[AI for Science] Syllabus

Instructor Information

Instructor	Email	Office Location & Hours
Ziming Liu	zmliu@mit.edu	TBA

General Information

Description

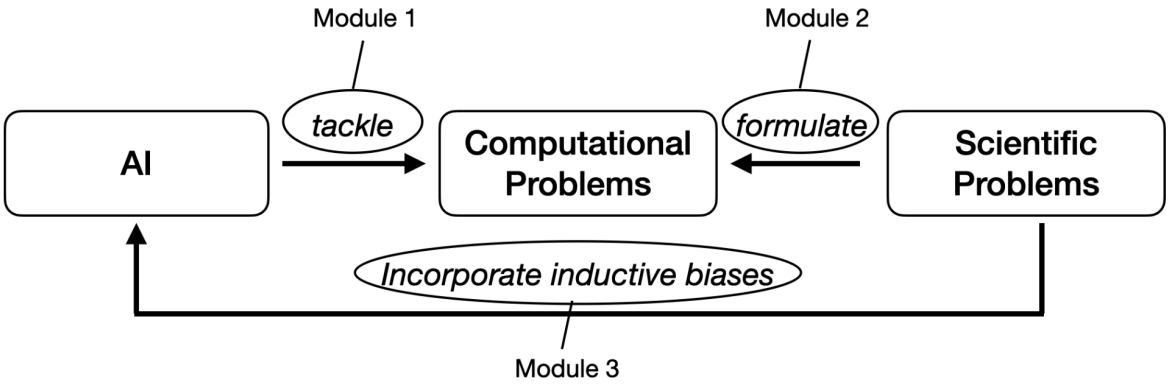

AI for science has become a big interdisciplinary field appealing to researchers from both AI and science. However, this is such an enormously big field that (1) any two people in the field barely understand each other's work; (2) researchers new to this field might be overwhelmed by the huge body of literature and do not know where to start. To make the field approachable to more researchers, and to foster collaboration between AI researchers and scientists, it is urgent to come up with (a) unified languages for people from different backgrounds to communicate and (b) global views of the whole field of AI for Science so that people can better connect with others by understanding the similarities and differences of the problems they are studying. This course attempts to provide such languages and views that connect people in AI for science.

This course has three modules:

Module 1: Introduction to AI, with a focus on the question "How to use AI to tackle well-formulated computational problems?"

Module 2: Overview of scientific fields with the central question "How to formulate scientific problems into computational problems that AI readily solves?"

Module 3: Building inductive biases in science to AI systems to facilitate learning

Prerequisites

This course targets graduate students and competitive undergraduates interested in AI for Science, regardless of their backgrounds – computer science, natural sciences, mathematical sciences, etc. Familiarity with either AI **or** Science is required. Familiarity with both AI and Science is preferred but not mandatory. Familiarity is defined as (1) AI: having experience in programming and training neural networks. (2) Science: finished undergraduate-level courses (or equivalent knowledge level) and/or having research experience in one or more scientific disciplines.

Expectations and Goals

(1) Understand how to formulate a scientific problem into a computational problem
(2) Know which AI models are suitable to solve which computational problems
(3) Be able to more effectively communicate with fellow AI researchers and scientists who work on different problems in AI for Science.

Course Schedule (Tentative)

Week	Topic
[Week 1 (Intro)]	Introduction to AI for Science
[Week 2 (M1)]	AI for property prediction: supervised learning
[Week 3 (M1)]	AI for concept discovery: unsupervised learning, self-supervised learning
[Week 4 (M1)]	AI for generation: Generative models
[Week 5 (M1)]	AI for search: Reinforcement learning, differentiable programming, Monte Carlo tree search
[Week 6 (M1)]	AI for planning & control: Reinforcement learning, LLM
[Week 7 (M2)]	Mathematics: AlphaGeometry, Knot theory, and Repr theory
[Week 8 (M2)]	Physics: Physics-informed learning, Fermi-net
[Week 9 (M2)]	Chemistry: NN force field, molecular synthesis
[Week 10 (M2)]	Biology: AlphaFold
[Week 11 (M2)]	Climate: Weather forecasting
[Week 12 (M3)]	Teaching AI Symmetries: Equivariant neural networks
[Week 13 (M3)]	Teaching AI Conservation laws: Lagrangian/Hamiltonian NN

[Week 14 (M3)]	Teaching AI qualitative insights: Representation matters, conditional generation, soft constraints
[Week 15 (M3)]	What can AI teach us? mechanistic interpretability
[Week 16 (End)]	Final exam or course project presentations