Introduction

Philipp Koehn

21 January 2025

Administrative

- **Instructor:** Philipp Koehn (phi@jhu.edu)
- **TA:** Steven Tan
- **Class:** Tuesdays and Thursdays 1:30-2:45, Maryland Hall 110
- **Textbook:** "Artificial Intelligence — A Modern Approach", by Russell and Norvig, 4th edition, 2020
- **Course web site:** `http://www.cs.jhu.edu/~phi/ai/`
- **Grading**
  - 5 assignments (10% each)
  - midterm exam (25%)
  - final exam (25%)

Fair Warning

- **This is not a Deep Learning class**

- We take a very broad view of Artificial Intelligence

- Some topics we cover have even fallen out of favor
  - they may just be waiting to be re-discovered
  - it is important to know why certain ideas failed

Main Topic Areas

- Artificial Intelligence in Context (5 lectures)
- Intelligent Agents, Heuristic Search, and Game Playing (5 lectures)
- Logic and Knowledge Representation (5 lectures)
- Uncertainty (3 lectures)
- Machine Learning (4 lectures)
- Generative AI (2 lectures)

ai in context

Artificial Intelligence in Context

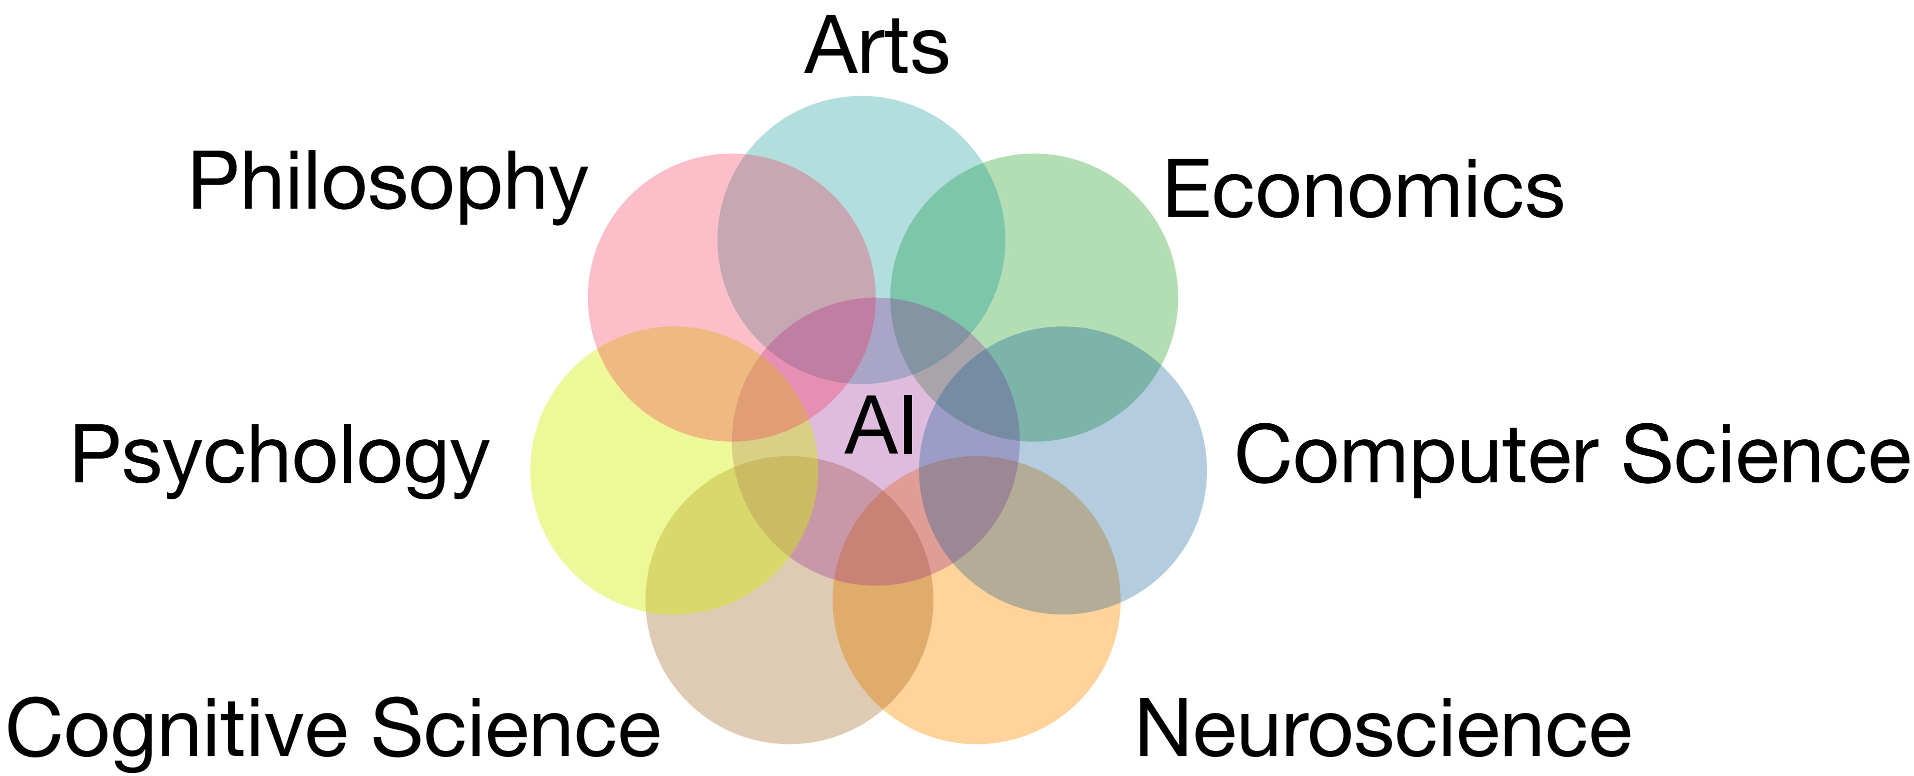

Artificial Intelligence

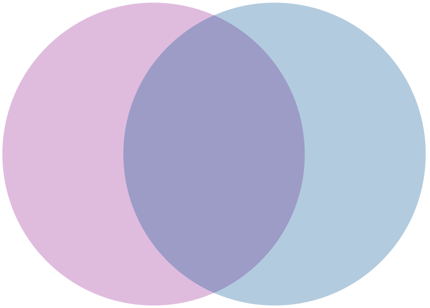

Computer Science

- In a way, all of computer science develops methods that replace thinking
- Artificial Intelligence is the Computer Science that does not work yet
- Algorithm (CS) = Automatic procedure that produces intended result
- Heuristic (AI) = May fail, does not work for all cases, is an estimate
- Modeling of the world: knowledge representation, common sense inference

Artificial Intelligence

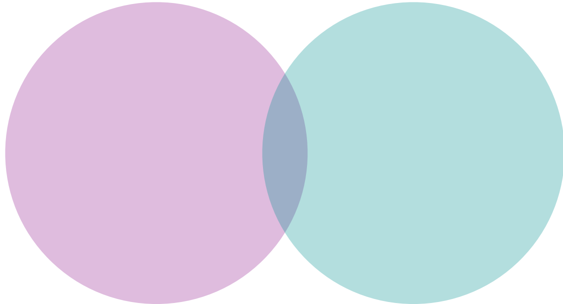

Arts

- Having intelligent machines a perennial topic in popular culture
- How do we interact with intelligent machines?
- How does artificial intelligence change interactions between humans?

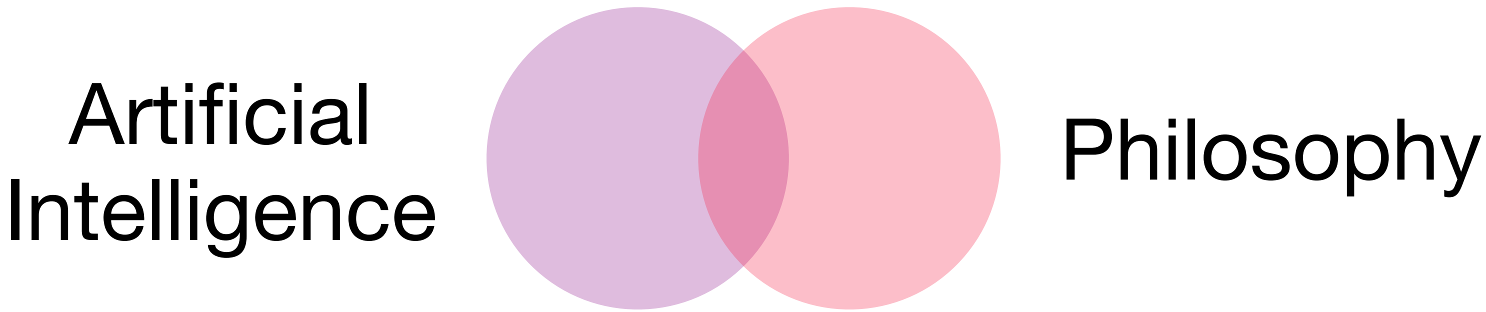

- Can computers ”think”?
- Can computers have ”consciousness”?
- What is the difference between simulation of a mind and a real mind?
- If we had a thinking robot, what would that say about us humans?

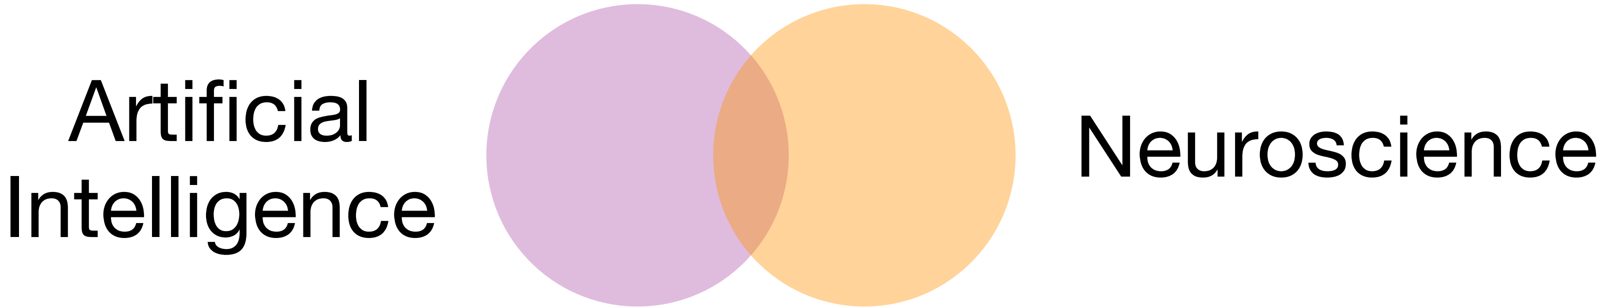

- We all already have our personal thinking machine: the brain
- How does it work?
- Can we build ”neural” computers that work the same way?
- If we better understand ”neural” computing, do we learn about our brain?

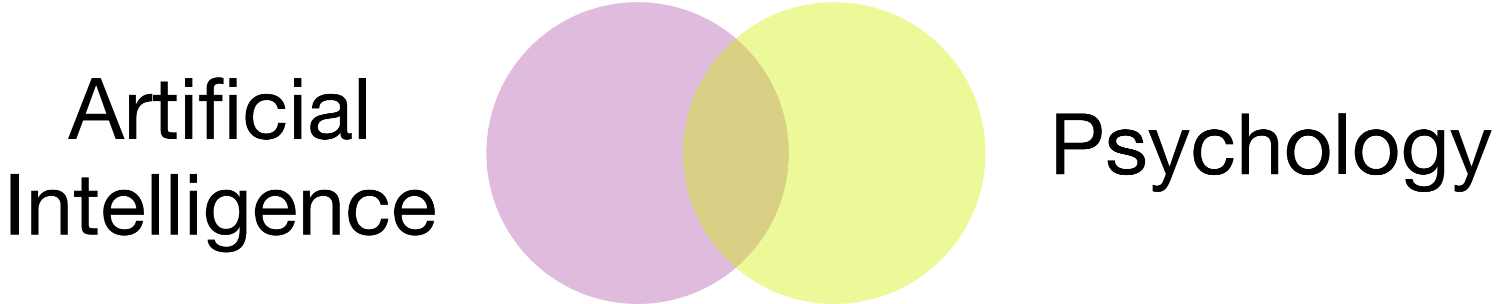

- How does the human mind process information?
- What are key properties of the mind?
- Besides rational thinking, what is the point of emotions, being embodied?

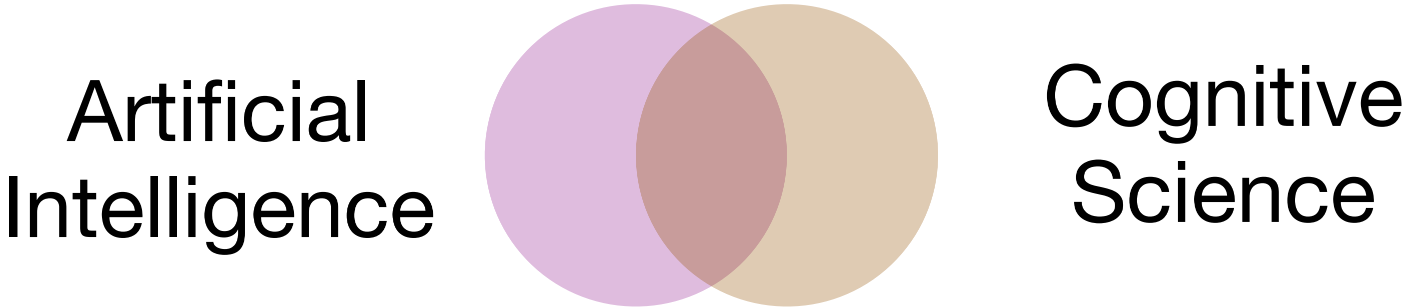

- Cognitive Science is an umbrella term for the previously mentioned disciplines
- Interdisciplinary: how do insights from one discipline inform the others?
- Role of artificial intelligence: implement grounded models of cognition

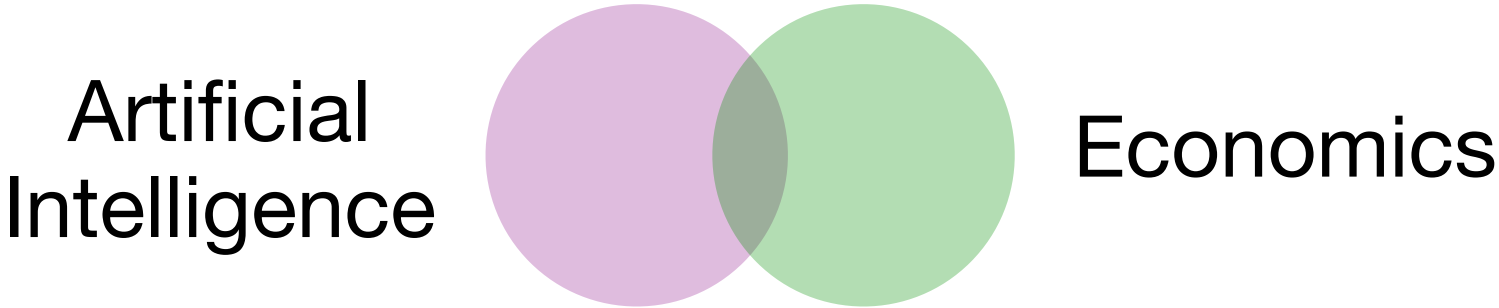

- Economic modeling inspired agent-based approach in artificial intelligence (operation research, decision theory, game theory)
- Both disciplines attempt to model human behaviour

Division of Sciences

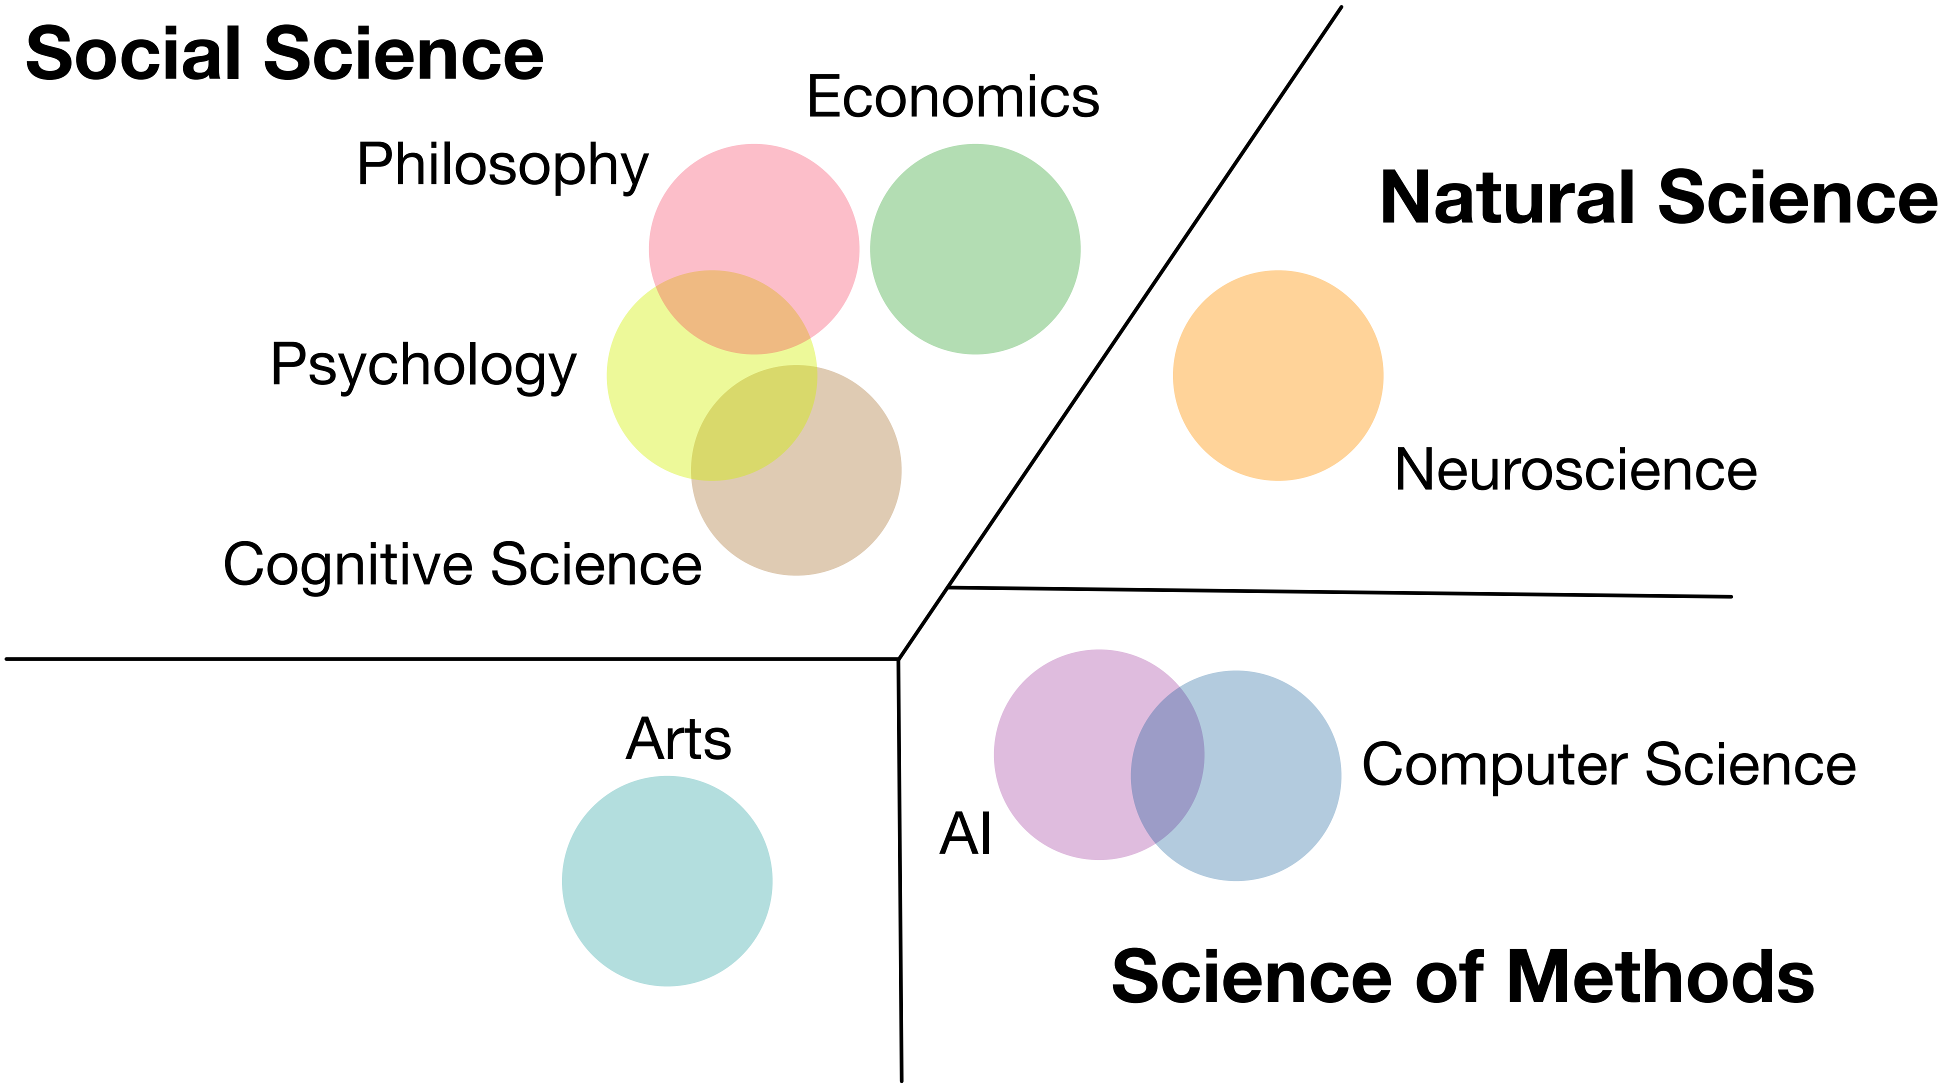

history

Dartmouth Conference

- Dartmouth Summer Research Project on Artificial Intelligence, 1956

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

Dartmouth Conference

- Dartmouth Summer Research Project on Artificial Intelligence, 1956

> *We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of* **the conjecture that** *every aspect of learning or* **any** *other* **feature of intelligence can in principle be so precisely described that a machine can be made to simulate it**. *An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.*

- Basic premise: intelligence can be described in formal rules

Dartmouth Conference

- Dartmouth Summer Research Project on Artificial Intelligence, 1956

> *We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines* **use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves***. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.*

- Problems to work on — a bit vague: language, abstraction, learning

Dartmouth Conference

- Dartmouth Summer Research Project on Artificial Intelligence, 1956

> *We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a* **significant advance can be made** *in one or more of these problems if a carefully* **selected group of scientists work on it together for a summer**.

- Great expectations: one summer.

Logic Theorist

- Program to perform mathematical proofs, Newell and Simon, 1955-1956

- Proved 38 of the first 52 theorems in Principia Mathematica

- Logic Theorist introduced several central AI concepts
  - Reasoning as search: consider exponential expansion of possible steps
  - Heuristics: rules of thumb to prune search tree
  - List processing: led eventually to development of Lisp

- Followed up by work on General Problem Solver

Boom and Bust

- Early successes
  - computers were winning at checkers
  - solving word problems in algebra
  - proving logical theorems

- Great promises

  ... within ten years a digital computer will be the world's chess champion.
  Herbert Simon and Allen Newell, 1958

  In from three to eight years we will have a machine with the general intelligence of an average human being.
  Marvin Minsky, 1970

- Late 1970s: AI Winter, funding stopped

Hyperbole and Disappointment

- Some hard problems (for humans) are easy to solve for computers

  recall: early success with mathematical proofs

- Some easy problems are hard for computers

  e.g., understanding language, recognizing objects, walking

- We tend to underestimate the ”easy” problems

  claims of (”will be solved in 10 years”)

- Maybe an intelligent computer needs to ”live” in the world to understand it

Expert Systems (Early 1980s)

- Idea
  - focus on a specific subject
  - consult with an expert to write down all facts and rules
  - build a computational system that applies rules to test cases

- Example: Medical Diagnosis
  - Collect set of symptons, diseases, and elements of treatment plans
  - Write rules that predict further testing steps
  - Write rules that predict disease
  - Define treatment plan from template, given state and severity of disease

- Not very successful
  - hard to formalize all aspects of expert knowledge
  - systems get quickly too complex to manage

Encoding Commonsense Knowledge

- For instance: Cyc project, started in 1986 (available as OpenCyc)
- Encode facts about the world
  - Barack Obama is a US President

    `(#$isa #$BarackObama #$UnitedStatesPresident)`
  - *all trees are plants*

    `(#$genls #$Tree-ThePlant #$Plant)`
- Inference engine that can answer queries
- Challenges: uncertainty, interface to natural language

Intelligent Agents (since late 1980s)

- Formal definition of intelligent agent (inspired by rational agent in economics)
  - perceives the environment
  - may have a model of the environment
  - has goals or a utility function
  - decides on an action
  - changes environment
  - may learn from environment

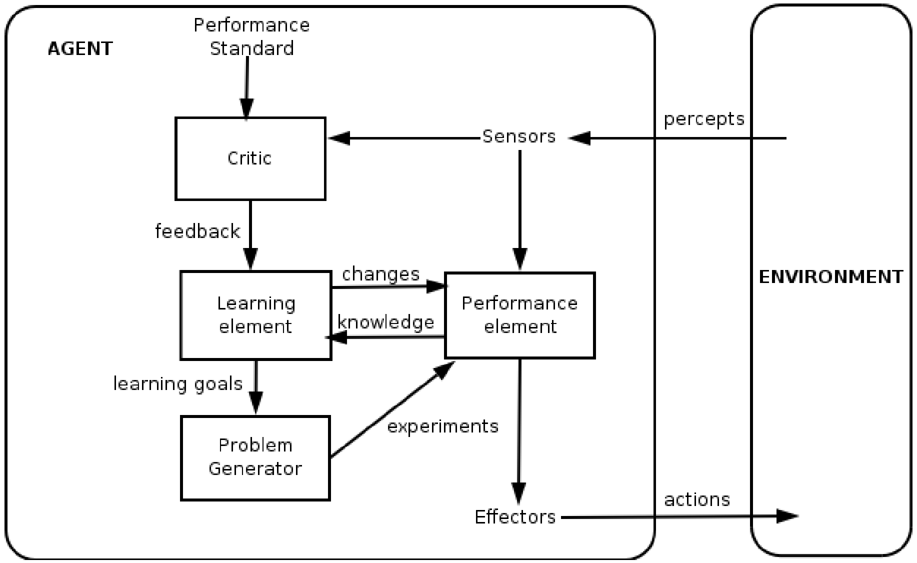

- Inclusion of uncertainty and probabilistic inference

- Requirement of empirical validation

⇒ AI a more rigorous "scientific" discipline

Machine Learning (since 1990s)

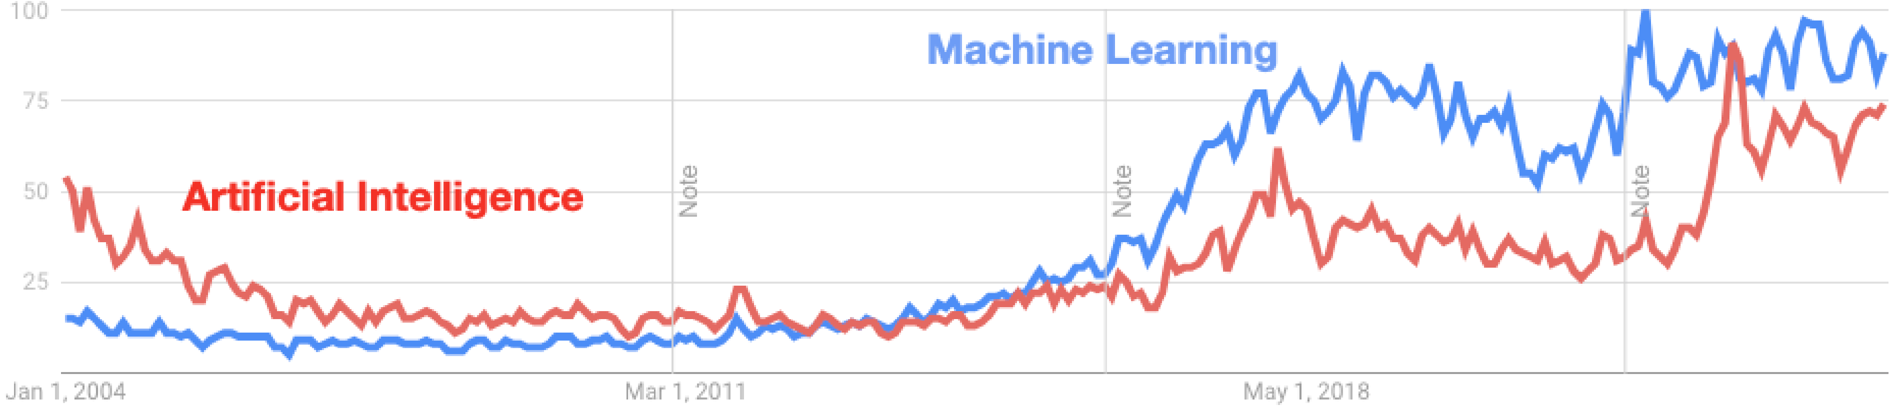

- Idea
  - collect data, maybe annotate data
  - learn patterns automatically
- Many approaches
  - is the truth known? maybe delayed? partially?
  - are we predicting a class or complex structure?
  - is the input/output continuous or discrete?
  - how much of the structure of the problem is known and can be used?
- Dominant paradigm in language and speech processing and many other fields

Big Data (since 2000s)

- Computers have became bigger
- Vast amounts of stored data available (e.g., the Internet)
- Better sensor systems allow collection of rich information about the environment

⇒ AI is big business

Deep Learning (since 2010s)

- Rediscovery of neural networks
  - initially proposed in the 1960s
  - very popular in 1980s / early 1990s ("connectionism")
- Why now? Better hardware
  - typically run on GPUs
  - 5000+ compute cores per processor
- Many toolkits available (Tensorflow, pyTorch, ...)

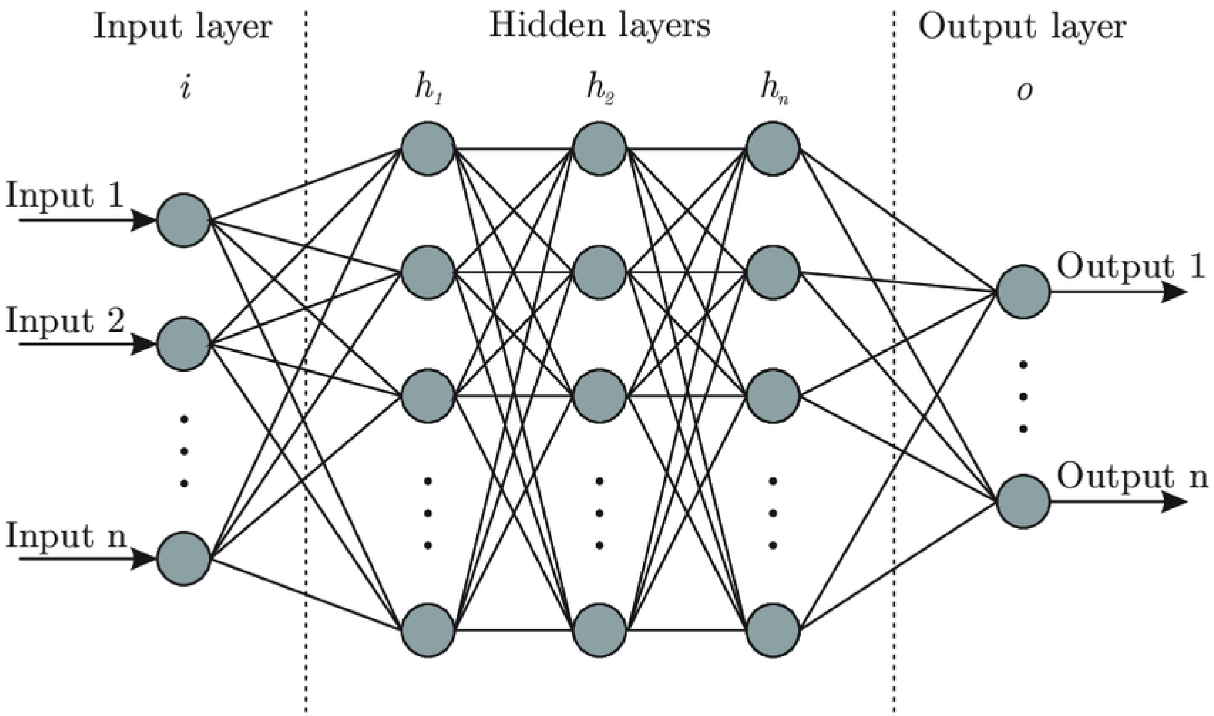

success stories

Chess

- Long considered too hard for machines: too complex, strategic thinking required
- 1997 first win of a chess computer, IBM's "Deep Blue", against reigning World Champion Garry Kasparov
- 2006, World Champion Vladimir Kramnik played Deep Fritz, lost 2–4

Machine Translation

La recette était simple. L'austérité budgétaire allait ramener la confiance, et celle-ci la croissance. On aurait l'équilibre des comptes publics et le retour de l'activité. Bon vieux bon sens britannique ! Ainsi parlait George Osborne, le péremptoire et sémillant chancelier de l'Echiquier. Il n'en n'a rien été. Depuis l'arrivée au pouvoir de David Cameron en 2010, l'impitoyable politique d'austérité budgétaire conduite à Londres n'a pas plus rétabli la croissance qu'elle n'a assaini les finances de l'Etat.

The recipe was simple. The austerity budget would restore confidence, and this growth. It would balance the public accounts and the return of the activity. Good old British common sense! Thus spoke George Osborne, the peremptory and spirited Chancellor of the Exchequer. It does not happen. Since coming to power in 2010 David Cameron, the relentless austerity policy London to conduct no more restored the growth it has cleaned up the finances of the State.

Speech Recognition

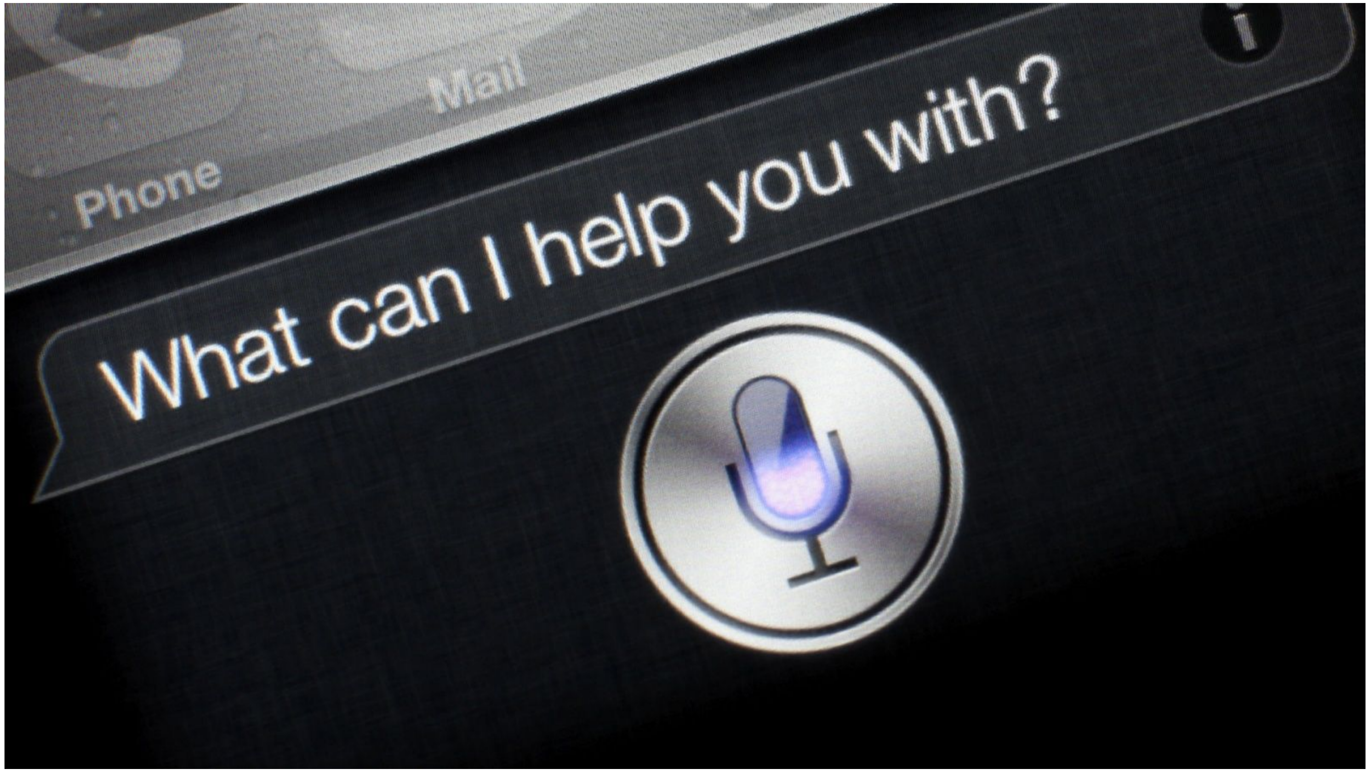

[video]

- Speech recognition has become useful enough to be a convenient input modality
- Most advanced use in personal assistants: Apple's Siri, Google Gemini, Amazon Echo, Microsoft Cartana, OpenAI's GPT

Face Recognition

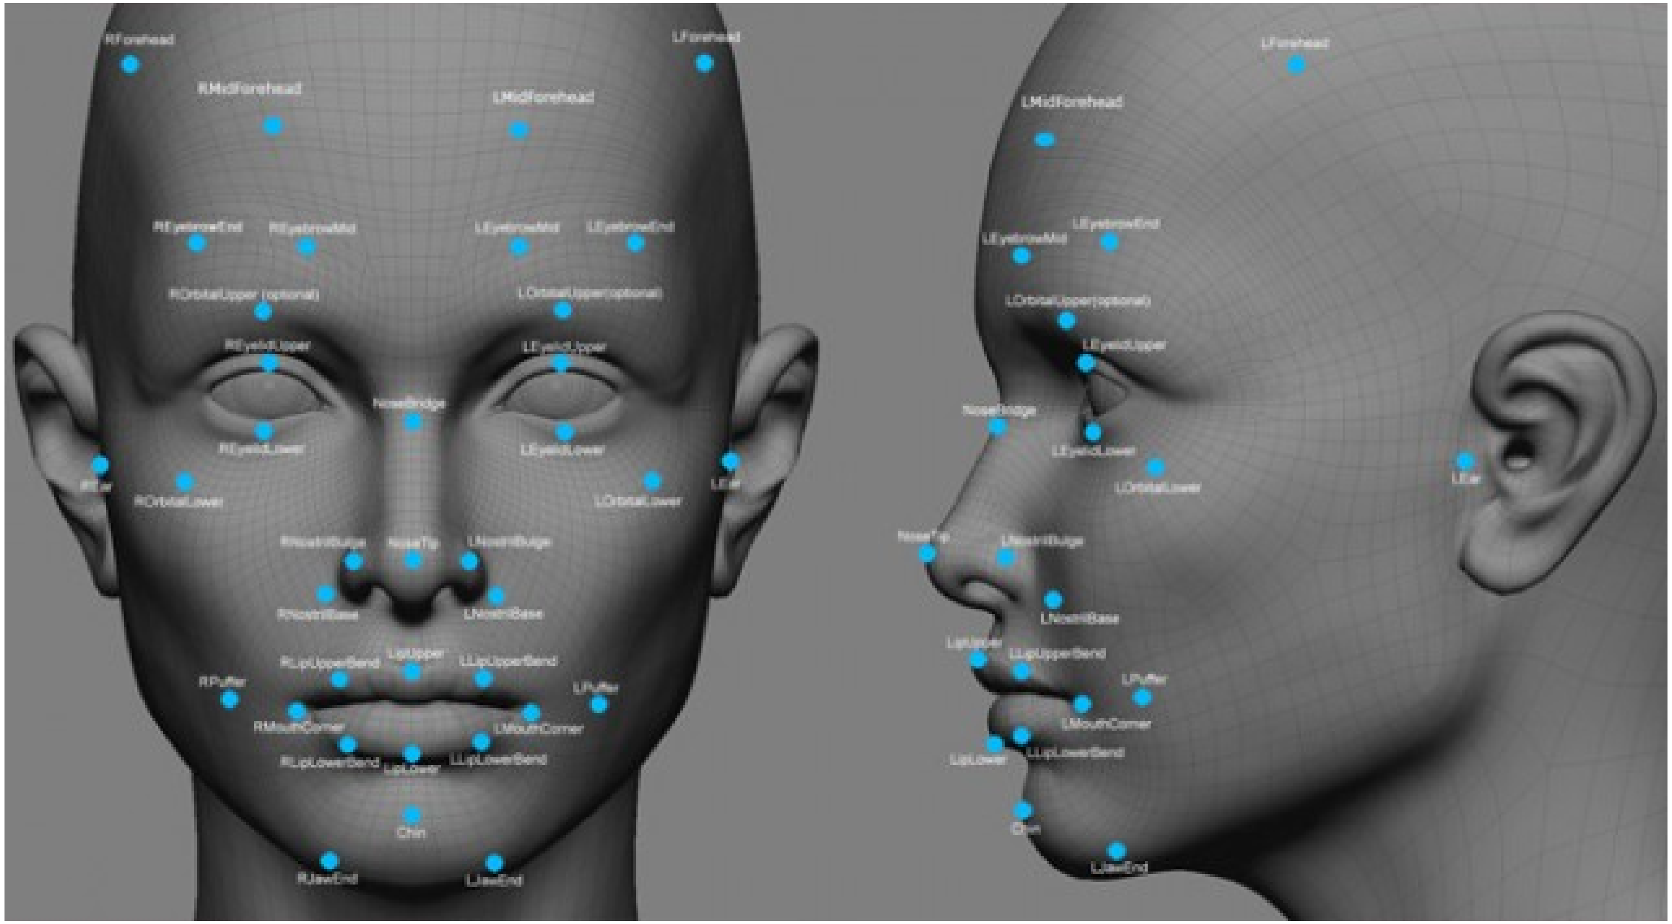

- Used by Facebook to tag photos
- Also used in border control automation

IBM Watson

- IBM built a computer that won Jeopardy in 2011
- Question answering technology built on 200 million text pages, encyclopedias, dictionaries, thesauri, taxonomies, ontologies, and other databases

Self-Driving Cars

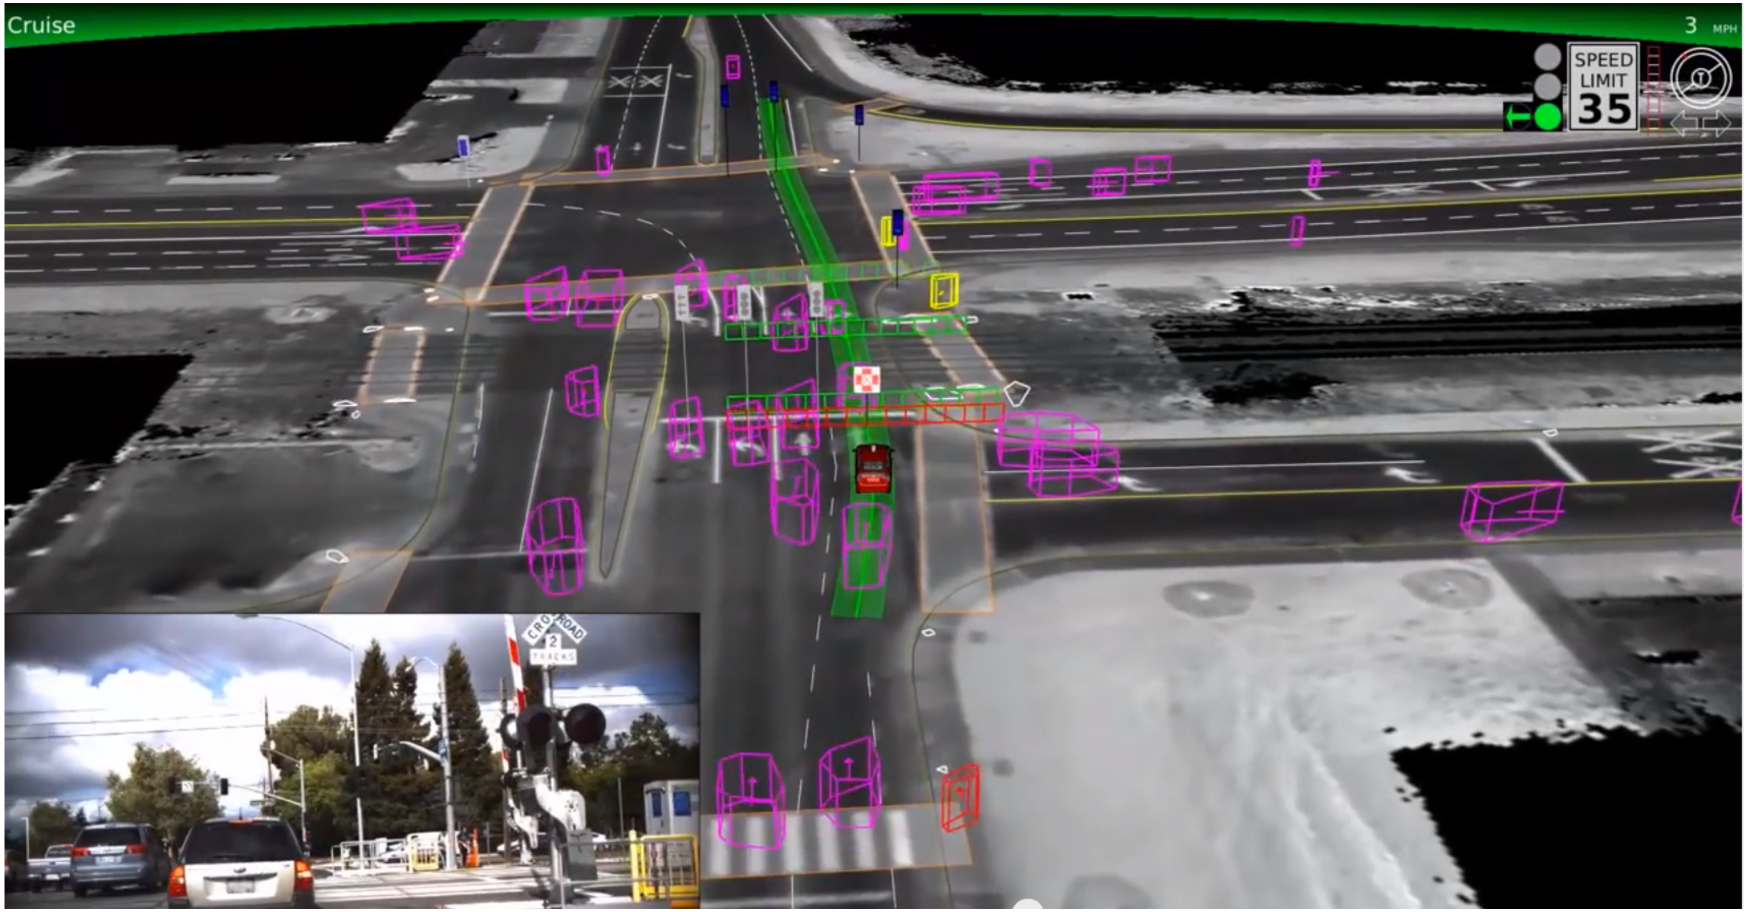

[video]

- Google self-driving car (Waymo) available as taxi service in selected cities
- Currently uses special high-precision maps

Walking Robots

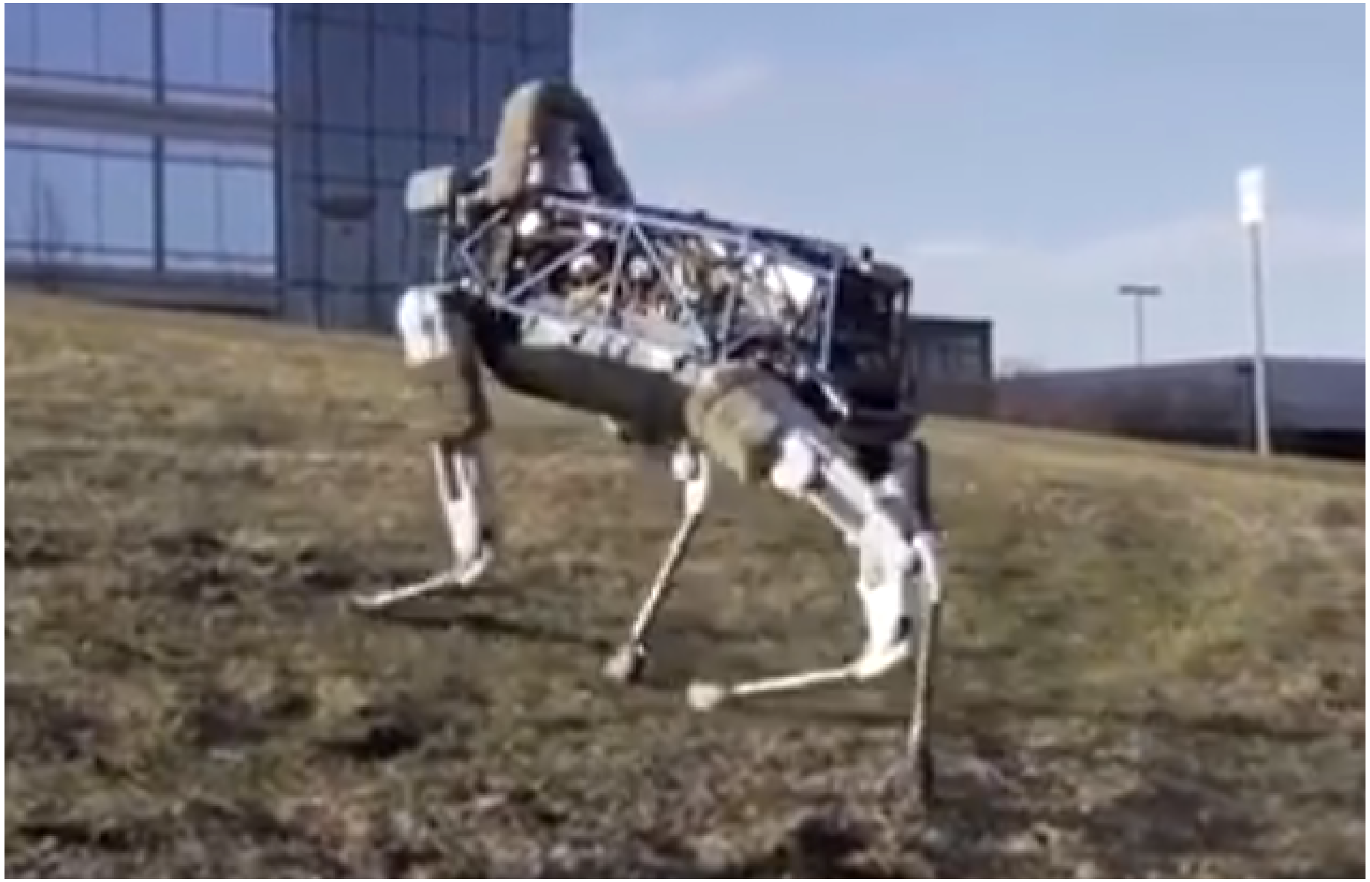

[video]

- Boston Dynamics walking robot: Spot, 2015
- Can walk across difficult terrain, stabilizes itself

Generative AI: Images

- Image manipulation
  - converting style
  - swapping faces
- Stable Diffusion:
  image from text
- Video?

Generative AI: Language

- November 2022: ChatGPT
  - explosion of use of large language models
  - bigger and better
  - still hard to separate hype from reality

- Cutting edge: generative speech and language models
  - speech synthesis
  - speech translation
  - speech chat

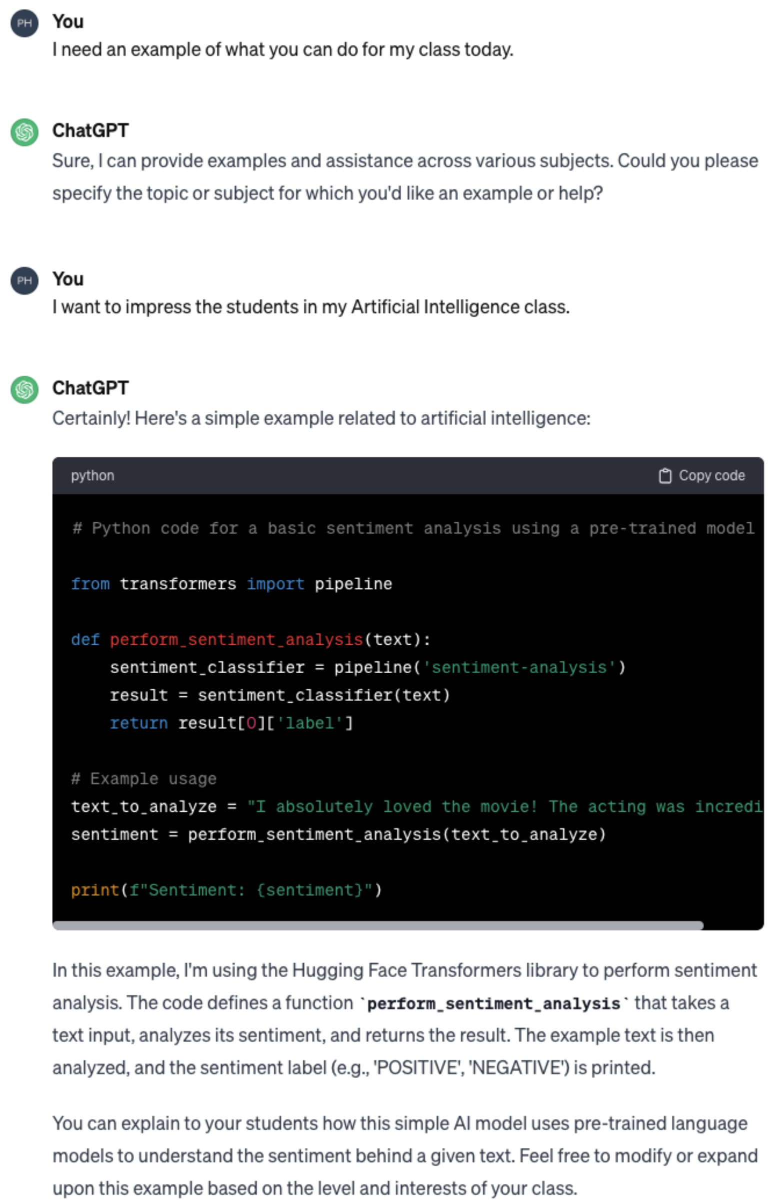

You
I need an example of what you can do for my class today.

ChatGPT
Sure, I can provide examples and assistance across various subjects. Could you please specify the topic or subject for which you'd like an example or help?

You
I want to impress the students in my Artificial Intelligence class.

ChatGPT
Certainly! Here's a simple example related to artificial intelligence:

python   Copy code

```python
# Python code for a basic sentiment analysis using a pre-trained model

from transformers import pipeline

def perform_sentiment_analysis(text):
    sentiment_classifier = pipeline('sentiment-analysis')
    result = sentiment_classifier(text)
    return result[0]['label']

# Example usage
text_to_analyze = "I absolutely loved the movie! The acting was incredi
sentiment = perform_sentiment_analysis(text_to_analyze)

print(f"Sentiment: {sentiment}")
```

In this example, I'm using the Hugging Face Transformers library to perform sentiment analysis. The code defines a function `perform_sentiment_analysis` that takes a text input, analyzes its sentiment, and returns the result. The example text is then analyzed, and the sentiment label (e.g., 'POSITIVE', 'NEGATIVE') is printed.

You can explain to your students how this simple AI model uses pre-trained language models to understand the sentiment behind a given text. Feel free to modify or expand upon this example based on the level and interests of your class.