# CS839: AI for Scientific Computing
## **Course Overview**

Misha Khodak

University of Wisconsin-Madison

20 January 2026

# Outline for today

- Logistics
- Course topics
- Goals of the class

# Logistics

# **Logistics:** Class times

- **Location**: Morgridge Hall 2538

- **Time**: Tuesdays & Thursdays from 1PM – 2:15PM

# **Logistics**: Enrollment

- Currently at capacity of 30 students
  - Typically, there is a lot of churn with 839s, so you may get in
  - But hard to predict for any given 839 topic

# **Logistics**: Teaching Team

Instructor: **Misha Khodak**

- Office Hours: by appointment
- Contact: khodak@wisc.edu
- Website: https://pages.cs.wisc.edu/~khodak/

# **Logistics**: Content (Two locations)

[Info + Schedule] **Course website**:
https://pages.cs.wisc.edu/~khodak/cs839spring2026

[Graded content] **Canvas:**
https://canvas.wisc.edu/courses/500358

# **Logistics**: Lecture Format

There will be four types of lectures:

1. **background lectures** (first quarter of the class)
2. **research lectures** (second quarter of the class)
3. **paper presentations** (third quarter of the class)
4. **project presentations** (fourth quarter of the class)

# Background lectures

Content:
- background on machine learning
- basics of AI for scientific computing

Format:
- presented by me
- questions encouraged
- no explicit participation required

# Research lectures

Content: cutting-edge research on AI for scientific computing

Format:

- presented by **experts at UW-Madison**
- questions encouraged
- participation graded (to be explained)

# Paper presentations

Content: cutting-edge research on AI for scientific computing

Format:

- presented by **student reading groups**
- questions encouraged
- participation graded (to be explained)

# Project presentations

Content: cutting-edge research on AI for scientific computing

Format:

- presented by **student project groups**
- questions encouraged
- participation graded (to be explained)

# **Logistics**: Grading

**Participation**:

- 30% of total grade
- assessed during research lectures and paper/project presentations
- comes in the form of questions submitted to a Canvas quiz in class

**Presentations**:

- 40% of total grade
- given in groups of 2-3

**Final project**:

- 30% of total grade
- done in groups of 1-4
- assessed via presentation and a project report

# **Logistics**: Participation

**During research lectures and paper/project presentations (but not background lectures):**

- students will think of and submit questions to the speaker(s) via a Canvas quiz
- students do not need to ask these questions out loud if they do not wish to do so
- **all questions must be thought of independently;** if someone else asks the same or similar question out loud *before* you can ask it, submit the speaker's answer in addition to your own version of the question
- **questions can only be submitted to Canvas during class time**

**Grading**

- questions will be assigned a binary grade based on relevance and significance (e.g. typo or basic clarification questions do not count)
- a third of the participation grade will be determined by taking the seven top-graded questions from across all research lectures, but only the first two questions during each lecture will be counted; this allows flexibility in case a student is not able to attend some of the classes, but also means that requests to make up these points in some manner are extremely unlikely to be granted
- the remaining thirds of the participation grade will be determined in the same manner based on the paper presentations and the project presentations, respectively

# **Logistics**: Presentations

Two presentations, each making up half the presentation grade.

The first will be of an existing paper or a cohesive line of work

- groups of 2-3
- assessed based on clarity of presentation, understanding of the material, and analysis of the consequences of the research
- topics should be determined in consultation with the instructor but can draw upon the students' own research backgrounds; I will suggest some topics, but you should start thinking of your own as well

The second is a progress report on the final project

- groups of 1-4
- assessed based on clarity of presentation, assessment of research progress and roadblocks, and analysis of next steps

# **Logistics**: Final projects

- groups of 3-4
- topics should be determined in consultation with the instructor but can draw upon the students' own research backgrounds, so long as they are sufficiently self-contained and relevant to the course topic
- I will send out some ideas, but you should start thinking of project topics now as well

**Deliverables:**

- **a two-page project proposal is due by midnight on March 27$^{th}$**
- **an eight-page project report is due by midnight on May 4$^{th}$**
- both may use unlimited room for references and appendices, should be typeset in LaTeX with 11pt font and 1-inch margins, and must be emailed to the instructor with all team members cc'd.

# **Logistics**: Recommended reading

No textbook

For (others') paper presentations, you are encouraged but not required to read the paper(s) before class.

# Q&A break on logistics

# Course topics

# **Course topics**: Required background

For success in this course you will need **one of the following:**

- working knowledge of machine learning (ML)
- ability to get up to speed quickly as needed to read papers and perform research (typically a strong knowledge of **linear algebra, optimization, and statistics suffices**)

There will be one review lecture on machine learning; for deeper understanding take CS 540, CS 760, ME 737, etc.

**Topics in scientific computing will be introduced as needed.**

# **Course topics**: What is this class about?

**How tools from AI are being used to replace or augment traditional scientific computing**

Here "AI" is defined broadly to include machine learning, optimization, statistics, etc.

This class has overlap with whatever "AI for Science" is, but

- that term has no agreed-upon definition
- the course was not called that on-purpose

# **Course topics**: Overview

Paradigms in which AI tools are applied in the natural sciences:

1. replacing, correcting, or augmenting traditional simulators
    - neural operators
    - physics-informed neural networks
    - simulation-based inference

2. processing large-scale experimental data
    - symbolic regression
    - property prediction

3. optimization and control

**This class will focus on 1** but also discuss 2 and 3.

# **Course topics**: Neural PDE solvers

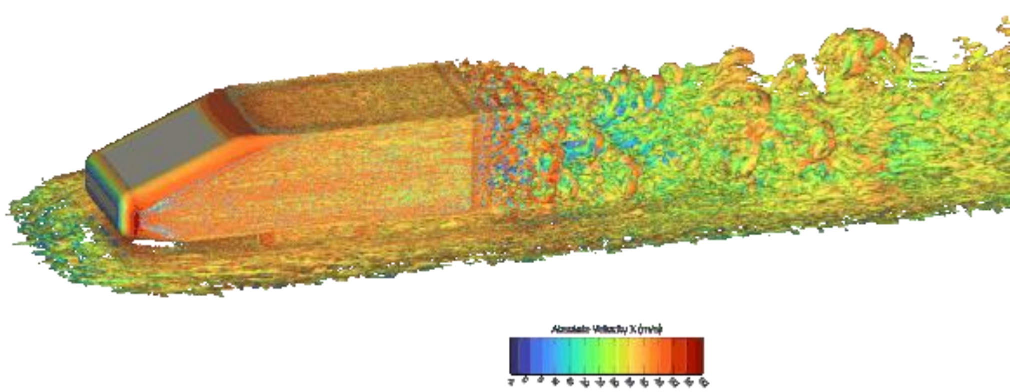

**Goal:** predict the evolution of a physical system from some initial conditions

Usually governed by a PDE relating different quantities across time and space

$$\nabla \cdot \overline{u} = 0$$

$$\rho \frac{D\overline{u}}{Dt} = -\nabla p + \mu \nabla^2 \overline{u} + \rho \overline{F}$$

**Why?**

- Scientists use simulation to gain understanding without running experiments
- Engineers use computer-assisted engineering (CAE) software for development
- Weather forecasting
- Physics visualization in movies/games
- Data-generation

**The Complicated Computer Modeling Behind an Animated Snowball**

BY AISHA HARRIS NOV 26, 2013 • 12:02 PM

# **Course topics**: Neural PDE solvers

Traditional approach:

- discretize in time
- discretize in space
- pose the solution at timestep $t+1$ as the solution of a linear system
- solve the linear system and advance to the next timestep

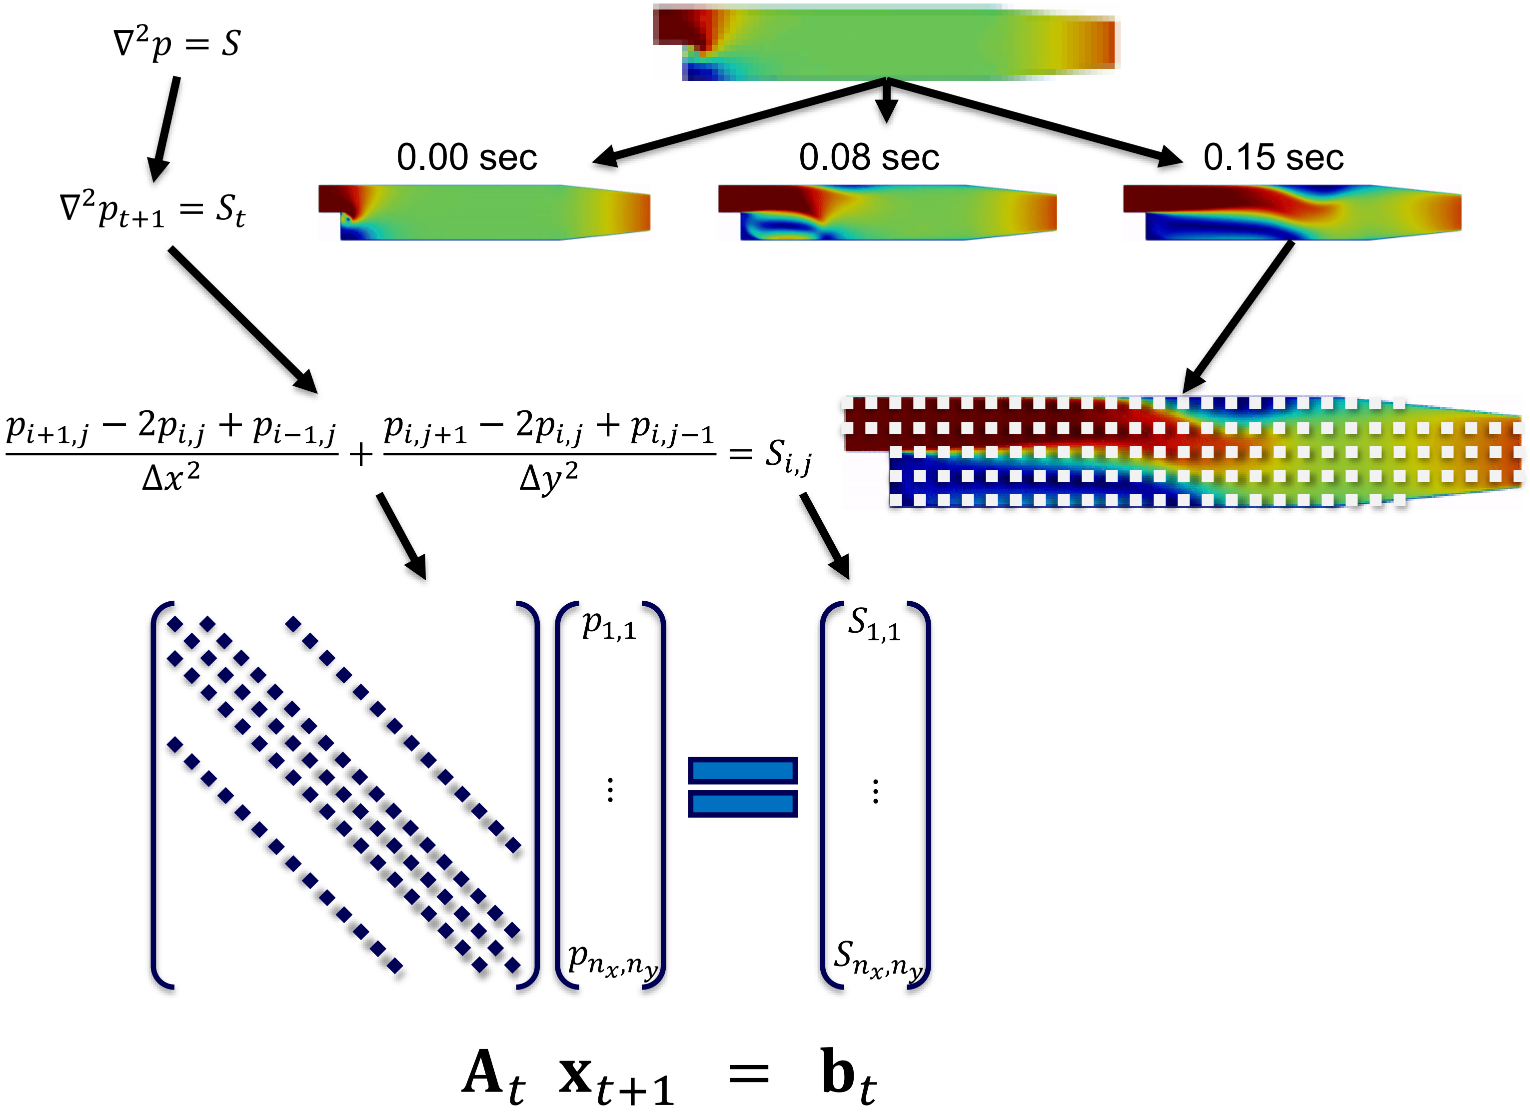

# **Course topics**: Neural PDE solvers

**Many other approaches to solving PDEs, but the pattern is**

- discretize space and time
- solve an expensive numerical problem (linear system, nonlinear system, Riemann problem) to advance in time

**This can be too slow for many applications of interest**

- real-time prediction and control
- solving inverse problems

# **Course topics**: Neural PDE solvers

One data-driven approach: **neural operators**

- use classical solvers to generate solutions for many different related systems
  - same PDE, different initial conditions
  - different PDEs in the same family
  - …
- train a neural network to output the solutions instead

Pros: very fast after training

Cons: expensive to generate data, typically low accuracy

# **Course topics**: Neural PDE solvers

Neural operator success story: **Weather**

NVIDIA's FourCastNet:

- neural network trained on historical weather data
- 60x faster 15-day forecast than traditional simulation

Key reasons for success:

- availability of historical data
- imperfect PDE model
- CPU-bound competitors

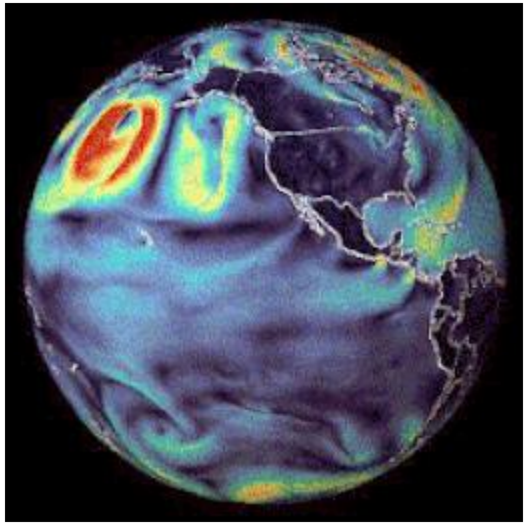

# **Course topics**: Neural PDE solvers

Another way: **physics-informed neural networks (PINNs)**

- set up a neural network that takes in grid points and outputs physical quantities at each one
- minimize the sum of squared residuals at those grid points w.r.t. the neural network parameters

Pros: can incorporate experimental data, no data generation

Cons:

- optimization doesn't always work
- requires higher-order derivatives

# **Course topics**: Application domains

Background lectures will typically have physics as the application domain

AI has been very successful in other domains as well

Paper presentations and final projects are welcome to consider data-driven research in any natural science or engineering field

# Q&A break on topics and application domains

# Goals of the class

# Goals of the class

1. **understand** the uses and limitations of machine learning applied to scientific computing

2. **develop** skills in conducting and presenting research

# **Goals of the class:** Key questions

We will see a lot of recent and ongoing research in an overhyped domain

**With any claimed advance, our goal will be to understand two things:**

1. What are we able to do now with AI that we couldn't do before?

2. Was it worth the cost?

# Thanks Everyone!