drake.edu/access-success. The same process will be followed as above, however, returning students do not need to meet with Access & Success unless their accommodations need to be changed.

Holiday Observance: If you miss class because of a holiday or observance, you can fill out the form to automatically notify me (faculty). You can find the form on the `https://www.drake.edu/diversity/initiatives` for Initiatives and Programs, or click here to directly access the `https://drake.qualtrics.com/jfe/form/SV_d5qfVUKtuTQdg7b`.