

CSCI 1420: Machine Learning

Course Information and Syllabus - Fall 2025

Instructor	Lorenzo De Stefani (lorenzo_destefani at brown.edu)
Course Description	How can artificial systems learn from examples and discover information buried in data? We explore the theory and practice of statistical machine learning, focusing on computational methods for supervised and unsupervised learning. Specific topics include empirical risk minimization, probably approximately correct learning, kernel methods, neural networks, maximum likelihood estimation, the expectation maximization algorithm, and principal component analysis. This course also aims to expose students to relevant ethical and societal considerations related to machine learning that may arise in practice.
Learning objectives	This course will cover both historic and modern machine learning algorithms viewed through the lens of various models, loss functions, and optimizers. Throughout this course, you will develop a toolbox of techniques used to address classic machine learning problems such as regression, classification, clustering, and dimensionality reduction. For each algorithm you work with, you will become familiar with the theoretical underpinnings that justify both its functionality and performance, as well as its limitations. Additionally, you will learn to recognize or discover which techniques are best to address a given learning problem.
Course Website	For more in-depth information about the course, refer to the course website: https://cs1420-f25.vercel.app/
Prerequisites	Students should have completed an intro sequence and taken at least two semesters of calculus, basic calculus-based probability, and linear algebra to be prepared for this course. We recommend comfort with multivariable calculus (particularly partial derivatives and gradients). We also recommend familiarity with writing formal proofs. Formally, the requirements are (CSCI 0160, 0180, 0190 or 0200) and (CSCI 1450, 0450, APMA 1650 or 1655) and (CSCI 0530, MATH 0520 or 0540). Per the recently announced departmental policy on CSCI 410/1411 and CSCI 1420, the former will become a prerequisite for the latter. This policy will be relaxed for senior students (and junior, to some extent) and students who have completed an equivalent 1000-level course. Students are also required to have some confidence in writing and debugging Python code. This is not an intro to programming/debugging code.
Lectures	Tuesdays and Thursdays. Time: 2:30 pm -3:50 pm. Room MacMillan Hall 115

	The course is offered in person. Attendance will be collected during class after the end of the shopping period and will count towards the final score in the course. Recordings of the lectures will be published within 48 hrs after the class on Panopto.
Lecture Slides	Lecture slides will be available on the Ed platform of the course.
Textbook	“Understanding Machine Learning: From Theory to Algorithms” by Shai Shalev-Shwartz and Shai Ben-David, Cambridge University Press, 2014. ISBN: 9781107298019. Print copies are available in the bookstore, and a free PDF is available online: cs.huji.ac.il/~shais/UnderstandingMachineLearning/copy.html
Enrollment Cap and Waitlist Policy	The course enrollment is currently capped. While we reserve to revise this cap upwards, we cannot guarantee this will happen. If you could not register for the course, to be added to the waitlist, please send an override request through CAB. No further action is necessary. Whenever slots are available, codes are sent out in the evening of every workday. Starting during the shopping period, codes that are unredeemed for 24 hours will be revoked. Please be mindful of your enrollment request. Students who drop out of the course but later wish to re-enroll will be placed at the end of the waitlist. CAB may display available seats, but those are likely allocated to override codes that have been sent but not yet redeemed. Inquiries regarding “position on the waitlist” and/or “likelihood of joining the course” will not be addressed. Due dates on HW assignments also apply to students currently on the waitlist, with no exception.
Enrolling in courses with overlapping meeting times	Petitions to enroll in CSCI 1420 with overlap with other courses will not be approved in accordance with the university policy on overlapping courses .
Professor’s Office Hours	Fridays 14:15 - 16:14 - Remote Office hours Please reserve a slot using the Official course calendar. The link to the Zoom meeting is provided in the description of the calendar slots.
Head TA	Jaideep Naik
UTAs	Peter Popescu Andrew Gao Arjan Chakravarthy Eric Kim Samuel Liu
TA Office Hours	See the course website.

Communications	All email communications regarding the course addressed to Lorenzo or the TA staff must have “[CSCI 1420]” as the first term of the subject.
Time Requirements	In addition to 3 hours of lecture each week, each student will spend approximately 8 to 12 hours per week on reading, homework, and the exam, for a total estimated time commitment of 180+ hours.
Course credit and Grading	Final grades will be determined using the following weighting of assignments: • 12 Homeworks (6% each, the lowest grade of the 13 assigned will be dropped) • Final Exam (23%) • Participation (5%, see Participation Policy below) Regrade requests on assignments (except for the Final) will be enabled via Gradescope for a week after the grades are released, and disabled past that. Regrade requests for the Final will be available for three days after the grades are released. Final Grade cutoffs: The cutoffs for an A/B/C grade are set to 90%/80%/70% of the available credit, including homework assignments and in-class midterms. Such thresholds may only be adjusted down to Lorenzo’s discretion. To receive an “S with distinction,” students need to complete the class with a score well within the “A” range (i.e., > 92%).
Homework policy	Thirteen take-home homework assignments will be released during the course. The schedule of homework releases and due dates is presented in the official course calendar. Homework assignments will be released on the same page. Any modifications will be communicated by the instructor, and such modifications will be reflected in the calendar. Credit: The credit tied to homework will be assigned based on the average score of all submitted homework, minus the one with the lowest score, all weighted equally. Gradescope policy: Homework assignments must be submitted using Gradescope. When registering, use your @brown.edu email account. Solutions and feedback on assignments will be delivered there as well. Collaboration policy: We allow (and encourage!) discussion of the material presented in lectures, help sessions, and the textbook, as well as discussion of concepts involved in homework assignments. Written work is another matter. You must write up your homework solutions by yourself. This ensures that you understand the material even if you collaborated on the problem. You should not take notes away from collaboration sessions and should be by yourself when writing the solution. Rule of thumb: If you cannot independently reproduce what you hand in, don’t hand it in. We reserve the

	right to request that you explain your solution to a problem, so we can ensure that the solution was your own work. We do not tolerate academic dishonesty. This includes cheating, lying about course matters, plagiarism, or helping others commit a violation. Keep in mind the obligations and expectations associated with the Academic Code of Conduct, Academic Code, Graduate Student Edition, and Code of Student Conduct. Late submission policy: At most two homework assignments can be handed in up to one day late. Other assignments submitted late will not receive any credit. Additional late-day passes may be granted by Lorenzo in consideration of extenuating circumstances (e.g., illness). Those requests must be accompanied by a Dean's Note and sent directly to Lorenzo. While we sympathize with other possible occurrences, extensions will not be given for random special events (e.g., I/my roommate/pet spilled water on my laptop, I lost my laptop/power adapter). You are free to update your submission up to the deadline date (or past it if you use late days). We recommend doing so to avoid losing potential partial credit, and/or keep a version of your work backed up online. Double-check your work before submitting it. Submitting blank pages and/or forgetting to upload materials or uploading the wrong version of an assignment will not be granted "special" late deadlines.
Final	The final will be a take-home exam. More details will follow. No collaboration is allowed. You should not discuss the problems and your solutions with other students. Violations of such policy will result in the inability to complete the course and will be reported as violations of the Academic Code of Conduct.
Auditing	Lecture attendance: Attending all of the planned lectures is required. To check in please talk with Lorenzo at the end of the class. Homework: Auditors must complete at least one homework assignment for each half of the semester. These must be completed within the assignment deadline, and the score must be at least 80%.
Incomplete grades	In extraordinary situations, students may request to receive an incomplete grade. A Dean's note must accompany such cases and will be assessed on a case-by-case basis.
Ed discussion policy	Students are encouraged to interact through the Ed Discussion platform for the course under the following policy:

	• Public posts cannot be anonymous. The TA staff will moderate these posts to make them visible. • Private posts may be anonymous. • We will observe quiet hours 8 pm-9 am • Ed is NOT meant to be the place to request extensions. Any such posts will be deleted. Ed will be placed in read-only mode during university-mandated holidays.
Diversity All are Welcome	The course intends to provide a welcoming environment for all students. We especially welcome diverse ideas and perspectives during class discussions. All members of the CS community, including faculty and staff, are expected to treat one another professionally. Toward this goal, TAs have undergone training in diversity and inclusion. However, despite our best efforts, we may accidentally slip up, so please feel free to speak to any member of the course staff with any concerns you have during the semester, and do not hesitate to contact Lorenzo directly. We will take your concerns very seriously.
Accommodations	If you feel you have physical, psychological, or learning disabilities that could affect your performance in the course, we urge you to contact SAS/SEAS (https://www.brown.edu/campus-life/support/accessibility-services/). We will do whatever we can to support accommodations recommended by SAS/SEAS. If you require SAS/SEAS accommodations, please let Lorenzo know at the beginning of the semester. Students for whom SAS/SEAS accommodations are granted for extended time on assignments and/or exams will have access to a budget of four late days, with the possibility of applying up to two late days to a single assignment.
Mental Health	Being a student can be very stressful. If you feel you are under too much pressure or there are psychological issues that are keeping you from performing well at Brown, we encourage you to contact Brown's Counseling and Psychological Services (CAPS: https://www.brown.edu/campus-life/support/counseling-and-psychological-services/). They provide confidential counseling.
Coping with Unforeseen Events	If there are events that are upsetting to you, whether political, family-related, weather-related, etc., that affect your ability to do well in class, we are happy to take them into account concerning our late and incomplete policies. Please feel free to talk to Lorenzo about this.
Agreeing to class rules	At the beginning of the course, all students are required to formally indicate they understand and agree to abide by the course collaboration policy by downloading the PDF of the syllabus, reading it thoroughly, and submitting it on Gradescope as HW 0 by 10/1/2025.

	Failure to do so in a timely fashion will not allow students to complete the course.
--	---