Intro to AI

CSE 473 24wi
Rob Minneker

Content adapted from many others including:
Hanna Hajishirzi, Dan Weld, Luke Zettlemoyer
Dan Klein, Pieter Abbeel @ai.berkeley.edu

Brief Introduction

Please call me Rob!

I went to UW: Studied Bioengineering and Computer Science and did the BS/MS in CSE. Graduated 2021.

- i.e. I was in your shoes just a few short years ago!

Been in industry ever since graduation working as an ML Engineer. Previously Microsoft, Currently Apple.

I have TA'd or been the instructor for numerous BioE and CSE courses (even did a CSE Education MS Thesis!)

Think of me as an instructor/mentor that will guide you through the high-level concepts of AI

TAs

- Ali Al-Shimari
- Emmanuel Azuh Mensah
- Kevin Milad Farhat
- Benlin Liu
- Entong Su
- Stefan Alex Todoran
- Weitian Xia

Office hours will be posted on the course website home page

Think of TAs as your friends that can help you understand concepts and debug your HW

Websites and tools

Course Website: cs.uw.edu/473

Q&A: Ed

HW: Gradescope

Grades: Canvas

Course Format

Programming Assignments (50% of Grade)

- 4 assignments
- Python
- Autograded
- More hands-on experience with algorithms

Written Assignments (50% of Grade)

- 8 assignments
- More conceptual understanding of the content

Grading Breakdown

Your course grade, X, will be determined as follows

90% → X >= 3.5

80% → X >= 2.5

70% → X >= 1.5

60% → X >= 0.5

*Subject to change based on class performance

What do you need to be successful? Prerequisites

CSE 332 and 312

- DSA, Probability

Programming

- All projects are in Python, we have an optional 0th project to help you get familiar if not already

Excitement to learn about a rapidly evolving field!

Optional Reading

It really is optional!

Check course website “reading” tab to follow along if you’d like to dive deeper into a particular topic

Policies

Assignments should be completed individually

You may collaborate with up to 3 other students, please list all collaborators on each assignment

Late policy: 6 penalty-free late days for the whole quarter; maximum of 4 days per assignment

Agenda

What is Artificial Intelligence?

What can AI do?

What is this course about?

AI is a *broad* term which encapsulates ML, DL, NLP, RL, GenAI, etc.

We are going to start broad in this course and dive deeper into each of these subgroups

GenAI is where all the hype is right now
These are things like ChatGPT, LLMs, Diffusion Models, etc.

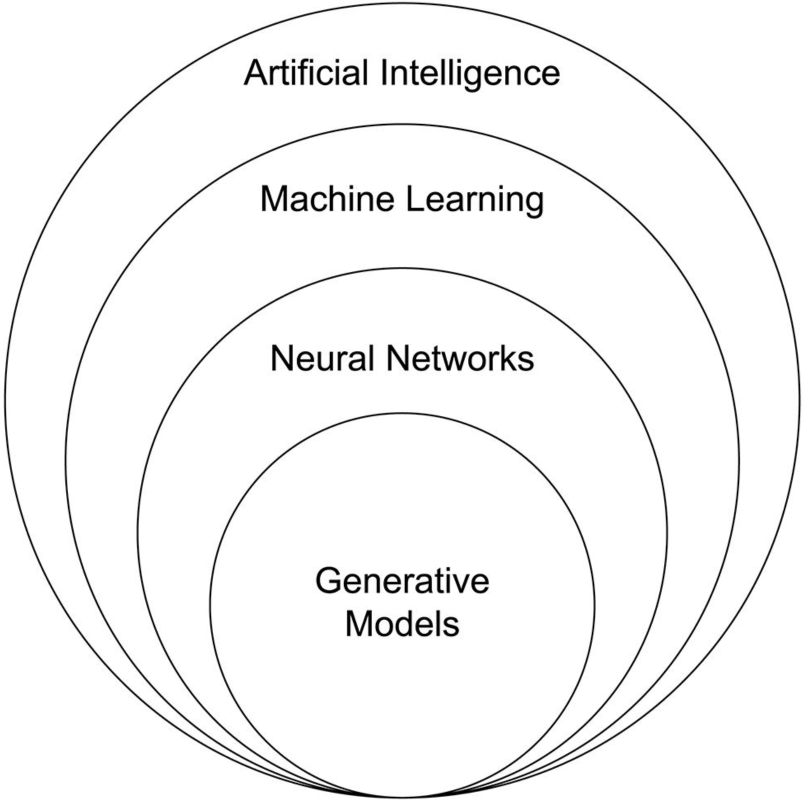

AI is all around us

ChatGPT

Visual Look Up
is available.

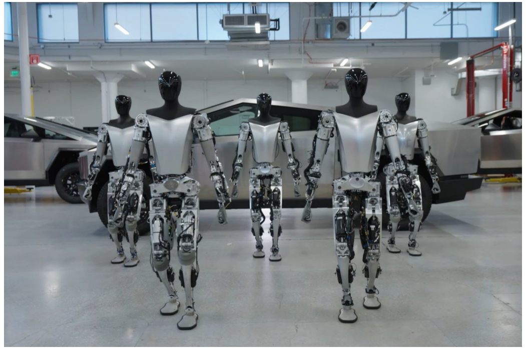

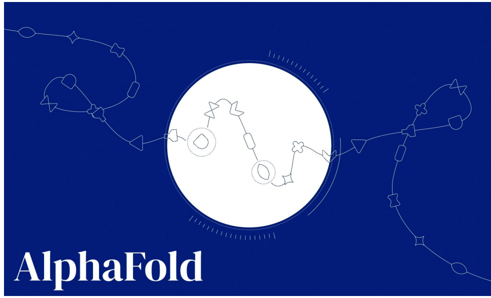

How ChatGPT defines AI

You

Explain AI like you're talking to a kindergartener

ChatGPT

Alright, imagine AI is like a really smart robot in a computer. This robot doesn't have arms or legs, but it has a super big brain. It can read lots of books really fast, remember everything it sees, and answer lots of questions. It's like having a magical friend in the computer who can help you learn new things and solve tricky puzzles! 🤖✨

How ChatGPT defines AI

You
Explain AI to a computer science major

ChatGPT
Artificial Intelligence (AI) is a field of computer science focused on creating systems capable of performing tasks that normally require human intelligence. These tasks include learning, decision-making, problem-solving, language understanding, and visual perception.

AI systems are built on algorithms and data. Machine Learning (ML), a subset of AI, involves training models on large datasets to recognize patterns and make predictions. This training can involve supervised learning (with labeled data), unsupervised learning (with unlabeled data), or reinforcement learning (where the system learns by trial and error).

Deep Learning, a further subset of ML, uses neural networks with many layers (hence 'deep') to process complex data like images and speech. These networks are inspired by the structure and function of the human brain.

AI applications are diverse, ranging from simple tasks like filtering spam emails to complex ones like autonomous driving. The field is evolving rapidly, with ongoing research in areas like natural language processing, computer vision, robotics, and AI ethics. The challenge is not just technical; it involves considering the societal impacts of AI as well.

What can AI do?

Talk to your neighbors for the next couple minutes

Please involve folks who don’t have a group!

Be prepared to share out after

What can AI do?

Play games

Generate text, images, audio, video

Understand multimodal data and act on it (LMM)

- LMM article by Chip Huyen

Operate machines: robots, cars, other agents acting with their environment

Everyday it seems a new capability emerges!

What can AI not do?

Every Model Learned by Gradient Descent Is Approximately a Kernel Machine

Pedro Domingos PEDROD@CS.WASHINGTON.EDU
Paul G. Allen School of Computer Science & Engineering
University of Washington
Seattle, WA 98195-2350, USA

Abstract

Deep learning's successes are often attributed to its ability to automatically discover new representations of the data, rather than relying on handcrafted features like other learning methods. We show, however, that deep networks learned by the standard gradient descent algorithm are in fact mathematically approximately equivalent to kernel machines, a learning method that simply memorizes the data and uses it directly for prediction via a similarity function (the kernel). This greatly enhances the interpretability of deep network weights, by elucidating that they are effectively a superposition of the training examples. The network architecture incorporates knowledge of the target function into the kernel. This improved understanding should lead to better learning algorithms.

Keywords: gradient descent, kernel machines, deep learning, representation learning, neural tangent kernel

Understanding Hyperdimensional Computing for Parallel Single-Pass Learning

Tao Yu*
Cornell University
tyu@cs.cornell.edu

Yichi Zhang*
Cornell University
yz2499@cornell.edu

Zhiru Zhang
Cornell University
zhiruz@cornell.edu

Christopher De Sa
Cornell University
cdesa@cs.cornell.edu

Abstract

Hyperdimensional computing (HDC) is an emerging learning paradigm that computes with high dimensional binary vectors. There is an active line of research on HDC in the community of emerging hardware because of its energy efficiency and ultra-low latency—but HDC suffers from low model accuracy, with little theoretical understanding of what limits its performance. We propose a new theoretical analysis of the limits of HDC via a consideration of what similarity matrices can be "expressed" by binary vectors, and we show how the limits of HDC can be approached using random Fourier features (RFF). We extend our analysis to the more general class of vector symbolic architectures (VSA), which compute with high-dimensional vectors (hypervectors) that are not necessarily binary. We propose a new class of VSAs, finite group VSAs, which surpass the limits of HDC. Using representation theory, we characterize which similarity matrices can be "expressed" by finite group VSA hypervectors, and we show how these VSAs can be constructed. Experimental results show that our RFF method and group VSA can both outperform the state-of-the-art HDC model by up to 7.6% while maintaining hardware efficiency. This work aims to inspire a future interest on HDC in the ML community and connect to the hardware community.

What can AI not do?

Protein Folding in the Hydrophobic-Hydrophilic (*HP*) Model is NP-Complete

BONNIE BERGER and TOM LEIGHTON

Published Online: 13 Mar 2009 | https://doi.org/10.1089/cmb.1998.5.27

PDF/EPUB Permissions & Citations Share

Abstract

One of the simplest and most popular biophysical models of protein folding is the hydrophobic-hydrophilic (*HP*) model. The *HP* model abstracts the hydrophobic interaction in protein folding by labeling the amino acids as hydrophobic (*H* for nonpolar) or hydrophilic (*P* for polar). Chains of amino acids are configured as self-avoiding walks on the 3D cubic lattice, where an optimal conformation maximizes the number of adjacencies between *H*'s. In this paper, the protein folding problem under the *HP* model on the cubic lattice is shown to be NP-complete. This means that the protein folding problem belongs to a large set of problems that are believed to be computationally intractable.

What can AI not do? Or can it?

SCIENCEINSIDER | TECHNOLOGY

Google's DeepMind aces protein folding

Artificial intelligence firm takes crown in biannual contest

6 DEC 2018 · BY ROBERT F. SERVICE

OCTOBER 31, 2022 | 9 MIN READ

One of the Biggest Problems in Biology Has Finally Been Solved

Google DeepMind CEO Demis Hassabis explains how its AlphaFold AI program predicted the 3-D structure of every known protein

BY TANYA LEWIS

How should AI do what it does?

UW NEWS

ENGINEERING | NEWS RELEASES | TECHNOLOGY

November 29, 2023

AI image generator Stable Diffusion perpetuates racial and gendered stereotypes, study finds

Stefan Milne

UW News

AI should be responsible, fair, and equitable

Responsibility: https://ai.google/responsibility/responsible-ai-practices/

Fairness: https://courses.cs.washington.edu/courses/cse599j/22wi/

Equity: https://www.gatesfoundation.org/ideas/science-innovation-technology/artificial-intelligence

These fields fall under the umbrella of Ethical AI

What are you curious to learn about in AI?

https://pollev.com/minneker

TODO

- Check out the course website!
- Post any questions you have on Ed
- Review search algorithms:
  - DFS, BFS, Dijkstra's, UCS, A*, etc.
- Work on Project 0 (or another pet project) to get familiar with Python
  - Reminder that Project 0 is NOT graded