Applied Machine Learning

COSC 410
Section and Lab A

Instructor Info

Forrest Davis (he/him)

Office Hours T: 3:00-6:00PM

Office Hours Bernstein 331

fdavis@colgate.edu

Course Info

Class: T, R | Lab: W

9:55-11:10AM | 12:45-2:35PM

212 Bernstein | 118 Bernstein

Overview

This course provides a practical introduction to applied machine learning. Students engage in supervised and unsupervised machine learning algorithms, including regression, support vector machines, decision trees, nearest neighbors, clustering, and ensemble methods. Students also learn deep learning techniques, including feed-forward, convolutional, and recurrent neural networks. Emphasis is placed on understanding and gaining hands-on experience with machine learning for practical use.

Material

Required Text: Dive into Deep Learning by Aston Zhang, Zachary C. Lipton, Mu Li, and Alexander J. Smola. It is available online here.

Course Website: You can find the course website where course materials, codelets, labs, and the schedule will be posted at `https://forrestdavis.github.io/cosc410/`.

Coursework

Codelets

Codelets are smaller assignments focused on a central problem, skill, or type of task, building on the course materials. There will be roughly 8 codelets graded on a satisfactory/unsatisfactory basis, where satisfactory demonstrates you meet the expectations of the codelet (roughly an 85% in standard grading). *Note: you might be requested to orally explain your answers to me to validate your assessment.*

Pre-Class Quizzes

There will be pre-class quizzes based on the course readings and content. Pre-class quizzes must be completed before the start of each relevant class period. Answering these questions will help ensure you are ready to engage in deeper learning and problem solving for the class meeting. There will be more pre-class quizzes offered than are required for full credit in this grading category, so you can still earn full credit even if you do not complete a few. They are graded on a satisfactory/unsatisfactory basis, where satisfactory is 60% in standard grading.

Labs

Lab assignments apply and extend the concepts from class. There will be around 9 lab assignments throughout the semester. They are graded on a satisfactory / unsatisfactory basis, where satisfactory demonstrates you meet the expectations of the labs (roughly an 85% in standard grading). Note: you might be requested to orally explain your answers to me to validate your assessment. Labs without assignments will focus on review and capstone projects.

Exams

The exams are designed to test your understanding of the core concepts covered in class, and your ability to apply them in different contexts. They will not involve any programming component. The scope of each exam is indicated on the course schedule. No discussion of exam questions or possible solutions is permitted from the time the exam is first administered until exams have been completed by all students and returned.

- Midterm Exam 1: October 2
- Midterm Exam 2: November 13

Please let me know in advance if you will be unable to take one of the exams on the scheduled date.

Capstone Project

A core component of the course is the capstone project completed in a small team (no more than 3 people) and related to machine learning. There are 6 components of the capstone project:

- Proposal (Friday, November 7)
- Feedback Discussions (Tuesday, November 11)
- Project Milestone I (Friday, November 21)
- Project Milestone II & Preliminary Presentation (Thursday, December 4)
- Poster Presentation Gala (Finals Slot TBD)
- Final artifact (Finals Slot TBD)

Completion of the proposal, feedback discussion meeting, project milestones, and preliminary presentation are graded on a satisfactory/unsatisfactory basis. The poster presentation and final artifact are individual grades, graded traditionally.

Society Reflection

Machine learning and artificial intelligence are playing an increasingly impactful role in our society. As part of a liberal arts education, broadly, and this class, narrowly, you are expected to build connections between the content of the course and the state of the world. I have identified the following events on campus as opportunities to engage in this process in a contained way:

- Robots as characters?: October 1 4:15pm Golden Auditorium
  - Do robots belong on stage? How do we perceive them as live characters inside a story? How might a robot endure the artistic rigor of developing a new play alongside live actors?
    Richard Maxwell will investigate the idea of robots as characters on stage as a means to explore empathy and how it is generated for characters in the minds of spectators. He will share his thoughts, experiments, and experiences working these crude machines into his productions.
- Ryan Family Film Screens: Films by Lawrence Lek: October 2 TBA
  - Lawrence Lek is a London-based Malaysian Chinese artist who unites film, video games, and electronic soundscapes in a singular cinematic universe. He is known for advancing the concept of Sinofuturism with immersive installations that explore spiritual and existential themes through the lens of science fiction. Featuring a recurring cast of wandering characters, his works are noted for their dreamlike narratives, evocative imagery, and preoccupation with technology and memory. Winner of the 2024 Frieze Artist Award, he was named one of Time's 100 most influential people in AI.
- Maria Antoniak Lecture: October 7 4:30pm Olin 350
  - Professor Antoniak's research focuses on natural language processing and cultural analytics. She develops and critically evaluates computational methods to analyze how language reflects culture and society, particularly in online communities and healthcare settings.

You must attend at least one of the following events and complete a reflection based on your experience. The reflection is a two page, single-spaced essay that addresses the following points:

- How does the event you attended present artificial intelligence or machine learning?
- How does this presentation relate (or not relate) to the technical aspects of such systems that you have learned in the semester?
- What role does the work suggest AI or machine learning should play in our society? Do you agree with this? Why or why not?

Grading Scheme

Your course grade is based on your codelets, exams, and capstone project. The table below defines specific criteria for each letter grade. A plus (+) or minus (-) will be added to your grade based on your attendance and completion of a society reflection. You must regularly attend class, pass 80% of pre-class quizzes, and complete the reflection for a plus (+). A minus (-) will be added to your grade if you fail to regularly attend class, fail more than 20% of quizzes, and fail to complete the reflection.

Grade	Codelets (S/U)	Exams (2)	Capstone
A	8/8	> 90% avg	> 90% avg and 4/4 S
B	6/8	80–89% avg	80–89% avg and 3/4 S
C	4/8	70–79% avg	70–79% avg and 2/4 S
D	2/8	60–69% avg	60–69% avg and 1/4 S
F	Failure to meet expectations for D results in an F		

The lab grade is based on lab assignments and attendance. The table below defines specific criteria.

Grade	Labs (S/U)	Attendance
A	9/9	0 unexcused absences
B	7/9	1 unexcused absences
C	5/9	2 unexcused absences
D	3/6	2 unexcused absences
F	Failure to meet expectations for D results in an F	

The final letter grade is determined by the average across the criteria. For example, an A for codelets and Bs for exams and capstone is a B. A grade of of A+ is awarded when the student demonstrates truly exceptional performance and is not simply determined by having a high final grade. I reserve the right to make adjustments; any such adjustments will only raise your grade, never lower it.

Policies

Attendance and Engagement

Attendance to class is expected and contributes to your final grade. However, I do not expect you to attend if you are feeling unwell. Additionally, if you have athletics, scheduled commitments, or other issues please let me know. If you miss a class, you can make up the credit by performing a check-in email. For a check-in, consult with a classmate and/or review the materials, and within 48 hours of the class send me an email containing:

1. a brief summary (2-4 sentences) of what was covered
2. any thoughts/questions you have

Check-ins are not a substitute for attending class regularly. While in the class, you are expected to engage with the material and the other students in the course. You should aim to be a good participant: raising your hand, respecting others, actively listening, and making sure to leave space for others to speak. There are no bad questions, and I would always rather you contribute than avoid doing so.

Deadlines

I will be reasonably flexible on deadlines. If you need some extra time due to illness, your workload in other classes, and/or personal matters, please let me know. As long as you have made a good faith effort to complete learning activities by the original deadline, I am willing to offer a reasonable extension. I will be less willing to grant an extension if you wait to start a assignment until the day before it is due (when you've had a week to work on it), repeatedly ask for extensions, etc.

Any assignment that is turned in late without prior approval will be marked as unsatisfactory.

Academic Honesty & Collaboration

You are expected to abide by Colgate's academic honor code. The overarching policy of this course is that the work you turn in should represent your own thinking. You should be accountable to its content (e.g., be able to explain what your code does). I may ask to meet with you to discuss your work. Your grade may change based on the conversation.

Collaborating with peers in the class

You may discuss course concepts, generic aspects of python, and work through the logic of something you don't understand with your peers. However, you should not share code (including psuedo code). Your submitted work should be your own. Here are some concrete rules that exemplify this (but are not intended to be comprehensive):

Do NOT:

- Ask a peer in either this section of the course or others (past or present) to debug your code.
- Ask a peer for pseudo code for an algorithm needed for a homework.

You CAN:

- Ask clarification questions about the fundamentals of programming (e.g., "How do I create a class in Python?")
- Ask for conceptual clarifications (e.g., "What is the difference between regression and classification?")
- Try to work through the logic of something you don't understand (e.g., "How is data formatted for a neural network?")
- Ask for help to pointers about libraries (e.g., "Where did you find information about using scikit-learn for clustring?")

Using generative AI tools

Generative AI systems, if used correctly, can serve as powerful tools. In this course, you can use generative AI systems to learn about concepts iteratively through a interaction, to help you debug parts of your code, to help direct you to resources, and to help with small amounts of code. However, you cannot ask these systems to directly give you answers or write all the code for you for an assignment. Ultimately, I want you to be able to work through an ML pipeline on your own (using tools in ways that extend your abilities rather than supplant them). When you include code from some generative AI system in your work you must include attribution and a brief – no more than 15 words – description of what you wanted it to do. Note, attribution is not sufficient to avoid claims of cheating. Here are some concrete rules that exemplify this (but are not intended to be comprehensive):

Do NOT:

- Give the model a problem description and ask it to sketch an algorithm for you or write you pseudo code
- Give the model the homework description and ask it to organize the code for you (e.g., generate the necessary function headers, write the main functions etc)
- Give the model the homework an ask it to produce a solution

You CAN:

- Ask clarification questions about the fundamentals of programming (e.g., "How do you create classes in Python?")
- Ask for conceptual clarifications (e.g., "What is the difference between recall, precision and AUC?")
- Try to work through the logic of something you don't understand (e.g., "How does the gradient tell the model which direction to move in during optimization?")
- Ask the model for help creating a helper function (e.g., "Write a python function that takes in a string and removes all non-character letters"; you should include an attribution in your code to the AI you used and briefly – no more than 15 words – what you wanted it to do)

Remember: Policies around the use of Generative AI tools, like any other course policies, vary across different courses both within and outside the department.

Anonymous Feedback

Your feedback on this course is important for helping me improve the learning environment. You can provide anonymous feedback at any point in the semester via this form. *Note, you must be accessing the form from a Colgate account, but your email is not recorded.*

Getting Help

A key to your success at Colgate is figuring out what resources are available and using them to help you achieve your goals. There are several options for getting help with this course:

1. Drop in during my office hours (noted at the top of the syllabus) or if no office hours times work, arrange an appointment with me – just send me an email with a few times you are available, and we will find a time that works well for both of us. If my door is open to my office (322 Bernstein), you may (occasionally) knock and ask a quick question.
2. Form a study group with other students in the class and work together on a regular basis (note the Academic Honesty & Collaboration policy above).

I also encourage you to reach out to many great resources at Colgate that can assist you with academic, personal, or other needs, including:

- Administrative Deans (https://www.colgate.edu/about/offices-centers-institutes/dean-college/administrative-advising) help you understand policies and procedures, navigate personal challenges, work with faculty, and engage with parents.
- Counseling Center (https://colgate.edu/counseling) staff are trained to help students manage a wide array of emotions. The counseling center meets with over half the student body for clinical services at some point during their four years at Colgate. You can arrange an appointment online or by phone (315-228-7385). For emergencies, a counselor is available 24/7 by calling campus safety at 315-228-7333 and asking for the counselor on call.
- Haven (https://colgate.edu/haven) is a sexual violence response center that provides confidential care, support, advocacy, and trauma-informed clinical services for survivors of sexual assault, intimate partner violence, child/family abuse, stalking, and/or harassment. You can call (315-228-7385) or visit during business hours. You can also contact the Help Restore Hope Center (855-966-9723).
- Student Health Services (https://colgate.edu/offices-and-services/studenthealthservice) provides accessible, convenient, cost-effective, non-judgmental, and confidential care for all students.
- Information Technology Services (https://colgate.edu/its) help desk consultants assist all students with problems concerning email, Portal, Moodle, or your personal laptop. Contact me if problems with your personal computer are affecting your ability to get your work done.
- Chaplains (https://colgate.edu/campus-life/religious-life/officeofthechaplains) provide the community with a dynamic, friendly, and supportive place.