AMCS 6045: Topics in Numerical Analysis and Scientific Computing: Machine learning approaches for inverse problems

Course Syllabus

Instructor: Shanyin Tong

Email: tong3@sas.upenn.edu

Lectures: Tuesdays & Thursdays 1:45 PM-3:14 PM, DRLB 3C6

Office hours: TBD

Course Description. This course aims to understand the mathematics behind machine learning approaches and to discuss their ideas and potential to impact the field of inverse problems. It is a research-focused course, offering students valuable opportunities, open questions, and side projects to conduct mathematical and computational research. The course will include learning frameworks (both supervised and unsupervised), optimization methods, Bayesian and variational inference perspectives, and their use in inverse problems. If time permits, we will also cover topics in data assimilation and rare events.

The main objective is to prepare students with mathematical, computational, and machine learning tools and theory to conduct further research in applied mathematics or related fields, and to connect these tools to various applications in applied mathematics.

Prerequisite. A good understanding of probability theory (especially conditional probability and Bayes' theorem), linear algebra, elementary analysis, and multivariable calculus are required. Basic knowledge on optimization, partial differential equations, numerical methods, and basic programming skills are desirable.

Website and Communication. The chief means of communication for this course will be the course Canvas site, accessed through `https://canvas.upenn.edu/`. Students are expected to check this for up-to-date assignments, including material separate from the text and announcements. Ed Discussions is an online discussion board for all issues related to content and logistics for the course. The communications and discussions will be conducted using this platform. You can access it through Canvas.

Homework. We employ an online grading system called Gradescope `https://gradescope.com/`. This should expedite the grading process and keep all assignments well-organized. You will be automatically enrolled, and the link can be accessed through Canvas. The assignments will be released and submitted through Gradescope.

There will be 3-4 homework assignments, and will be a mixture of theoretical questions, computational questions, and reading assignment and a short reading report. You do not need to have experience coding beforehand, but you must be willing to learn basic programming during the class. A high-level software package that lets you easily plot things, such as Python or Matlab, is recommended.

Please pay attention to the deadline of each homework, it may vary during the semester. We do not grant extensions for homework except for emergencies (need be documented and contact beforehand). For late submission, a penalty of 20% will apply for each day after the assigned deadline.

Final project. This research-focused course will require each student to complete a final project. Students may work individually or in groups (number of group members ≤ 2. If number of group members ≥ 2, the final report must include an additional section describing each member's contributions to the collaboration). The project topic must be related to the course content. You will be asked to submit a midterm report (Due: Oct 14) outlining the project topic, problem statement, and proposed methods and a 5-min midterm presentation on class to summarize the report. At the end of the semester, you need to submit a full final report (Due: Dec 10) in the format of a scientific paper. A final presentation (on Dec 4) of your project will also be required.

Academic Honesty. You are encouraged to work with others, on the homework problems and to study. However, you must write up your own solutions. Students who violate university rules on academic dishonesty are subject to disciplinary penalties, including the possibility of failing the course and/or dismissal from the University. Detailed information on academic integrity at Penn is available here: `https://catalog.upenn.edu/pennbook/code-of-academic-integrity/`.

The appropriate use of AI in this course is for brainstorming ideas (similar to how you would use Google as a search engine), for editing/revising code, or for proofreading your drafts. Using such tools to generate an entire assignment or to copy any other person's assignment will be considered a violation of Penn's Code of Academic Integrity, and suspected use will be reported to the Center for Community Standards & Accountability.

If you discuss homework with others or use AI tools in your homework or project, please include a short paragraph at the beginning of your submission describing this usage.

Attendance. Regular attendance is expected. Please notify me in advance if you must be absent due to an emergency.

Grading. 30 % Homework + 20 % Midterm report + 50% Final project & presentation

Textbooks. The theoretical part of this course is mainly based on the following notes:

- *Inverse Problems and Data Assimilation: A Machine Learning Approach*, Eviatar Bach, Ricardo Baptista, Daniel Sanz-Alonso, Andrew Stuart (2024): `https://arxiv.org/abs/2410.10523`. You can also find a list of references at the end of each chapter of these notes for different topics. These provide a rich set of resources if you wish to read further or generate ideas for your final project.

There are also other related literature for further reading (more to come):

- *Computational methods for inverse problems*, Vogel, Curtis R., Society for Industrial and Applied Mathematics (2002).
- *Inverse Problems and Data Assimilation*, Daniel Sanz-Alonso, Andrew M. Stuart, Armeen Taeb (2018): `https://arxiv.org/abs/1810.06191`
- *Solving inverse problems using data-driven models*, Arridge, Simon, et al., Acta Numerica 28 (2019): 1-174.

Tentative Schedule & Topics. Here is a rough schedule for the lectures:

- Introduction to inverse problems
  - Mathematical formulation
  - Application examples
  - Challenges, ill-posedness
  - Other related problems: data assimilation, rare event studies
- Inverse problems – deterministic & statistical perspectives
  - Optimization framework, regularization
  - Bayesian inversion
  - Connections through MAP
  - Error analysis
- Machine learning in inverse problems:
  - Variational inference, distribution learning for prior, posterior, or their maps
  - Forward surrogate
- Learning frameworks:
  - Different metrics, divergence
  - Unsupervised learning: density estimation, transport methods, normalizing flows, diffusion and generative models
  - Supervised learning: neural networks, PINNs, random features, Gaussian process
  - Optimization: algorithms for unconstrained/constrained optimization, auto-differentiations, PDE-constrained optimization, adjoint method and variational frameworks.

This schedule is only tentative. Changes of the schedule will be announced in class.

A detailed and updated syllabus (with class resources, communication and submission tools) will be given to students before classes begin, and will be reviewed together with students on the first day of the class (August 26).

If you have difficulty enrolling in the class, please contact me or the Math/AMCS administration.