CS 7545 Machine Learning Theory

Fall 2026

Basic Info

Class Hours: 2:00pm-3:15pm on Tuesdays and Thursdays
Class Room: 1447 in Klaus Advanced Computing Building

Course Staff:

- Steve Mussmann
  - Instructor
  - mussmann@gatech.edu
  - Office Hour Location: Klaus 3320
  - Office Hour Time: TBD
- TA(s) TBD

Office hours are held from the first week of class (August 24) through December 11, except for holidays: Labor Day, Sep 7; Fall Break, Oct 5-6; and Thanksgiving, Nov 25-27.

Course Description

This course provides a basic arsenal of powerful mathematical tools for the analysis of learning algorithms, focusing on both statistical and computational aspects.

Learning Goals

The goal of this course is for students to:

- Digest a set of common topics and techniques within theoretical ML:
  - Common analytical tools (e.g., concentration inequalities)
  - Learning frameworks (e.g., PAC learning)
  - Generalization and bias-variance trade-offs (e.g., Rademacher complexity, VC dimension, parametric convergence, benign overfitting, inductive bias and NFL theorems)
- Theoretically analyze ML settings and algorithms through proving mathematical results.
- Test and analyze ML algorithms in idealized settings with computational tools.

- Analyze and appreciate how the area of theoretical ML relates to ML in general

Required Prior Knowledge

- Ability to create rigorous proofs
- Background in probability, algorithms, and linear algebra
- General familiarity with ML at the level of CS 4641
- Ability to implement simple ML settings, algorithms, and analyses in Python

Course Tools

- Canvas for announcements, assignments, and gradebook
- Piazza for questions
- GradeScope for submitting and grading assignments
- PointSolutions for in-class polls

Assessments and Grading

Scribing [5%]

As part of the course, each student will sign up to be part of a 2 or 3 student team that scribes a lecture. Each student should create a separate write-up (using the provided LaTex template) and then the team creates a "merged" write-up. The team should submit the tex and pdf files for all individual write-ups as well as the merged write-up. The scribed notes are due one week (7 days) after lecture and the team must address subsequent feedback.

Written Homeworks [25%]

The written homework problems are intended for practice writing proofs and exploring extensions of results covered in class.

- [5%] HW 1, due TBD
- [5%] HW 2, due TBD
- [5%] HW 3, due TBD
- [5%] HW 4, due TBD
- [5%] HW 5, due TBD

Computational Mini-Projects [20%]

The computational mini-projects involve running numerical simulations (given Python starter code) and answering questions analyzing and interpreting the results.

- [5%] CMP 1, due TBD
- [5%] CMP 2, due TBD
- [5%] CMP 3, due TBD
- [5%] CMP 4, due TBD

Exams [50%]

The exam questions will reflect simplified and more conceptual versions of the types of problems on the homework, though a few lectures will not have homework questions but may have corresponding exam questions. The final exam will cover the union of the topics from the first two exams. The worst exam score will be dropped, or equivalently, the final is optional but can replace the worst exam score.

- [25%] Exam 1, in-class TBD
- [25%] Exam 2, in-class TBD
- [25%] Final Exam (optional), TBD

Final Grade

The final grade is a weighted average (using percentages above) of the scores on the assignments. The conversion to a letter grade is based on a standard conversion (without any rounding), though the score thresholds may be decreased.

Grade	Score threshold:
A	90%
B	80%
C	70%
D	60%
F	0%

For the pass/fail option, the score threshold for pass is 70%.

Mid-Semester Check-in

There will be a mid-semester check-in survey (due TBD) with the goal of students giving the course staff feedback on what is working well and what is not. Students completing the survey will be given bonus points.

Student Well-Being

At Georgia Tech, we are concerned about your overall physical, social, and mental well-being. A comprehensive list of wellness related resources has been compiled and maintained by the Office of the Vice President for Student Engagement and Well-being and can be found here: https://students.gatech.edu/student-resource-guide.

A list of resources for graduate students is given on the Office of Graduate and Postdoctoral Education website. Specific information includes:

- Academic Resources such as the Communications Center, Language Institute, Library, Catalog, Registrar, resources for conducting research, Advocacy and Conflict Resolution resources, and how to manage unexpected situations that may impact your academic performance
- Student Resources such as Campus Services, Child Care/Family programs, Health & Wellness, Career Services, and the Student Resource Guide
- Professional Development such as the programming from the Career Center and other professional development resources and events

Course Policies

Academic Integrity

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. Review Georgia Tech's Honor Code and the student Code of Conduct.

Any student suspected of cheating or plagiarizing on an exam or assignment will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations

AI Policy

We recognize that generative AI tools (e.g., ChatGPT, Copilot) can support learning when used responsibly. In this course, you may use AI tools to brainstorm, explore ideas, or clarify concepts in a similar way to how you might collaborate with peers. However, all submitted work must reflect your own understanding and original expression.

You are responsible for ensuring that all submitted work is your own. Submitting content generated by AI tools will be treated as academic misconduct. If you're unsure whether your use of AI is appropriate, please ask.

For the mini-projects, any submitted code is ancillary to the main submission of the mini-project report. You may use AI tools for writing code, but not for writing the report.

Guidelines for Ethical AI Use:

- Use AI for learning, not for writing your submission. You may consult AI tools to help you understand a topic or generate ideas. However, do not copy and paste AI-generated text into your assignment. Instead, reflect on what you've learned and write a response in your own words.
- Separate your writing from AI interactions. Do not work on an assignment and use a Generative AI tool simultaneously. Treat AI interaction as a preparatory step, similar to reading a source or discussing with a peer. After using AI, close the tool and write your assignment independently reflecting your revised knowledge.

Collaboration Policy

In this class, it is considered plagiarism to use written content from a classmate. In particular, do not share your written homework, mini-project report, or mini-project code with your classmates. However, you are encouraged to brainstorm and "whiteboard" ideas and general approaches with other students.

For the scribed lecture notes, you must write your own individual write-up (which is submitted) before viewing others' write-ups or the merged write-up.

For exams, it is considered cheating to view other student's exams or discuss the exam in any capacity before you've submitted your own exam.

Accommodations for Students with Disabilities

If you are a student with learning needs that require special accommodation, contact the Office of Disability Services (404-894-2563) as soon as possible to make an appointment to discuss your special needs and to obtain an accommodations letter. Please also e-mail me as soon as possible to discuss your learning needs.

Student-Faculty Expectations Agreement

At Georgia Tech, we believe that it is important to strive for an atmosphere of mutual respect, acknowledgement, and responsibility between faculty members and the student body. The Student-Faculty Expectations articulate some basic expectations that you can have of me and that I have of you. In the end, simple respect for knowledge, hard work, and cordial interactions will help build the environment we seek. Therefore, I encourage you to remain committed to the ideals of Georgia Tech while in this class.

Attendance Policy

Attendance is not required though completing and submitting exams requires in-person attendance on exam days. If you will miss an exam date, please let the course staff know as soon as possible. Lectures will not be recorded nor have a remote option. If you miss a lecture, I recommend that you review the scribed notes, attend office hours, and/or learn from a classmate.

When you attend class, please be respectful of other students and avoid behaviors that could be distracting. Please use discretion for eating in class, using mobile devices, and use of electronic devices for activities unrelated to lecture.

Exam Notes Policy

While taking the exam, you may bring and refer to a handwritten sheet of paper (double-sided is fine) no larger than 8.5” x 11” (standard printer paper size).

Assignment Late Days

Each student has a total of 5 late days that they can use during the semester. Additional late days will each incur a 0.5% penalty on the course grade (10% of the assignment grade).

At most 2 late days can be spent on a homework assignment.

At most 2 late days can be spent on a computational mini-project assignment.

Regrade Requests

Any regrade requests are due by the third day after receiving the grade.

Office Hours

To encourage discussion and interaction outside of class, each student receives a bonus for attending (before November) a TA office hour session and for attending an instructor office hour session. You don’t need to stay for the entire time. Please contact course staff if you have scheduling conflicts for the entire semester.

Inclement Weather and Digital Learning Days

For a Digital Learning Day, all class activities will be held remotely via a video conferencing tool.

Tentative Semester Schedule

This schedule will be modified later

Week	Day	Date	Content
1	T	Aug 25	Lecture 1: Mathematical Basics
1	R	Aug 27	Lecture 2: Deviation Bounds
2	T	Sep 1	Lecture 3: Deviation Bounds
2	R	Sep 3	Lecture 4: Union Bounds, Learning Theory
3	T	Sep 8	Lecture 5: VC-Dimension
3	R	Sep 10	Lecture 6: Generalization Bounds for VC-dimension
4	T	Sep 15	Lecture 7: Fundamental Theorem of Learning Theory
4	R	Sep 17	Lecture 8: Rademacher Complexity
5	T	Sep 22	Lecture 9: Rademacher Complexity Continued
5	R	Sep 24	Lecture 10: Parametric Convergence
6	T	Sep 29	Lecture 11: Kernels
6	R	Oct 1	Lecture 12: Kernels continued
7	T	Oct 6	Lecture 13: Summary
7	R	Oct 8	Exam 1
8	T	Oct 13	*Fall Break*
8	R	Oct 15	Lecture 14: Bayesian Learning
9	T	Oct 20	Lecture 15: Bayesian Inference
9	R	Oct 22	Lecture 16: PAC-Bayes
10	T	Oct 27	Lecture 17: No Free Lunch, and Inductive Bias
10	R	Oct 29	Lecture 18: Benign Overfitting
11	T	Nov 3	Lecture 19: L1 regularization and sparsity
11	R	Nov 5	Lecture 20: Uniform Stability
12	T	Nov 10	Lecture 21: From labels to actions (contextual bandits)
12	R	Nov 12	Lecture 22: From iid to ood (transfer learning)
13	T	Nov 17	Lecture 23: Summary
13	R	Nov 19	Exam 2
14	T	Nov 24	Lecture 24: Bonus topic (TBD)
14	R	Nov 26	*Thanksgiving*
15	T	Dec 1	Lecture 25: Bonus topic (TBD)
15	R	Dec 3	Lecture 26: Bonus topic (TBD)
16	T	Dec 8	Summary and feedback