# Artificial Intelligence

CSE 120 Spring 2017

Slide credits: Pieter Abbeel, Dan Klein, Stuart Russell, **Pat Virtue** & http://csillustrated.berkeley.edu

**Instructor:** Justin Hsia

**Teaching Assistants:** Anupam Gupta, Braydon Hall, Eugene Oh, Savanna Yee

## Candy Grab Game

1) Grab a pack of "game pieces" and a suggested Agent card
2) Play the following game
   a) 10 pieces on the table
   b) Take turns taking either 1 or 2 pieces
   c) Player that takes the last piece(s) wins ☺
3) How do humans learn to play this game?
4) How would a computer learn to play this game?

```
int takeTurn(int numPiecesAvailable)
```

# Administrivia

- Assignments:
  - Tic-Tac-Toe due tonight (5/19)
  - Project Proposal due Saturday (5/20)
  - Innovation Exploration post due Sunday (5/21)

- Project:
  - Dream big! It's okay to scale back later
    - You might be surprised at what you can accomplish
  - Feel free to brainstorm with a staff member
  - Make it something that you're excited about

# Outline

- **What is AI?**
- AI History
  - AI winter and the resurgence
- AI Today
  - Applications and how they work
- AI Tomorrow
  - Ethics and the singularity

# Artificial Intelligence

- Why learn about AI?

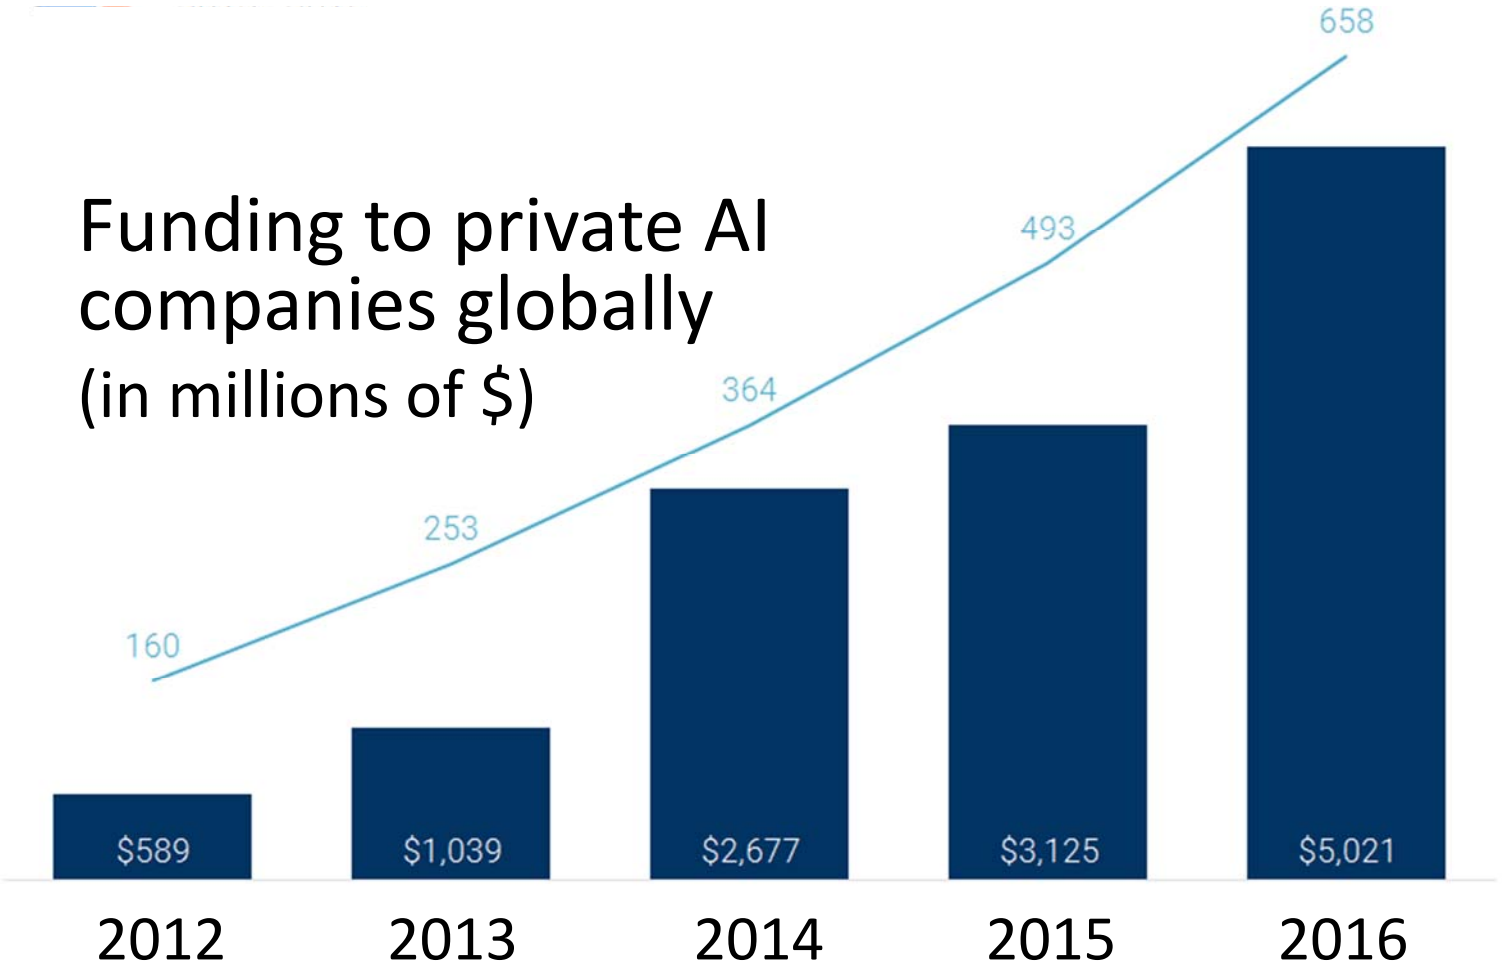


https://www.cbinsights.com/blog/artificial-intelligence-startup-funding/

# Candy Grab Game

```
int takeTurn(int numPiecesAvailable) {
    return ?;



}
```

# Candy Grab Game

- Agent 001 – always choose 1

```
int takeTurn(int numPiecesAvailable) {
    return 1;



}
```

# Candy Grab Game

- Agent 002 – always choose 2

```
int takeTurn(int numPiecesAvailable) {
    return 2;



}
```

# Candy Grab Game

- Agent 007 – whatever you think is best

```
int takeTurn(int numPiecesAvailable) {
    return ?;



}
```

# Candy Grab Game

- Agent 007 – whatever you think is best

```
int takeTurn(int numPiecesAvailable) {
    if(numPiecesAvailable%3 == 2) {
        return 2;
    } else {
        return 1;
    }

}
```

# Candy Grab Game

- Agent 007 – whatever you think is best

```
int takeTurn(int numPiecesAvailable) {
    return ?;
}
```

  - What if we had:

```
boolean willIWin(int numPiecesAvailable,
                 boolean isMyTurn,
                 int action)
```

# Candy Grab Game

- Agent 007 – whatever you think is best

```
int takeTurn(int numPiecesAvailable) {
    return ?;
}
```

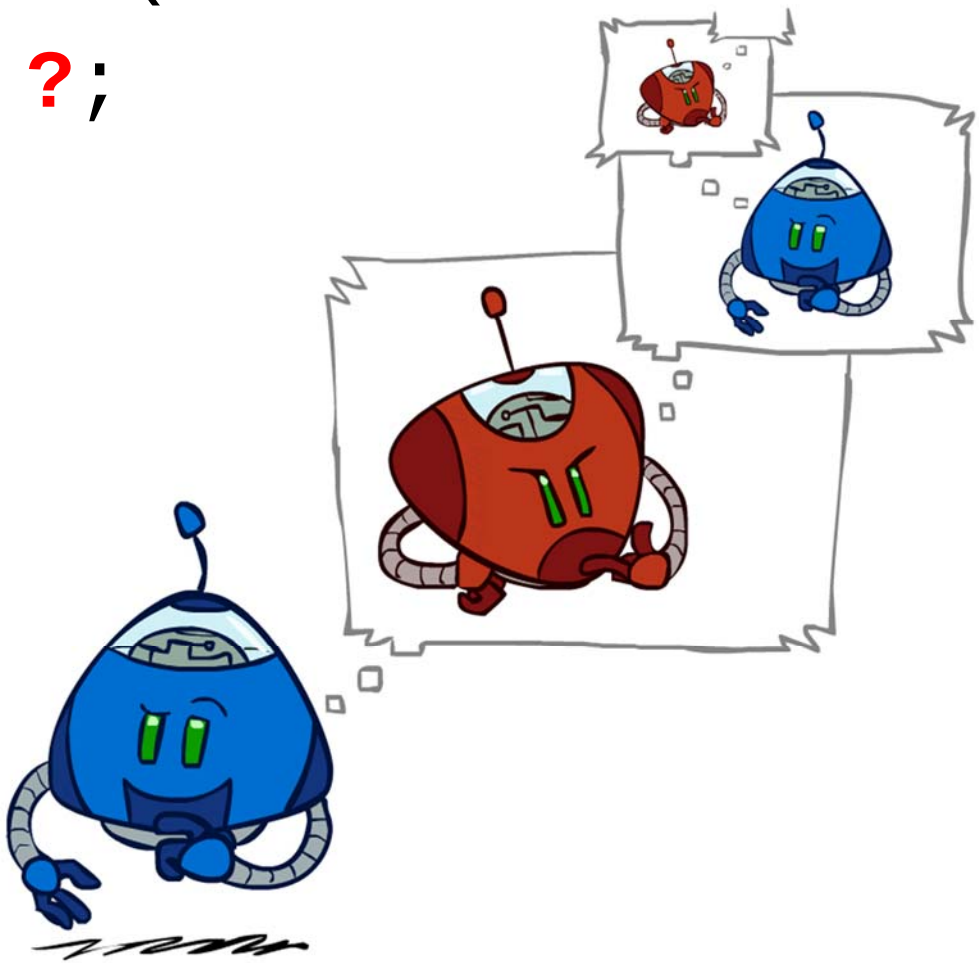

# How: Search & Planning

- I take an action… then what? … then what?

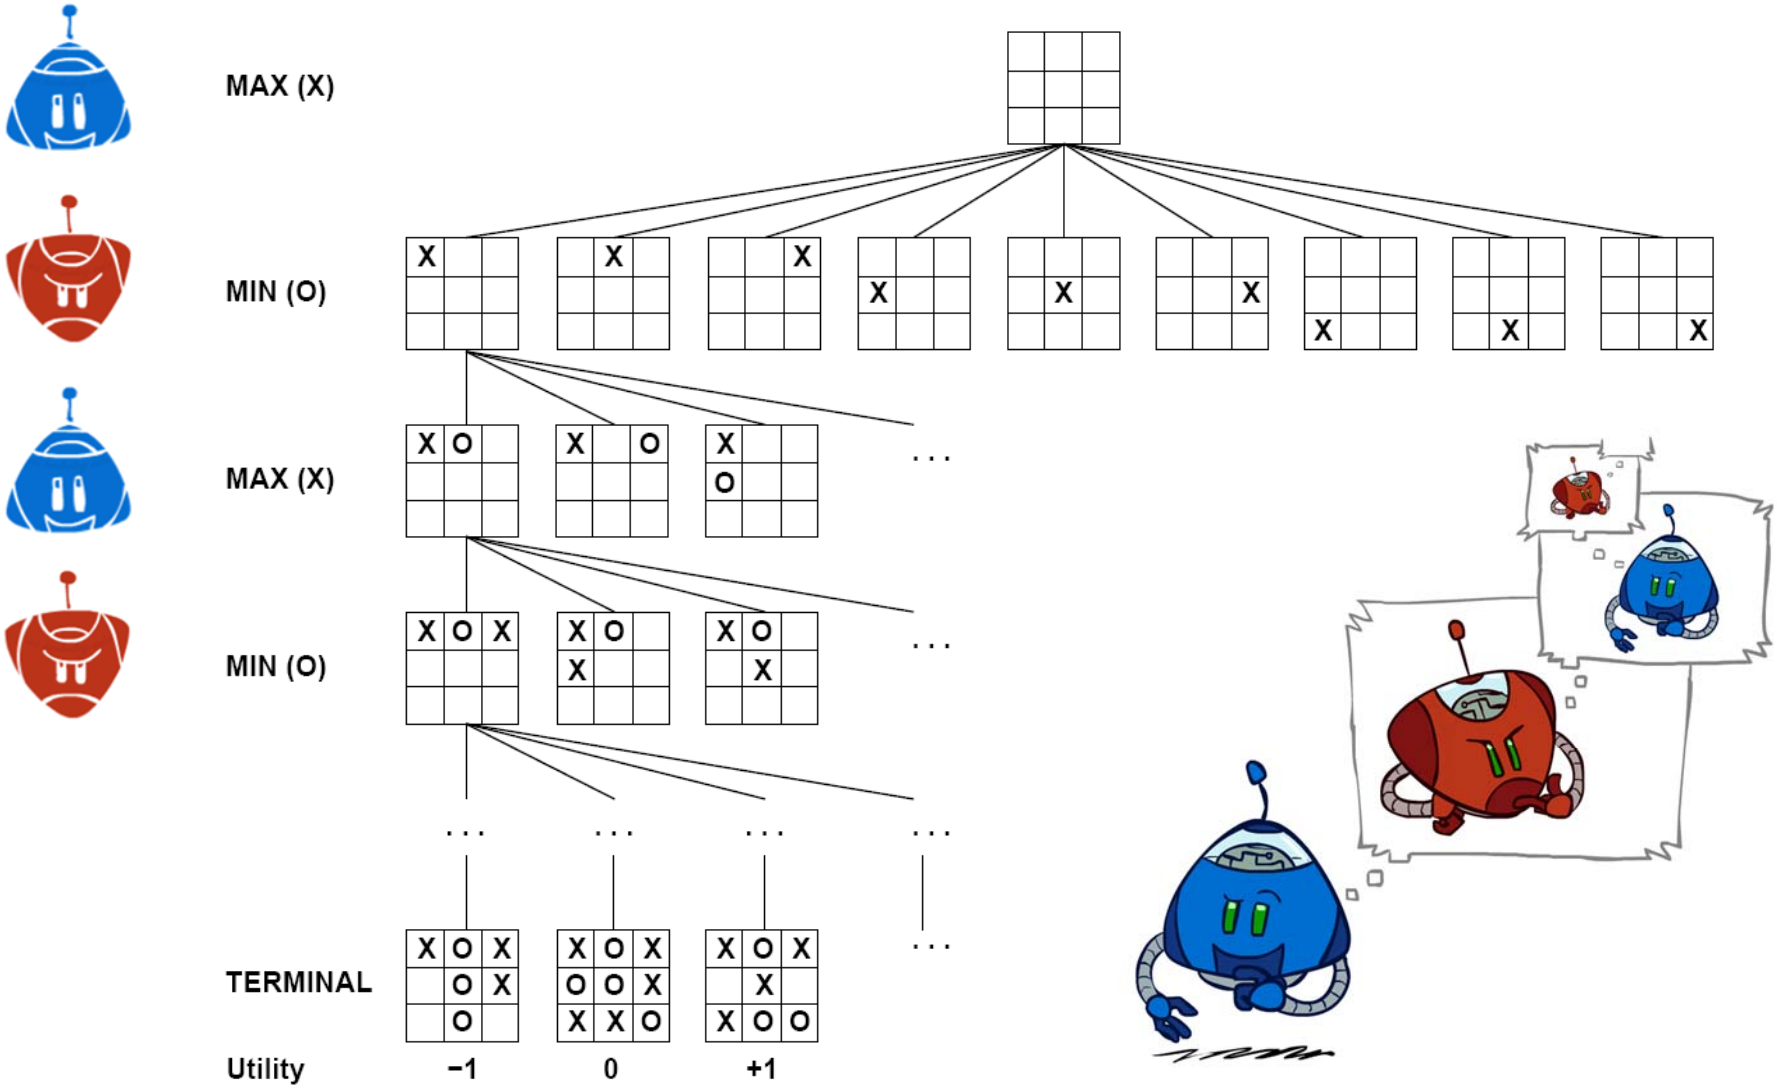

# Candy Grab Game (Min-Max)

- Agent 007 – whatever you think is best

```
int takeTurn(int numPiecesAvailable)
```

# Candy Grab Game (Statistical)

- Agent 007 – whatever you think is best

```
int takeTurn(int numPiecesAvailable)
```

| Pieces Available | Take 1 | Take 2 |
|---|---|---|
| 2 | 0% | 100% |
| 3 | 2% | 1% |
| 4 | 75% | 2% |
| 5 | 4% | 68% |
| 6 | 5% | 6% |

# AI Games in the News


- https://www.youtube.com/watch?v=EfGD2qveGdQ
- https://www.youtube.com/watch?v=cjpEIotvwFY

# So What is AI?

# Sci-Fi AI?

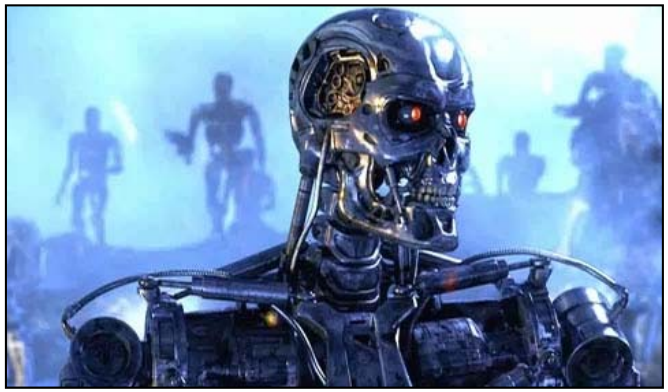

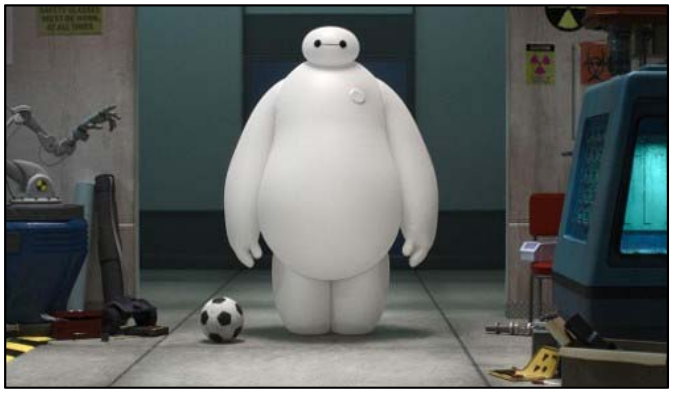

# So What is AI?

- The science of making machines that:

**Think like people**

**Act like people**

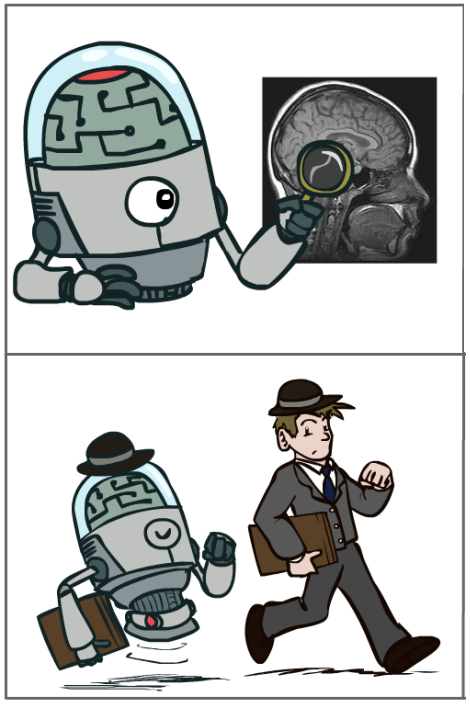

# Turing Test for Intelligence

- In 1950, Turing defined a test of whether a machine could "think":
  - "A human judge engages in a natural language conversation with one human and one machine, each of which tries to appear human. If judge can't tell, machine passes the Turing test."

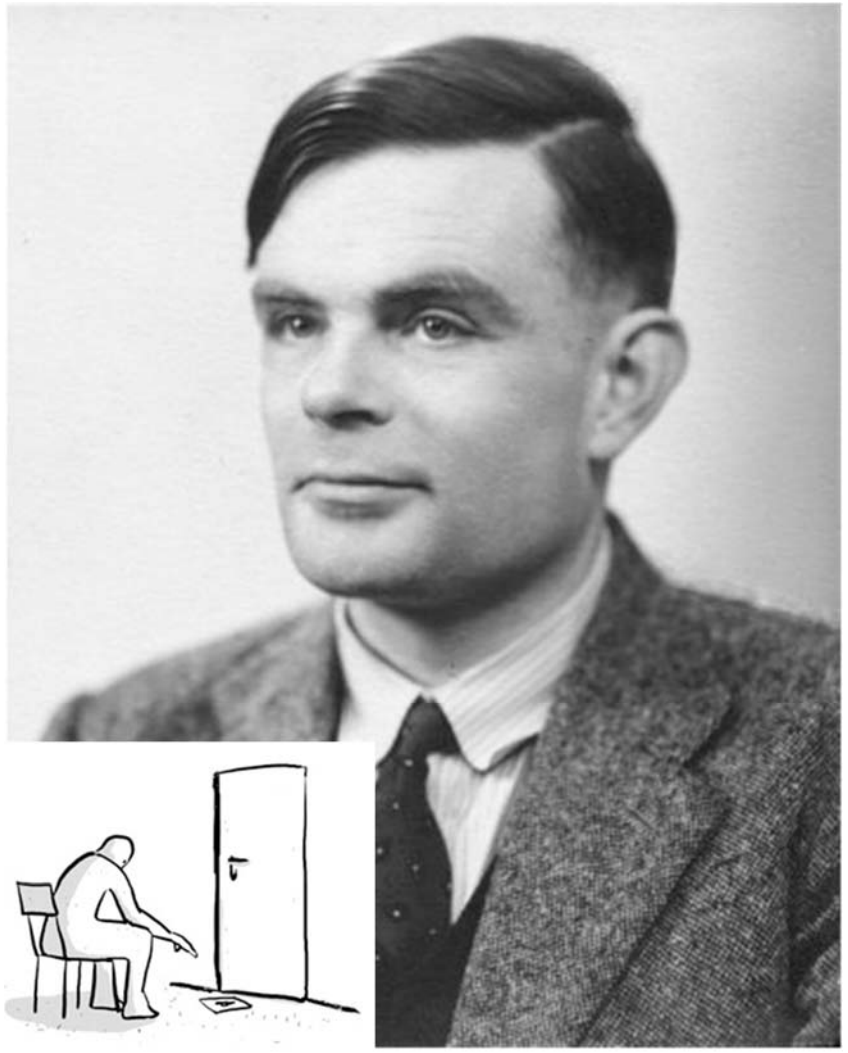

http://en.wikipedia.org/wiki/Turing_test

# So What is AI?

- The science of making machines that:

Think like people

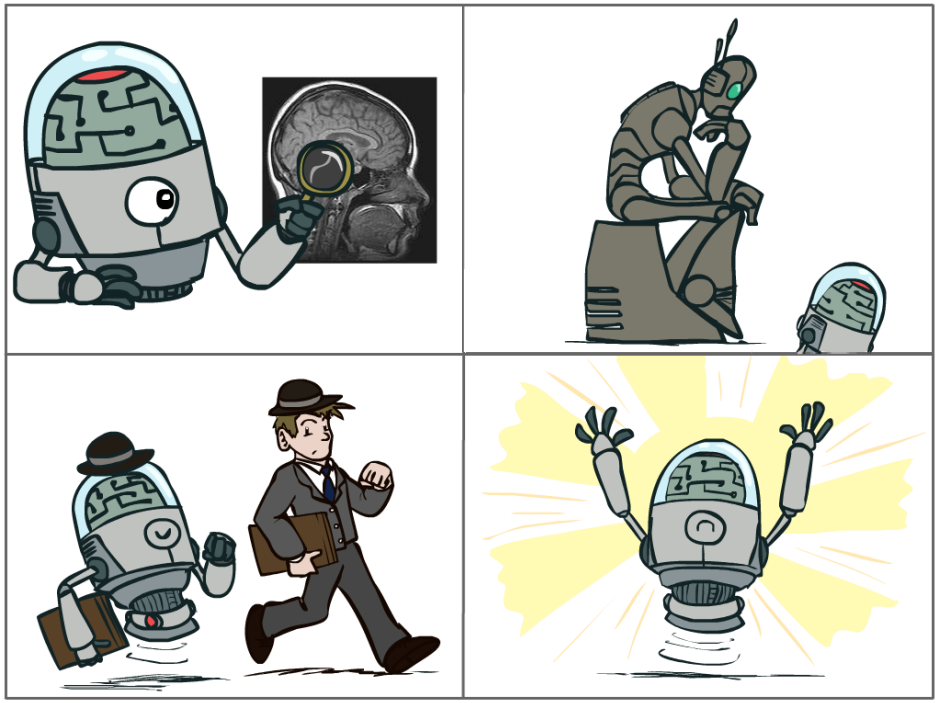

Think rationally

Act like people

Act rationally

# AI Definition by John McCarthy

- "Getting a computer to do things which, when done by people, are said to involve intelligence"
  - Finesses the idea of whether a computer has consciousness, whether they have rights, etc.

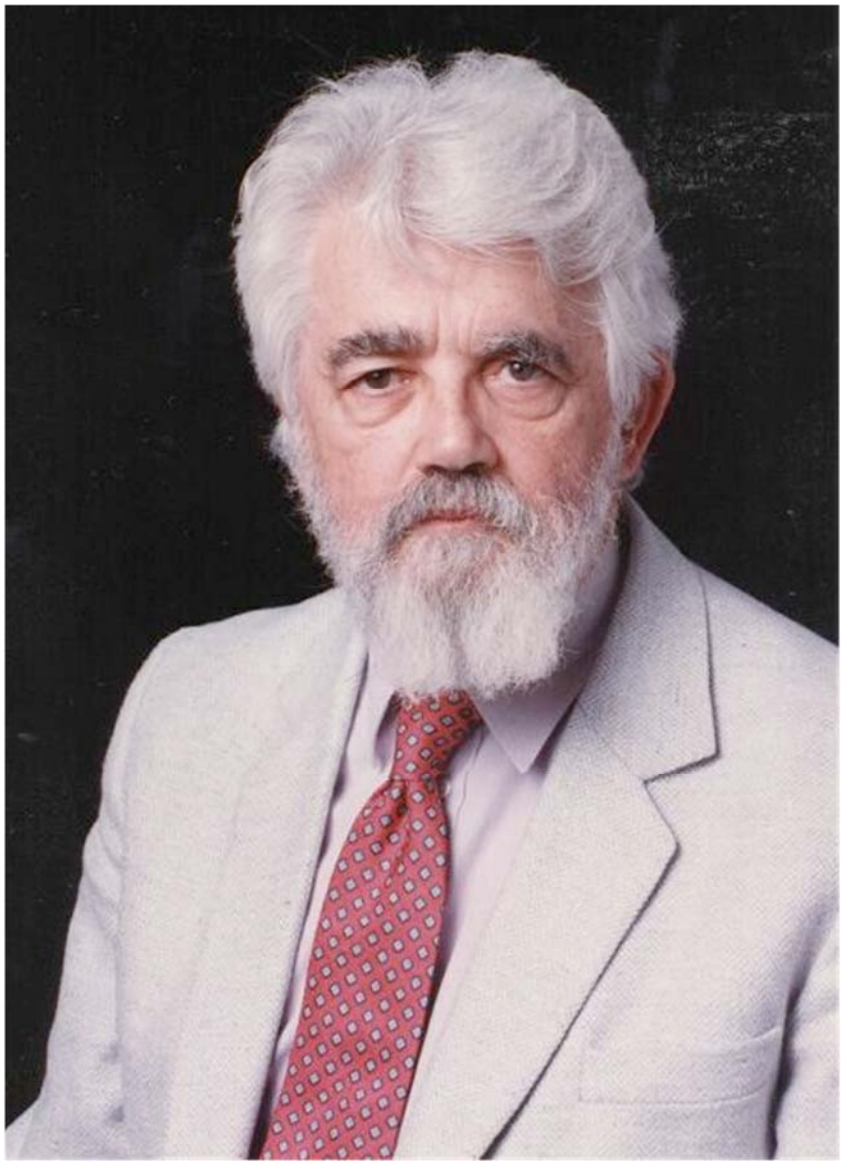

# Outline

- What is AI?
- **AI History**
  - **AI winter and the resurgence**
- AI Today
  - Applications and how they work
- AI Tomorrow
  - Ethics and the singularity

# A Brief History of AI

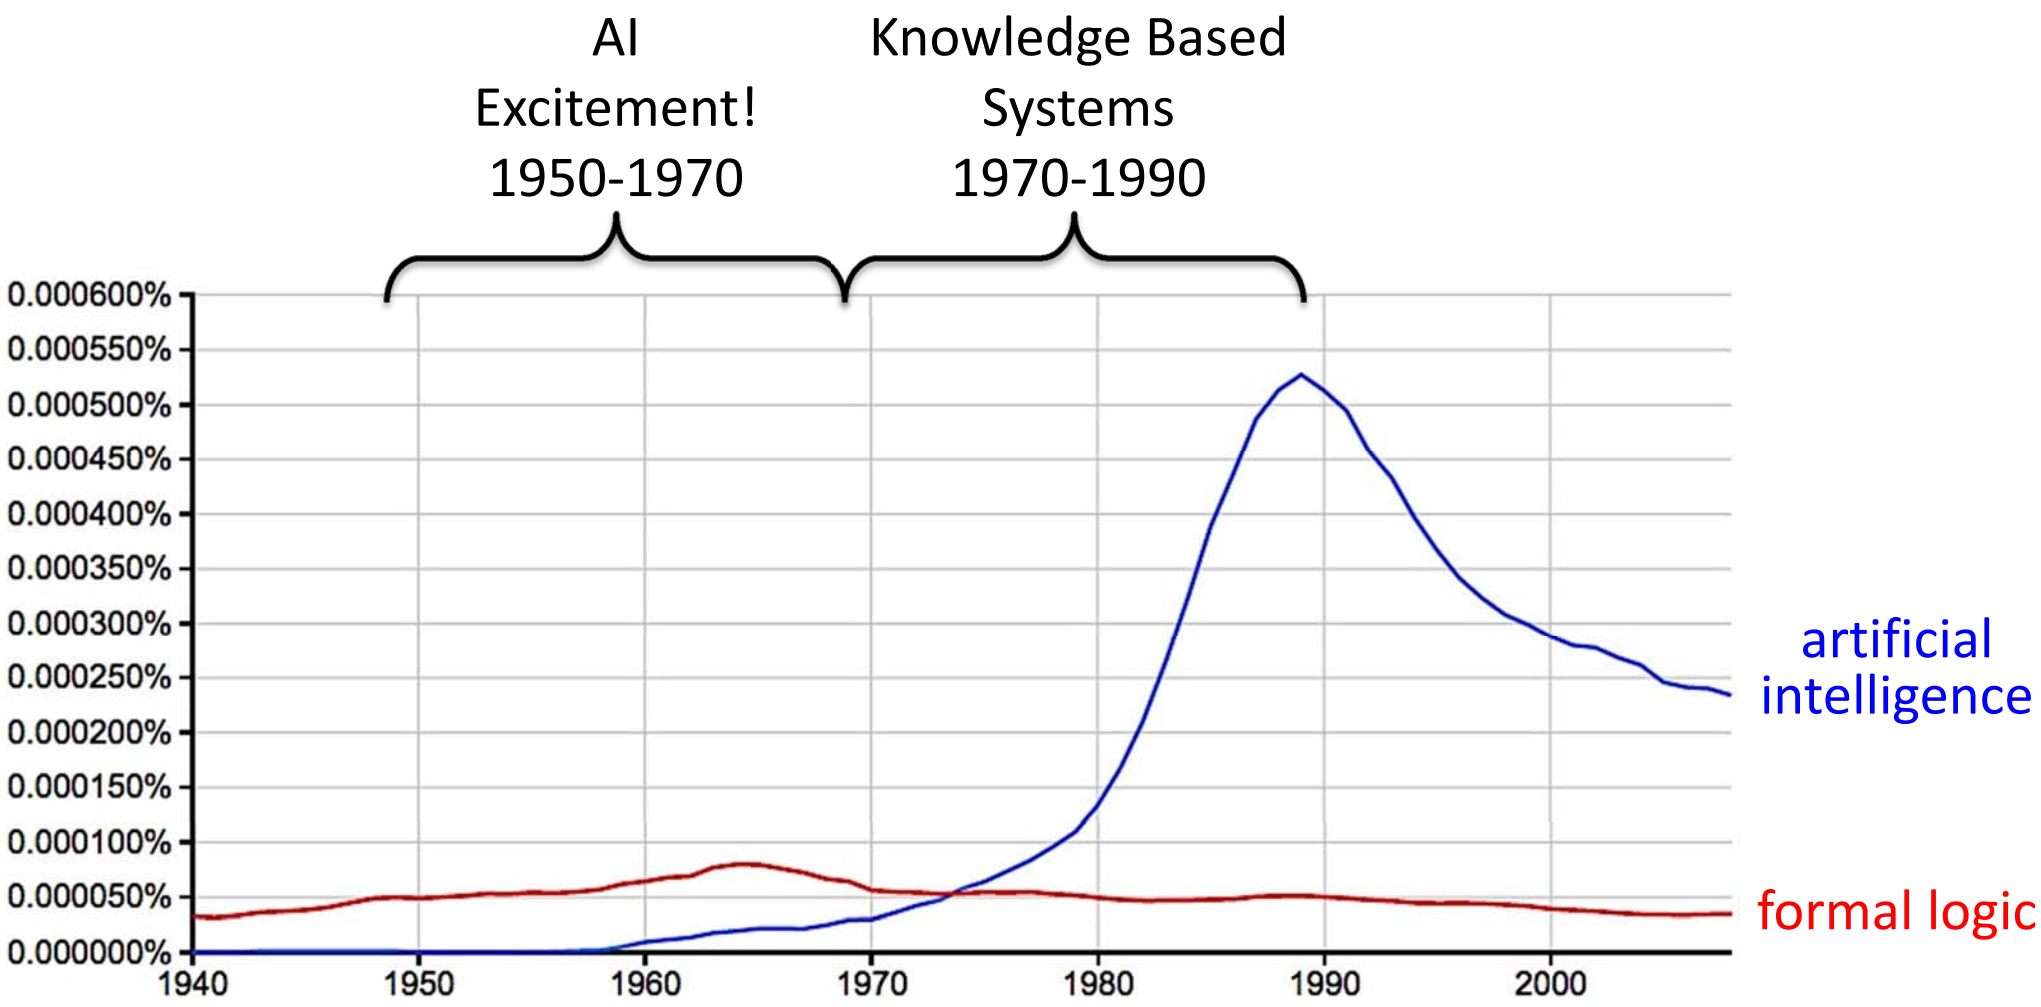


Google Books Ngram Viewer: https://books.google.com/ngrams

# Logic

- A formal representation of our knowledge of the world
- Use **knowledge base** and **perception** to infer new knowledge

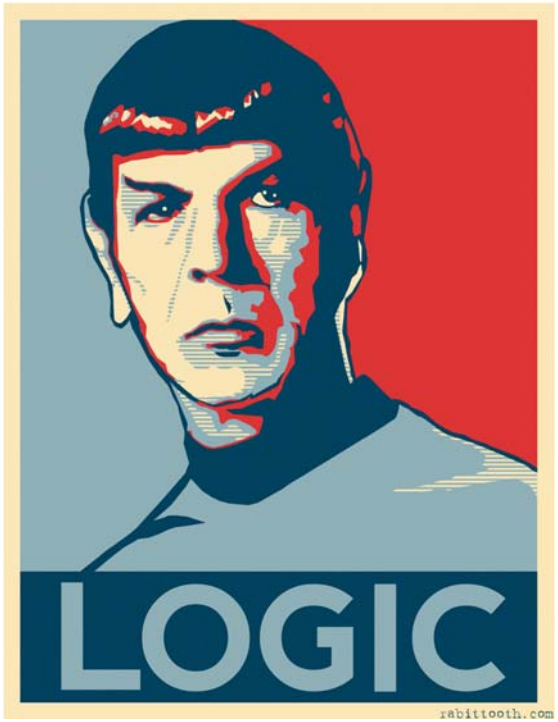


http://www.rabittooth.com/gallery/spock_in_2012_by_rabittooth.jpg

# isDog(animal)

- $barks \ \& \ fur \ \& \ fourLegs \Leftrightarrow dog$

```
boolean isDog(animal A) {
  if(!barks(A)) {
    return false;
  }
  if(!hasFur(A)) {
    return false;
  }
  if(!hasFourLegs(A)) {
    return false;
  }
  return true;
}
```

What's the problem?

Dog

- Barks
- Has Fur
- Has four legs

# isDog(animal)

- $barks \ \& \ fur \ \& \ fourLegs \Leftrightarrow dog$

What's the problem?

Dog — Sheila

- Barks ✓
- Has Fur ✓
- Has four legs ✗

# A Brief History of AI

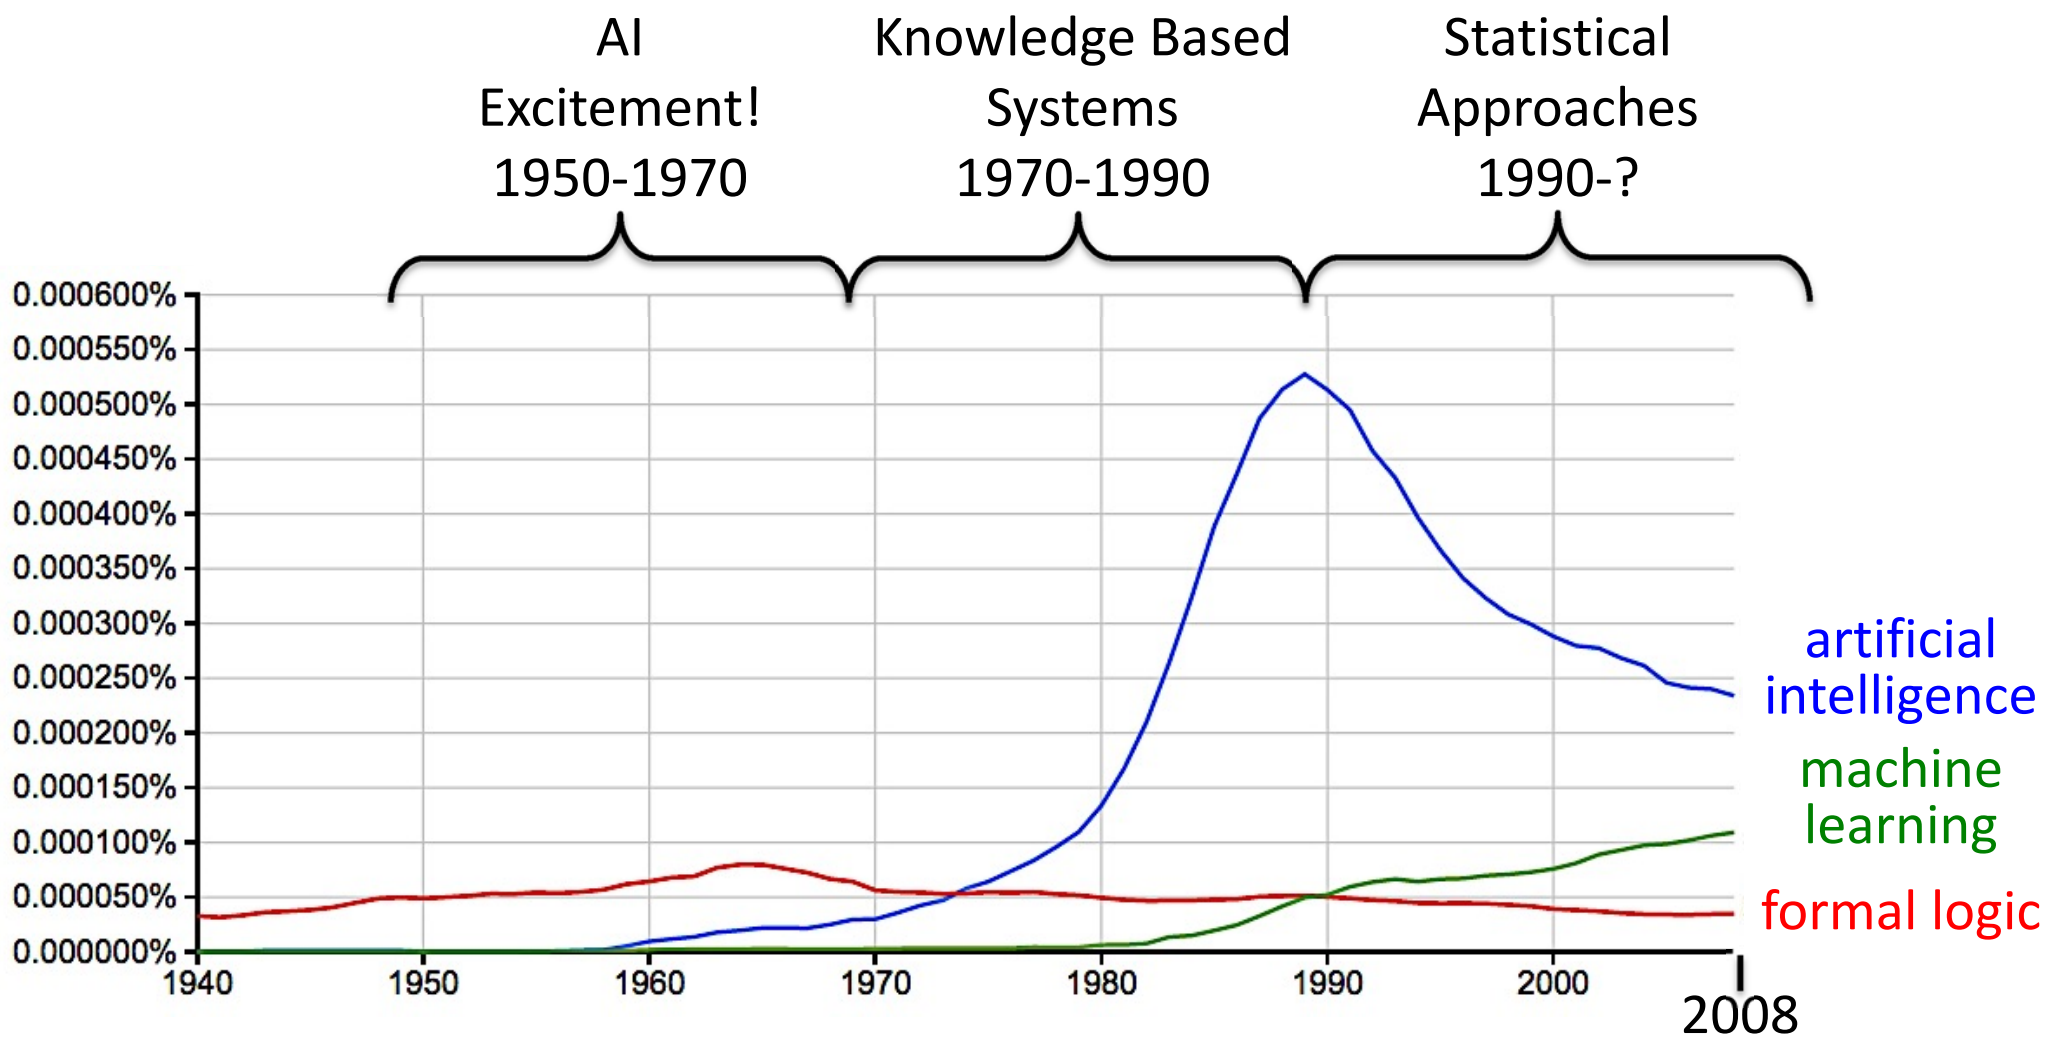


Google Books Ngram Viewer: https://books.google.com/ngrams

# Outline

- What is AI?
- AI History
  - AI winter and the resurgence
- **AI Today**
  - **Applications and how they work**
- AI Tomorrow
  - Ethics and the singularity

# "Intelligent" Applications

- Discuss in pairs/groups:
  - List at least four existing applications that seem intelligent
- Audience responses:

# Applications: Natural Language

- Speech technologies (*e.g.* Siri)
  - Automatic speech recognition (ASR)
  - Text-to-speech synthesis (TTS)
  - Dialog systems
- Language processing technologies
  - Google translation
  - Web search
  - Spam filter

# How: Probability

**Notation:**

$P(\text{limb} \mid \text{artificial}, \text{audio})$

"Probability of 'limb' given 'artificial' and audio"

- Example: speech recognition of "artificial …"
  - Find most probable next word given "artificial" and the audio for the second word

Which second word gives the highest probability?

Break down problem

n-gram probability * audio probability

$P(\textbf{limb} \mid \text{artificial}, \text{audio})$

$P(\textbf{limb} \mid \text{artificial}) * P(\textbf{limb} \mid \text{audio})$

$P(\textbf{intelligence} \mid \text{artificial}, \text{audio})$

$P(\textbf{intelligence} \mid \text{artificial}) * P(\textbf{intelligence} \mid \text{audio})$

$P(\textbf{flavoring} \mid \text{artificial}, \text{audio})$

$P(\textbf{flavoring} \mid \text{artificial}) * P(\textbf{flavoring} \mid \text{audio})$

# Applications: Vision (Perception)

- Tasks related to understanding images/camera input

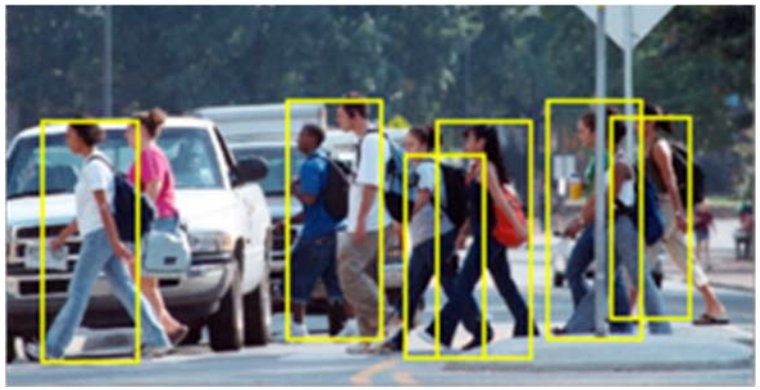

Pedestrian Detection

Segmentation

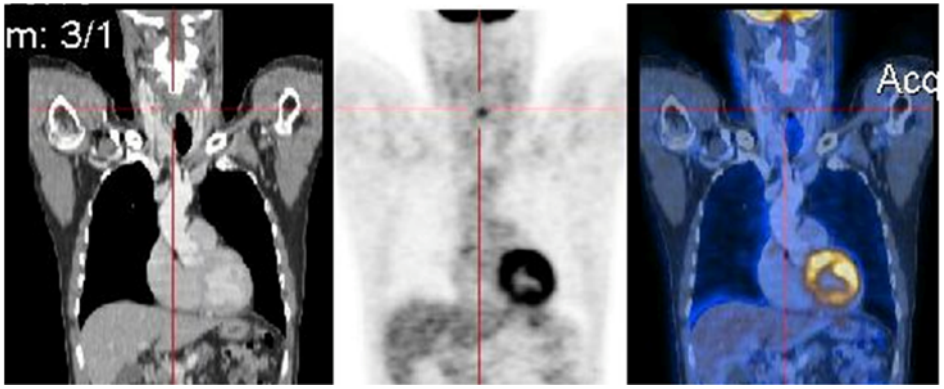


Alignment/Registration

Long, Shelhamer, Darrell. arXiv preprint arXiv:1411.4038 (2014).

https://en.wikipedia.org/wiki/Medical_image_computing#/media/File:CT-PET.jpg

# How: Neural Networks

Input Signal

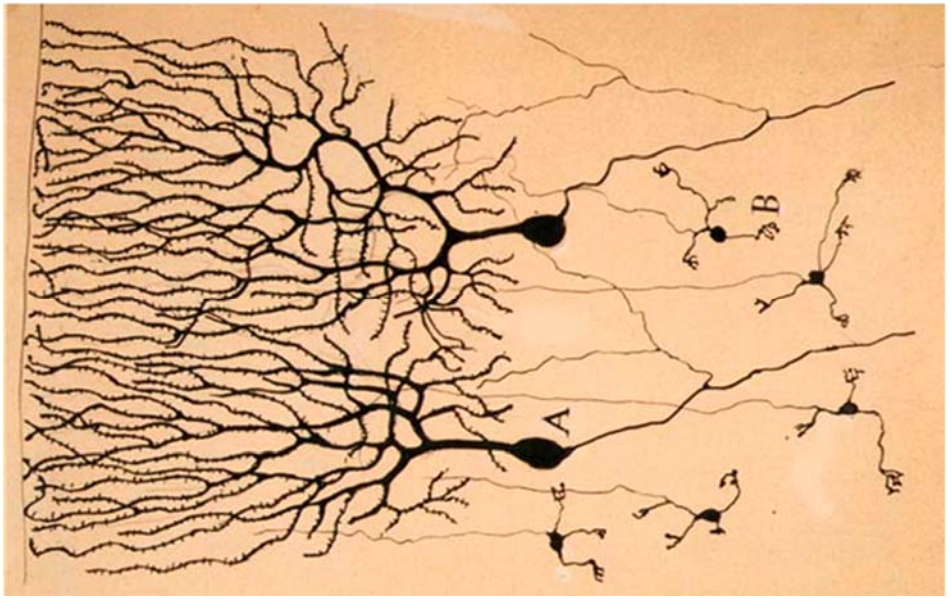


Output Signal

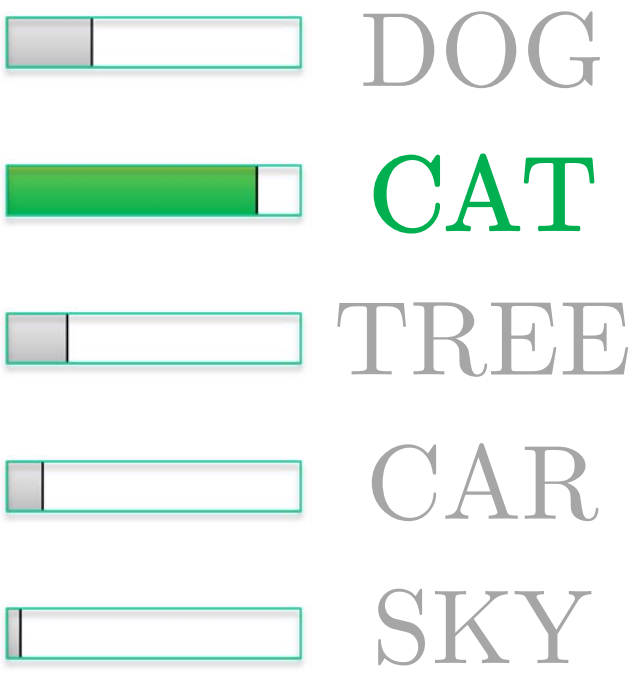


Image: https://en.wikipedia.org/wiki/Neuron

# How: Neural Networks

Input Signal

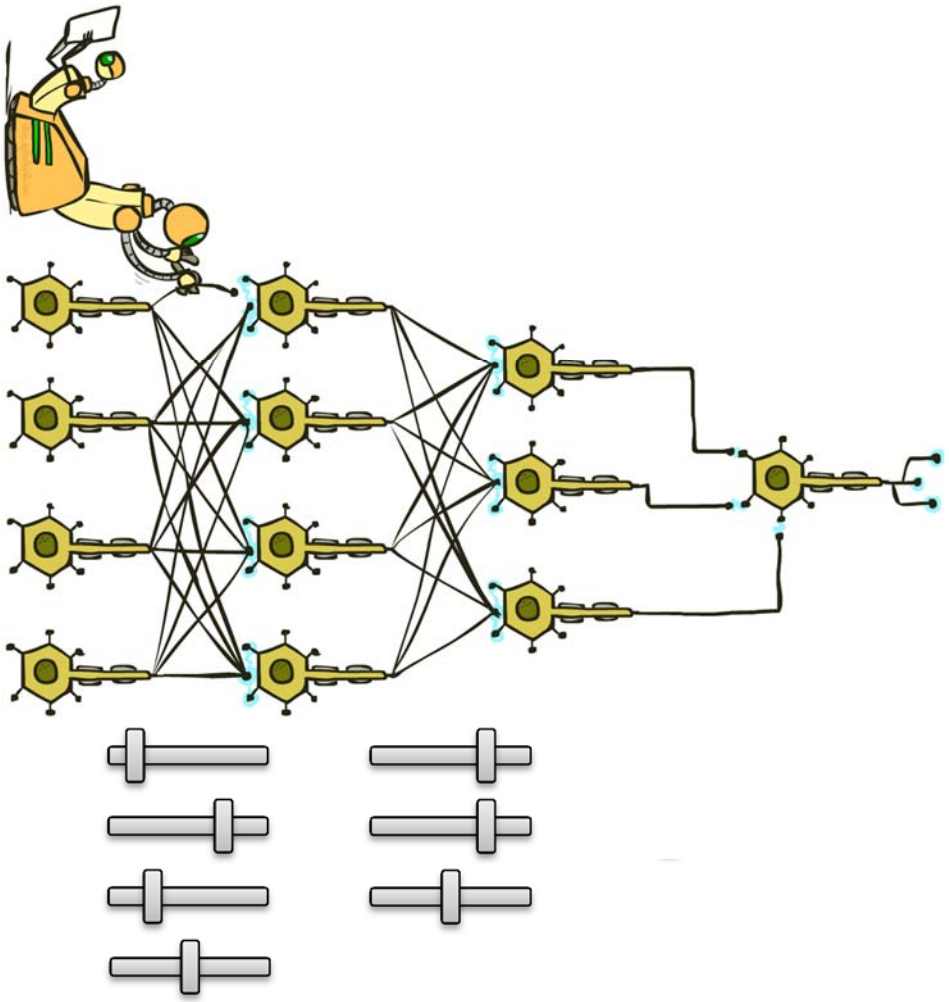

Output Signal

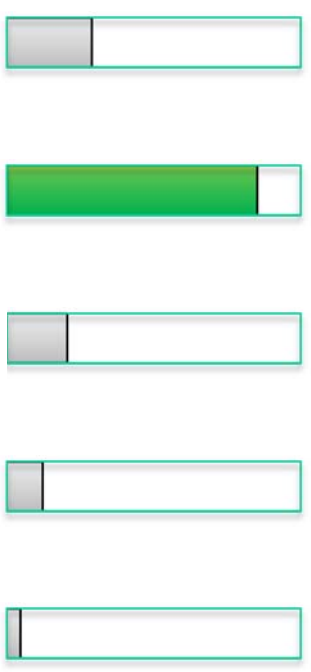


Image: https://en.wikipedia.org/wiki/Neuron

# Applications: Games

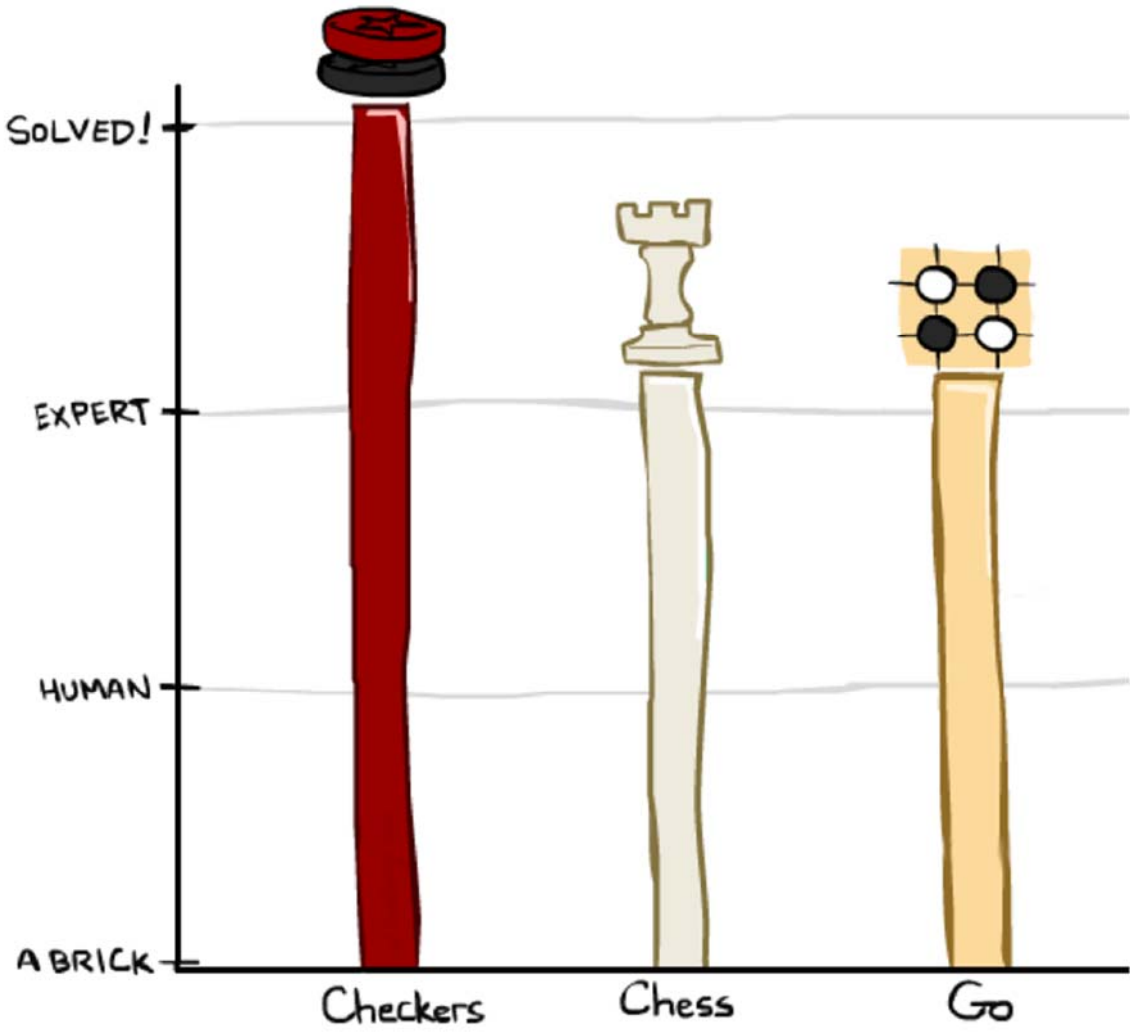

# How: Search & Planning

- I take an action… then what? … then what?

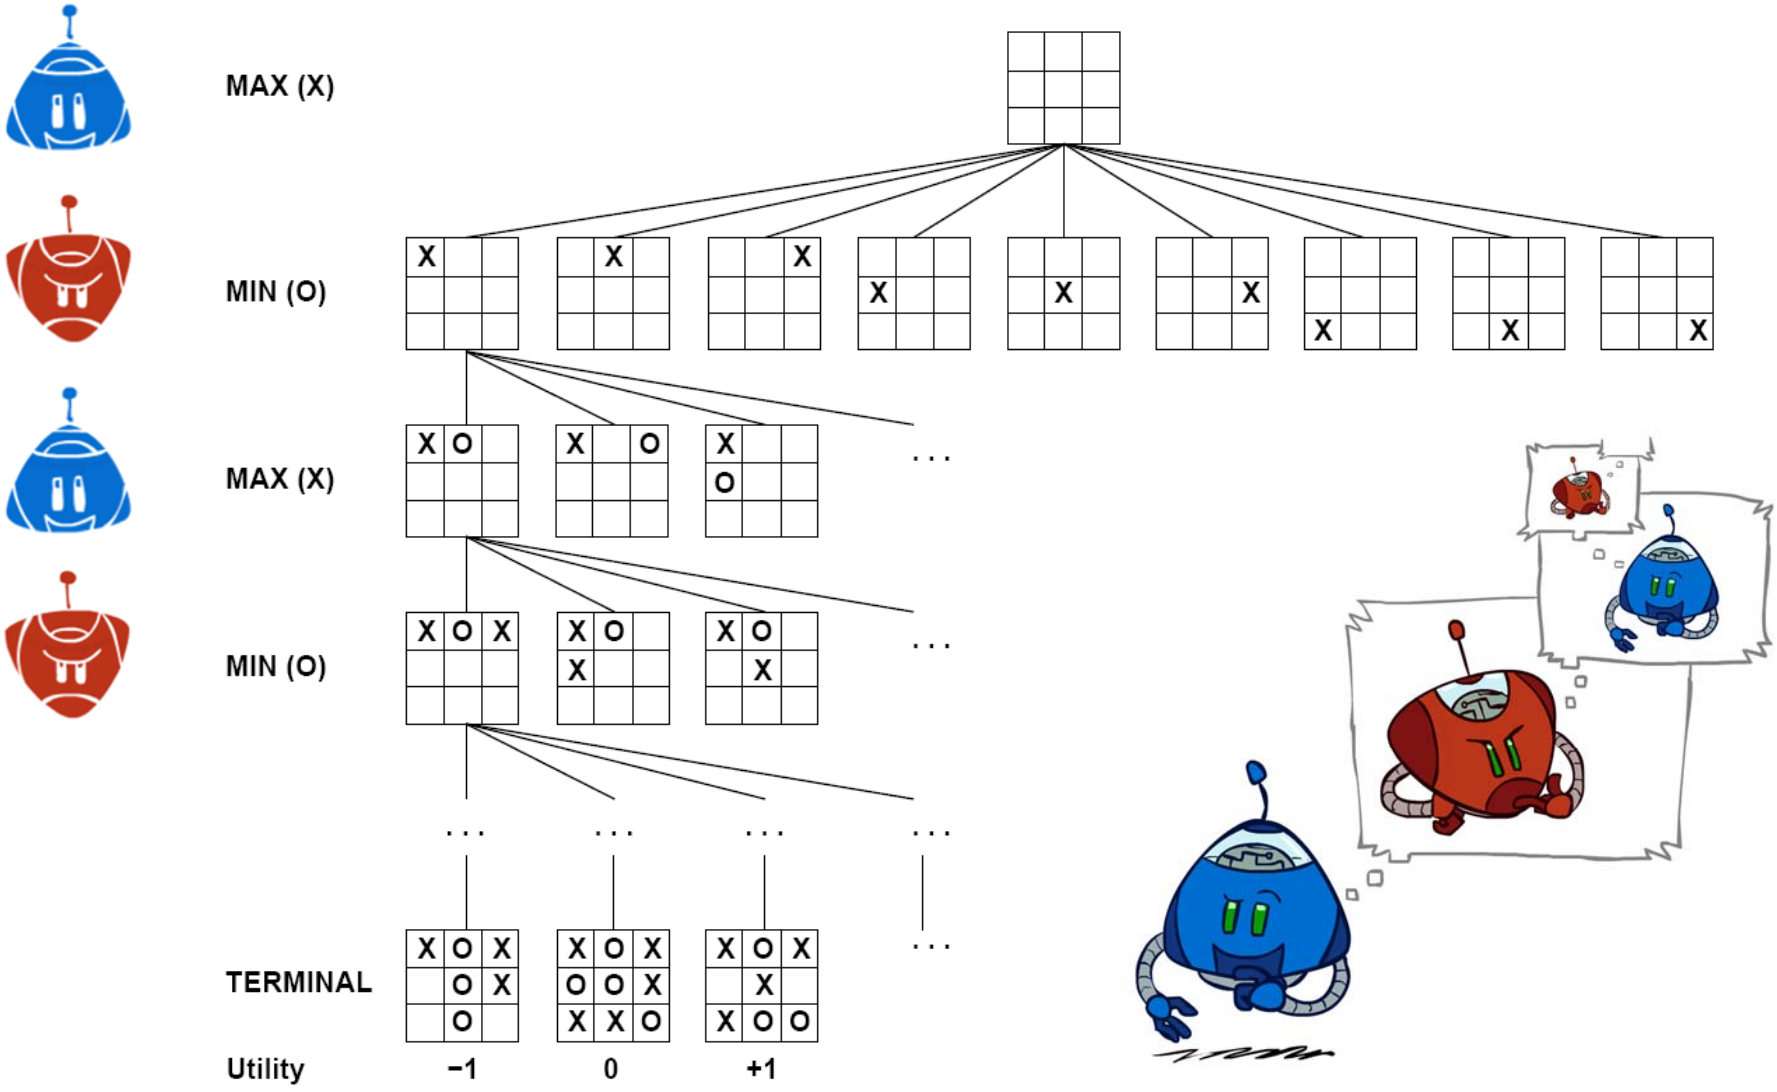

# Applications: Robotics

- For many, the coolest and scariest part of AI
- Combines fields of AI/CS
  - Speech recognition
  - Synthetic voice
  - Machine vision
  - Planning
  - HCI

Surgical robots

Autonomous helicopter

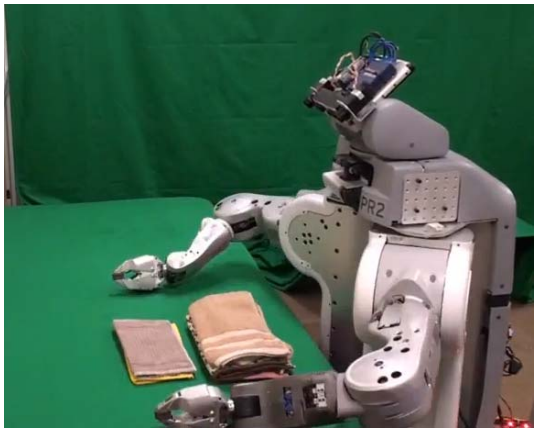

Towel-folding!

TOPIO, the ping-pong playing robot

# Applications: Robotics

- Video: Bipedal Robot – Boston Dynamics
  - https://www.youtube.com/watch?v=rVlhMGQgDkY
- Video: Robot Preschool – UC Berkeley
  - http://www.bloomberg.com/features/2015-preschool-for-robots/

# Applications: Driving

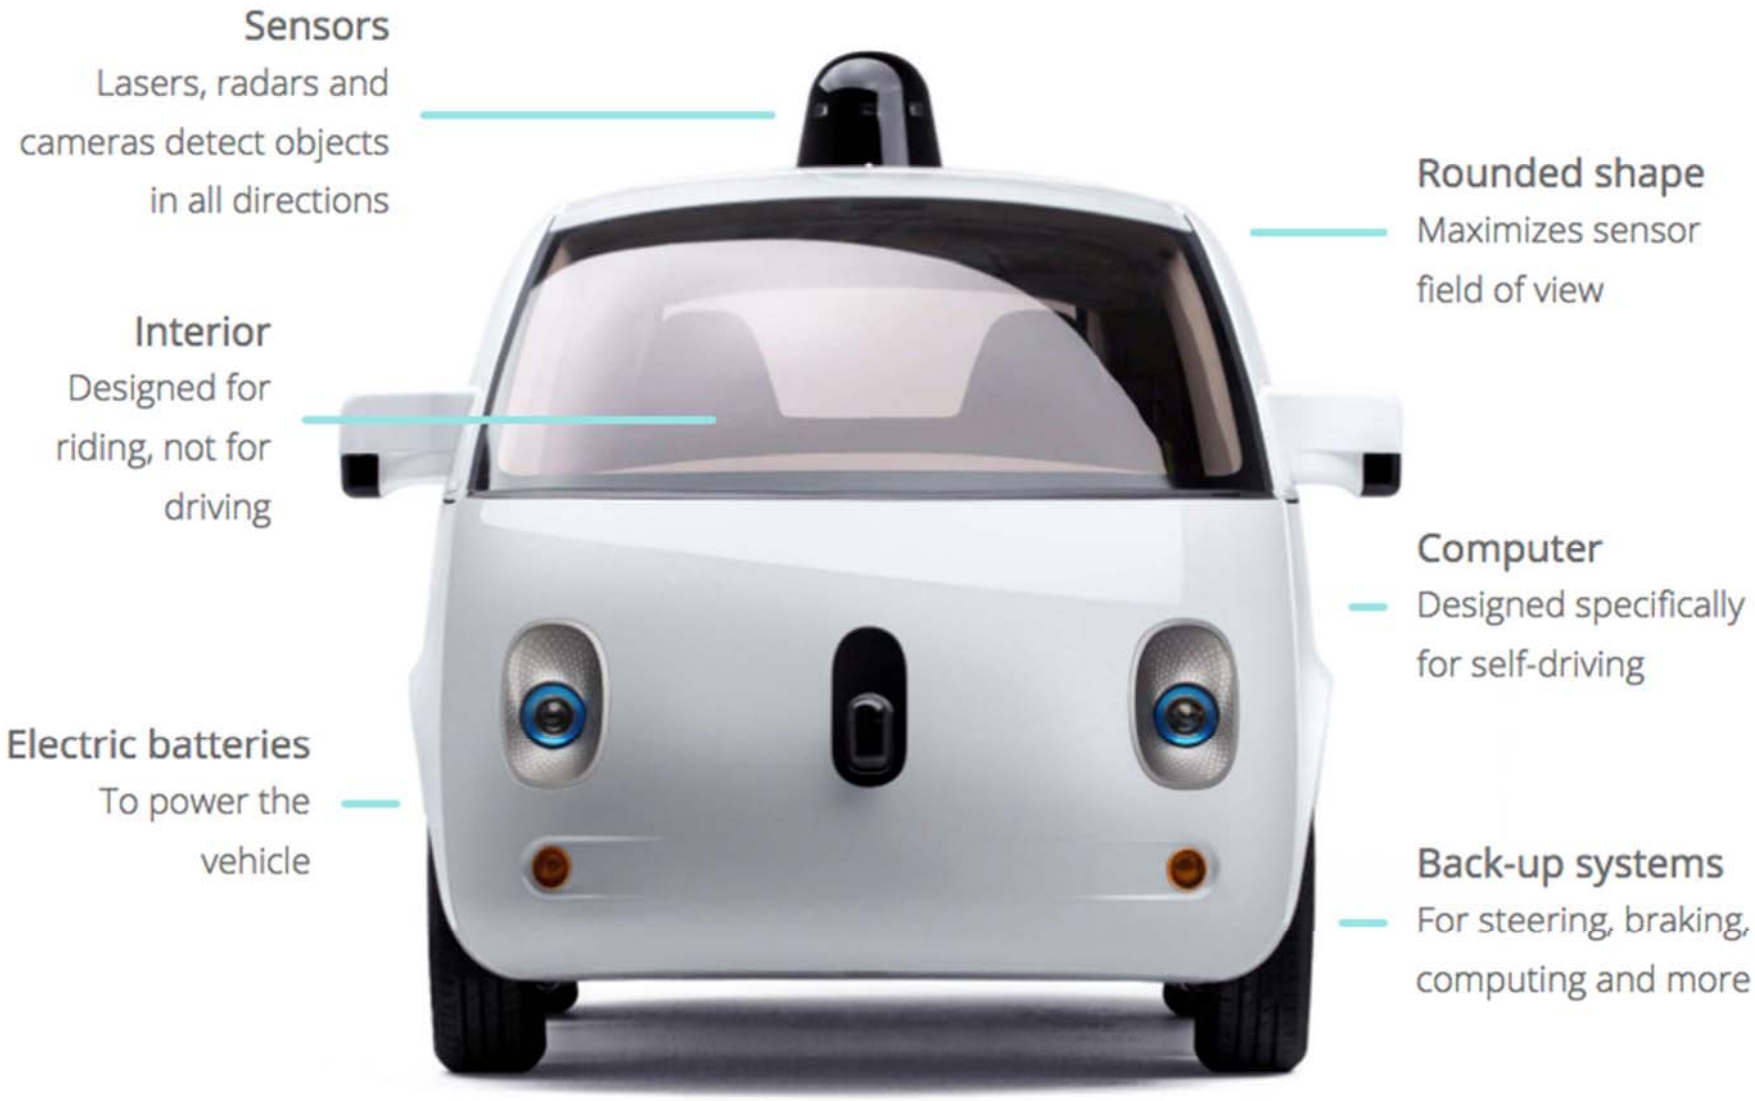


Image: https://www.google.com/selfdrivingcar/how/

# Applications: Much, Much More

- Scheduling, *e.g.* airline routing, military
- Route planning, *e.g.* Google maps
- Medical diagnosis
- Web search engines
- Spam classifiers
- Automated help desks
- Fraud detection
- Product recommendations
- … Lots more!

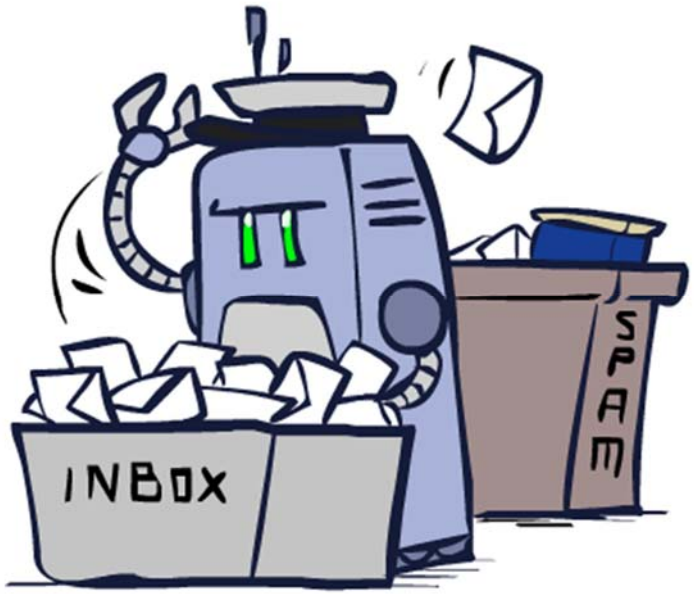

# Outline

- What is AI?
- AI History
  - AI winter and the resurgence
- AI Today
  - Applications and how they work
- **AI Tomorrow**
  - **Ethics and the singularity**

# What's In a Training Set?

- Just like humans, machines can only learn what they are taught (or can read about on their own)
  - Biased training set = biased behavior
  - https://techcrunch.com/2016/12/10/5-unexpected-sources-of-bias-in-artificial-intelligence/

- Examples:
  - In March 2016, Microsoft released AI Twitter bot Tay
    - http://www.complex.com/life/2016/03/microsoft-racist-ai
  - Software used to predict future criminals is biased against African Americans
    - https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

# AI: What Should We Worry About?

- AI Ethics: Immediate concerns

Jobs

Liability

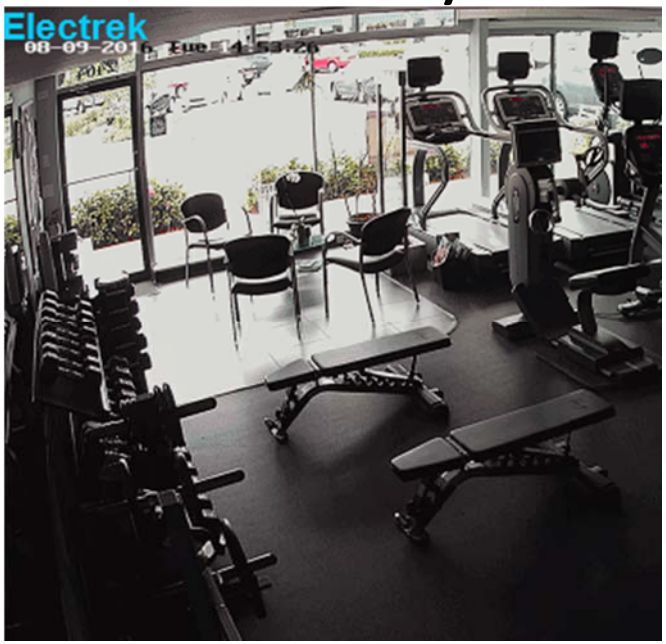


Weapons

Images:
http://ot.to/
https://electrek.co/2016/09/25/tesla-model-s-crashes-into-gym-driver-claims-autonomous-accel
http://futureoflife.org/2016/09/20/podcast-what-is-nuclear-risk/

# AI: Superintelligence

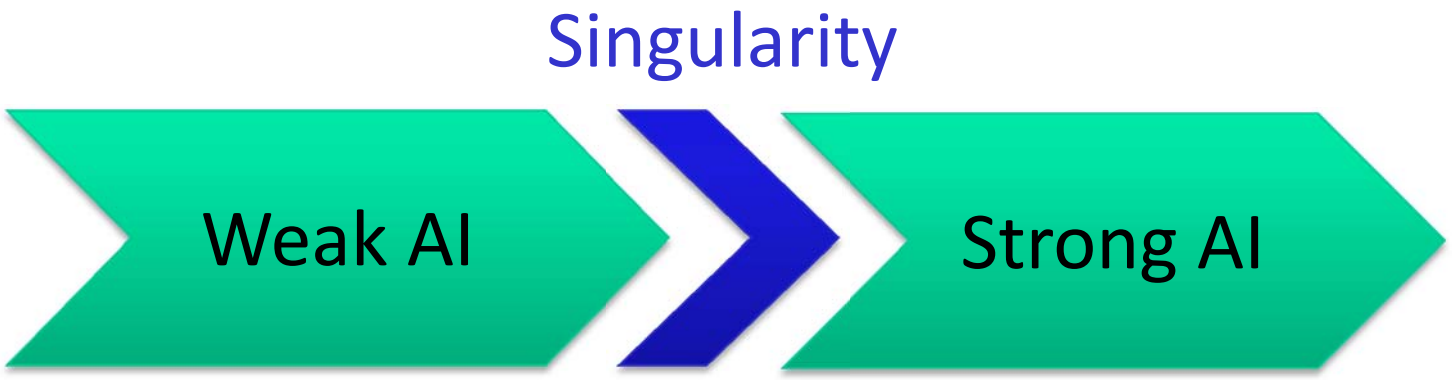


- Narrow AI
- Limited number of applications

- Artificial General Intelligence
- Recursive self-improvement
- Beyond human control