

Fundamentals of AI/ML

Logistics

Ashley Y. Gao

William & Mary

January 21, 2026

About this Course

- Broad introduction to machine learning
 - Algorithms and principles for supervised learning
 - Nearest neighbors, decision trees, ensembles, linear regression, logistic regression, SVMs
 - Algorithms and principles for unsupervised learning
 - Unsupervised Domain Adaptation
 - (Unsupervised) Domain Generalization
- Coursework is aimed at grads and advanced undergrads. We will use multivariate calculus, probability, and linear algebra.

Course Information

- Course website
 - <https://lindagaw.github.io/courses/CSCI416-Spring26/CSCI416.html>
 - Main source of information is the course webpage; check regularly!
 - Lecture slides will be posted on the website as topics are covered during class.
- Announcements, grades
 - Announcements will be sent as group emails to students enrolled in this class
 - Grades will be posted on Blackboard
- Discussions
 - Your grade does not depend on your participation on Piazza. It's just a good way to ask questions, discuss with your instructor, TAs, and peers. We will only allow questions that are related to the course materials, assignments, and exams.

Office Hours

- Instructor: Ashley Y. Gao
 - Email: ygao18@wm.edu
 - OH Time: Tu/Th 9:00-10:30, and by appointment
 - Location: Integrated Science Center 4 Room 2381, or on Zoom (by appointment)
- Teaching Assistant: TBD
 - Email: TBD
 - Time and Location of OHs: TBD
- You can also find the most up-to-date information regarding OHs on the course website.

Textbooks are Optional

- All materials that will be covered in exams and homeworks are in the lecture slides, but you are encouraged to read the following optional reading materials on ML and AI!
 - Goodfellow et al: “Deep Learning”, 2015.
 - Christopher Bishop: “Pattern Recognition and Machine Learning”, 2006.
 - Kevin Murphy: “Machine Learning: a Probabilistic Perspective”, 2012.
 - David Mackay: “Information Theory, Inference, and Learning Algorithms”, 2003.
 - Shai Shalev-Shwartz & Shai Ben-David: “Understanding Machine Learning: From Theory to Algorithms”, 2014.
 - David Barber: ”Bayesian Reasoning and Machine Learning”, 2012
 - Richard S. Sutton and Andrew G. Barto: ”Reinforcement Learning: An Introduction” (2nd ed.), 2018.
- There are also a lot of high-quality resources on ML and AI available on YouTube!

Grading

- Homeworks: 40 pts
 - Homework #1: 10 pts
 - Homework #2: 10 pts
 - Homework #3: 10 pts
 - Homework #4: 10 pts
- Exams
 - Midterm: 20 pts
 - Final Exam: 20 pts
 - Exams will be closed-book and held during class hours but you are allowed a one-sided US letter sized cheat sheet.
 - If you are a grad student, you are required to complete additional questions in each exam in order to receive the full mark. If an undergrad correctly answers these questions, they will receive extra credits.
- Final Project: 20 pts
 - Will require you to apply several algorithms to a challenging problem and to write a short report analyzing the results
 - More details TBA

More on Homeworks

- There will be 4 homeworks in this course. The assignments will be released on the course webpage. Homeworks are due on Blackboard.
- Late submissions
 - Homeworks will be accepted up to 3 days late, but 10% will be deducted for each day late, rounded up to the nearest day. No credit will be given for assignments submitted after 3 days. Extensions will be granted only in special situations with valid proof (e.g. Doctor's note).
- Collaboration policy
 - Each student is responsible for their own work. Discussions on homeworks are allowed, but students should write their own submissions solely by themselves. Write down the name and email of every student you've discussed with on each homework.

Honor Code

- Presenting someone else's ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in our campus policies. Please read the Honor Code at William and Mary.
- If you violate this rule, your final exam will be failed. When you refer to some source codes on GitHub, please cite it with a URL in your report.
- Please do not copy the answers from the Internet directly without any references. You should rephrase your answers based on your own understanding.

Accommodations

- William & Mary accommodates students with disabilities in accordance with federal laws and university policy. Any student who feels they may need an accommodation based on the impact of a learning, psychiatric, physical, or chronic health diagnosis should contact Student Accessibility Services staff at 757-221-2512 or at sas@wm.edu to determine if accommodations are warranted and to obtain an official letter of accommodation. For more information, please see www.wm.edu/sas.
- As per the university's guidance, if you have a religious observance that conflicts with a deadline, please notify me as soon as possible so that I can attempt to make an appropriate adjustment.

- You'll be expected to synthesize and apply critical analysis, solve problems in an applied and/or academic setting, create original material or original scholarship, and communicate effectively with diverse audiences.
- Since this is a COLL 400 class with a significant portion being writing, **you will not receive the final letter grade unless you finish the final report!** If you haven't submitted the final report when it is due, you will receive an Incomplete and the Incomplete will remain there until you submit your final report!
- **Reflected in your final project!**

About the Instructor

- B.S. in Computer Science
University of California, San
Diego (UCSD)
- B.A. in Literature/English
University of California, San
Diego (UCSD)
- M.S. in Computer Science
University of Virginia (UVa)
- PhD in Computer Science
University of Virginia (UVa)

- Large Audio Language Models
(Audio LLMs)
- Natural Language Processing
(NLP)
- Affective Computing