Course overview and logistics

CS861: Theoretical Foundations of Machine Learning

Kirthi Kandasamy
University of Wisconsin - Madison
Fall 2024

September 4, 2023

AI/ML is popular nowadays!

- *“A breakthrough in ML will be worth 10 Microsofts”*
  - Bill Gates

- *“ML is the new internet”*
  - Tony Tether, Director, DARPA

- *“AI will be the best or worst thing ever for humanity”*
  - Elon Musk

Object detection & segmentation

Weather forecasting & Climatology

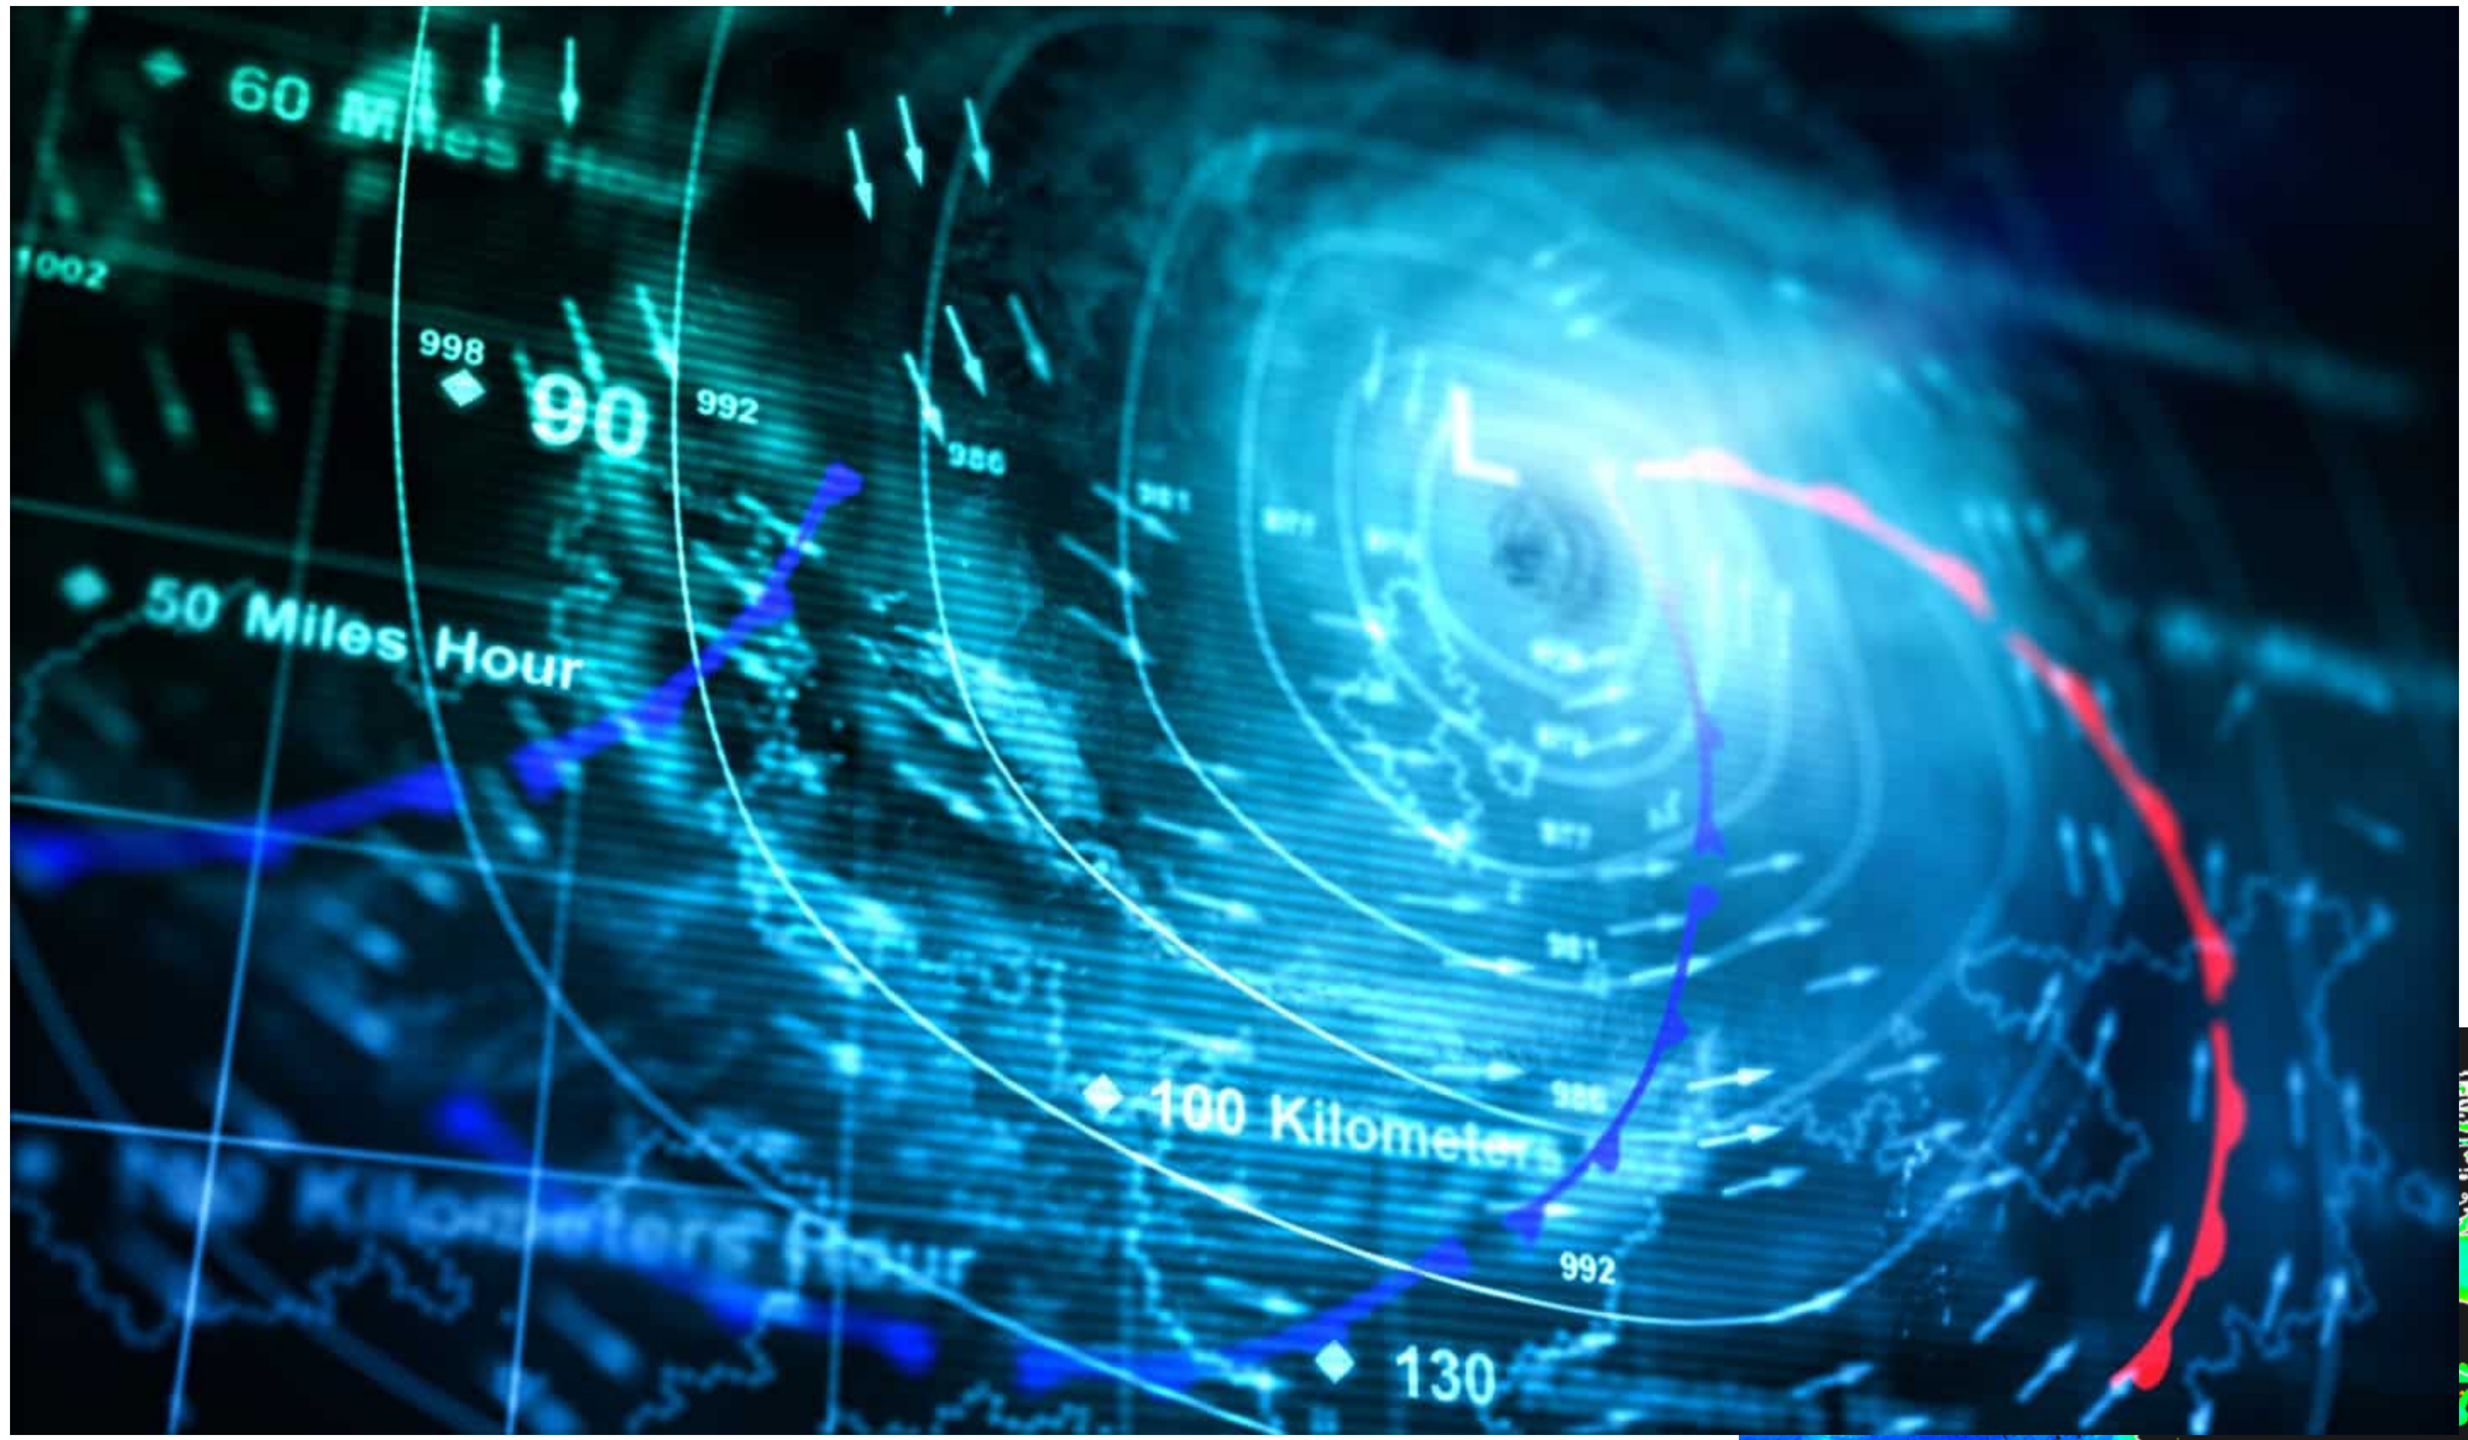

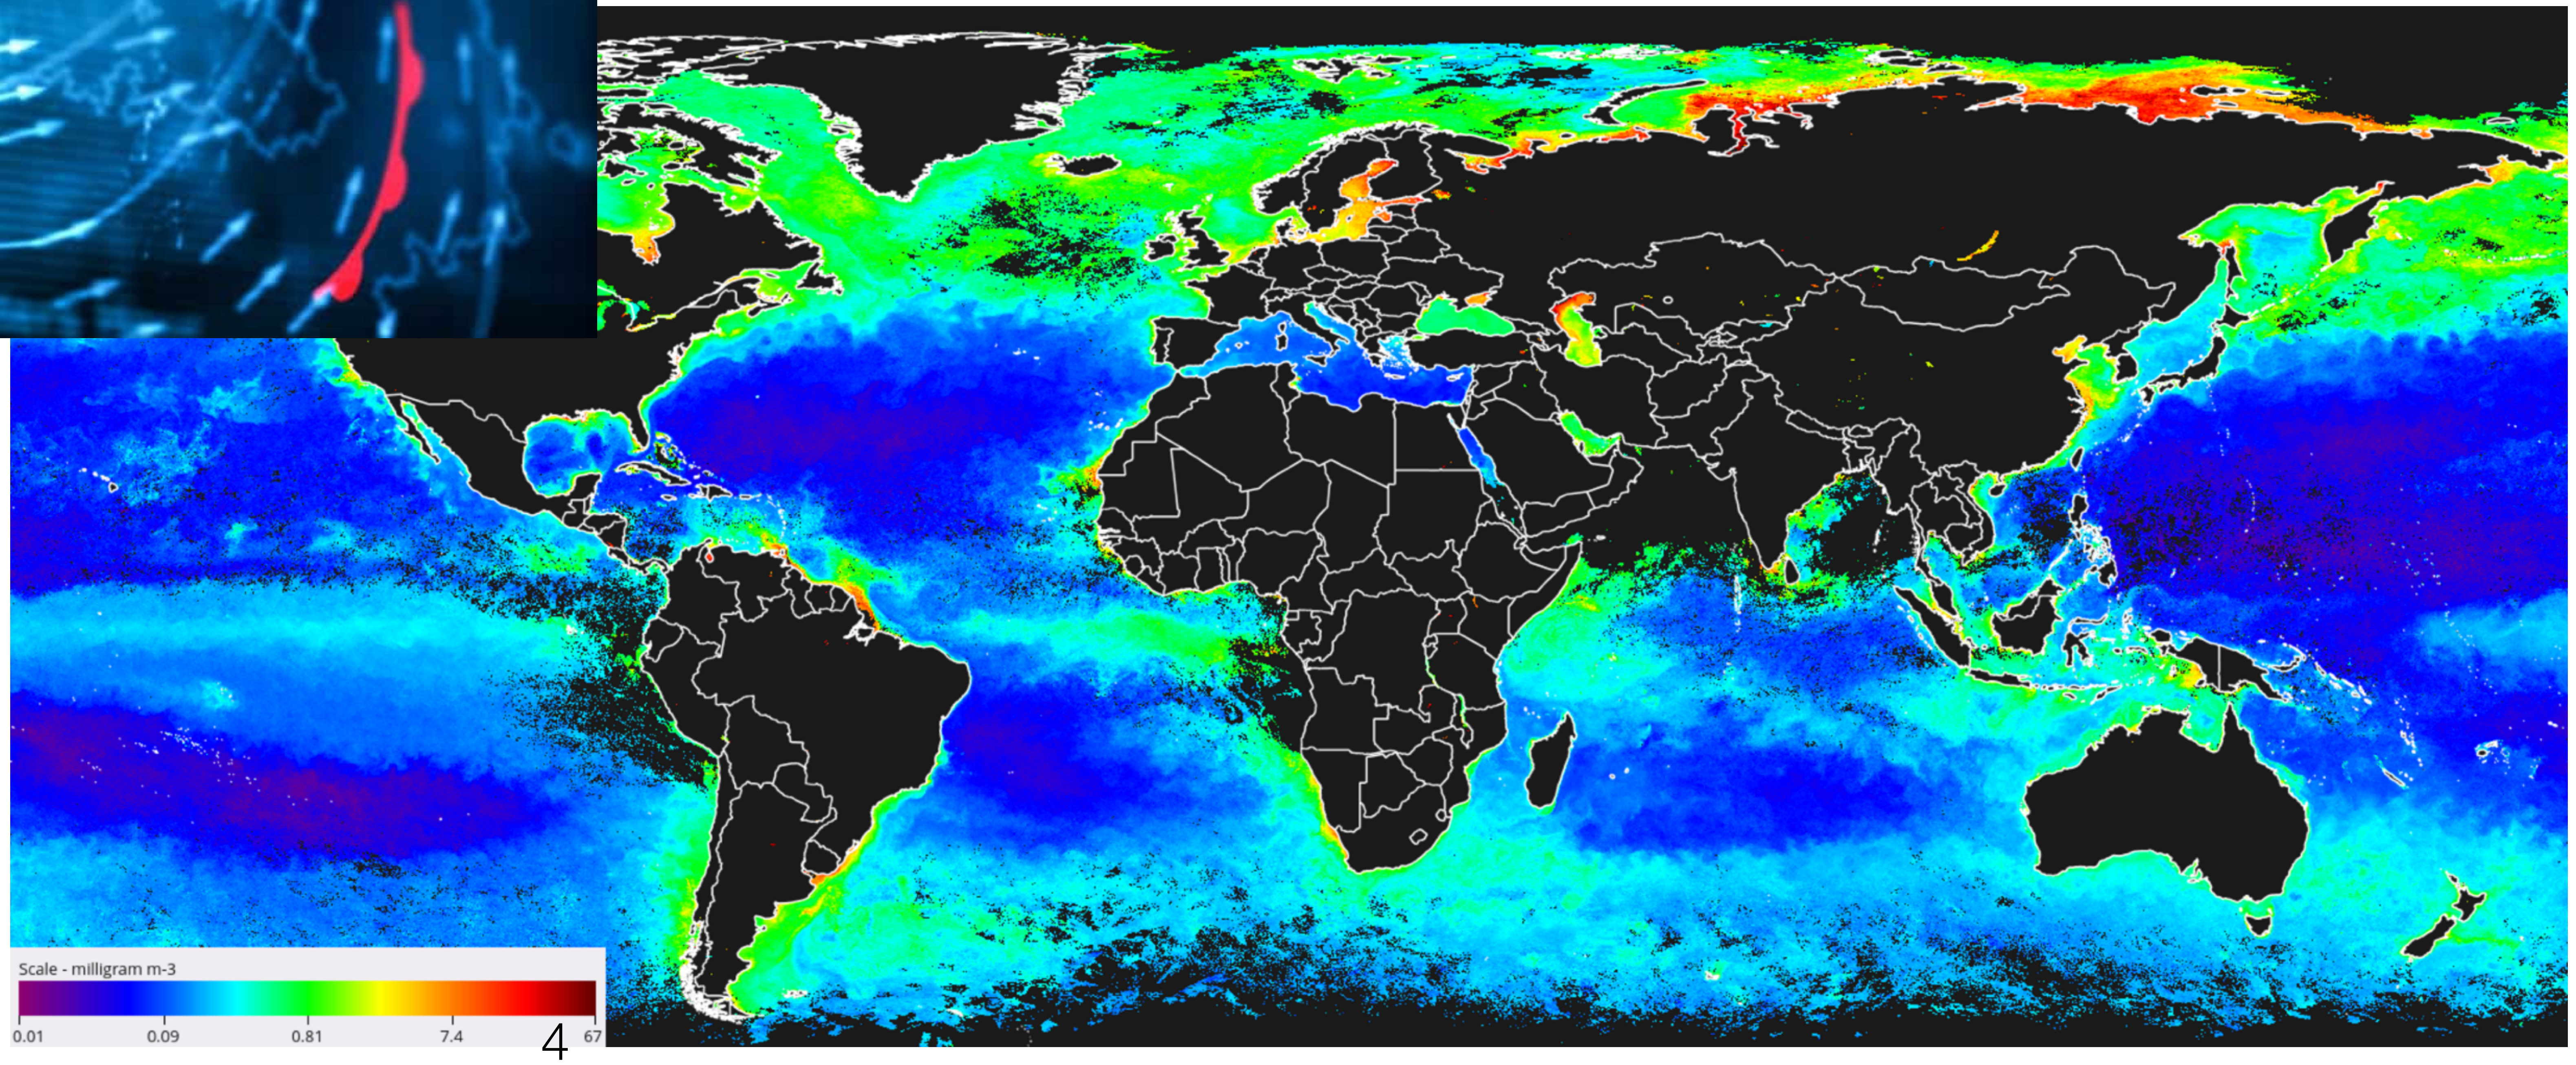

Optimizing drugs and materials

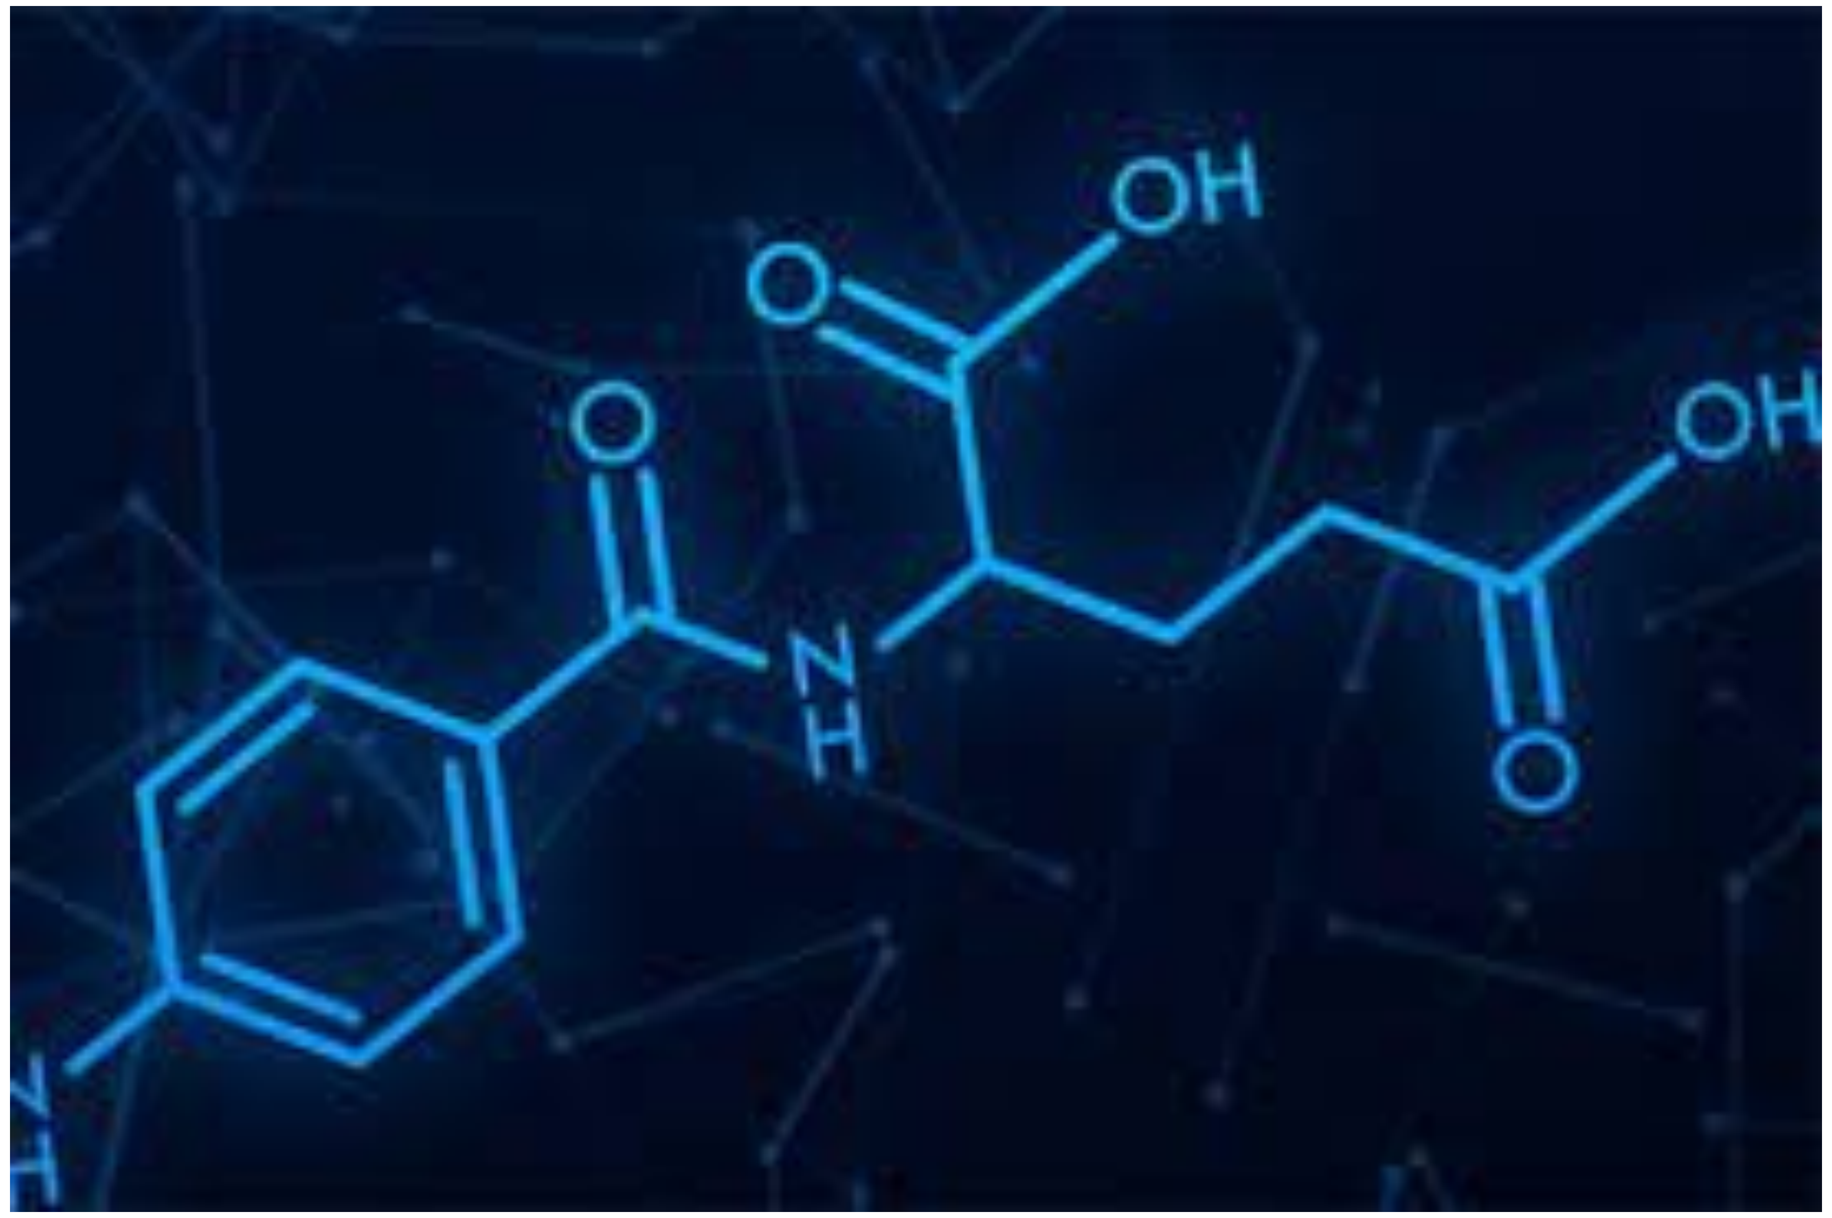

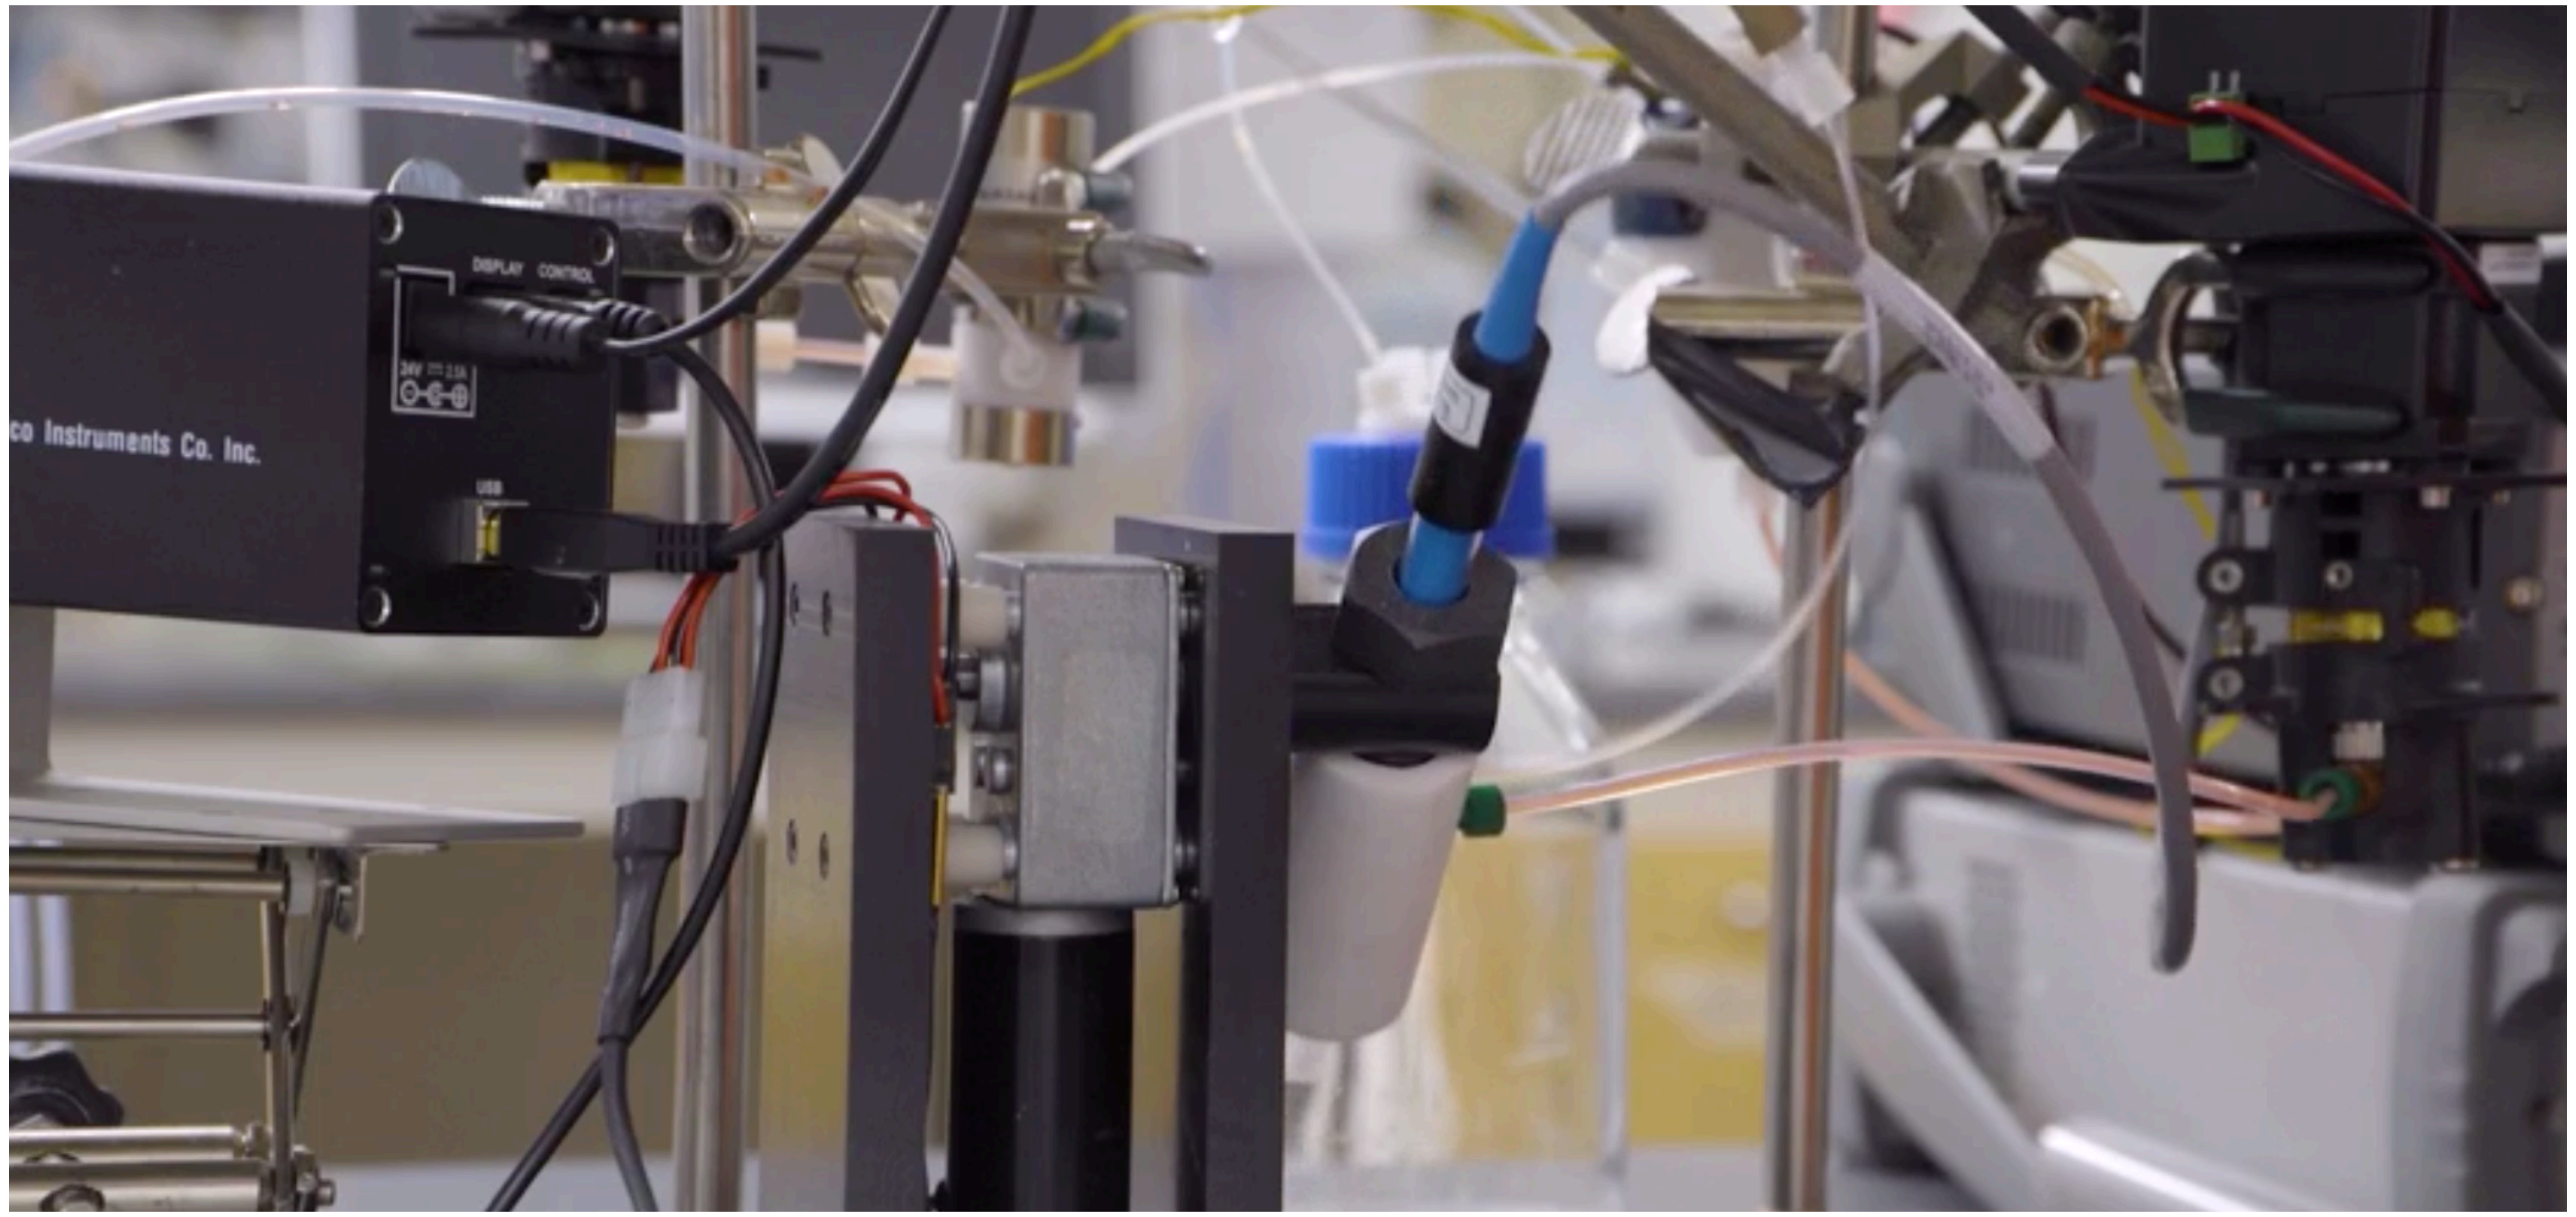

Autonomous vehicles

Language generation & LLMs

Q: What is your favorite animal?
A: My favorite animal is a dog.

Q: Why?
A: Because dogs are loyal and friendly.

Q: What are two reasons that a dog might be in a bad mood?
A: Two reasons that a dog might be in a bad mood are if it is hungry or if it is hot.

Q: How many eyes does a giraffe have?
A: A giraffe has two eyes.

Q: How many legs does a frog have?
A: A frog has four legs.

Q: Are there any animals with three legs?
A: No, there are no animals with three legs.

Q: Why don't animals have three legs?
A: Animals don't have three legs because they would fall over.

Image generation

Image to text generation

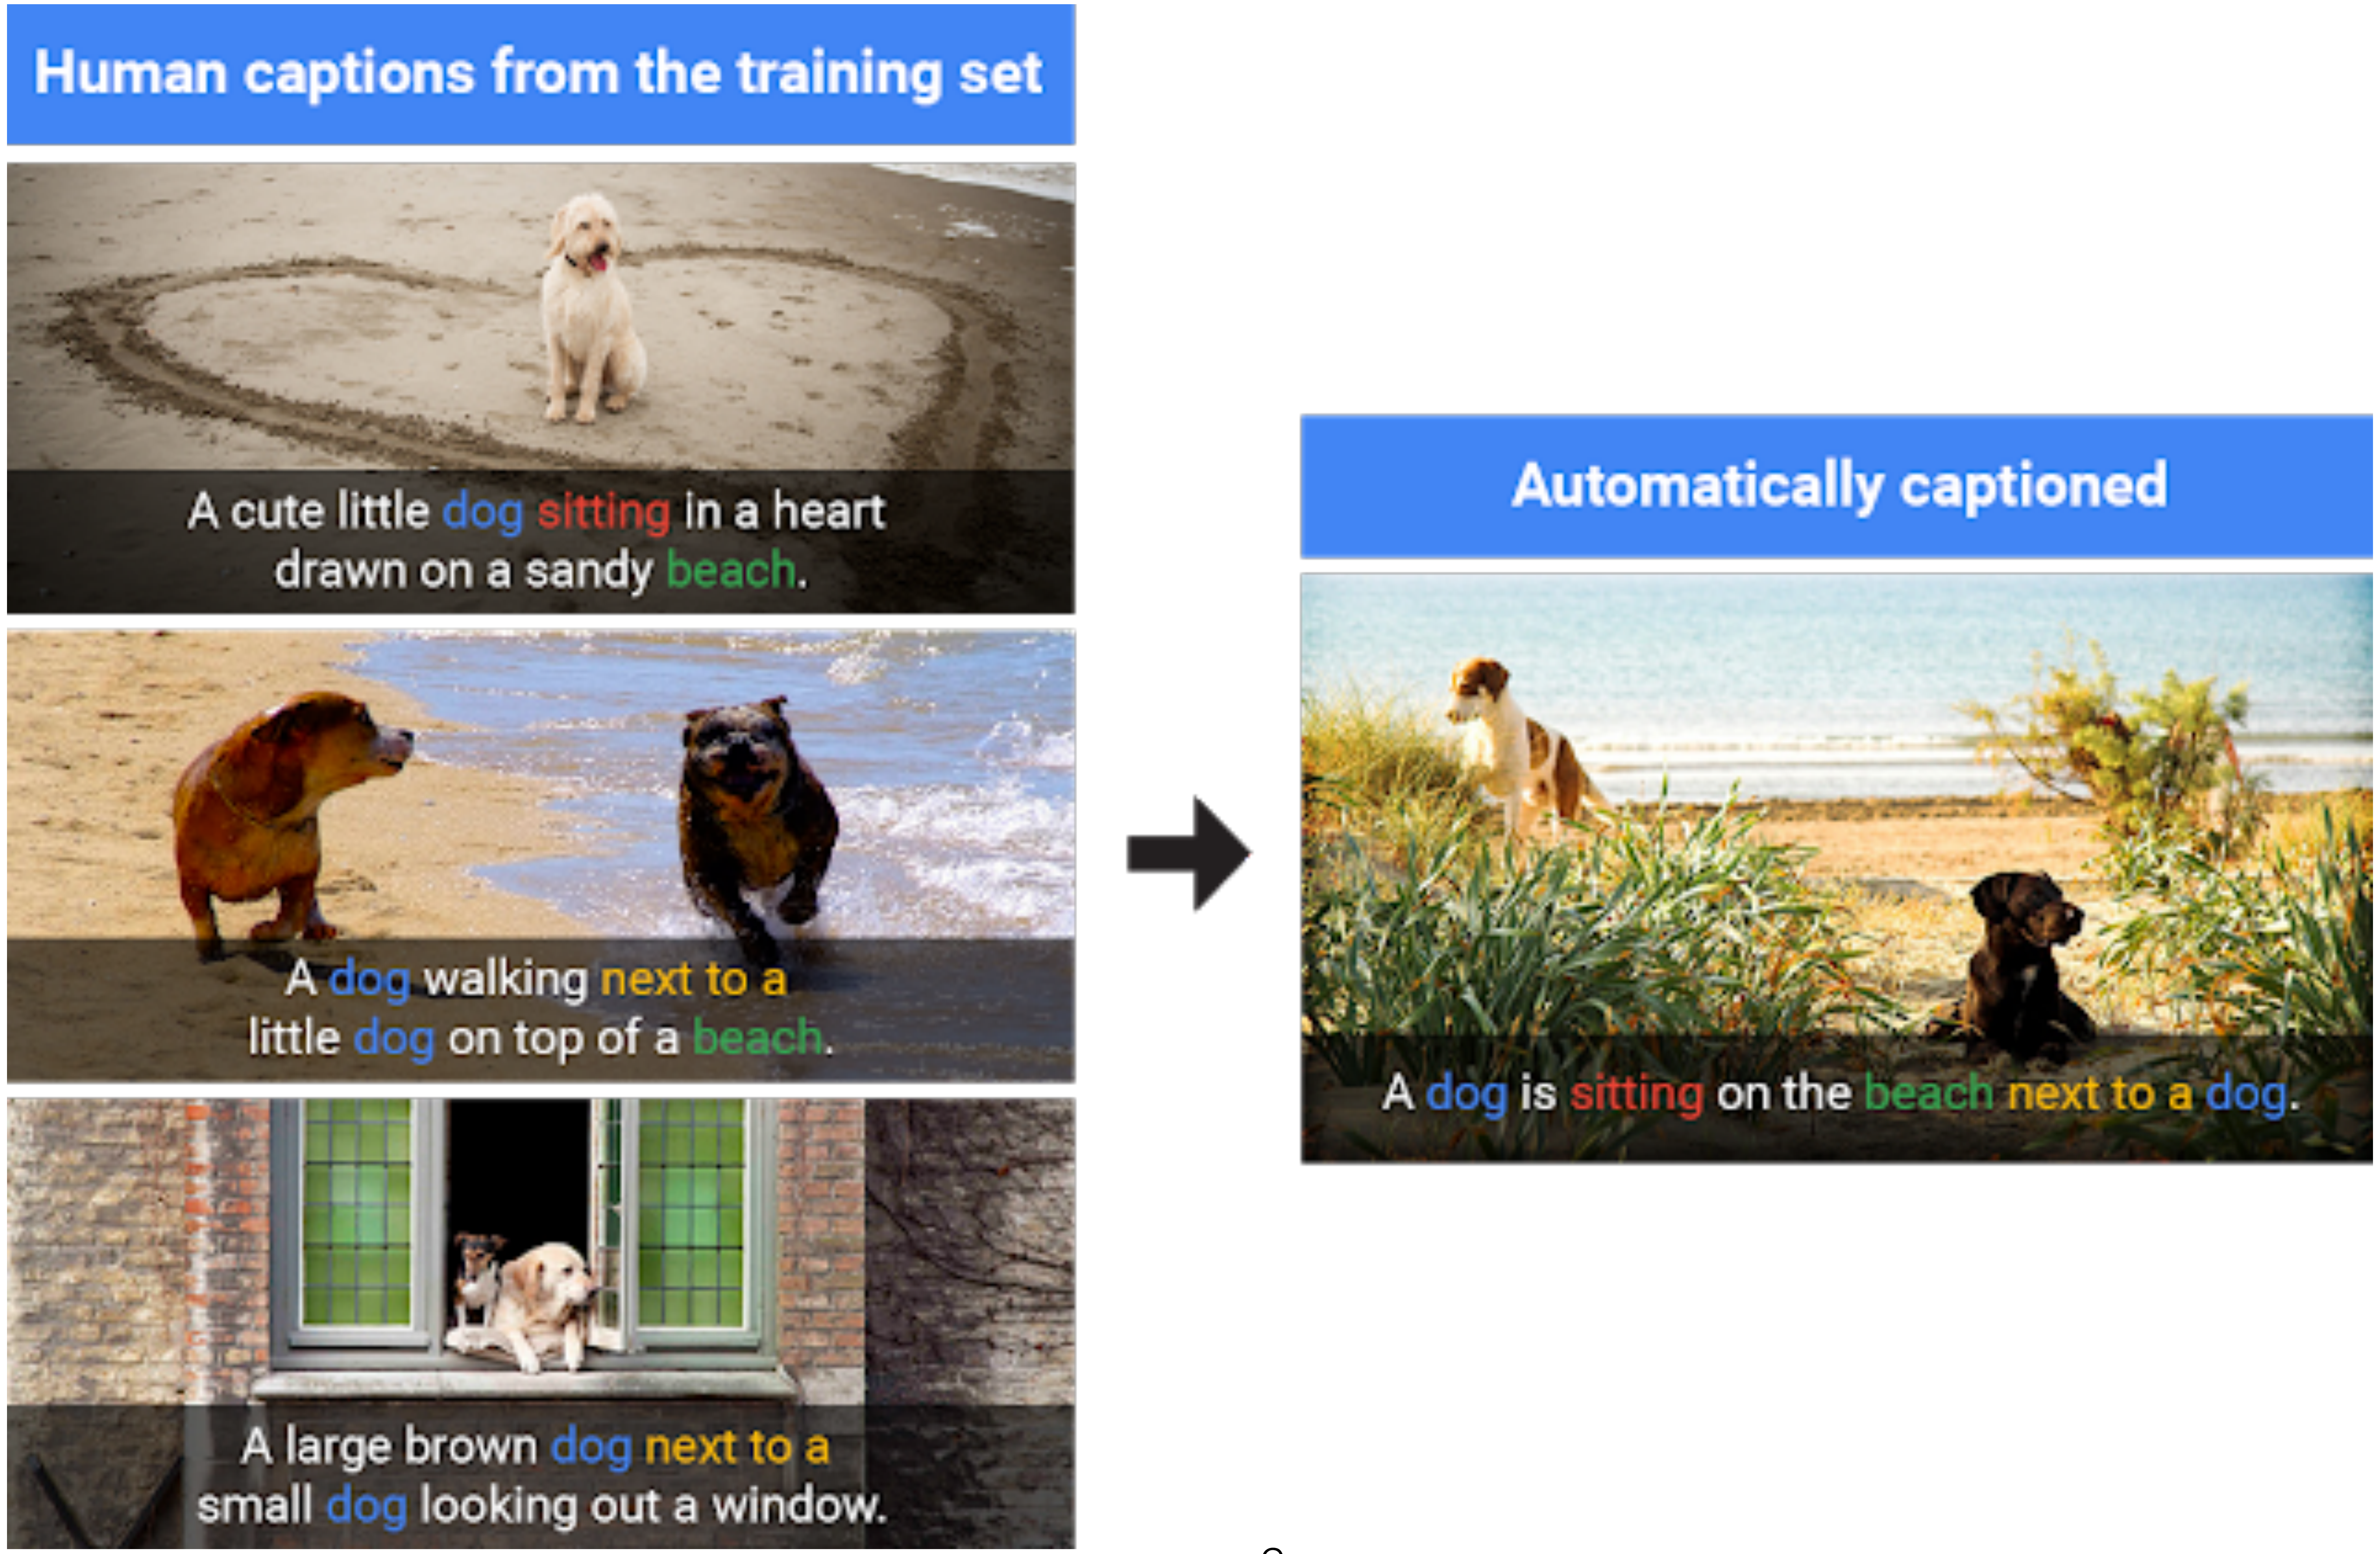

This class: *Theoretical Foundations* of ML

Why take this class? Why study ML theory?

1. Understand fundamental limitations about a learning problem.
    - When is it possible to learn?
    - What are the primary challenges we need to solve?
    - How much data do we need to learn?

This class: *theoretical foundations* of ML

Why take this class? Why study ML theory?

2. Develop fundamental intuitions for designing learning algorithm
   - What is the "correct" way to think about these challenges?
   - How do we trade-off between multiple challenges?

Will focus on simple (as opposed to *"realistic/practical"*) settings

3. It is fun!

Outline

1. Course logistics
2. Syllabus
3. Who should take this class? Prerequisites and expectations

Outline

1. **Course logistics**
2. Syllabus
3. Who should take this class? Prerequisites and expectations

Logistics: Lectures, OHs, Grader, Enrollment

- **Lectures**
  - MWF, 9.30-10.45am at ENGR HALL 2540
  - Will be mainly on the whiteboard. (Some lectures on I-Pad)
  - 27-30 lectures.

- **My office hours:** Wed 1.30 - 3pm at CS5375
- **Grader:** Albert Dorador-Chalar
- **Enrollment**
  - If you cannot enroll due to pre-requisites, please speak to me.

Logistics: Webpages

- **Course website**
  - https://pages.cs.wisc.edu/~kandasamy/courses/24fall-cs861
  - Logistics, syllabus, schedule, homeworks, and grading

- **Piazza**
  - https://piazza.com/wisc/fall2024/csece861 (**access code: wisc86124**)
  - Ask public questions whenever possible.
  - Announcements, peer discussions on lectures, homework clarifications.

- **Canvas**
  - Homeworks, exams, and some announcements

Grading

- **Scribing: 10%**
- **Homeworks: 40%**
- **Exam: 30%**
- **Course project (setting a homework question): 20%**

Scribing

- Each student will scribe ~2 lectures. Two students per lecture.
  - This may change if enrollment drops.
  - Sign up for scribing via the sign-up spreadsheet (see course website for link).

- **Instructions (see course website as well)**
  - Written in *full prose*, proof steps written in detail, intuitions explained well.
  - Prepare in Overleaf, and add me as a collaborator within 2 days
  - If you are unsure about taking the class, sign up for after Oct 4.
  - If you decide to drop, *delete your name **and** email me*.
  - ***Please use notes from last year***, but you will be graded on your submission.

Homework

- 5 Homeworks (including homework 0)
- Typically due every other Saturday at 11.59pm on Canvas.

- A total of 3 late days for the entire course. Extensions only for documented emergencies.

- 5 percent extra credit if you LaTeX your solutions.

- Homeworks will be *difficult*.
  - Expect to spend multiple hours/days on some problems.
  - Unless otherwise specified, you *are allowed* to collaborate with up to 2 classmates.

Take-home Exam

- Take-home exam from Mon 11/18 12.01 AM – Fri 11/22 11.59 PM.
- You have 48 hours to complete the exam.
- No collaboration allowed on the exam.

Course project (setting a homework problem)

- You will work in groups of size up to 3, to design a homework problem.
- Your peers will attempt and evaluate your question.
- See webpage for guidelines. I will point to examples from the homework questions, which you can use as a model for depth and level of difficulty.

Key dates

- A preliminary draft of the problem (with solutions) due on Sat 10/19.
- Final problems due on Sat 11/16.
- I will assign 1-2 questions to each of you. Solutions and evaluation of the questions due on 12/07.
- If you set a good question, I will include it in future courses and acknowledge your contribution!

Outline

1. Course logistics
2. **Syllabus**
3. Who should take this class? Prerequisites and expectations

Syllabus: Overview

1. PAC Learning
2. Statistical lower bounds
3. Nonparametric methods
4. Online learning & bandits
5. Online convex optimization

1. PAC Learning (4-6 lectures)

- Loss, risk, Empirical risk minimization
- PAC Learning: realizable vs agnostic
- Rademacher complexity & VC dimension
- Sauer’s lemma

2. Statistical lower bounds (7-10 lectures)

- Average-risk optimality vs minimax optimality
- Minimax optimal estimators for point estimation
- From estimation to testing: Le Cam & Fano methods
- Applications
    - regression, classification, density estimation

3. Nonparametric methods (2-3 lectures)

- Nonparametric regression
- Nonparametric density estimation

4. Online learning & bandits (7-10 lectures)

- Stochastic bandits and the UCB algorithm
- Learning from experts and the Hedge algorithm
- Adversarial bandits and the EXP-3 algorithm
- Contextual bandits and the EXP-4 algorithm
- Lower bounds for online learning and bandits

5. Online Convex Optimization (3-4 lectures)

- Follow the leader
- Follow the regularized/perturbed leader
- Online gradient descent

Outline

1. Course logistics
2. Syllabus
3. **Who should take this class? Prerequisites and expectations**

Target audience for the class

- Ph.D students doing research in theoretical **(statistical)** machine learning.

- **Background knowledge**
  - **Formal prerequisite:** CS761 or equivalent.
  - Strong background (intermediate-level graduate course) in calculus, statistics, and probability.

- **Who should not take this class.**
  - *"I want to learn about ML/AI"* (Take 540, 532)
  - *"I want to apply ML in an applied area of research"* (Take 760)
  - *"I want to learn take an introductory ML theory class"* (Take 761)

Homework 0

Two questions:

1. Normal mean estimation
2. A simple bandit model and algorithm

Three Objectives

I. A preview of some topics

II. Calibrate my teaching/expectation

III. Lets you assess if you are ready to take this class

General advice when taking this class

1. Focus on learning, and not on grades.
    - Class will be challenging. But if you are able to keep up, you will get a good grade.

2. Give me feedback about the course.

3. Be good citizens: attend class, ask questions, answer questions, let others answer/ask questions, respond to questions on piazza.