Machine Learning: Introduction

Hao Tang

January 14, 2026

https://homepages.inf.ed.ac.uk/htang2/mlg2025/

What is machine learning?

What is machine learning?

Name a few **examples** of machine learning systems

Large language models (LLMs)

Podcast generation

https://notebooklm.google.com/

Voice cloning

https://inworld-ai.github.io/tts/

Image super resolution

https://store.google.com/gb/product/pixel_10_pro

Face recognition

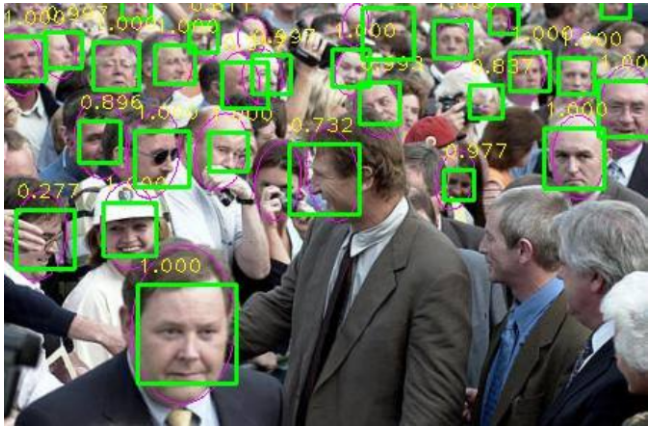

Image credit: (Triantafyllidou and Tefas, 2016)

Photo editing

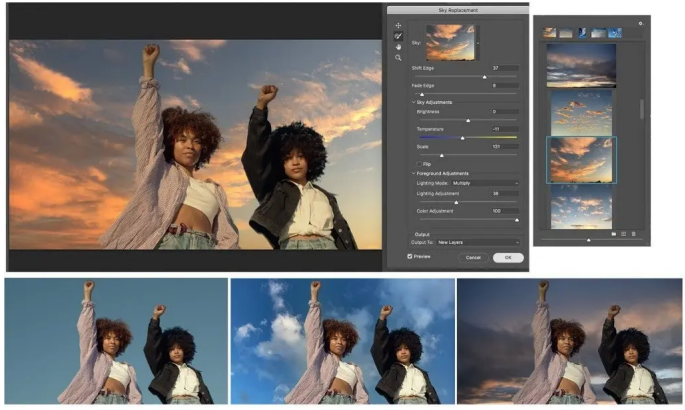

Image credit: https://blog.adobe.com/en/publish/2020/10/20/photoshop-the-worlds-most-advanced-ai-application-for-creatives

Photo editing

Image credit: https://blog.adobe.com/en/publish/2020/10/20/photoshop-the-worlds-most-advanced-ai-application-for-creatives

Video editing

Source | Driver | Ours | Driver | Ours

Image credit: https://neeek2303.github.io/MegaPortraits/

AlphaGo

Image credit: https://www.theguardian.com/technology/2016/mar/15/alphago-what-does-google-advanced-software-go-next (left)
https://twitter.com/demishassabis/status/692430224382902272 (right)

Recommender systems

NETFLIX

Spotify®

TikTok

YouTube

Speech recognition

Image credit: https://www.bbc.co.uk/news/technology-47893082

Speech verification

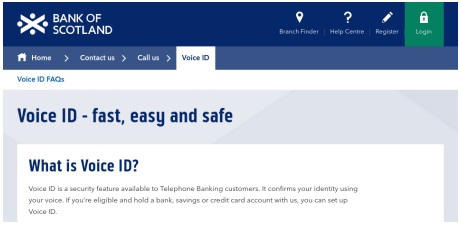

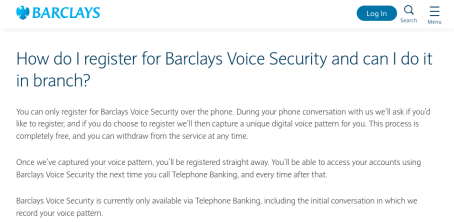

Robotic vacuum cleaner

Image credit: https://www.irobot.co.uk/en_GB/roomba.html

Autonomous driving

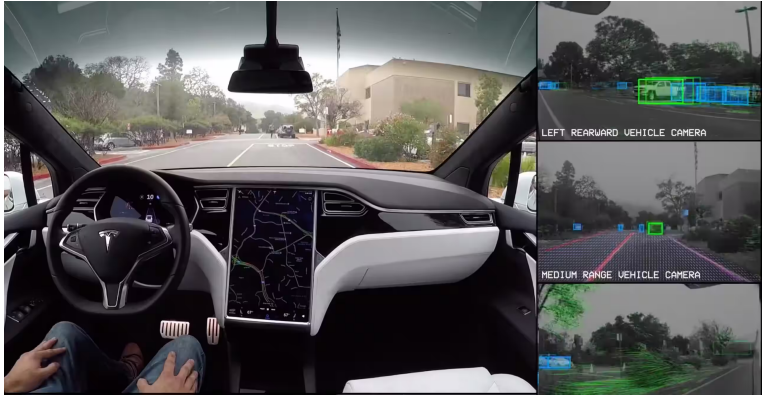

Image credit: https://www.tesla.com/autopilot

What is machine learning?

What is machine learning?

How would you write a program to recognize pasta?

What is machine learning?

How would you write a program to recognize hand-written 2s?

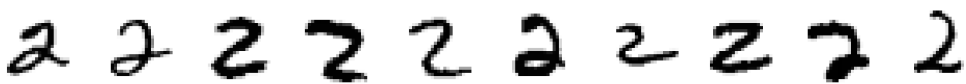

What is machine learning?

How would you write a program to recognize hand-written 2s?

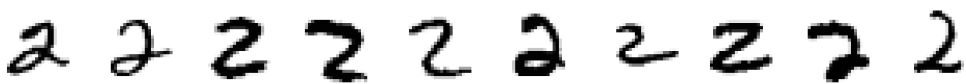

What is machine learning?

How would you write a program to recognize hand-written 2s?

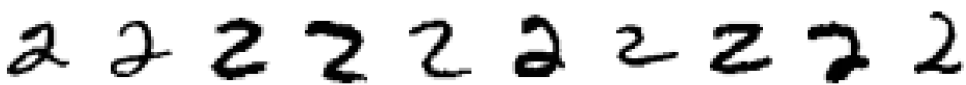

We don't write a program to enumerate all the possible ways of writing 2s.
We "produce" a program using the examples of 2s.

What is machine learning?

- There are programs that are hard to describe algorithmically and hence hard to write.
- However, in certain cases, the input-output pairs might be cheap to collect.
- We can alternatively describe a program in terms of data points.

Programming with Data

Programming with Data

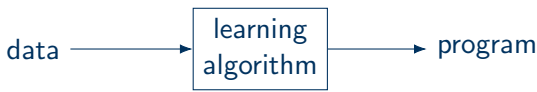

Programming with data

- This is a paradigm shift, but not all problems need to be solved with machine learning.
- Machine learning targets programs that are hard to describe algorithmically and hence hard to write.
- Machine learning assumes the input-output pairs are cheap to collect.

Pause and reflect

- We can infer a lot about big machine learning systems from small examples.
- In fact, finding a minimal example is usually the first step to understanding (and solving) a problem.
- This principle will be used a lot in this course.

Ingredients of machine learning systems

What **ingredients** do we need
to implement video recommendation?

To build a video recommendation system, we need ...

- videos
- people's preference
- ML engineers developing the learning algorithm
- GPUs and GPU cluster admin
- additional models that recognize speech in the videos
- data engineers
- UI/UX engineers, frontend and backend engineers
- databases and database admin
- video codec engineers
- website load balancing engineers
- users

Software cycle

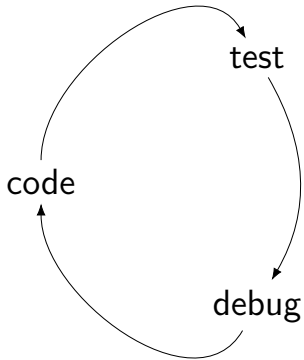

Software cycle

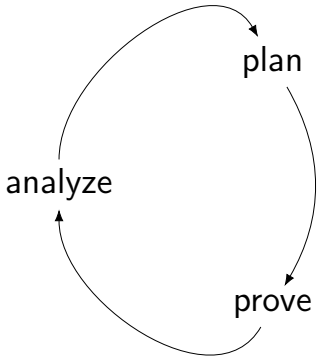

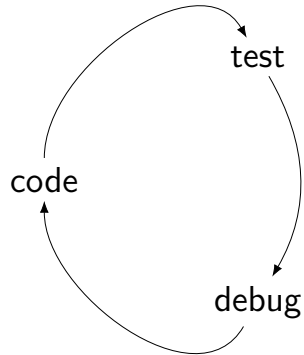

Machine learning software cycle

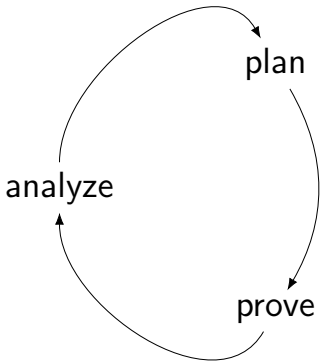

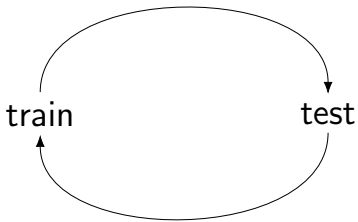

Machine learning software cycle

- Machine learning simplifies software programming.
- Writing software with machine learning now requires more analyses, planning, and verification.
- We will see some training and testing in this course, but the focus will be on analysis and verification.

Machine learning concepts

Name a few **concepts** in machine learning

Course calendar

https://homepages.inf.ed.ac.uk/htang2/mlg2025/calendar.html

Learning machine learning

- Stanford CS229 (2018)
  - https://cs229.stanford.edu/syllabus-autumn2018.html
  - https://www.youtube.com/playlist?list=PLoROMvodv4rMiGQp3WXShtMGgzqpfVfbU
- Hands-on books
- Sketchy sources
  - https://x.com/PinakiLaskar/status/1329748899347767296
  - https://x.com/PhDemetri/status/1723802334545813507

Hands-on books

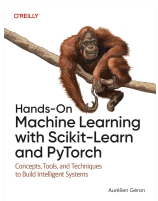

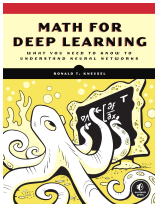

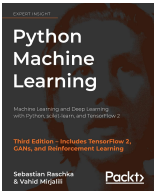

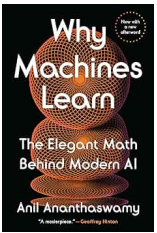

The goal of this course

Getting you prepared for year-4 courses

The goal of this course

- Getting you prepared for year-4 courses
- Getting you prepared for honors projects
- Getting you prepared for internships
- Getting you prepared for hackathons
- Getting you prepared for interviews

Our approach to learning machine learning

- Attend the lectures to get the high-level ideas
- Read the assigned reading after every lecture
- Summarize the assigned reading and lectures *without* looking at any notes or textbooks
- Work on the exercises (released on odd weeks)
- Join the discussion in tutorial sessions
  - Tuesday 1:10pm in AT 5.02
  - Tuesday 2:10pm in AT 5.02
  - Tuesday 3:10pm in AT 5.02

Textbooks

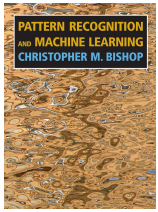
PATTERN RECOGNITION
AND MACHINE LEARNING
CHRISTOPHER M. BISHOP

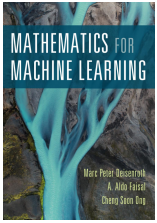
MATHEMATICS FOR
MACHINE LEARNING
Marc Peter Deisenroth
A. Aldo Faisal
Cheng Soon Ong

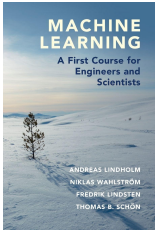
MACHINE
LEARNING
A First Course for
Engineers and
Scientists
ANDREAS LINDHOLM
NIKLAS WAHLSTRÖM
FREDRIK LINDSTEN
THOMAS B. SCHÖN

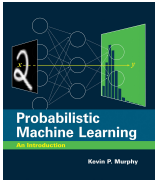
Probabilistic
Machine Learning
An Introduction
Kevin P. Murphy

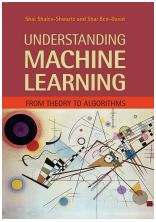
UNDERSTANDING
MACHINE
LEARNING
FROM THEORY TO ALGORITHMS

Assessment

- 20% from the coursework
- 80% from the final exam

MLG last year

- Average mark: 65.4
- Course feedback last year: MLG.pdf

Further reading

- Karpathy, Software 2.0,
  `https://karpathy.medium.com/software-2-0-a64152b37c35`, 2017
- Karpathy, Software is changing (again),
  `https://www.youtube.com/watch?v=LCEmiRjPEtQ`, 2025
- Togelius, Artificial general intelligence, MIT Press, 2024