



# CS 540 Introduction to Artificial Intelligence

## **Course Overview and Probability**

University of Wisconsin–Madison

Fall 2025, Section 1 & 2



# Today's outline

- What's in CS540
- Course logistics
- Software
- Probability



# Part I: Course overview

# What is artificial intelligence?

# Language models

ChatGPT 5

Give me a very short introduction to AI, to begin a lecture to college students.

Here's a concise way you could open your lecture:

"Artificial Intelligence, or AI, is the science of creating machines that can perform tasks requiring human-like intelligence—things like recognizing patterns, learning from data, making decisions, and even generating new ideas. At its core, AI is about building systems that can adapt and improve, not just follow fixed instructions. It powers technologies we use every day—from recommendation systems on streaming platforms to voice assistants and self-driving cars—and continues to shape how we live, work, and think about the future."

Would you like me to make it more **engaging and conversational** (to capture attention at the very start), or more **formal and academic** (to set a scholarly tone)?

+

Ask anything

ChatGPT can make mistakes. Check important info.

DeepMind and OpenAI achieve IMO Gold. What does it all mean?

What we know, what we would like to know, and what it may take years to know

ERNEST DAVIS AND GARY MARCUS


JUL 22, 2025

175

65

24

Share



Drawing (and typo) by ChatGPT

ADVENTURES IN AUGMENTATION

Google CEO says over 25% of new Google code is generated by AI

We've always used tools to build new tools, and developers are using AI to continue that tradition.

BENJ EDWARDS

– OCT 30, 2024 11:50 AM | 145

Conversation on 8/7/25

<https://garymarcus.substack.com/p/deepmind-and-openai-achieve-imo-gold>

<https://arstechnica.com/ai/2024/10/google-ceo-says-over-25-of-new-google-code-is-generated-by-ai/>

# Autonomous driving



<https://www.sfchronicle.com/sf/article/cruise-waymo-driverless-cars-in-s-f-18282902.php>

# Image and Video Generation



# Bioengineering

## The Nobel Prize in Chemistry 2024

The Royal Swedish Academy of Sciences has decided to award the Nobel Prize in Chemistry 2024 with one half to

**David Baker**

University of Washington, Seattle, WA, USA  
Howard Hughes Medical Institute, USA.

*“for computational protein design”*

**Demis Hassabis**

Google DeepMind, London, UK

*“for protein structure prediction”*

**John Jumper**

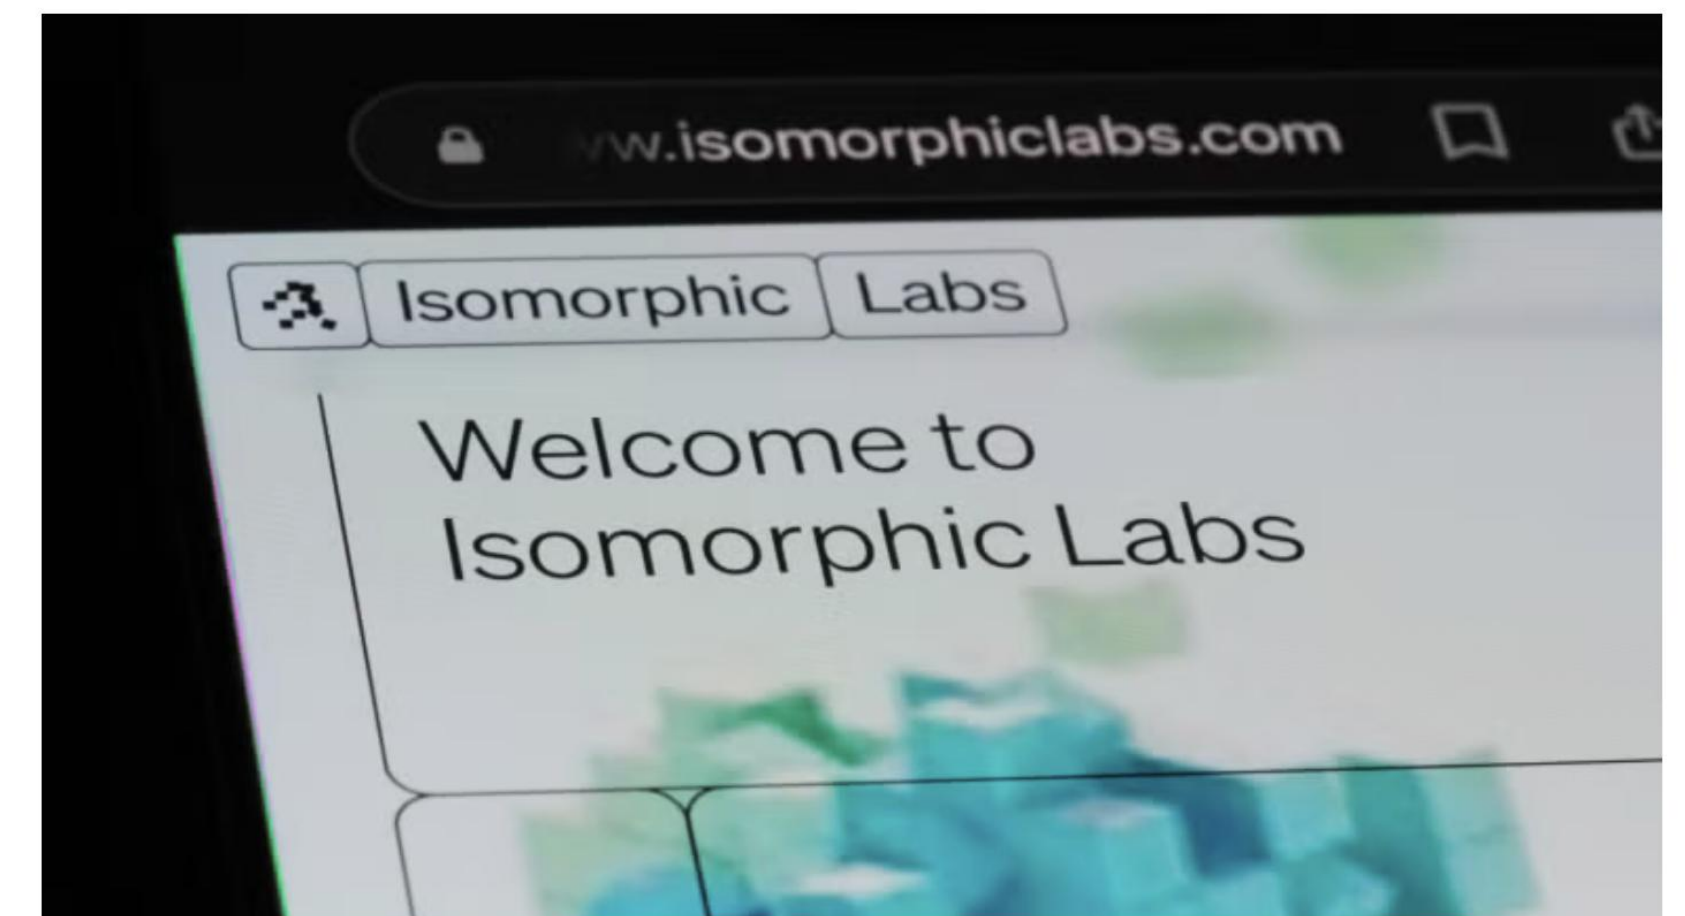
Google DeepMind, London, UK

## Isomorphic Labs prepares to launch trials for AI-designed drugs

The Google DeepMind spinout raised \$600m in its first financing round in March 2025.

Abigail Beaney | July 7, 2025

Share



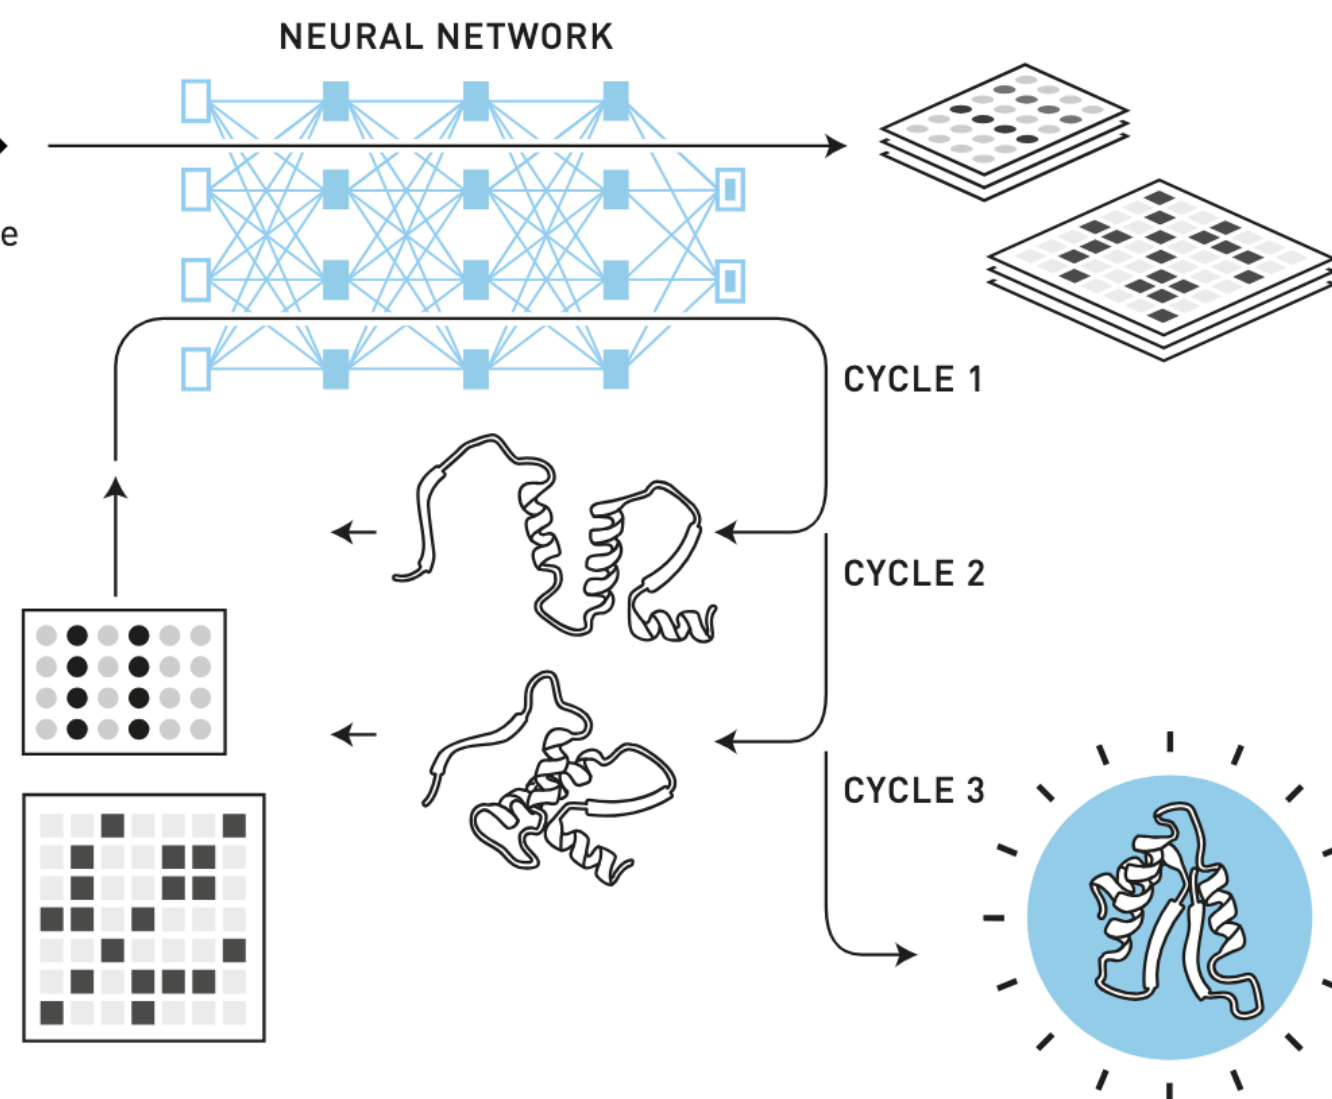
Isomorphic Lab's AlphaFold 3 system, developed with Google DeepMind, won its developers a Nobel Prize. Image credit: Shutterstock / Robert Way

### 3. AI ANALYSIS

Using an iterative process, AlphaFold2 refines the sequence analysis and distance map. The AI model uses neural networks called transformers, which have a great capacity to identify important elements to focus on. Data about other protein structures – if they were found in step 1 – is also utilised.

### 4. HYPOTHETICAL STRUCTURE

AlphaFold2 puts together a puzzle of all the amino acids and tests pathways to produce a hypothetical protein structure. This is re-run through step 3. After three cycles, AlphaFold2 arrives at a particular structure. The AI model calculates the probability that different parts of this structure correspond to reality.



<https://www.nobelprize.org/prizes/chemistry/2024/press-release/>

<https://www.clinicaltrialsarena.com/news/isomorphic-labs-prepares-trials-ai-designed-drugs/>

# Playing Games

## AlphaGo seals 4-1 victory over Go grandmaster Lee Sedol

DeepMind's artificial intelligence astonishes fans to defeat human opponent and offers evidence computer software has mastered a major challenge



📷 The world's top Go player, Lee Sedol, lost the final game of the Google DeepMind challenge match. Photograph: Yonhap/Reuters

## *Hold 'Em or Fold 'Em? This A.I. Bluffs With the Best*

Pluribus, a poker-playing algorithm, can beat the world's top human players, proving that machines, too, can master our mind games.

FUTURE PERFECT

## AI triumphs against the world's top pro team in strategy game Dota 2

It's the first time an AI has beat a world champion e-sports team.

by Kelsey Piper

Apr 13, 2019, 10:30 PM UTC



<https://www.vox.com/2019/4/13/18309418/open-ai-dota-triumph-og>

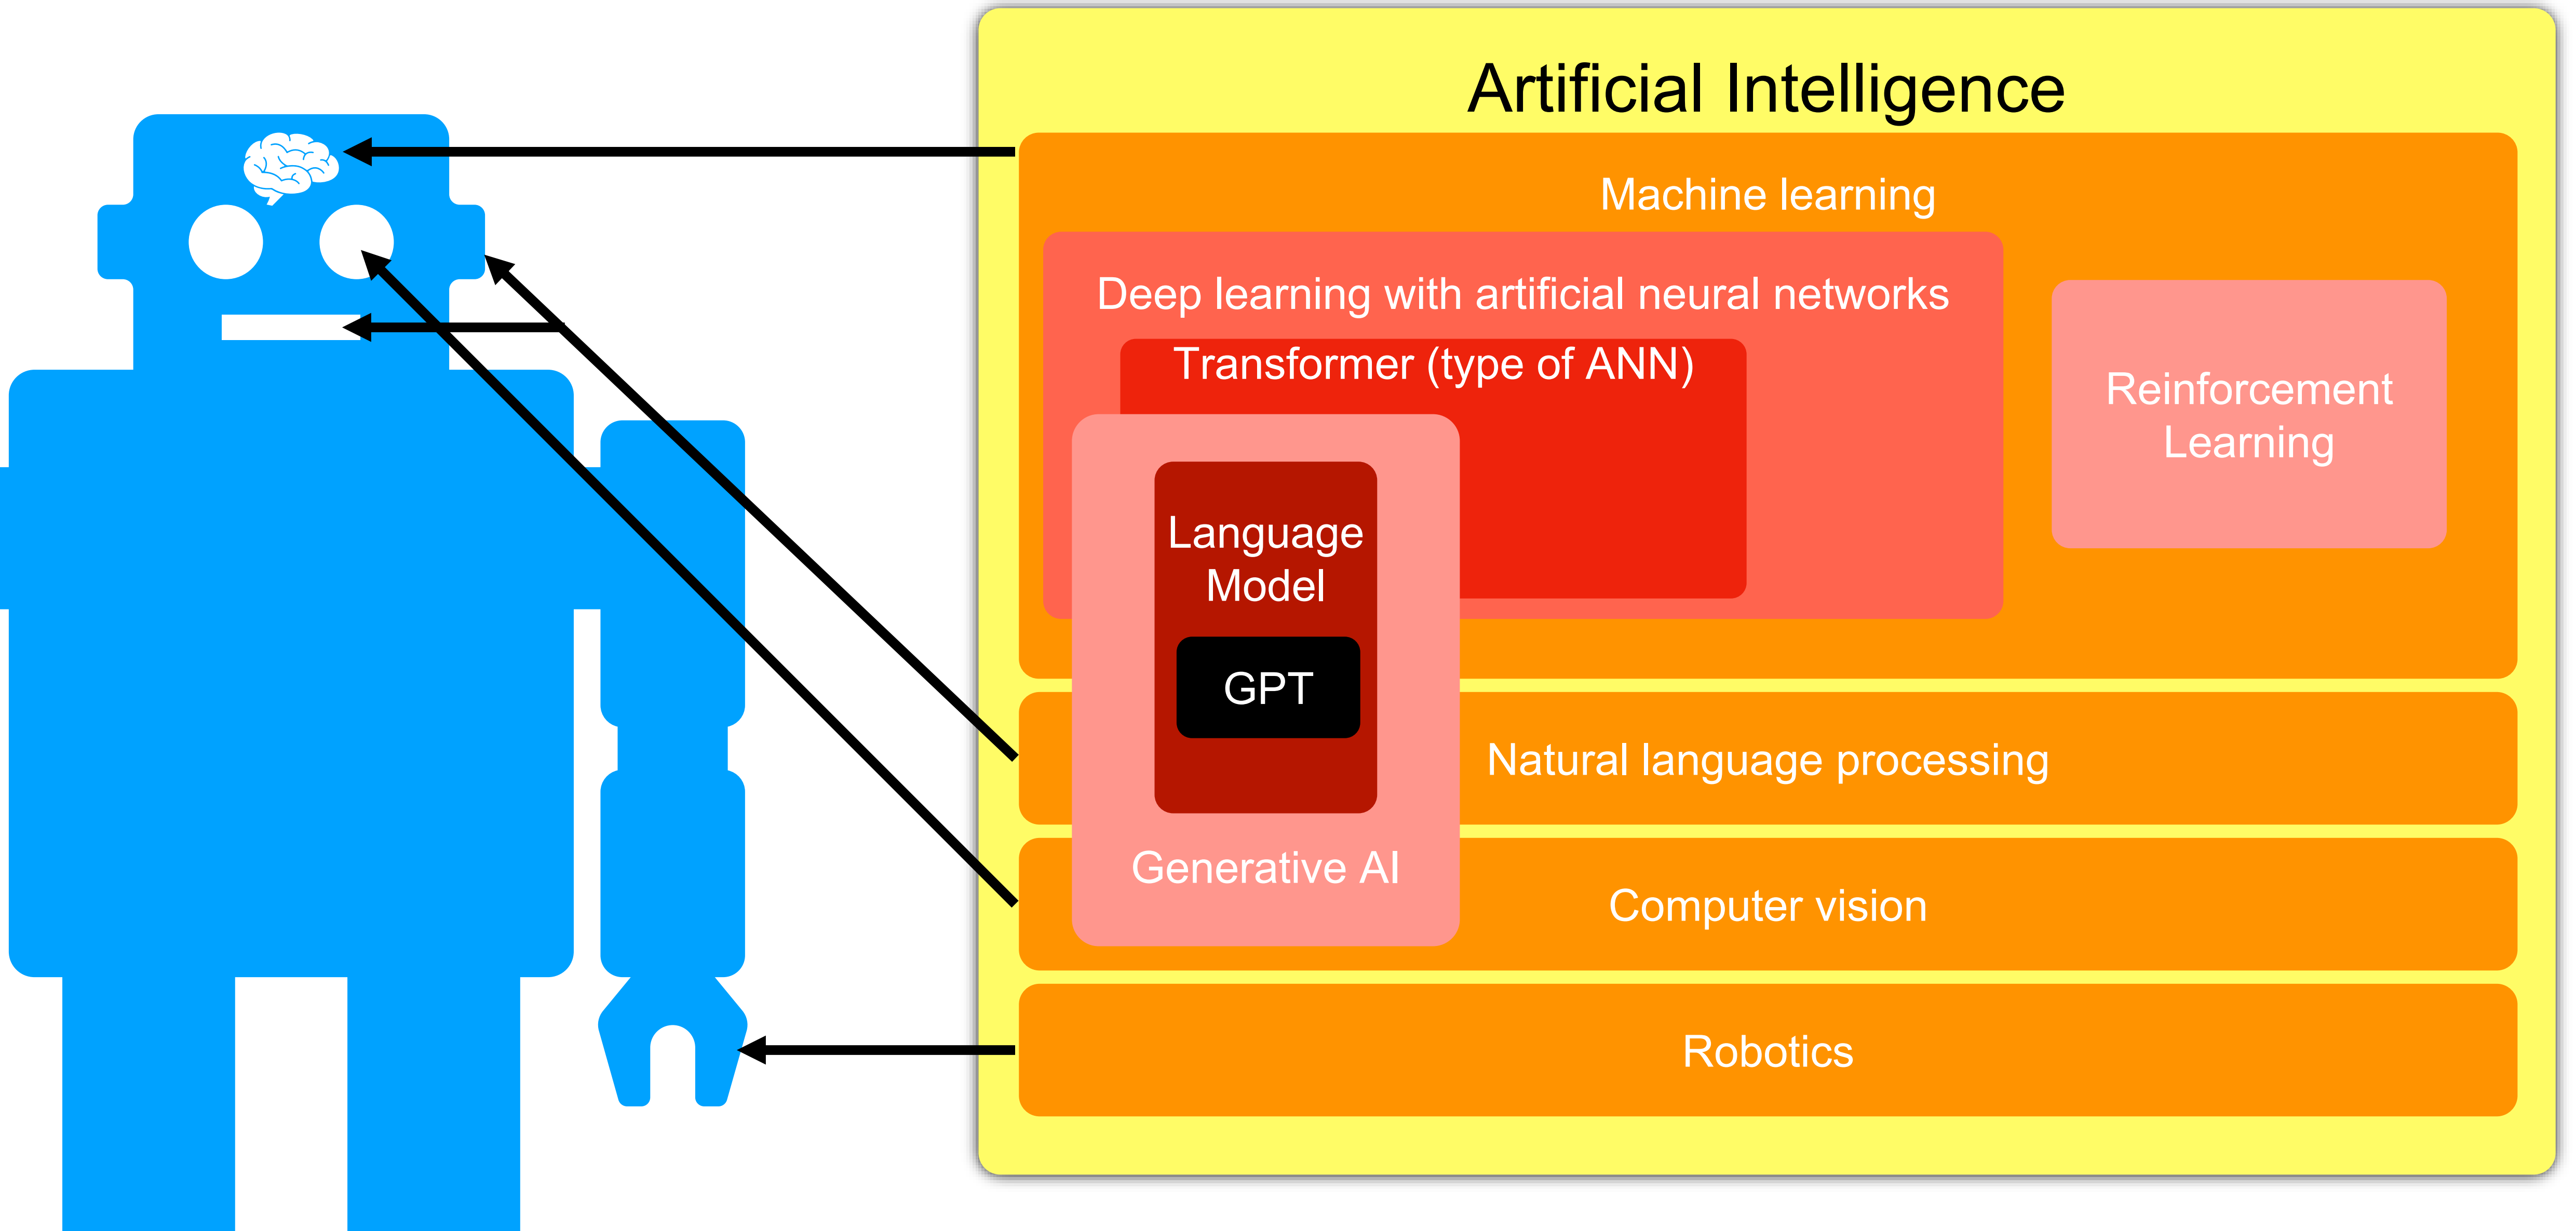
<https://www.nytimes.com/2019/07/11/science/poker-robot-ai-artificial-intelligence.html>

<https://www.theguardian.com/technology/2016/mar/15/googles-alphago-seals-4-1-victory-over-grandmaster-lee-sedol>

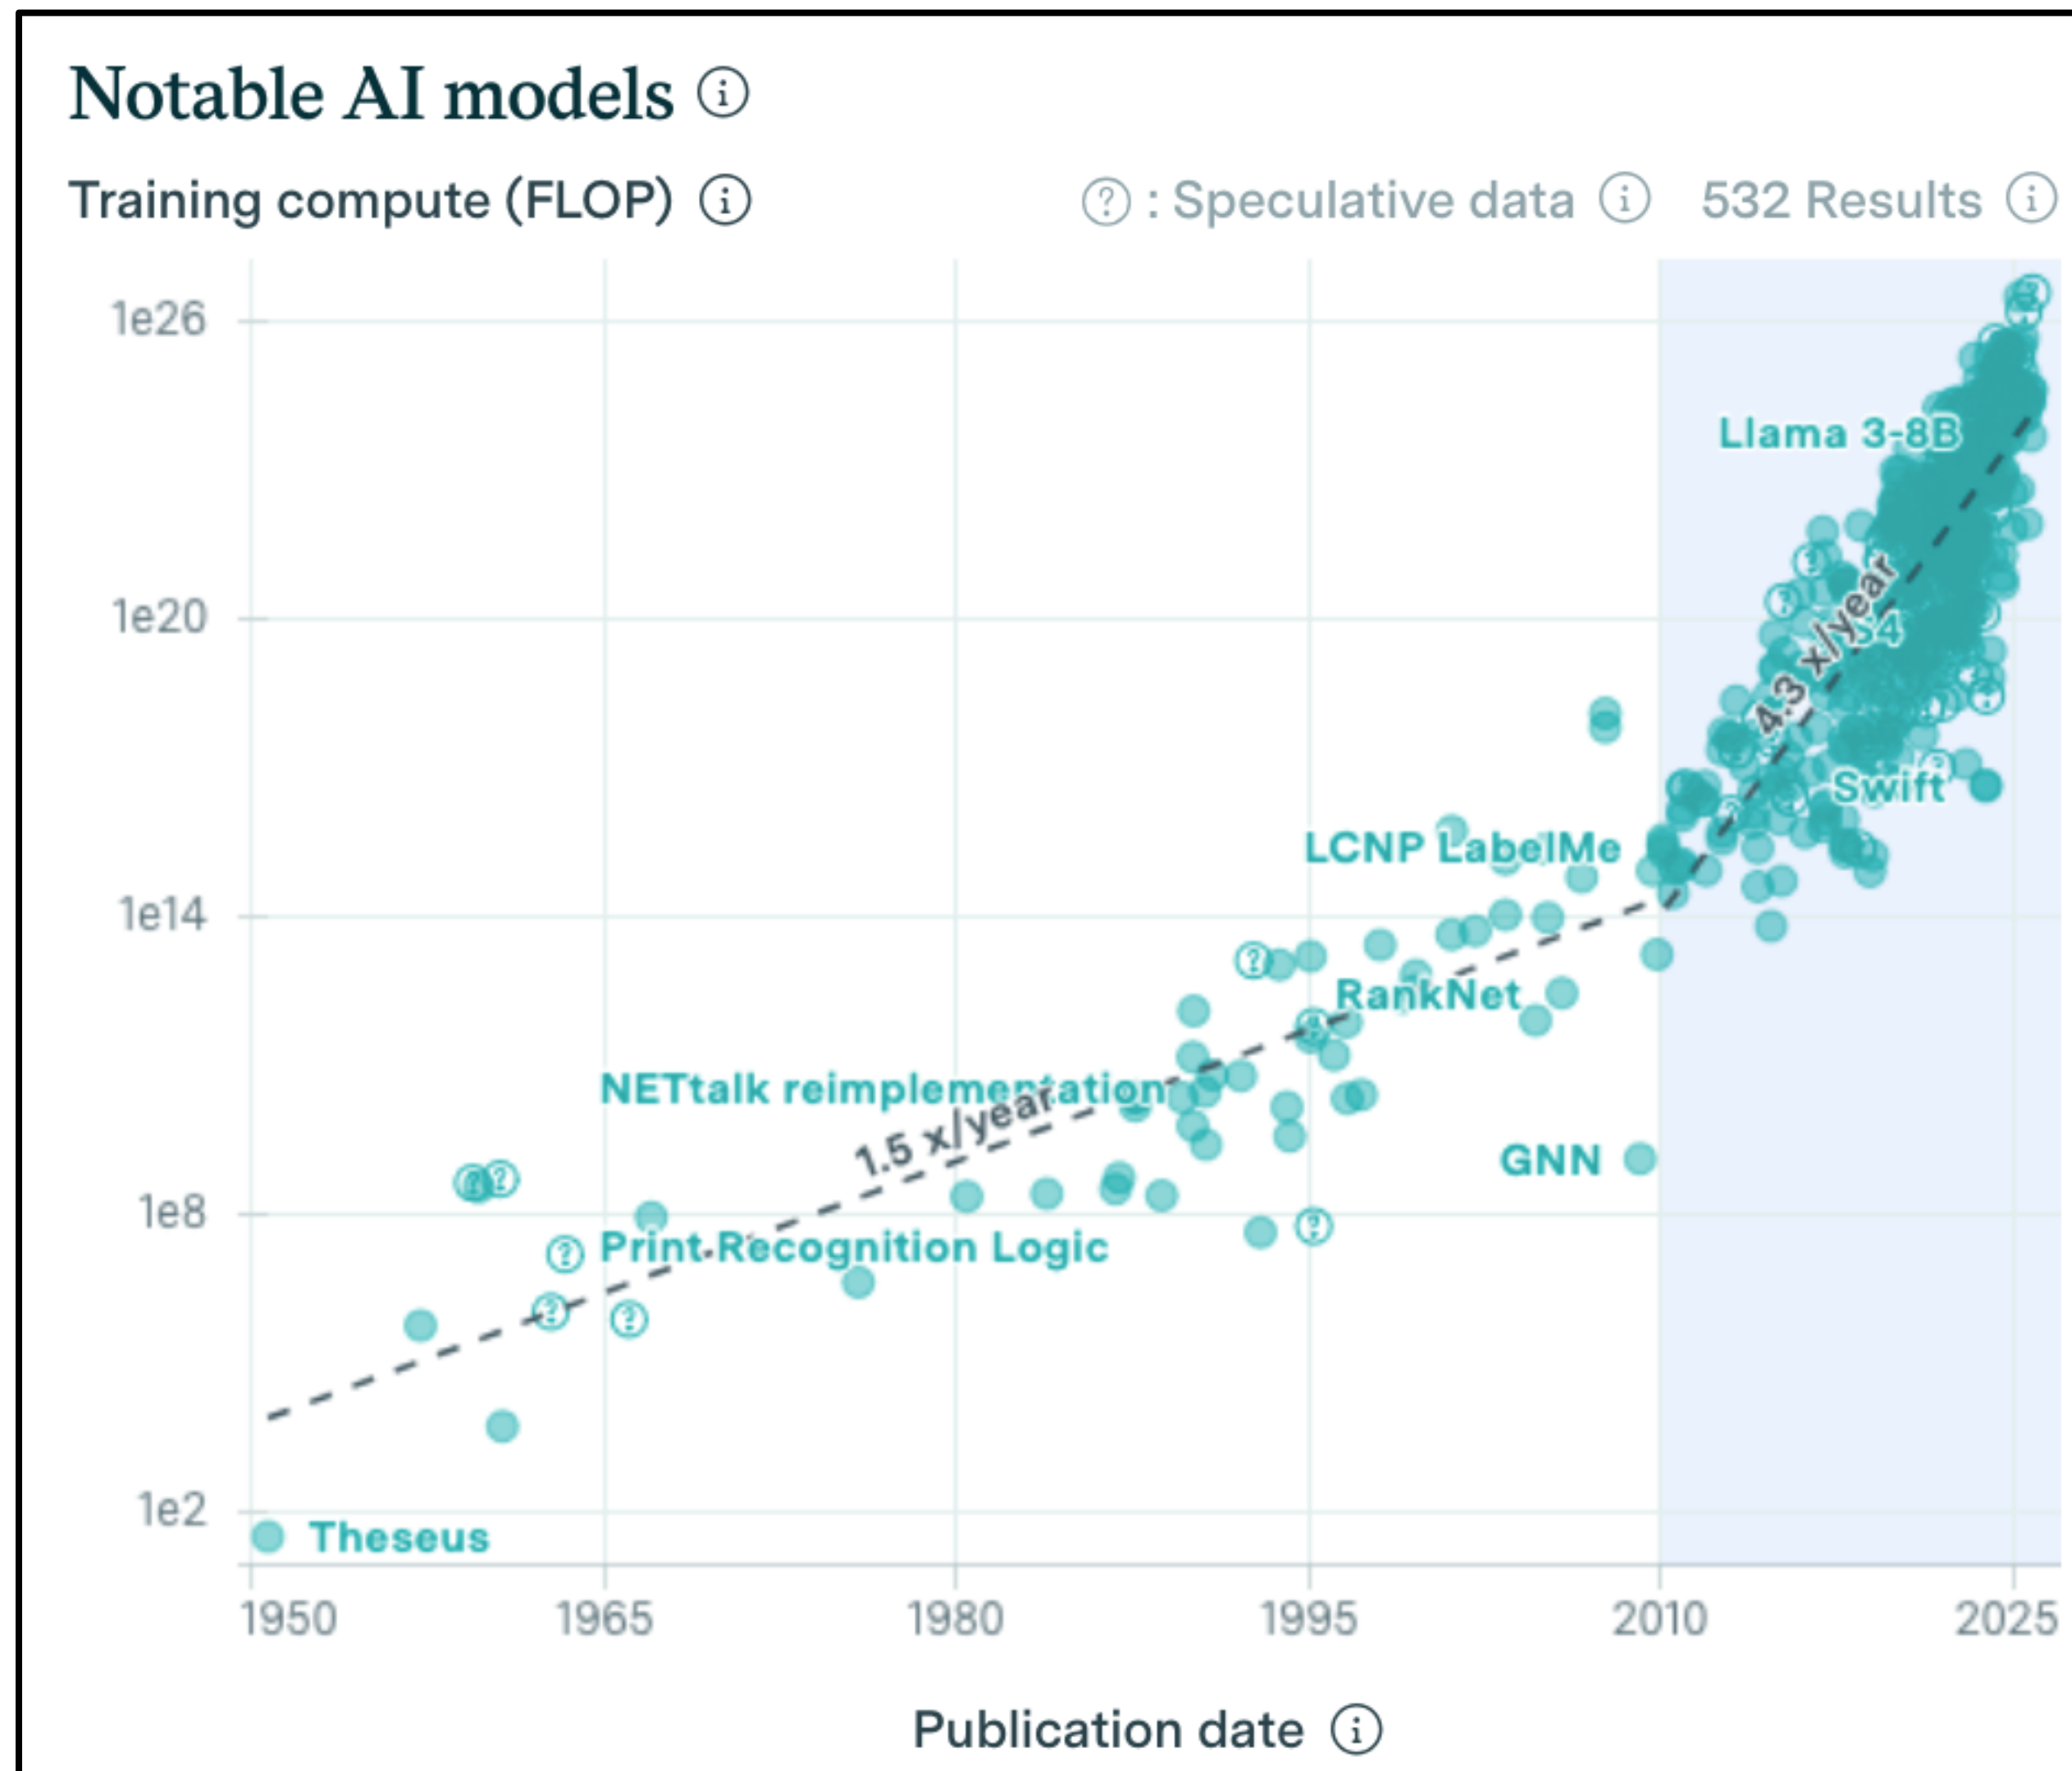
GPT, large language model, generative AI, deep learning...

Oh my!

# Artificial Intelligence Is a Broad Topic



# Artificial Intelligence Is Rapidly Changing



Standard textbook on AI is Russell & Norvig's  
***Artificial Intelligence: A Modern Approach***

- First edition in 1995
- Fourth edition in 2020
- Third edition from 2010 devotes just 10 pages to neural networks!

# Artificial Intelligence Is Not Magic

It relies on **fundamental** techniques in:

- Algorithms
- Mathematics
- Logic
- Probability and Statistics
- Optimization

# UW–Madison Conducts World-Class AI Research

- Computer Sciences
- ECE
- Statistics
- ISyE
- Math
- iSchool
- Almost all other departments on AI applications
- Data Science Institute, Data Science Hub

# What Will You Learn in CS540?

- Foundational mathematical tools
  - Linear algebra, probability, statistics
- Core techniques in natural language processing
  - Language modeling, n-gram models, transformers
- Basics of machine learning
  - Classification, regression, clustering
- Neural networks and deep learning
  - Multilayer perceptron, convolutions, ResNets
- Search and reinforcement learning
  - Game theory, Q-learning, A\*
- Ethics in real-world applications

# What Will You Learn in CS540?

- Foundational mathematical tools
  - Linear algebra, probability, statistics
- Core techniques in natural language processing
  - Language modeling, n-gram models, transformers
- Basics of machine learning

**TL;DR Lots of useful stuff, theory and practice in AI**

- Neural networks and deep learning
  - Multilayer perceptron, convolutions, ResNets
- Search and reinforcement learning
  - Game theory, Q-learning, A\*
- Ethics in real-world applications



## Part II: Course Logistics

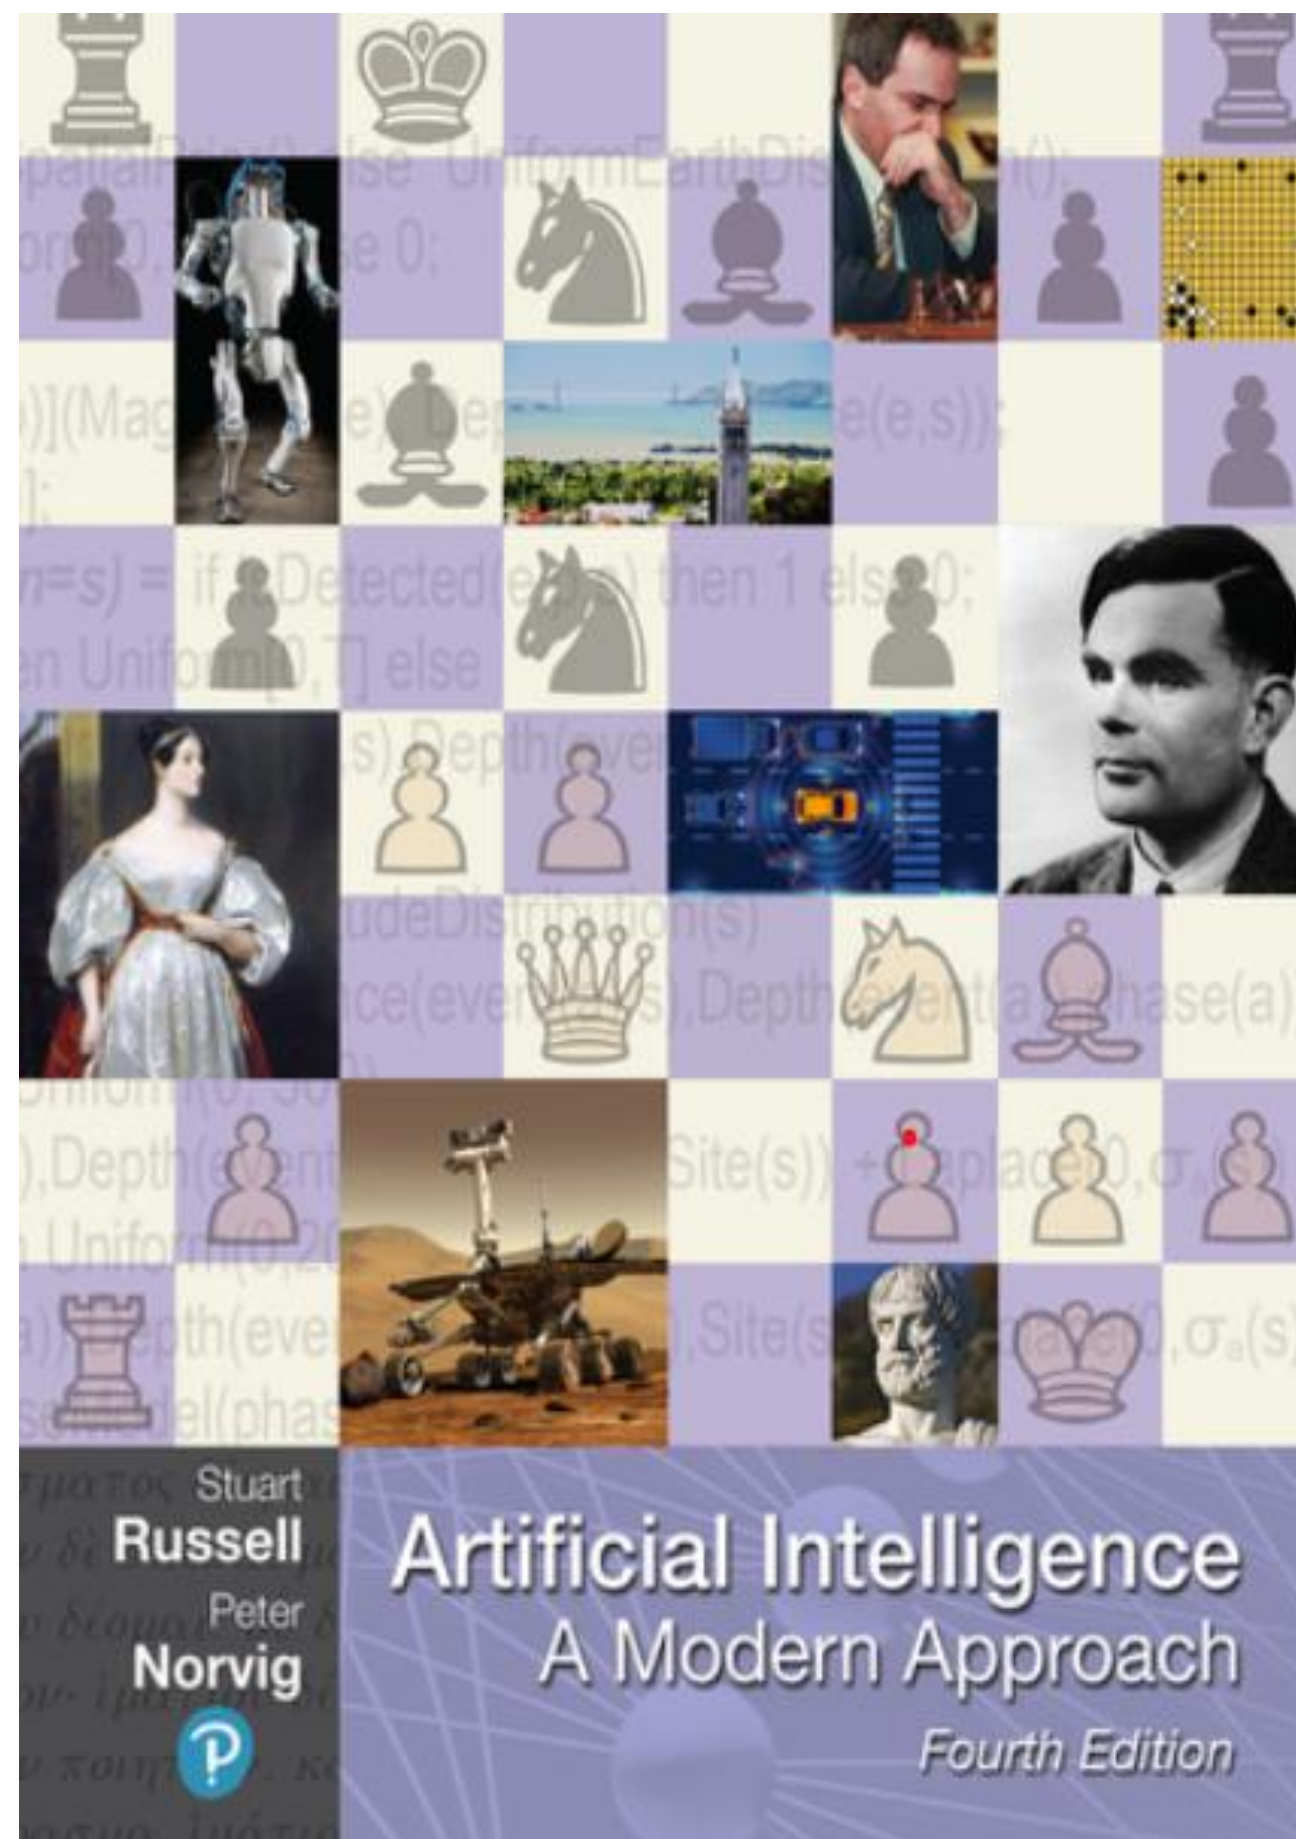
# Where to find everything?

- **Canvas** - *Pointer to everything*
  - grades, other private materials *that should not be shared*
- **Course website** - public materials
  - <https://pages.cs.wisc.edu/~gkotse/cs540f25/index.html>
  - Slides, schedule, policies
- **Piazza**
  - <https://piazza.com/wisc/fall2025/fa25compsci540/info>
  - Discussion, questions, announcements

# Textbook

Artificial Intelligence: A Modern Approach (4th edition).

Stuart Russell and Peter Norvig. Pearson, 2020. ISBN 978-0134610993.  
**(textbook is optional, but may be useful)**



# Instruction Team

(See course [webpage](#))

Merged across 3 sections:

- Teaching Assistants (TAs): hold office hours, grade your homework
- Peer Mentors: hold office hours

# Meet your classmates!

- Introduce yourself
- Share something that excites you about AI
- Find a commonality between you

# Office Hours

- Available on the course website
- Morgridge Hall Room B2582 (garden level)
- All office hours are **merged** across sections, you can go to anyone
- Use TA and Peer Mentor hours for detailed-level questions (e.g. coding related), and use professor office hours for conceptual level questions

# Grading scheme

- **Midterm Exam: 15%**
- **Final Exam: 15%**
- **Homework Assignments: 70%**
  - Ten assignments, submit via **Gradescope by 11:59 pm on due date**
  - **Two lowest homework scores dropped**
- **Top Hat:** up to 1% bonus for participation (20% of questions dropped)

# Integrity

Just don't cheat. It's not worth it.

You are encouraged to discuss with your peers, the TA or the instructors ideas, approaches and techniques broadly. However, all examinations, programming assignments, and written homeworks must be written up individually. For example, code for programming assignments must not be developed in groups, nor should code be shared. Make sure you work through all problems yourself, and that your final write-up is your own. If you feel your peer discussions are too deep for comfort, declare it in the homework solution: “I discussed with X,Y,Z the following specific ideas: A, B, C; therefore our solutions may have similarities on D, E, F...”.

You may use books or legit online resources to help solve homework problems, but you must always credit all such sources in your writeup and you must never copy material verbatim.

We are aware that certain websites host previous years' CS540 homework assignments and solutions against the wish of instructors. Do not be tempted to use them: the solutions may contain “poisonous berries” previous instructors planted intentionally to catch cheating. If we catch you copy such solutions, you automatically fail.

Do not bother to obfuscate plagiarism (e.g. change variable names, code style, etc.) One application of AI is to develop sophisticated plagiarism detection techniques!

Cheating and plagiarism will be dealt with in accordance with University procedures (see the [UW-Madison Academic Misconduct Rules and Procedures](#))

# Quiz

Join Top Hat and reply to questions online: <https://app.tophat.com>

1. Where can I find all the 540 stuff, if I didn't write down the URL?
2. I feel sick, should I still show up to class?
3. I can't finish my homework because I was traveling, I was sick, my dog ate it, etc. Can I ask for an extension?
4. Can I do homework with a group?


# Answers

1. Where can I find all the 540 stuff, if I didn't write down the URL? **Your Canvas has the main link.**
2. I feel sick, should I still show up to class? **NO. Study materials online.**
3. I can't finish my homework because I was traveling, I was sick, my dog ate it, etc. Can I ask for an extension? **No. But we discard 2 lowest hw scores.**
4. Can I do homework with a group? **Yes (and encouraged) for high level discussions. No for exact solutions.**



## Part III: Software

# Tools



- **Python**

- Main tool for machine learning & data science
- Conda package manager (for simplicity)

- **Jupyter**

- So much easier to keep track of your experiments
- Obviously, you should put longer code into modules

# Python for Java Pros (slides available on website)


## A Crash Course in Python

---

1. Why are we doing this in Python?
2. Where do I write Python code? How do I run it?
  - a. Online
  - b. Offline
3. What are the big differences between Java and Python

# Colab

- Go to [colab.research.google.com](https://colab.research.google.com)
- Activate the GPU supported runtime (this is a T4 GPU)



# **ML loves MATH** **(Probability, Statistics and Linear Algebra)**



# Part IV: Probability

# Probability: What is it good for?

- Language to express **uncertainty**



# In AI/ML Context

- Quantify predictions

$$[p(\text{lion}), p(\text{tiger})] = [0.98, 0.02]$$



$$[p(\text{lion}), p(\text{tiger})] = [0.01, 0.99]$$



$$[p(\text{lion}), p(\text{tiger})] = [0.43, 0.57]$$

\* If we know for sure the photo must contain either a lion or a tiger

# Model Data Generation

- Model complex distributions



**StyleGAN2** (Kerras et al '20)

# Win At Poker

- Wisconsin Ph.D. student Ye Yuan 5<sup>th</sup> in WSOP

Not unusual: probability began  
as study of gambling techniques

Cardano

*Liber de ludo aleae*

Book on Games of Chance  
1564!



[pokernews.com](http://pokernews.com)

# Basics: Outcomes & Events


- **Outcomes:** possible results of an **experiment**

$$\Omega = \underbrace{\{1, 2, 3, 4, 5, 6\}}_{\text{outcomes}}$$

- **Events:** subsets of outcomes we're interested in

$$\underbrace{\emptyset, \{1\}, \{2\}, \dots, \{1, 2\}, \dots, \Omega}_{\text{events}}$$

- Always include  $\emptyset, \Omega$




# Basics: Probability Distribution

- We have outcomes and events
- Assign **probabilities**: for each event  $E$ ,  $P(E) \in [0,1]$
- Back our example

$$\underbrace{\emptyset, \{1\}, \{2\}, \dots, \{1, 2\}, \dots, \Omega}_{\text{events}}$$

$$P(\{1, 3, 5\}) = 0.2, P(\{2, 4, 6\}) = 0.8$$




# Basics: **Axioms**

- Rules for probability:
  - For all events  $E$ ,  $P(E) \geq 0$
  - Always,  $P(\emptyset) = 0, P(\Omega) = 1$
  - For disjoint events,  $P(E_1 \cup E_2) = P(E_1) + P(E_2)$
- Easy to derive other laws. Ex: non-disjoint events

$$P(E_1 \cup E_2) = P(E_1) + P(E_2) - P(E_1 \cap E_2)$$


# Visualizing the Axioms: I

- Axiom 1: for all events  $E$ ,  $P(E) \geq 0$



# Visualizing the Axioms: II

- Axiom 2:  $P(\emptyset) = 0, P(\Omega) = 1$




$$P(\Omega) = 1$$




# Visualizing the Axioms: III

- Axiom 3: disjoint  $P(E_1 \cup E_2) = P(E_1) + P(E_2)$



# Visualizing the Axioms

- Also, other laws:



# Break & Quiz

- **Q 1.1:** We toss a biased coin. If  $P(\text{heads}) = 0.7$ , then  $P(\text{tails}) = ?$
- A. 0.4
- B. 0.3
- C. 0.6
- D. 0.5

# Break & Quiz

- **Q 1.1:** We toss a biased coin. If  $P(\text{heads}) = 0.7$ , then  $P(\text{tails}) = ?$
- A. 0.4
- **B. 0.3**
- C. 0.6
- D. 0.5

# Break & Quiz

- **Q 1.2:** There are exactly 3 candidates for a presidential election. We know X has a 30% chance of winning, B has a 35% chance. What's the probability that C wins?
- A. 0.35
- B. 0.23
- C. 0.333
- D. 0.8

# Break & Quiz

- **Q 1.2:** There are exactly 3 candidates for a presidential election. We know X has a 30% chance of winning, B has a 35% chance. What's the probability that C wins?
- **A. 0.35**
- B. 0.23
- C. 0.333
- D. 0.8

# Break & Quiz


- **Q 1.3:** What's the probability of selecting a black card or a number 6 from a standard deck of 52 cards?
- A.  $26/52$
- B.  $4/52$
- C.  $30/52$
- D.  $28/52$

# Break & Quiz

- **Q 1.3:** What's the probability of selecting a black card or a number 6 from a standard deck of 52 cards?
- A.  $26/52$
- B.  $4/52$
- C.  $30/52$
- **D.  $28/52$**

# Coming up: Probability and Linear Algebra Review

$C = AB$  where  $C_{ik} = \sum_j A_{ij} B_{jk}$





# Recap

- What's in CS540
- Course logistics
- Software
- Probability



**Thanks!**