

Instructions

CSE 7752 Syllabus

Scientific ML

Instructor Information

Instructor

Raphael Pestourie

Email

rpestourie3@gatech.edu

General Course Information

Description

Welcome to scientific machine learning, a hot and promising topic that combines two of the most exciting fields of engineering: scientific modeling and machine learning!

Machine learning and modeling through ordinary differential equations and partial differential equations are prevalent across all the areas of engineering. Machine learning's paradigm is to see the world through data. While modeling's paradigm is to use knowledge of governing equations, the scientific machine learning approach is to take advantage of both, leveraging knowledge from both model and data. Scientific machine learning leverages computing technologies of the 21-st century, and it has already impacted the world in COVID pandemic modeling, pharmaceutical drug discovery, and climate research, to name a few. Scientific machine learning is a growing research community which encompass major tech companies, other industry, and academia. This class is an opportunity to join this community!

In this graduate-level course, you will discover scientific machine learning thinking through, reading, in-class conversation, and a team project on a complex real-world problem. You will also learn the peer review process, an important part of the nascent scientific machine learning literature.

This is not a ML class nor an ODE/PDE solving class, however, we will cover how to leverage and combine methodologies from ML and ODE/PDE solving to find solutions to real-world engineering problems. This class is not coding heavy, but rather, it will focus on methodology design and peer-review.

This document is meant to help you make the most of this course. I welcome feedback and hope you will let me know if I am meeting your expectations and needs. I am excited to work with you this semester!

Course Learning Outcomes

Upon successful completion of this course, you should be able to

- Combine machine learning and scientific computing to create a scientific machine learning model
- Design or implement a scientific machine learning model to solve a significant real-world problem

- Assess the differences between a mathematical model and its computational implementation
- Complete the peer-review process (of machine learning conferences) as a submitter or reviewer
- Operate in a team in a process-focused progression towards results

Required Course Materials

Linear algebra (MATH 1554 or similar), deep learning (CS 4803 or similar), and coding experience (i.e., comfortable working with python/Julia code and using git)

Grading Policy:

Assignment	Date	Weight (Percentage, points, etc)
Participation/Mini quizzes (individual)	All sessions starting week 2	20% (Each quiz is worth 1 point; your final grade is based on the sum of your 10 best quiz scores.)
Scribe paper discussion (individual)	1 per semester	10% (graded by provided rubric)
Initial project proposal (group)	Week 5	5% (graded SAT/UNSAT, submission more than 1 day late is UNSAT)
Project proposal review (individual)	Week 7	20% (graded via a given rubric, submission more than 2 days late is graded 0 automatically)
Revised project proposal with point-by-point response (group)	Week 11	10% (graded SAT/UNSAT, submission more than 3 days late is UNSAT)
Project design paper (group)	Week 13	25% (graded by provided rubric)
Project presentation (group)	Week 13+	10% (graded SAT/UNSAT, UNSAT if absent)

Note: SAT and UNSAT translate into 100% and 0%, respectively.

Grading Scale

Your final grade will be assigned as a letter grade according to the following scale:

A	90-100%
B	80-89%
C	70-79%
D	60-69%
F	0-59%

Description of Graded Components

Participation/Mini quizzes (individual)

Purpose: Most lectures in this course will have a discussion component. Mini quizzes about the reading will be due on Canvas before the Thursday class starts and help you prepare for in-class conversations. The idea is that while we are not all equal regarding spontaneous in-class participation, Prior Preparation Prevents Poor Performance (aka the five (5) Ps). These mini quizzes strive to create an equal footing among students towards in-class participation. They will also be an opportunity for me to ask for your feedback.

What you need to do: The quiz is about the assigned reading and/or prior class material and is due on Canvas before the class starts. In the case of a reading, you will be asked to summarize the reading in a few sentences. Then you will be guided to focus on select aspects of the reading as a formative exercise towards learning peer-review.

Evaluation: This class is more process-oriented than result-oriented. I trust that processes (prior preparation through the mini quizzes) will bring results (in-class participation). For that reason, the participation grade is solely evaluated on the mini quizzes, specifically the timeliness and concision of the quiz answers. Each quiz is worth 1 point; your final grade is based on the sum of your 10 best quiz scores..

Scribe paper discussion (individual)

Purpose: At conferences, a lot of insights can be gained from Questions and Answers (Q&A) following presentations. The ability to take notes and put in a broader context will maximize your learning from going to conferences.

What you need to do: 1. Take notes of the discussion of the article (this part can be shared with other group members). 2. Individually attempt to put into a wider context of the scientific machine learning literature. 3. Individually reflect on how the reading and the discussion may help you with your project.

Evaluation: You will be evaluated in the quality (not quantity) of your notes and insights. A rubric will also be provided to explain what success looks like.

Initial project proposal (group)

Purpose: The project will be your sounding board as you are exposed to different scientific machine learning methodologies in this class. This initial project proposal comes early in the semester to help you determine an engineering problem you want to study early on. You will produce a short paper that you will get feedback on and that you will expand for your project design paper.

What you need to do: Write a paragraph about your research question, a paragraph on the model input, a paragraph on the model output, a paragraph on the model. Please address risks related to your methodology, and potential mitigations when discussing each of your choices.

Evaluation: Your proposal will be evaluated on timeliness and teamwork.

Project proposal review (individual)

Purpose: Due to the quantity of submissions to ML conferences, it is more and more the case that graduate students will be asked to perform reviews. Reviewing papers is an important service to the community as well as a great way to learn about the state-of-the-art in your field. However, it is important to follow appropriate standards which comes with practice. In this class, you will experience the review process both as a submitter and a reviewer. This is one of the intended outcomes of this class because members of the scientific machine learning community are asked to review articles very early on in their research career.

What you need to do: You will perform a review of your assigned initial project proposal following a given rubric. Remember that the people who submitted the proposal are like you and are also part of your community; you want to help them improving their project so that it makes a great contribution to the field! Start with a summary of the paper you are reviewing to show your understanding, then detail your constructive feedback following our rubric given in class about the research question, the input, the output and the model. The review should take at least 1 page without counting references.

Evaluation: Your review will be evaluated on quality (thoughtfulness), collegiality, and timeliness. A rubric will also be provided to explain what success looks like.

Revised project proposal with point-by-point response (group)

Purpose: Responding to reviewers is an important step of the review process. Through thoughtful point-by-point responses to the reviewer comments and thorough manuscript revision, you can dramatically increase your likelihood to have your paper accepted by a conference or a journal.

What you need to do: You will update the manuscript of your initial project proposal and provide a point-by-point response document where you address all the points raised by the reviewers. Keep in mind that if you feel that the reviewer did not understand your point, it might be because it was not clear in the first iteration of the proposal.

Evaluation: Your project proposal and the responses to the reviewers will be evaluated on quality (thoughtfulness). A rubric will also be provided to explain what success looks like.

Project design paper (group)

Purpose: In this class, you will design a scientific machine learning methodology to solve a real-world engineering problem. Ideally, with this work in hand, you would be set up to start doing research as part of the scientific machine learning community and submit to a journal or a conference once you have the results.

What you need to do: Rearrange and expand content into a motivation section of your project, a section on how it can relate to a data problem or a modeling problem, a section on prior work, a section on what will you baseline be, a section on what results you hope to have to show the success of your methodology, a section on risks related to your methodology, and potential mitigations. As well as a carefully-written: "The abstract should start by outlining the background to your work (why the topic is important) and the main question you have addressed (the specific problem that initiated your research), before going on to describe your new observations, main conclusions and their general implications, introduced by the phrase "Here we show" or equivalent. Because researchers from other sub-disciplines will be interested in your results and their implications, it is important to explain essential but specialised terms concisely. Overall, the abstract should contain between 150 and 200 words." (from Author checklist of Nature Springer)

Evaluation: Your work will be evaluated on quality (thoughtfulness), clarity, timeliness, and teamwork. A survey will be sent to you to evaluate your contribution and that of your teammates confidentially. A rubric will also be provided to explain what success looks like.

Project presentation (group)

Purpose: The ability to present your work effectively and to articulate the motivation and results clearly is key to becoming a successful researcher both in academia and in the industry. If you aim for industry, working on the presentation could set you apart from others when talking about your project in an interview.

What you need to do: Prepare 5-6 slides about your project to present to the class: 1. Problem motivation, 2. Problem characteristics that relate to scientific machine learning, 3. Related work and limitations of prior art, 4. Main result: choice of the approach and tradeoffs, 5 (Encouraged but not mandatory) Early results, 6. Concrete next steps outline to get to a ML conference submission. The presentation will be followed by Q&A.

Evaluation: The presentation and Q&A will be evaluated on clarity and engagement with the broader class. A rubric will also be provided to explain what success looks like.

USG Required Course Policies

Attendance and/or Participation

This class will not be recorded. Attendance is key to help you learn and feel that you belong to the scientific machine learning community because we will dedicate some time to discussions during every lecture. Attendance is mandatory for graded assessments and strongly recommended otherwise, however I acknowledge that some students may have strenuous circumstances that prevent them to attend some of sessions. As mentioned in the assessment description, participation will be evaluated based on preparation of in-class discussions in the form of mini quizzes.

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 plagiarism on a quiz, exam, or assignment will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations.

Core IMPACTS

Not applicable

Additional Georgia Tech Required Policies

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 in order to set up a time 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.

Teamwork

Teamwork creates synergy. Because the combined effect of an effective team is significantly greater than the sum of individual efforts, teams can tackle problems that are too big to solve for any individual. In the professional world, effective teamwork is paramount. For this reason, this class uses a team-based approach that will help you develop collaborative skills, that will help you work effectively in a team, and that will maximize your learning. As in the professional world, three important features affect your productivity and success in a team: your own effort, the effort of people you depend on, and the way you work together.

Throughout the term, you will work closely with other classmates, as part of a project team. Research on teamwork suggests the following good team practices:

- Come to class prepared. Before working as a team, read any relevant material(s) and formulate your own approach to the task at hand.
- Actively participate and contribute to all activities when the team is together (both in and out of class). When even one team member checks out and starts working individually (or starts checking email, text messages, etc.) instead of engaging with the team, the overall performance of the team is adversely affected.
- In all team activities, be prepared to share three things with your teammates: (a) what approach you chose as an individual, (b) why you chose that approach, and (c) how confident you are about your approach.
- Be respectful and listen and evaluate other people's points of view.
- Deliberate as long as time permits. Regardless of the make-up of the team, teams that deliberate longer do better in team activities.

Peer Assessment

It is important to provide positive feedback to people who truly worked hard for the good of the team and to also make suggestions to those you perceived not to be working as effectively on team tasks. Twice times during the semester you will provide a confidential online assessment of the contributions of the members of your team (including yourself) to all the activities in class and to the project. The feedback you provide should reflect your judgment of each team members':

preparation - were they prepared when they came to class?

contribution - did they contribute productively to the team discussion and work?

respect for others' ideas - did they encourage others to contribute their ideas?

flexibility - were they flexible when disagreements occurred?

Your teammates' assessment of your contributions and the accuracy of your self- and peer-assessments play an important role in your final grade for the course.

Failure — the unavoidable price of success

Throughout your education, you have probably been led to view mistakes and failure as something that is unfavorable and that negatively affects you — something to be ashamed of. However, success is not possible without taking calculated risks, which inevitably means failing sometimes. The road to creativity and innovation is littered with failure. "If you haven't failed, you're not trying hard enough," goes a well-known saying. Failure is a problem only if it is your end point. On the way to finding a solution, failure can be very productive as it can teach you a lot (what doesn't work, what might work, and what you might want to explore in greater detail) and lead you to success.

In this class, we want to create a culture that encourages creativity and calculated risk-taking. Also, we design all the activities in this class so that they leave ample room for errors for anyone (including the staff). Only then can we guarantee that everyone's learning will be maximized and that you will learn to feel comfortable with the (productive) failures that go hand-in-hand with creativity. See them as learning opportunities, not negatives, as stepping-stones to success, not the end point. So be bold and take risks, both as an individual and as a team — failure, even repeated failure, is a healthy and necessary part of becoming successful. We trust the process rather than the results: most of the

assignments will be graded SAT-UNSAT, and some assignments will give you a chance to resubmit if you have an UNSAT on the first submission

Late policy

There is no late submission allowed for mini quizzes because they are meant to prepare the in-class discussion. However, up to five (5) UNSAT quizzes are allowed throughout the semester and will be dropped from the final grade. Due to its collaborative nature, it is not possible to make up any team project presentation. The assignment related to the peer-review process have a limited grace period because the submissions are needed for the assignments in the next step of the review process. Approved Institute activities (IV. Attendance < Georgia Tech.) will trump the present late policy.

Community norms

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. See XXII. Student-Faculty Expectations < Georgia Tech. for an articulation of some basic expectation 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.

We will also spend time in the first lecture to co-create community norms that create a classroom environment that is psychologically safe for everyone and conducive to learning.

Student Use of Mobile Devices in the Classroom

Mobile devices to enhance learning (laptop or tablet, NOT smartphones) are allowed in the classroom and will likely be needed in support of discussions. Although I won't be able to check, I would encourage you to stay away from emails, social media, and other multitasking inherited from the Zoom era to help you stay focused on the lecture.

Preparation for guest speakers

Preparation for guest lectures will be using the same process of mini quizzes as regular lecture. In the mini quizzes prior guest lectures, you will be asked to prepare 1-2 question based on the assigned reading authored by the speaker. In the mini quizzes following guest lectures, you will be asked to evaluate the presentation based on a given rubric. This same rubric will be used to evaluate your project presentation

Large Language Models (ChatGPT, GitHub Copilot)

We treat AI-based assistance, such as ChatGPT and Copilot, the same way we treat collaboration with other people: you are welcome to talk about your ideas and work with other people, both inside and outside the class, as well as with AI-based assistants.

However, all work you submit must be your own. You should never include in your assignment anything that was not written directly by you without proper citation (including quotation marks and in-line citation for direct quotes).

Including anything you did not write in your assignment without proper citation will be treated as an academic misconduct case. If you are unsure where the line is between collaborating with AI and copying AI, we recommend the following heuristics:

Heuristic 1: Never hit "Copy" within your conversation with an AI assistant. You can copy your own work into your own conversation, but do not copy anything from the conversation back into your assignment.

Instead, use your interaction with the AI assistant as a learning experience, then let your assignment reflect your improved understanding.

Heuristic 2: Do not have your assignment and the AI agent open at the same time. Similar to the above, use your conversation with the AI as a learning experience, then close the interaction down, open your assignment, and let your assignment reflect your revised knowledge.

This heuristic includes avoiding using AI directly integrated into your composition environment: just as you should not let a classmate write content or code directly into your submission, so also you should avoid using tools that directly add content to your submission.

Deviating from these heuristics does not automatically qualify as academic misconduct; however, following these heuristics essentially guarantees your collaboration will not cross the line into misconduct.

Course Schedule

The course schedule of this class will be populated on Canvas progressively during the semester at least one week ahead of the class. It will include class topics, assigned readings, and assignment due dates. You can refer to the “Course Requirements & Grading” for the timeline of the assignments.

Note: the sections on teamwork, peer assessment, and failure were adapted from the excellent AP50 Spring 2020 syllabus by Erik Mazur at Harvard University.

Graduate Student Academic and Professional Success Resources: For graduate A list of resources for graduate students is given on the [Office of Graduate and Postdoctoral Education](#) website. Specific information for [current graduate students](#) 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; and
- [Professional Development](#) such as the programming from the Career Center and other professional development resources and events”]

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 ([student-resource-guide \(gatech.edu\)](#))