CS 486/686
Introduction to Artificial Intelligence

Yuntian Deng

Lecture 1 (Part 1)

(Adapted from Prof. Wenhu Chen's CS486/686 slides)

Outline

Learning goals

- Get to know a bit about the instructor and one or more classmates.
- Understand the topics of the course
- Name an application of AI.

Outline

Who am I?

My name is Yuntian Deng. I am from China.

Appointments:

- Assistant Professor, University of Waterloo
- Visiting Professor, NVIDIA (Prof. Yejin Choi)
- Associate, Harvard CS
- Faculty Affiliate, Vector Institute

Previous:

- Postdoc at AI2 (Prof. Yejin Choi)
- PhD, CS, Harvard University (Profs. Alexander Rush, Stuart Shieber)
- MS, Language Technologies, CMU (Prof. Eric Xing)
- BE, Automation, Tsinghua University (Prof. Jie Zhou)

Research: natural language processing, machine learning

Who are the TAs?

The TAs will monitor Piazza, grade your homework, host TA office hours, and grade the final exam.

- Wentao Zhang (w564zhan@)
- Bihui Jin (b27jin@)
- Max Ku (m3ku@)
- Ruoxi Ning (r2ning@)
- Hala Sheta (hsheta@)
- Shuhui Zhu (s223zhu@)

Outline

Topics in CS 486/686 (1/2)

Introduction to AI and CS486/686

1: Introduction to AI.

Search:

2: Uninformed Search
3: Heuristic Search
4: Constraint Satisfaction Problem
5: Local Search

Uncertainty Estimation:

6: Uncertainty and Probability
7: Independence of Bayesian Networks
8: Bayesian Networks
9: Variable Elimination Algorithm
10: Hidden Markov Models 1
11: Hidden Markov Models 2

Topics in CS 486/686 (2/2)

Markov Decision Process:

12: Decision Theory
13: Decision Networks
14: Markov Decision Processes 1
15: Markov Decision Processes 2
16: Reinforcement Learning 1
17: Reinforcement Learning 2

Machine Learning & Deep Learning:

18: Machine Learning, Decision Trees 1
19: Decision Trees 2
20: Neural Networks 1
21: Neural Networks 2
22: TBA
23: Conclusion and Recap

What's included in Search

- Learn the generic search algorithm
- Discuss uninformed search algorithms like DFS, BFS
- Discuss heuristic search algorithms like the A* algorithm
- understand the complexity, and completeness of search algorithms
- Understand constraint satisfaction problem
- Understand how to solve CSP with the arc-consistency algorithm
- Learn about local search algorithms like annealing, genetic, etc

What's included in Uncertainty Estimation

- Sum Rule, Product Rule, Chain Rule, Bayes Rule
- Understand the random variable independence
- D-separation Principle for Testing independence
- Understand how to construct Bayesian Networks
- Learn how to answer query about Bayesian Networks
- Understand the concept of Hidden Markov Model
- Forward-Backward Algorithm to compute hidden states

What's included in Markov Decision Process

- Understand Decision Network, Actions, and Utility
- Variable Elimination in Decision Networks
- Markov Decision Network to Make Sequence of Decisions
- Value Iteration Algorithm to Compute Best Policy in Markov Decision Network
- Reinforcement Learning (Temporal Difference Learning + Q-Learning)

What's included in Machine Learning

- Understanding Supervised/Unsupervised Learning, Bias-Variance Trade-off
- Understand Decision Tree Algorithm, Learn to construct Decision Tree
- Basics of Neural Networks
- Backward Propgation Algorithm in Neural Networks

Outline

Main Website

https://bit.ly/cs486

- Frequently updated with learning materials and announcements.
- Includes all the deadlines, release date, slides and notes.

Piazza

Best way to reach me = office hour

Second-best way to reach me = private Piazza post

- Ask questions related to course
- You're encouraged to answer other students' questions
- We will try to respond in 2 business days

TA Support

If you have questions regarding assignments or learning materials, you can also consult TAs.

Office Hours

- Instructor Office Hour: Thursdays 4-5pm in DC2633
- TA office hours: TBA on Piazza

Outline

Assignments

There are 3 assignments – each has a written and programming part

Written part

- E.g. trace through the algorithm, prove something, analyze results from programming part
- Submit on Crowdmark

Programming part

- Implement functions for algorithms discussed
- Submit code to Marmoset
- Week 1, week 2, final test setups

Assignments

Each assignment will have roughly 20 days of completion time

- Assignment 1
- Assignment 2
- Assignment 3

Exam

Final: 2.5 hours, date/time (TBA)

- Every student must attend the final exam, no makeup exam
- CS 486 students must pass the final exam
- Final exam will cover all the materials

Project

Solve a problem with AI (or for AI) and write a report

- CS 686 students must pass the project to pass the course
- Optional for CS 486 students
- Done individually or in groups of up to 3

CS 686 students only

Where can I find free computation resources like GPUs

- Google Cloud Colab (https://colab.research.google.com)
- Math Faculty Computing (https://uwaterloo.ca/math-faculty-computing-facility/services/service-catalogue-teaching-linux/access-teaching-gpu-cluster)

Project

- Topics are not restricted, it could be reinforcement learning, image classification, natural language processing, human-computer interaction, biomedical, etc.
- It does not need to reach publication-level novelty.
- It cannot be trivial, like cloning some code from existing public GitHub.
- Students can form a team up to 3 members, but the report must specify the responsibility of each member.
- Proposal due around the middle of the term, final report due after the final exam.

Score Breakdown

- For CS486: 30% from Assignment, 20% from Quiz, 50% from Final Exam, 10% bonus from Project
- For CS686: 30% from Assignment, 30% from Project, 40% from Final Exam

Outline

The Goal of AI

Artificial Intelligence:

- Enable machines to sense, comprehend, act and learn with human-like levels of intelligence
- Combined with computer science to solve challenging problems
- Not confined to methods that are biologically observable

The History of AI: Alan Turing

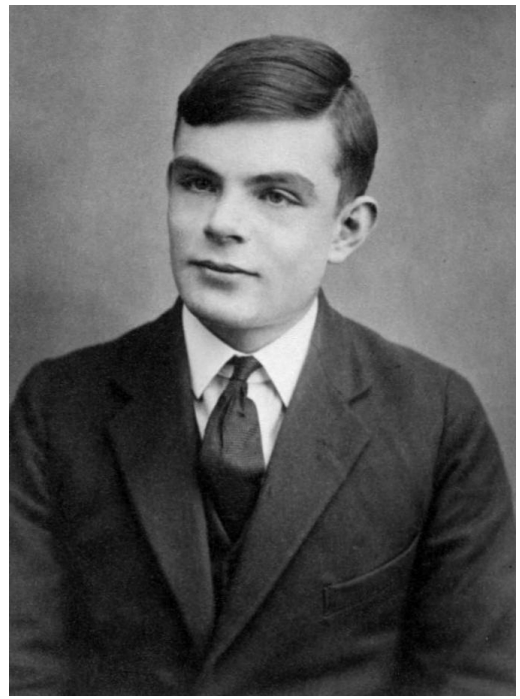

Alan Turing: 1912-1954 was primarily known for codebreaking in WW1 (https://www.imdb.com/title/tt2084970/) Also asked 'can machines think?', proposed the imitation game.

The History of AI: Dartmouth Conference

- During Dartmouth conference in mid-1950s, Jon MacCarthy proposed the term "Artificial Intelligence"
- McCarthy also created Lisp, the standard programming language used in robotics

The History of AI: the Turing Test

The interrogator is connected to one person and one machine via a terminal, therefore can't see her counterparts. Her task is to find out which of the two candidates is the machine, and which is the human only by asking them questions. If the machine can 'fool' the interrogator, it is considered intelligent.

The History of AI: the Turing Test

ELIZA - (1966 Joseph Weizembaum, Lisp)

ME: Hi
ELIZA: How do you do? What brings you to see me?
ME: I'm thinking of taking a course in artificial intelligence.
ELIZA: Why do you say that?
ME: I'm worried about it
ELIZA: Are you worried often?
ME: when I take AI courses
ELIZA: What happened when you take ai courses?
ME: I usually spend long nights doing assignments
ELIZA: Earlier you said you felt worried?
ME: about AI courses
ELIZA: Maybe your plans have something to do with this.
ME: no
ELIZA: Why not?
ME: because I think you are a computer
ELIZA: Is it because you think I am a computer that you came to me?

Loebner Prize - Turing test competition

- fool all judges: 1M prize
- no one has ever won before 2020
- example interaction:

CONFEDERATE: Hello, my name is Sam and I am the human
JUDGE: Hi Sam, tell me which is larger, a TV screen or Belgium.
CONFEDERATE: Belgium is a little bit larger.
JUDGE: Wljso lwjejdo wleje elwjeo wjeol, isn't it?
CONFEDERATE: Is that Polish?
JUDGE: Which is further away from this apartment, Queens or the Queen of England?
CONFEDERATE: If the Queen is visiting Manhattan, she is closer.

Blocks World - SHRDLU - Winograd 1968

- first autonomous physical interaction
- blocks world only

Pick up a big red block
OK
grasp the pyramid
I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN
What does the box contain?
THE BLUE PYRAMID AND THE BLUE BLOCK.
What is the pyramid supported by?
THE BOX.
How many blocks are not in the box?
FOUR OF THEM

Game Playing

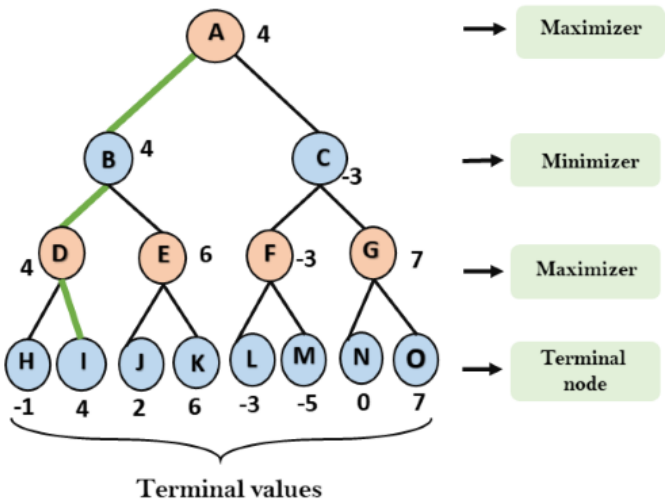

- Game playing is also a very important metric to evaluate intelligence
- Min-Max Game of two players, one wants to maximize the chance of wining, the other aims to minimize it
- Game playing can be formulated as a search problem with nodes representing the state

Chess

- game tree has more than 10^{100} nodes
- IBM - Deep Blue uses lookahead search and complex evaluation function to predict the outcome.
- In 1997, IBM Deep Blue defeated the world champion Gary Kasparov in a 6-match game.
- Deep Blue AI 2022 (https://www.scientificamerican.com/article/20-years-after-deep-blue-how-ai-has-advanced-since-conquering-chess/)

Go (Weiqi)

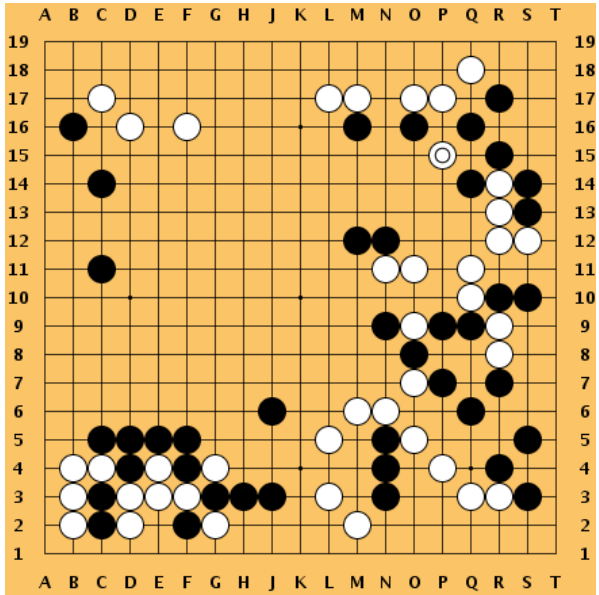

- game tree has more than 10^{360} nodes, which takes longer for AI to master
- AlphaGo (Deepmind): Silver et al. (2016). Mastering the game of Go with deep neural networks and tree search. Nature, 529(7587), 484-489.
- Deep neural networks; supervised learning; reinforcement learning.

AlphaGo

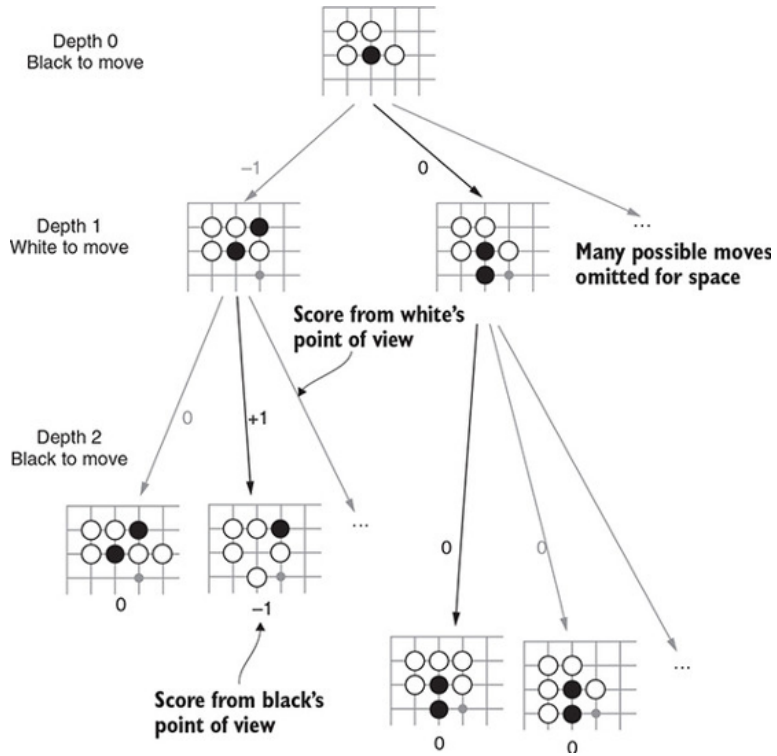

- AlphaGo combines machine learning and tree search techniques
- Alphago uses Monte-Carlo tree search algorithm, guided by a "value network" and "policy network" (covered in our lectures)
- AlphaGo learns by playing against both humans and itself

Poker

- Play with uncertainty. Must model opponent(s). Care about long-term payoff.
- Bowling, M., Burch, N., Johanson, M., & Tammelin, O. (2015). Heads-up limit hold'em poker is solved. Science, 347(6218), 145-149.
- Brown, N., & Sandholm, T. (2019). Superhuman AI for multiplayer poker. Science, 365(6456), 885-890.

Atari Games

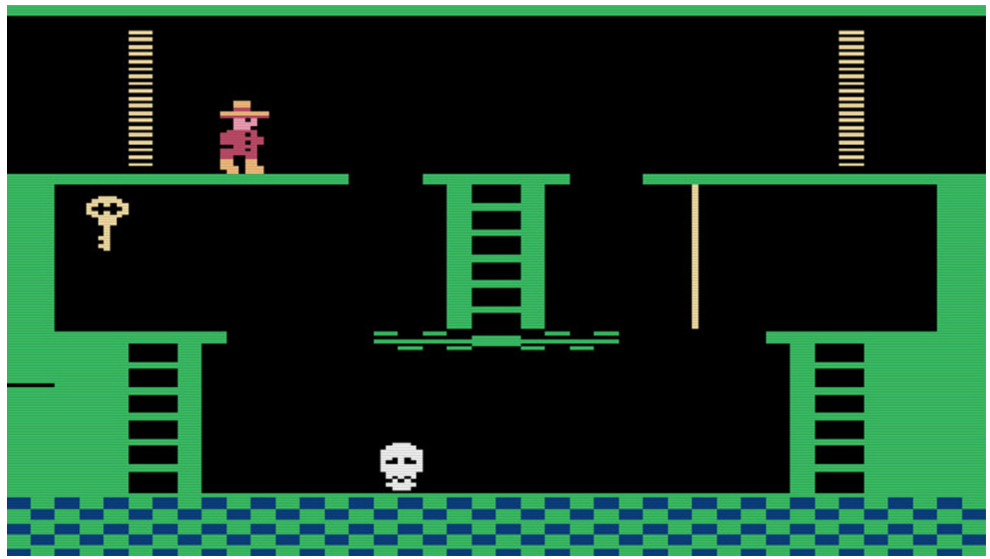

- Mnih et al. (2013). Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602.
- Reinforcement learning; Convolutional neural network; High-dimensional sensory input. Previous approaches use hand-crafted visual features.
- Outperforms previous approaches; Surpasses a human expert on 3/7 Atari 2600 games.

StarCraft II

Figure: https://snl.no/StarCraft_II

- Multi-agent problem; Imperfect information;
  Large action and state space; Delayed credit assignment.
- Vinyals et al. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. Nature, 575(7782), 350-354.
- Video: https://www.youtube.com/watch?v=jtlrWbl0yP4

AlphaFold

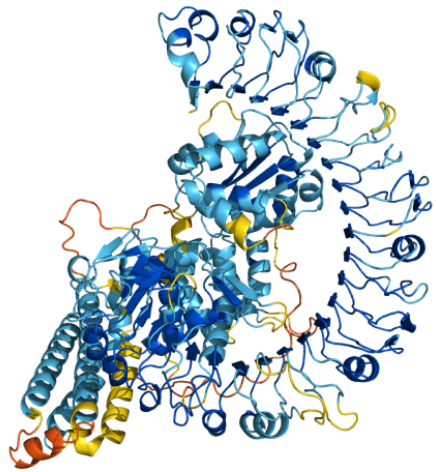

- Synthesize structures of unknown protein (previously requires painstaking months to years human labor)
- Jumper et al. (2021) Highly accurate protein structure prediction with AlphaFold
- Benefit greatly from the growth of the Protein Data Bank
- Uses a model called EvoFormer to make structure prediction

The surge Deep Neural Network

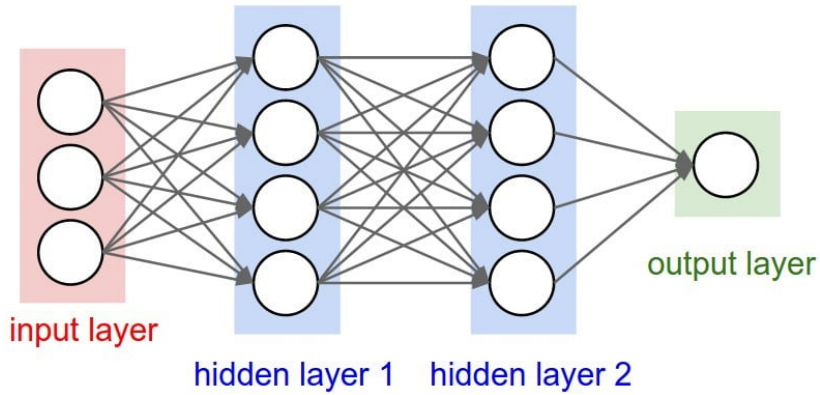

- Old-fashioned ML used decision tree, SVM, linear regression, boosting algorithms.
- Neural Network exist for quite a while. It really took off with AlexNet (Alex et al. ImageNet Classification with Deep Convolutional Neural Networks)
- Increasing the depth of neural networks to hundreds of layers.

Development of AI Models

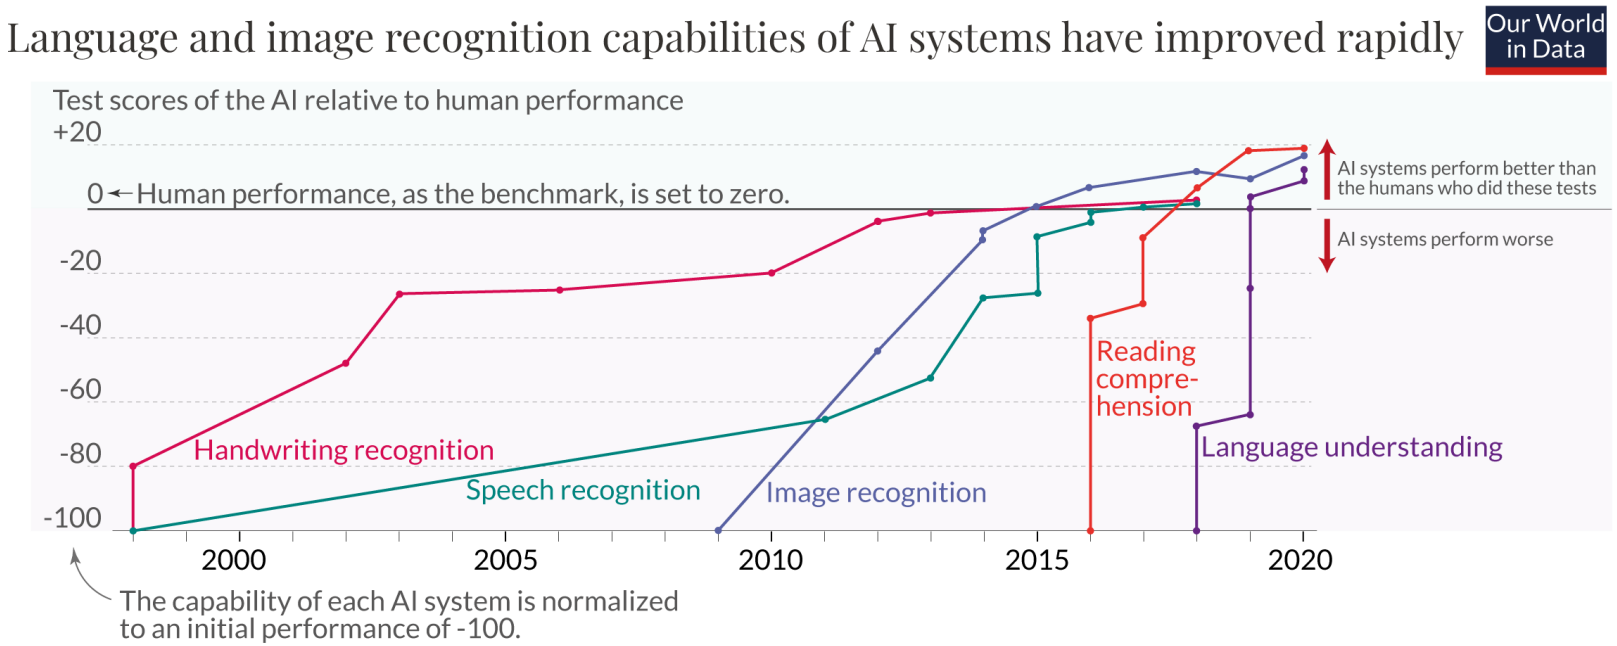

- The language and image recognition capabilities of AI systems have developed very rapidly.
- Just 10 years ago, no machine could reliably provide language or image recognition at a human level. But, as the chart shows, AI systems have become steadily more capable and are now beating humans in tests in all these domains.

ImageNet Challenge

IMAGENET

- 1,000 object classes (categories).
- Images:
  - 1.2 M train
  - 100k test.

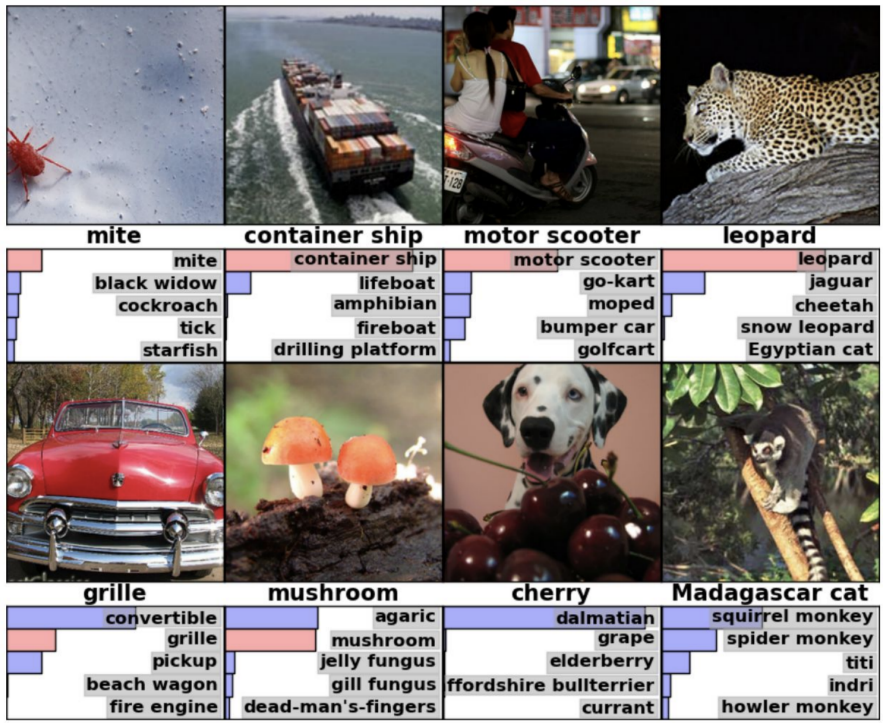

Figure: ImageNet Challenge

A standard classification problem to categorize an image into 1 of 1000 object classes.

Image Classification

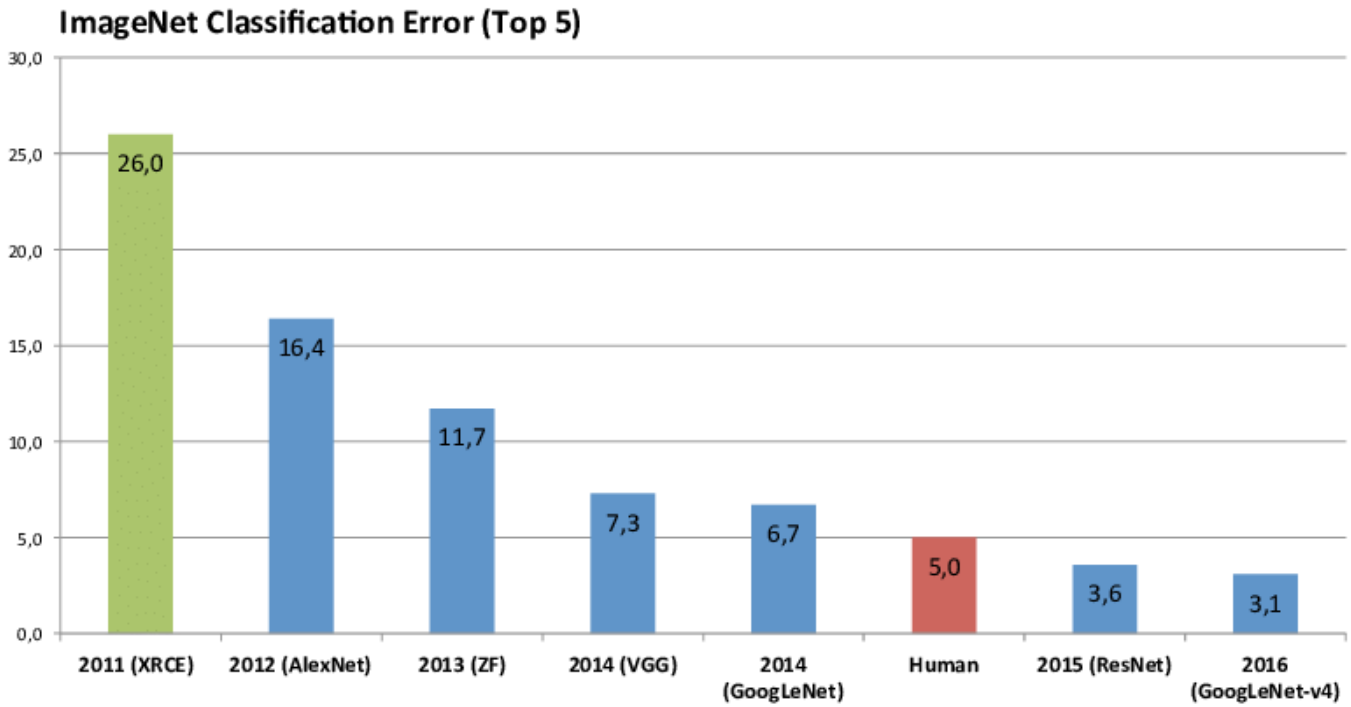

Figure: ImageNet winner over time

Before 2011, the curve was relatively flat. In 2012, AlexNet was proposed as the first deep neural network to address the long-standing problem. Already solved.

Image Generation

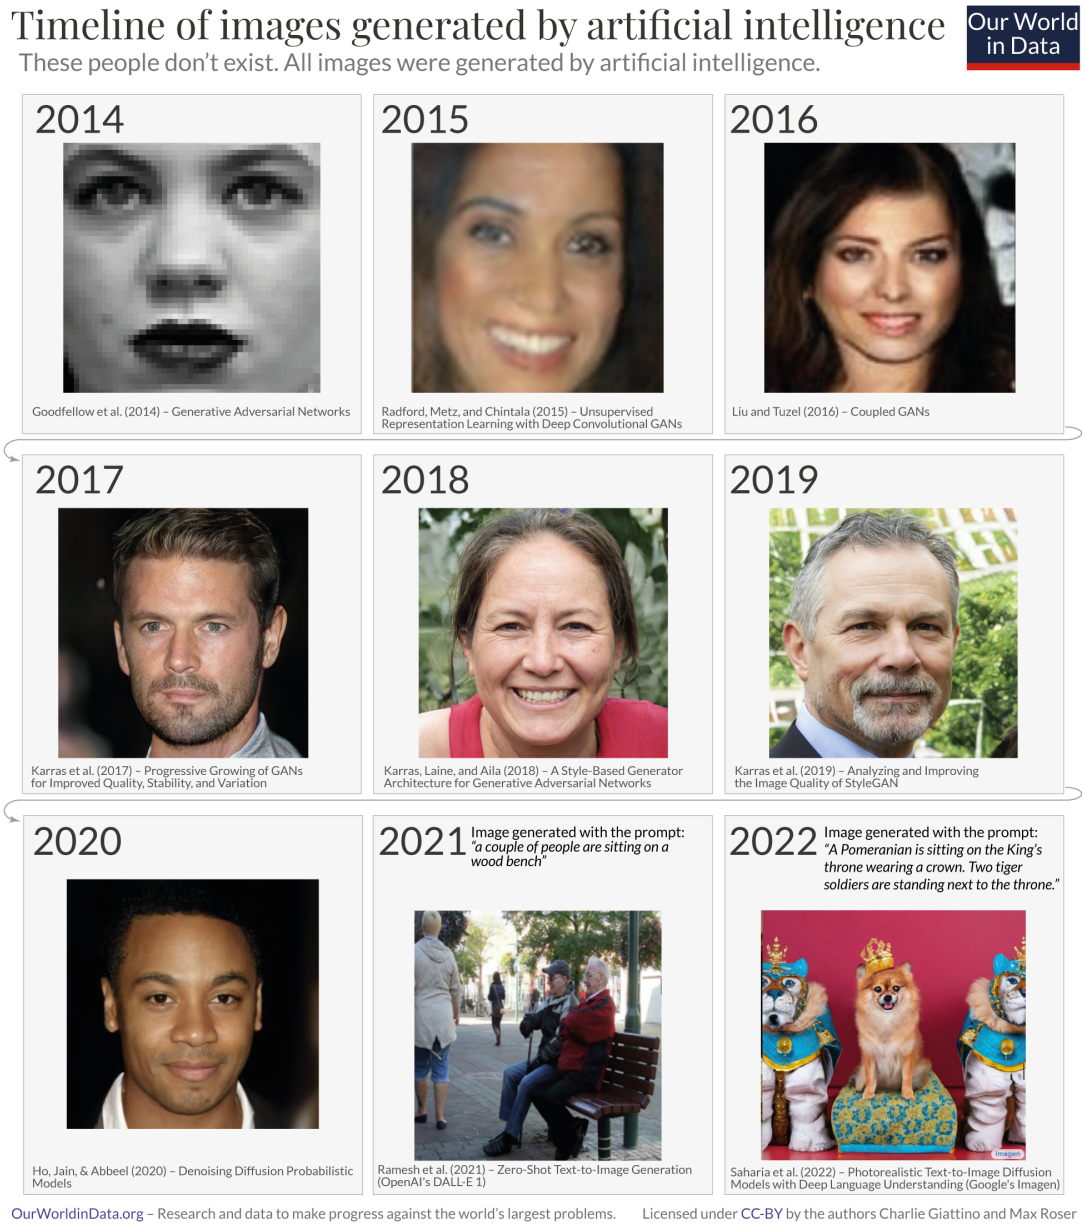

Supervised Learning (Human Annotation)

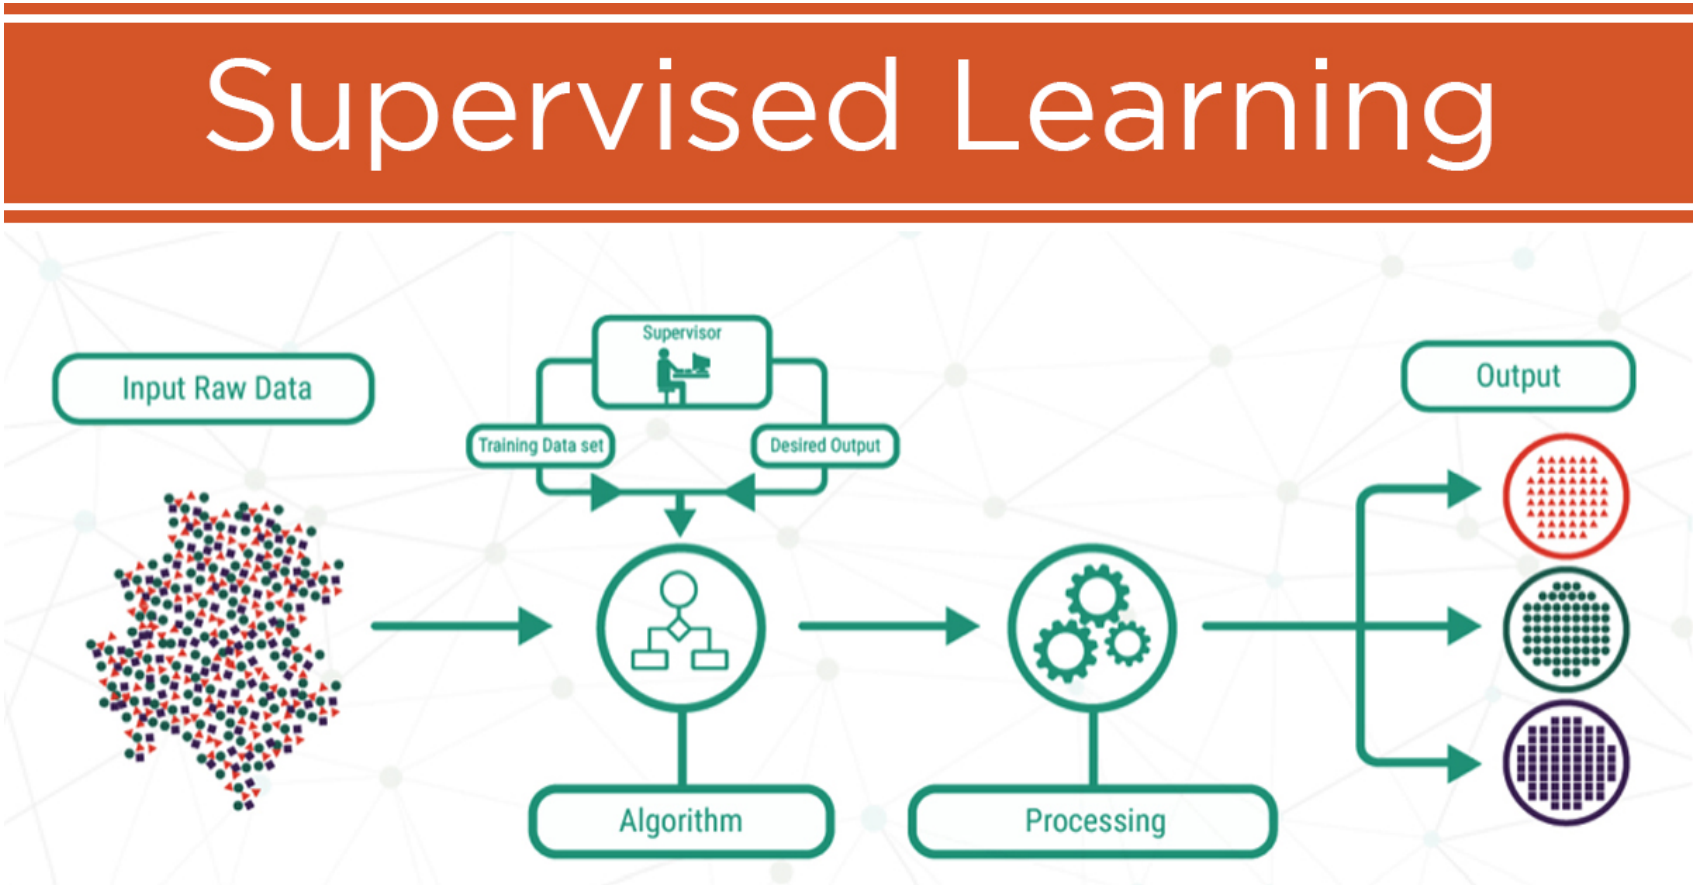

Self-Supervised Pre-training

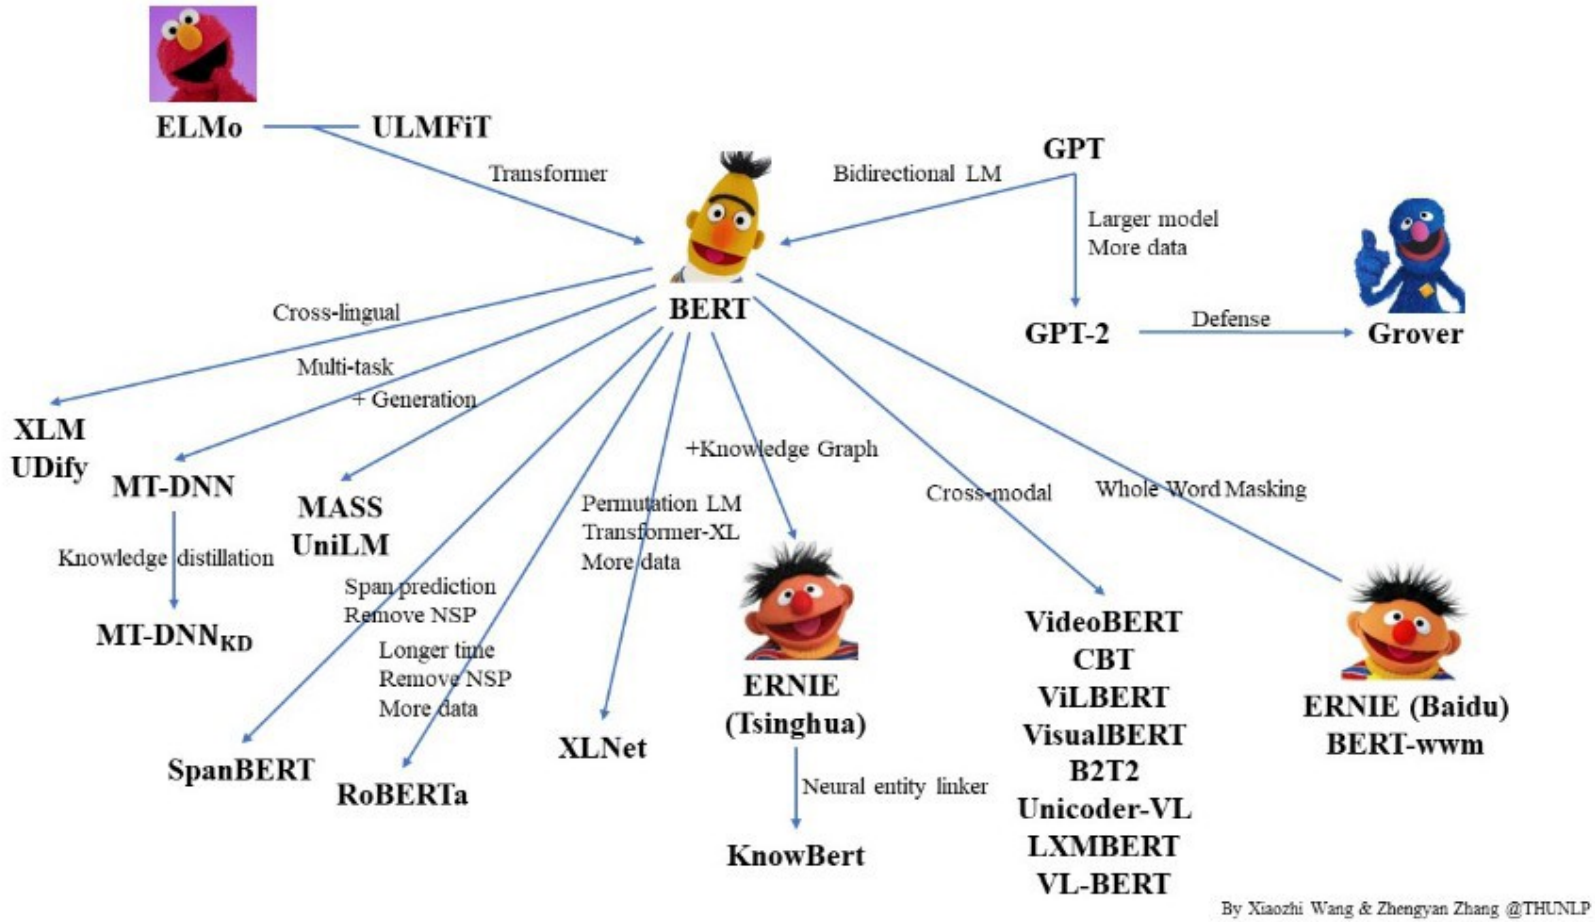

Instead of training on annotated data, the model trains on the crawled web with self-supervised objective

Large Language Models

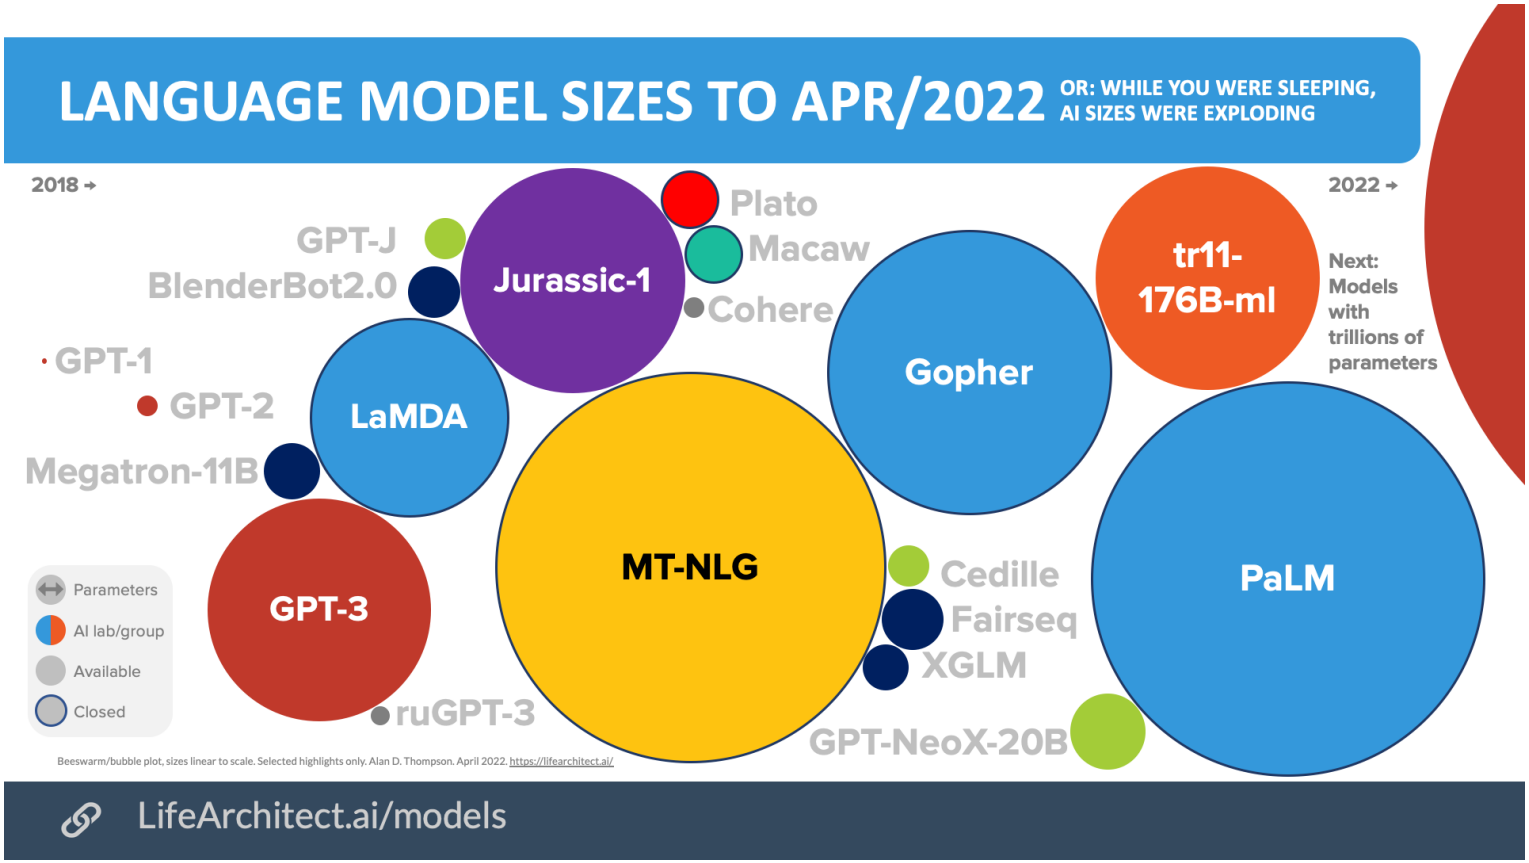

- the current text/language models grow larger and larger
- Google's PaLM model is already achieving 540B parameters

Large Language Models

The current large language models have shown striking performance in many aspects. The most recent ChatGPT model from OpenAI can give astonishingly accurate responses.

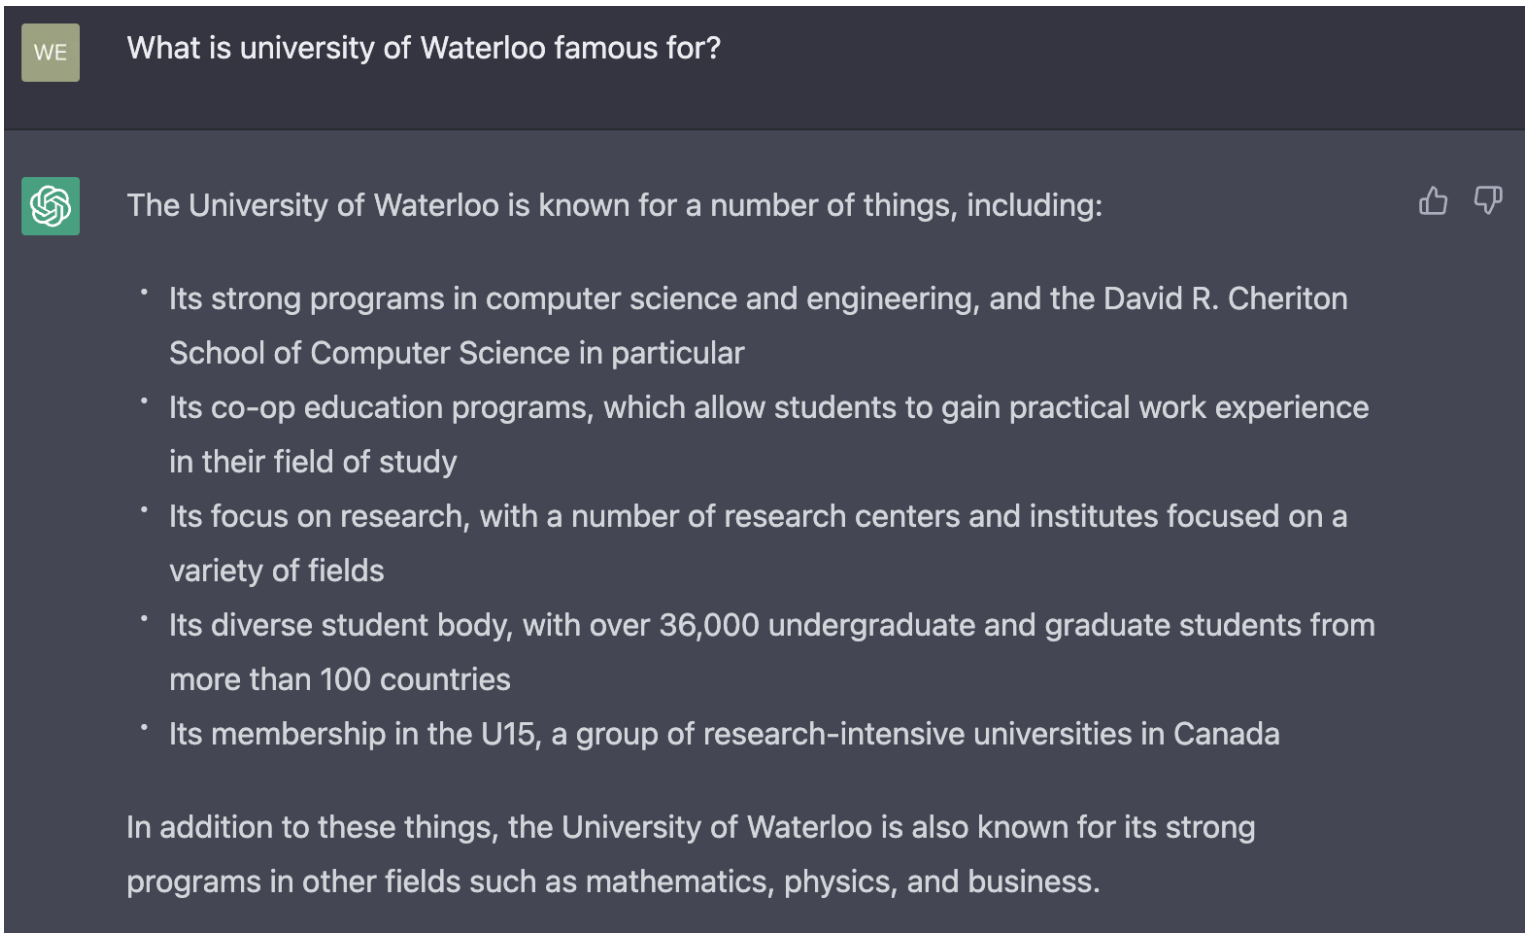

Large Language Models

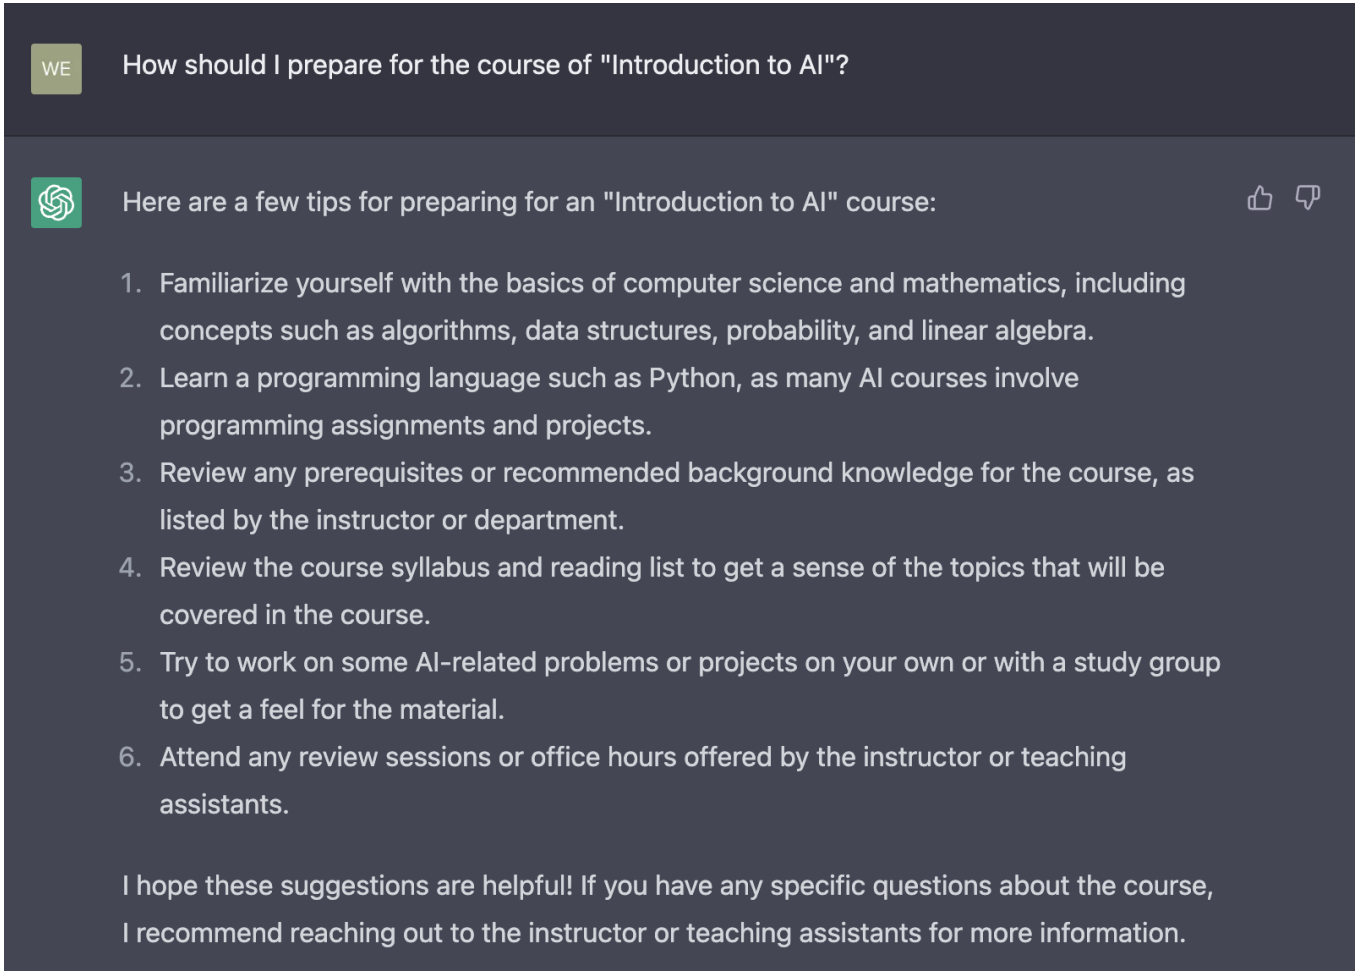

Learning goals

- Get to know a bit about the instructor and one or more classmates.
- Understand the topics of the course
- Name an application of AI.