

CSCI 1420: Machine Learning

Instructor: Stephen Bach

Spring 2024

Time: Tuesday and Thursday, 2:30-3:50p
Location: Metcalf Auditorium
Web: cs1420.vercel.app

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.

Required Materials

Understanding Machine Learning by Shai Shalev-Shwartz and Shai Ben-David.
Print copies are available in the bookstore, and a free PDF is available online:
cs.huji.ac.il/shais/UnderstandingMachineLearning/copy.html

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).

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 their functionality and performance, as well as their limitations. Additionally, you will learn to recognize or discover which techniques are best to address a given learning problem.

Estimated Time Commitment

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 total estimated time commitment of 180+ hours.

Teaching Staff

Instructor: Stephen Bach <stephen_bach@brown.edu>
Office: CIT 335

Head TAs:

Emily Ye

Kevin Lu

Matthew Meeker

<cs1420headtas@lists.brown.edu>

TAs:

Aditya Agashe

Alex Liang

Alex Lin

Andrew Yang

Jaideep Naik

Johnny Elias

Keitaro Nishijima

Krishi Saripalli

Luke Choi

Marco Ayala

Mason Lee

Nitin Sreekumar

Noah Foster

Reggie Zheng

Sarah Peters

Spencer Dellenbaugh

Taishi Nishizawa

Thomas Chang

Youjung Koo

Zeeshan Bhalwani

Socially Responsible Computing TAs:

Emma Huang

Julie Qian

Lectures

Lectures will be held live and recorded. Recordings will be available on [Canvas](#) shortly after lecture.

To reduce the spread of diseases, students are encouraged to wear masks.

In the event that the instructor is unable to teach a class session in person, it will either be taught remotely or recorded and posted on Canvas, depending on the circumstances.

Instructor Office Hours

Office hours are held from 10 am to 11 am on Tuesdays and by appointment. Visit [this link](#) to reserve a regular slot or email the instructor to make an appointment at another time.

TA Office Hours

See cs1420.vercel.app/#calendar for availability.

Edstem

Found at [this link](#). Students should be added automatically through Canvas. For guidelines on Edstem, see our collaboration policy.

Grading

Final grades will be determined using the following weighting of assignments:

12 Homeworks (7% each, lowest grade of the 13 assigned will be dropped)
Final Exam (11%)
Participation (5%, see Participation Policy below)

Final grades will be then be determined using the following cutoffs for those weighted averages: $\geq 90\% = A$, $\geq 80\% = B$, $\geq 70\% = C$, $< 70\% = NC$. The instructor reserves the right to lower (but not raise) these cutoffs. In other words, 90% is *always* an A, but the instructor may or may not decide that lower averages are also an A.

Exam Date

The final exam is on **Thursday–Friday, May 16–17**. We will release the final at 12:00 pm noon ET on Thursday, and it is due at 11:59 pm ET Friday. It will be designed to take at most 4 hours, so this 36 hour window is intended to accommodate students in different time zones, conflicting exams, etc.

Course Policies

Collaboration Policy

Students are required to familiarize themselves with and abide by this course's collaboration policy: cs1420.vercel.app/assets/docs/collaboration_policy.pdf.

Key elements include the requirement that all students complete their own homework assignments. Students may discuss the assignments with each other, but may not work directly on the deliverables (write ups, code, etc.) during those discussions. No student may look at another student's work in progress, nor permit another to look at theirs. The only exceptions are when asking for help on Ed or in Collab hours.

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 Brown [Academic Code](#), [Academic Code, Graduate Student Edition](#), and [Code of Student Conduct](#).

Participation Policy

We use **TopHat** during class to facilitate active learning. Participation in these questions will count for 5% of a student's final grade, as described above. Participation is counted on a per-vote basis, with students receiving credit for answering a question, regardless of whether the answer is correct. Students can miss participating in up to 15 votes without penalty. After those 15, each missed vote will reduce a student's participation grade. Additional free votes will not be allowed for any reason.

Please also note that a single question might include multiple votes, e.g., before and after discussing with a partner. A class therefore may contain as many as four to eight votes. Voting without being physically present in class will be considered academic dishonesty.

Hours Policy

Our philosophy for TA hours is that they are for students to receive guidance on concepts and assignments, as opposed to receiving answers. Therefore, please do not go into hours expecting that a TA will give you the answer to a problem. We will not provide answers directly and we will not tell you if your answers are correct. As a reminder, the [Code of Student Conduct](#) applies during hours, and we will ask anyone who does not uphold it to leave hours.

Individual hours will make use of hours.cs.brown.edu, and will follow department policy where each student has up to 15 minutes with a TA before they must sign up again. Before signing up for TA hours,

students must have a specific question in mind and state their question in the description when signing up. If you do not have a sufficiently specific question, the TA will skip over you and not see you until you clarify your question. In addition, students will not be able to ask more questions after the first five minutes with a TA so that more students can receive help from TAs during busy TA hours.

Collab hours are also available and described in the collaboration policy.

Incomplete Policy

We expect everyone to be able to complete the course on time. However, especially in these challenging times, we certainly understand that unexpected and uncontrollable situations may arise that prevent you from finishing the course on time, such as health problems and family crises. If you feel you cannot complete the course on time, please email Prof. Bach to discuss the possibility of being given a grade of Incomplete for the course and setting a schedule for completing the course in the upcoming year.

Due Dates and Turning in Homework

All assignments are due at **12:00 pm noon ET**. Homework is to be submitted through **Gradescope**. Students should be added automatically to Gradescope through Canvas. All grading and re-grade requests will be handled through Gradescope.

Late Policy

The late-day policy described below applies to all late days other than those due to illness, religious holidays, and other extensions granted by Prof. Bach in extreme circumstances. Thus days missed because of job interviews are included in the late-day policy. Late days are applied to an entire homework, and not to individual components of the homework. In other words, choosing to use a late day for the code handin automatically constitutes the use of a late day on the written handin and vice versa.

Everyone is allowed a total of four late days on homework, and no more than three late days may be applied to any one homework assignment. Beyond that, you will be penalized 25 percentage points for each day an assignment is late. We will apply late days to assignments in an optimal fashion at the end of the semester. If you are ill, you may get an extension without using late days if you get a note from either Health Services, a university dean, or other healthcare provider. Please email Prof. Bach regarding illness or other issues. If a religious holiday impacts your ability to meet an assignment deadline, you may also get an extension without using late days; please email Prof. Bach at the beginning of the semester. In addition, if other circumstances are affecting you, please email Prof. Bach.

Diversity and Inclusion

We strive to create and sustain a diverse and inclusive environment in which all students, faculty, and staff can thrive. It is everyone's responsibility to uphold this mission, and we ask students in the course to honor this. We take all complaints about discrimination, harassment, and other violations of **Brown's Nondiscrimination and Anti-Harassment Policy** seriously. If anybody involved in the course has failed to make it an inclusive space for you, please contact either Prof. Bach or the department chair (Prof. Tamasia). The department's **Diversity Advocates** are also available as a resource for members of historically underrepresented groups. **Additional resources** are listed on the department's website.

In addition, Brown welcomes students from all around the country and the world, and their unique perspectives enrich our learning community. To empower students whose first language is not English, an array of support is available on campus, including language and culture workshops and individual appointments. For more information, visit the Sheridan Center's **English Language Support** page.

Accommodations

Brown University is committed to full inclusion of all students. If you have physical, psychological, or learning disabilities that could affect your performance in the course, we urge you to contact [SAS](#). We will do whatever we can to support accommodations recommended by SAS. If you have an accommodation letter from SAS, please submit it to Prof. Bach through the SAS portal. Please also do not hesitate to contact Prof. Bach if you would like to discuss your accommodations. The departmental [Health and Wellness Advocates](#) are also available as a resource for you to discuss any concerns, and to guide you through options and next steps for accommodations.

Mental Health

We understand that health encompasses your physical, mental, and emotional well-being. If for any reason you are struggling to balance your health with your coursework, we recommend that you take advantage of the resources available to you. We commit to working with you and any supporting offices at Brown to provide reasonable accommodations they recommend.

Academic deans are accessible every weekday in University Hall, and [Student Support Services](#) deans are available over the phone 24/7 and in person by appointment. Brown's [Counseling and Psychological Services](#) (CAPS) provides confidential in-house counseling, outside mental health resources, and notes for health accommodations. The department also has [Health and Wellness Advocates](#) who can assist you.