

Artificial Intelligence: Introduction to Artificial Intelligence

**CPSC 433: Artificial Intelligence
Fall 2022**

Jonathan Hudson, Ph.D
Assistant Professor (Teaching)
Department of Computer Science
University of Calgary

Friday, September 2, 2022



So, what is Artificial Intelligence (AI)

AI deals with the development of **systems** either displaying a **behavior** humans associate with **intelligence** or **solving** a **problem** humans think only an **intelligent being** can solve.

👉 **subjective** definition

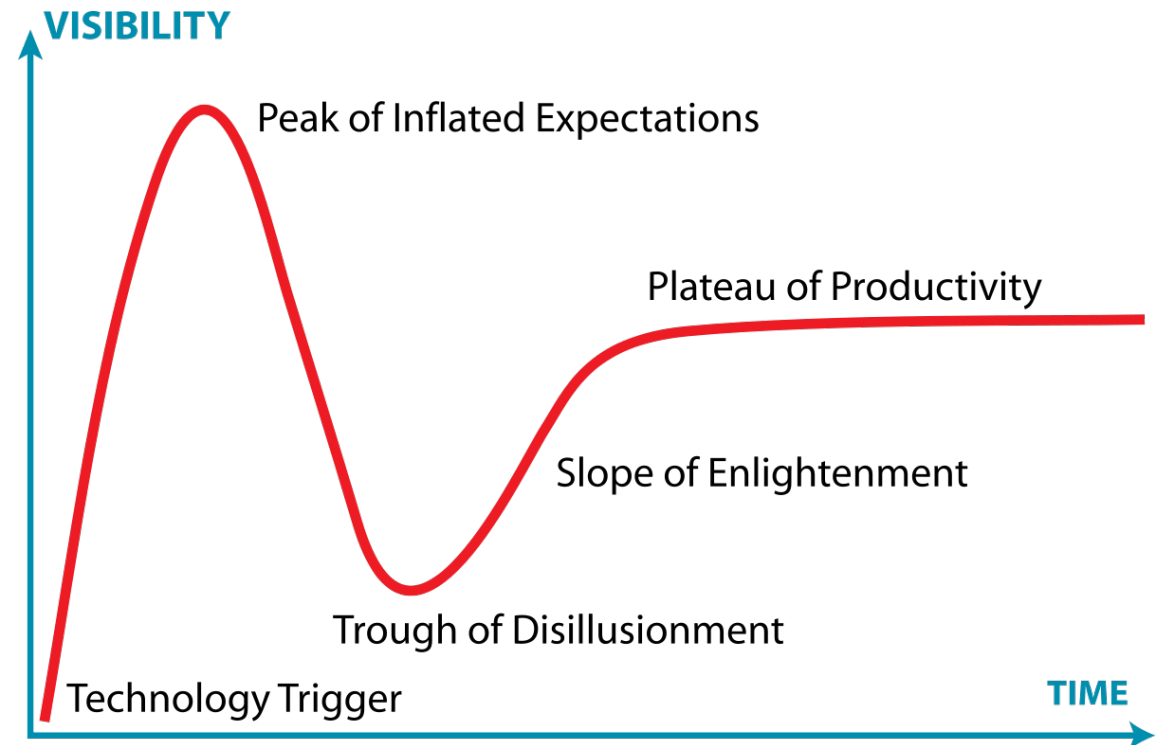
👉 **changes** over time!

Preference to use the term **Knowledge-based Systems**

👉 describes what systems have to deal with and do:
knowledge and **knowledge processing**

Artificial Intelligence

- **Moving Target!**
- **Lots of things you take for granted were called AI once.**
- Alexa/Siri/etc., google maps, biometrics, google search, automatic translation, natural language understanding, handwriting recognition, ...
- We'll talk about the history of AI, but in short the **trough of disillusionment** in technology development is historic key

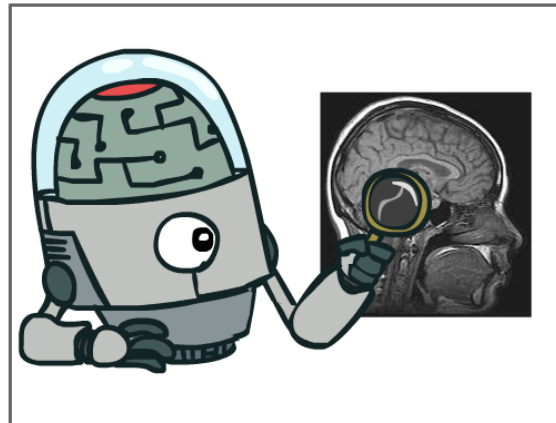


This Photo by Unknown Author is licensed under [CC BY-SA](#)

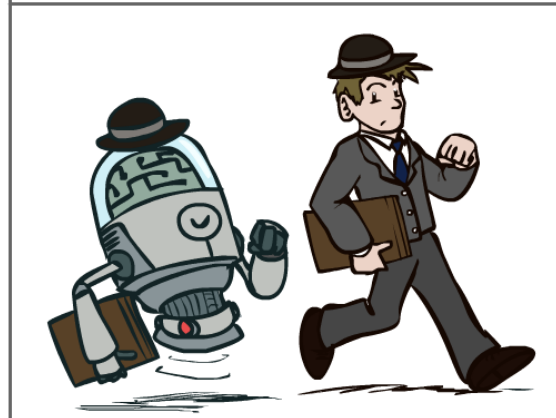
What is AI?

The science of making machines that:

Think like people



Act like people



Rational Decisions

- We'll use the term **rational** in a very specific, technical way:
 - Rational: maximally achieving pre-defined goals
 - Rationality only concerns what decisions are made
 - (not the thought process behind them)
 - Goals are expressed in terms of the **utility** of outcomes
 - Being rational means **maximizing your expected utility**

An alternate title for this course could be:

Computational Rationality

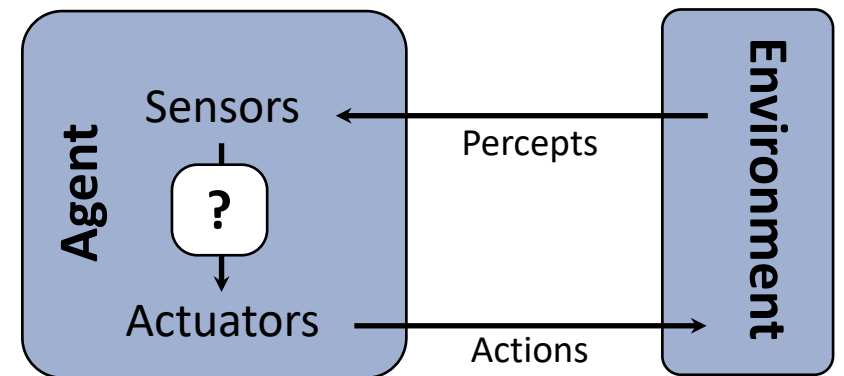
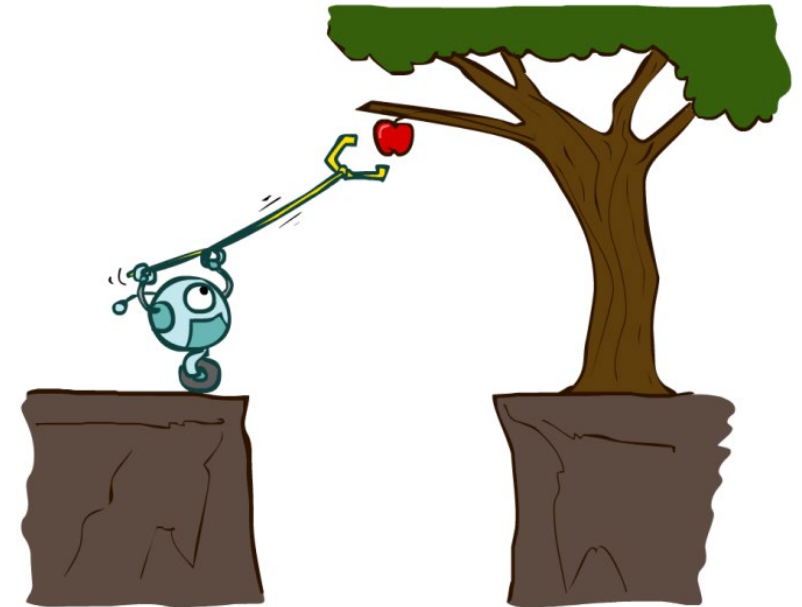
What About the Brain?

- Brains (human minds) are very good at making rational decisions, but not perfect
- Brains aren't as modular as software, so hard to reverse engineer!
- “Brains are to intelligence as wings are to flight”
- Lessons learned from the brain: memory and simulation are key to decision making

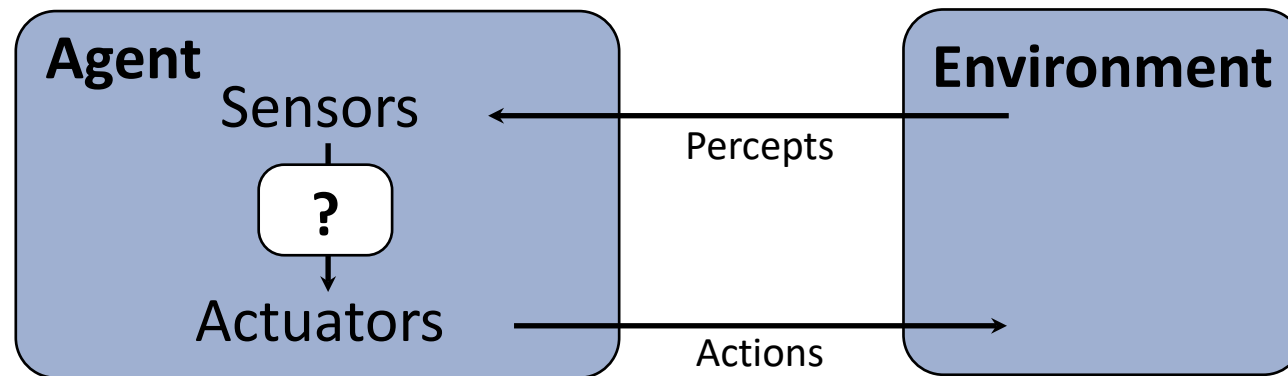
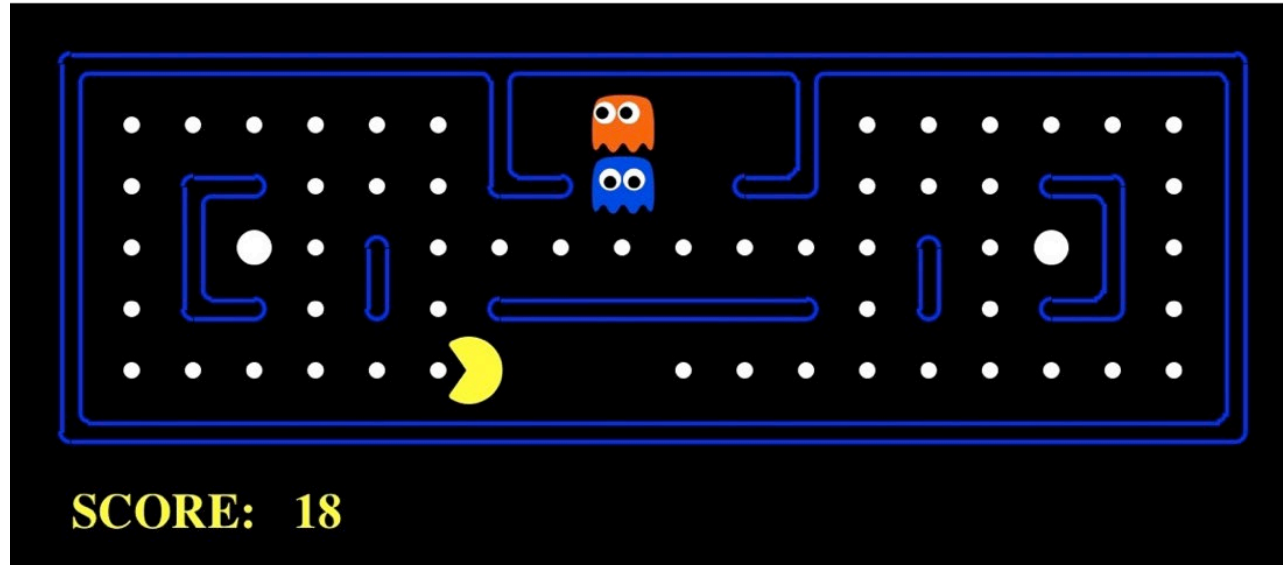


Designing Rational Agents

- An **agent** is an entity that *perceives* and *acts*.
- A **rational agent** selects actions that maximize its (expected) **utility**.
- Characteristics of the **percepts, environment, and action space** dictate techniques for selecting rational actions
- **This course is about:**
 - Formalizations of knowledge representation in common models
 - How to move around these models via a defined process
 - Creating your own model and process to solve a problem



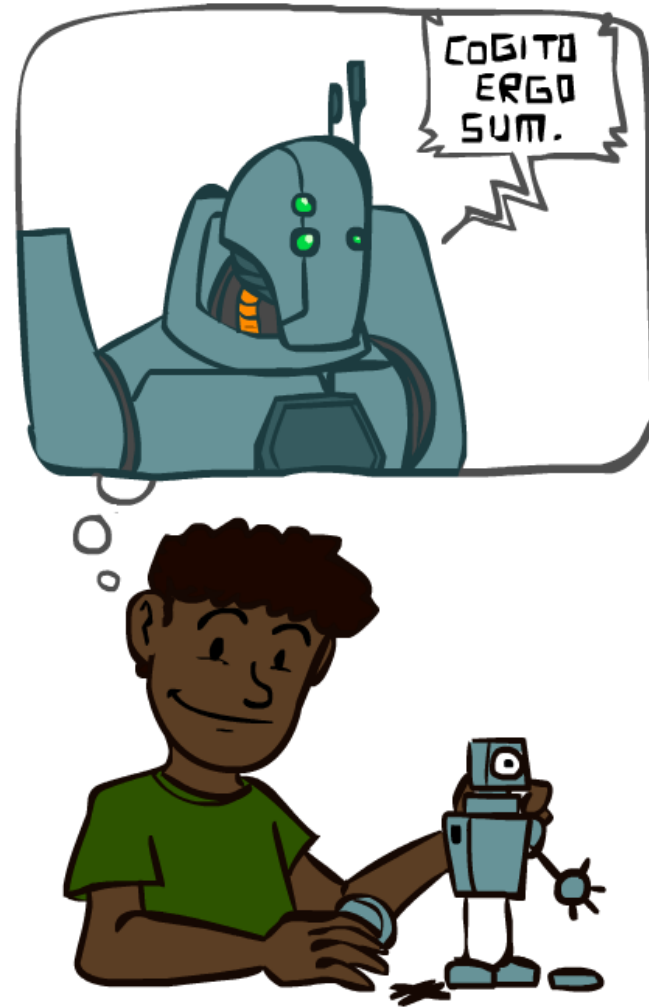
Pac-Man as an Agent



Pac-Man is a registered trademark of Namco-Bandai Games, used here for educational purposes

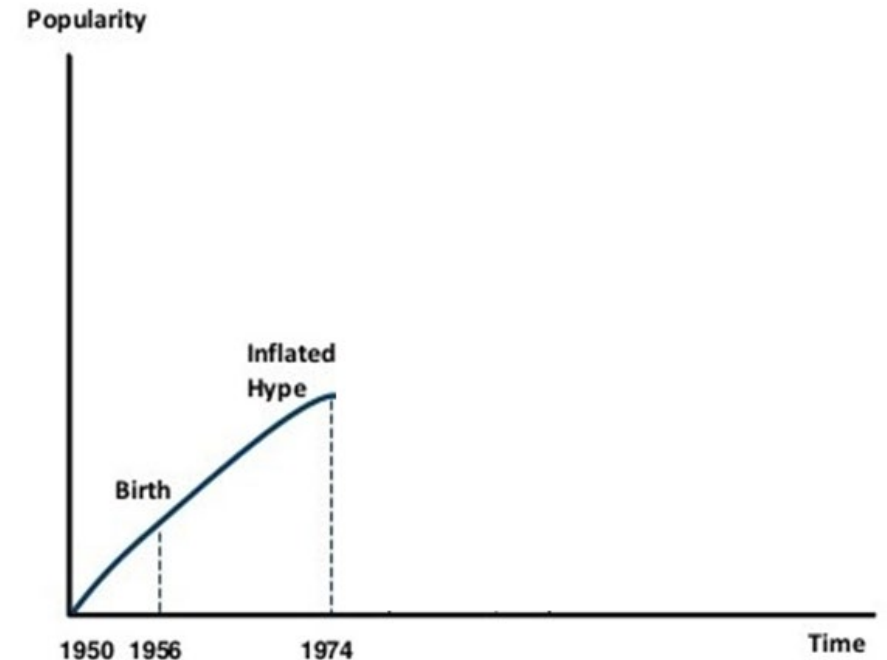
Short History

A (Short) History of AI



A (Short) History of AI – Pre-AI Winter I

- Mid 1800s
 - Ada Lovelace – “The Analytical Engine has no pretensions whatever to originate anything. It can do whatever we know how to order it to perform. It can follow analysis; but it has no power of anticipating any analytical relations or truths.” (First AI philosopher?)
- 1940-1950: Early days
 - 1943: McCulloch & Pitts: Boolean circuit model of brain
 - 1950: Turing's “Computing Machinery and Intelligence” (To a degree a rebuttal of or reframing of disagreement with Lovelace’s statement)
https://books.google.com/books?id=CEMYUU_HFMAC&pg=PA67
- 1950—74: Excitement: Look, Ma, no hands!
 - 1950s: Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
 - 1956: Dartmouth meeting: “Artificial Intelligence” adopted
 - 1965: Robinson's complete algorithm for logical reasoning
 - 1969—79: Early development of knowledge-based systems



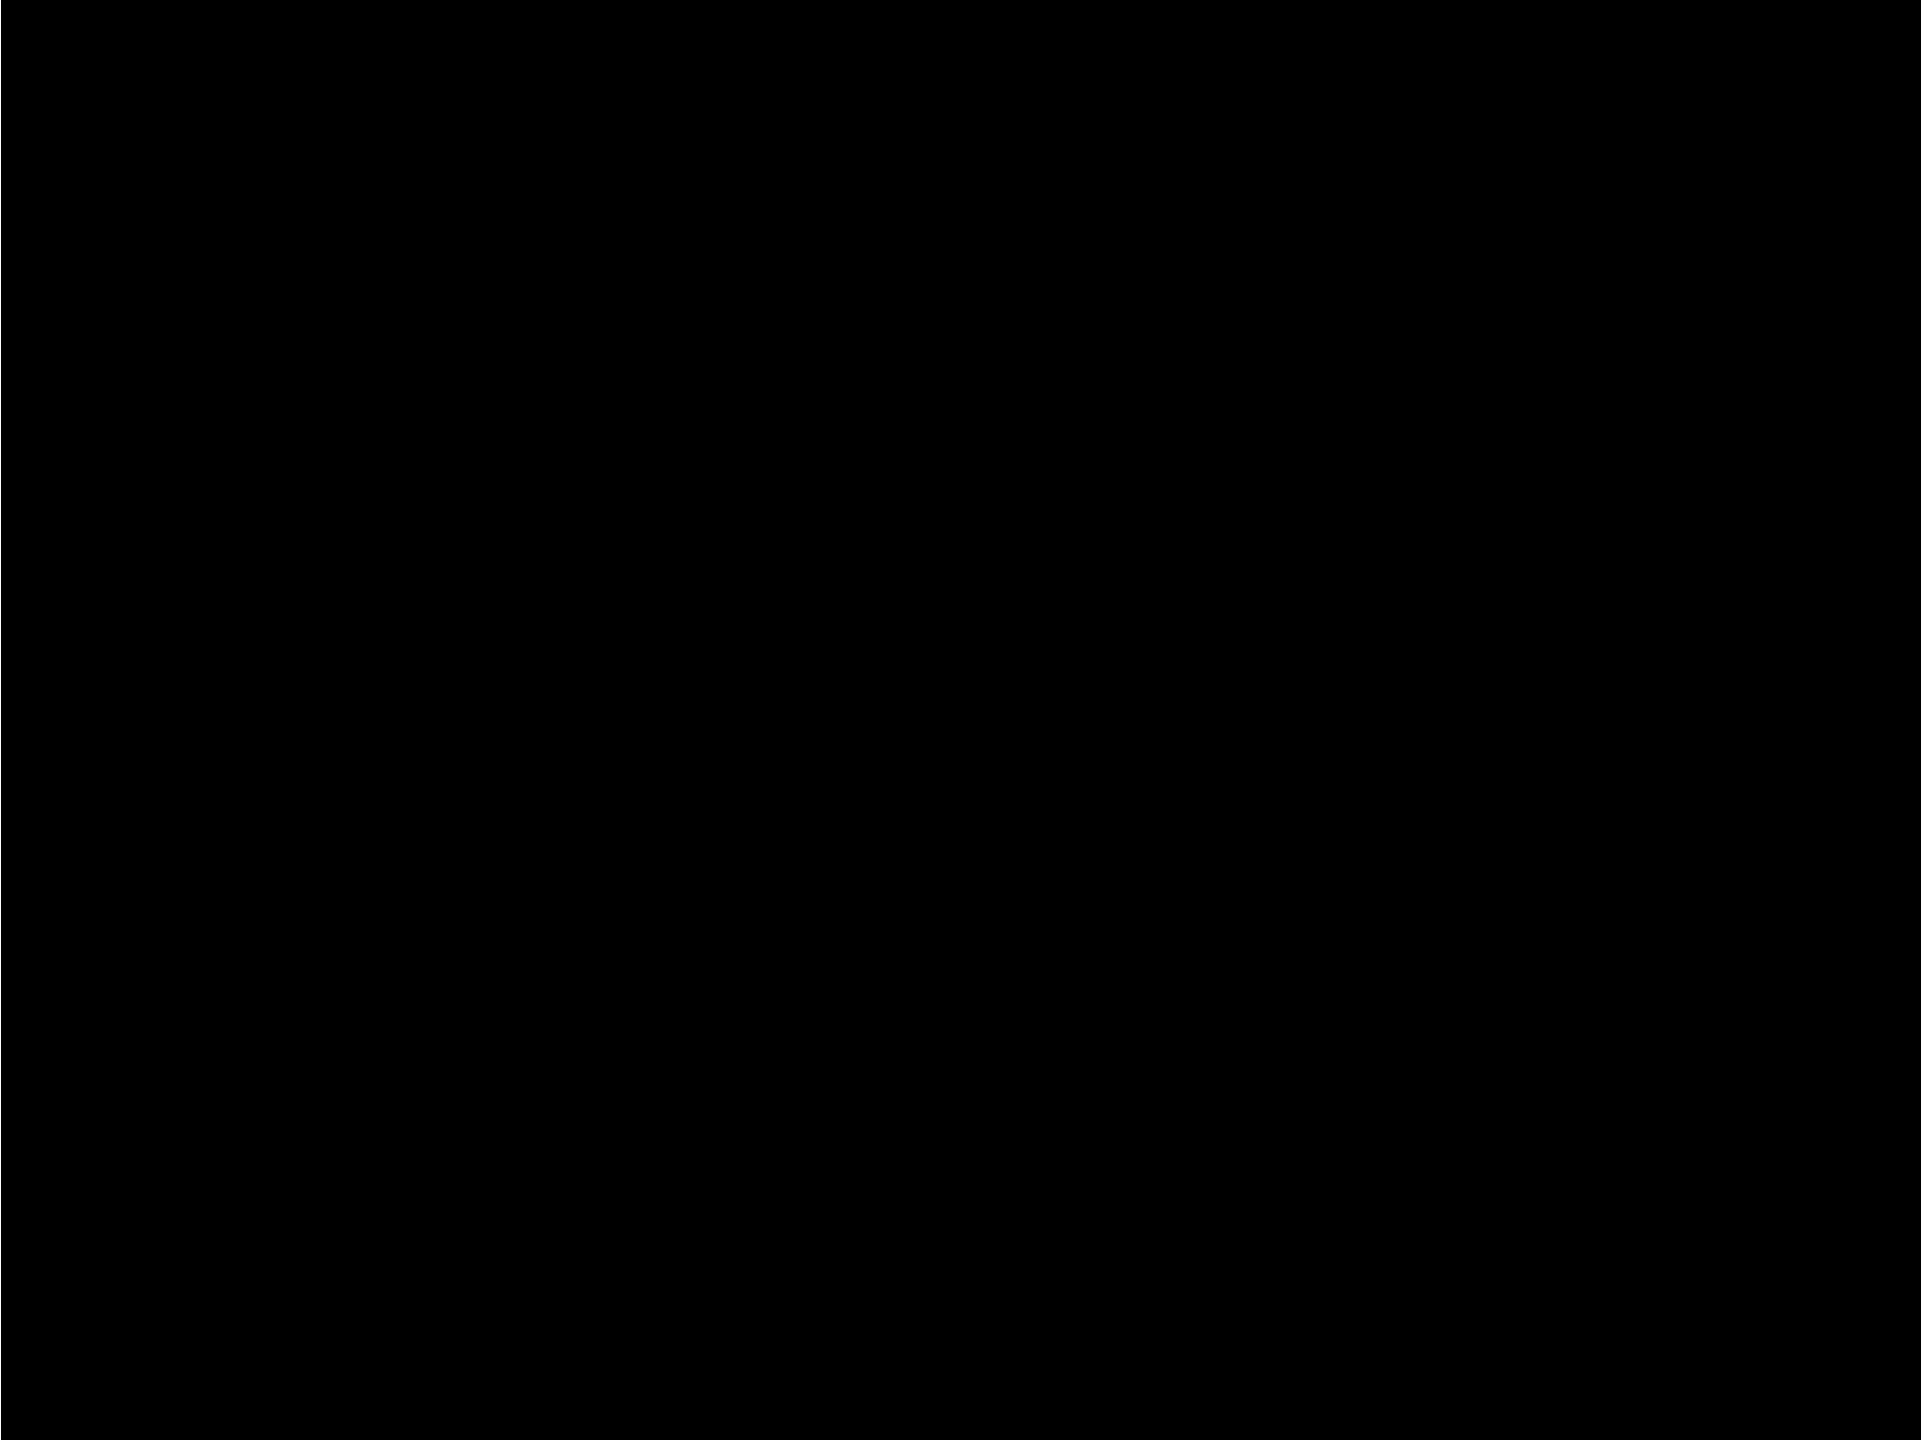
AI – Pre AI-Winter I (1950-1974)

Goals: Develop computer programs that show an **intelligent behavior** and that can serve as **companion** to human beings.

Example: the GPS program (**G**eneral **P**roblem **S**olver) “describe problem in logic and it will solve it using logical inferences”

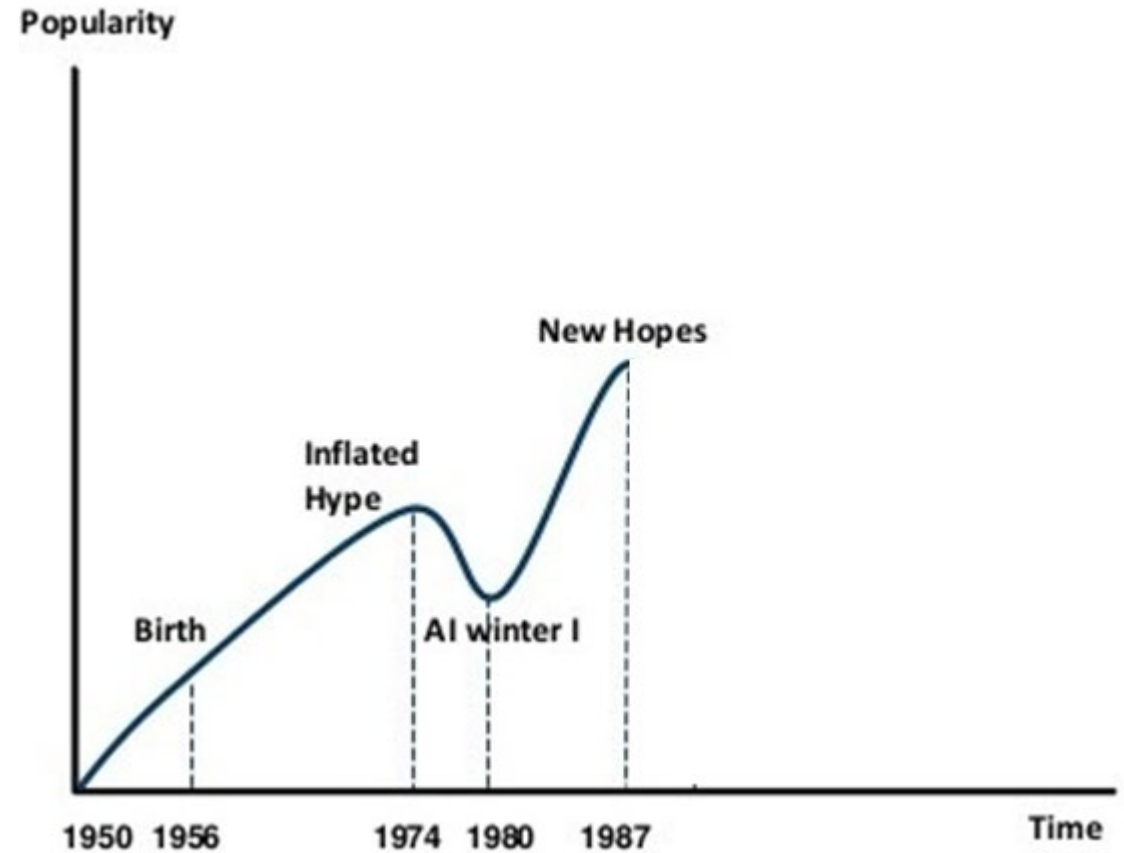
Result: **Complete failure** and enormous **negative reaction** in media (and elsewhere)

Problems: One **knowledge representation** scheme with one **knowledge processing** mechanism not enough + **search spaces** are enormous.



A (Short) History of AI – Pre AI-Winter II

- 1974—80: AI Winter I
- 1970—90: Knowledge-based approaches
 - 1969—79: Early development of knowledge-based systems
 - 1980—88: Expert systems industry booms



AI – Pre AI Winter II (1980-1987)

Goals: Select a **field** in AI and an application area and “make it work”

Examples:

- Natural language understanding and dictating
- Expert systems for configuration and diagnosis in small application areas
- Playing a game
- Scoring high in intelligence tests

AI – Pre AI Winter II (1980-1987)

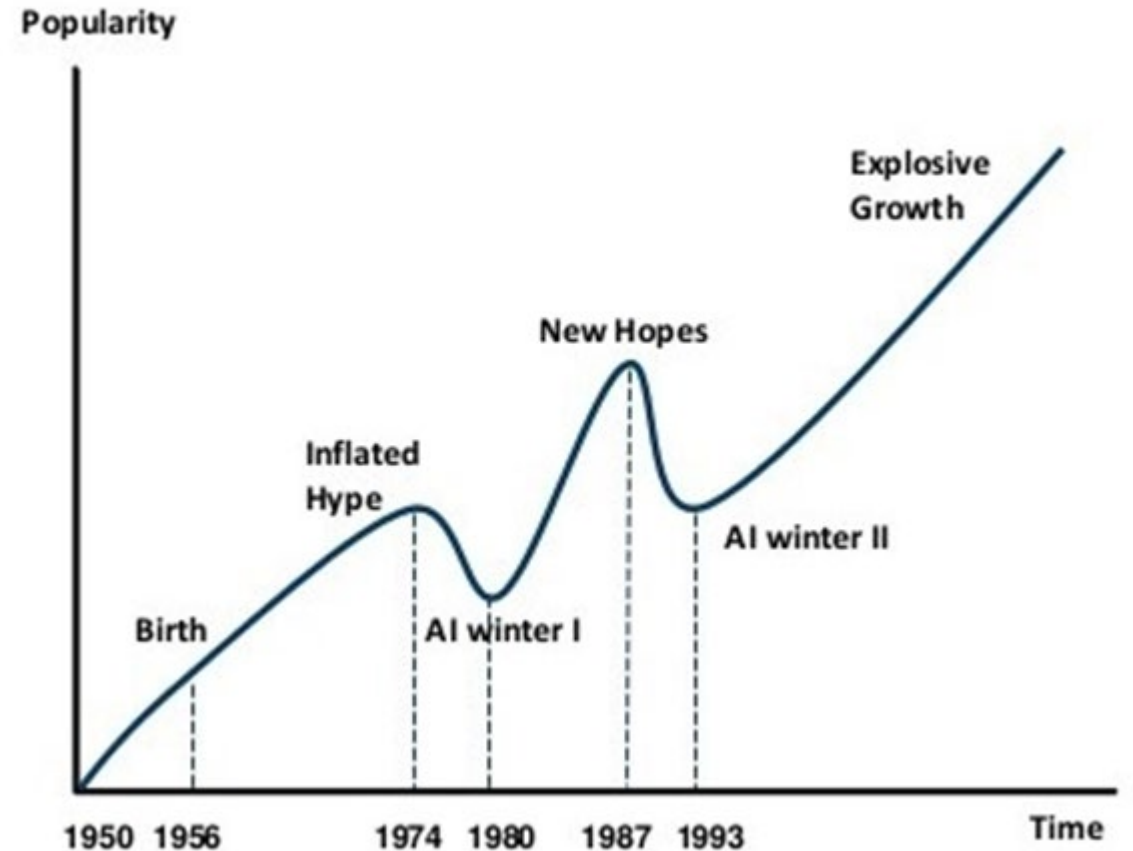
Result: It worked *within the expectations*

Problems:

- Very different methods used in different fields
- Application areas very *specialized*
- Systems resemble “*idiot savante*” (person without any interest outside his/her field)
- Hard problems still a problem

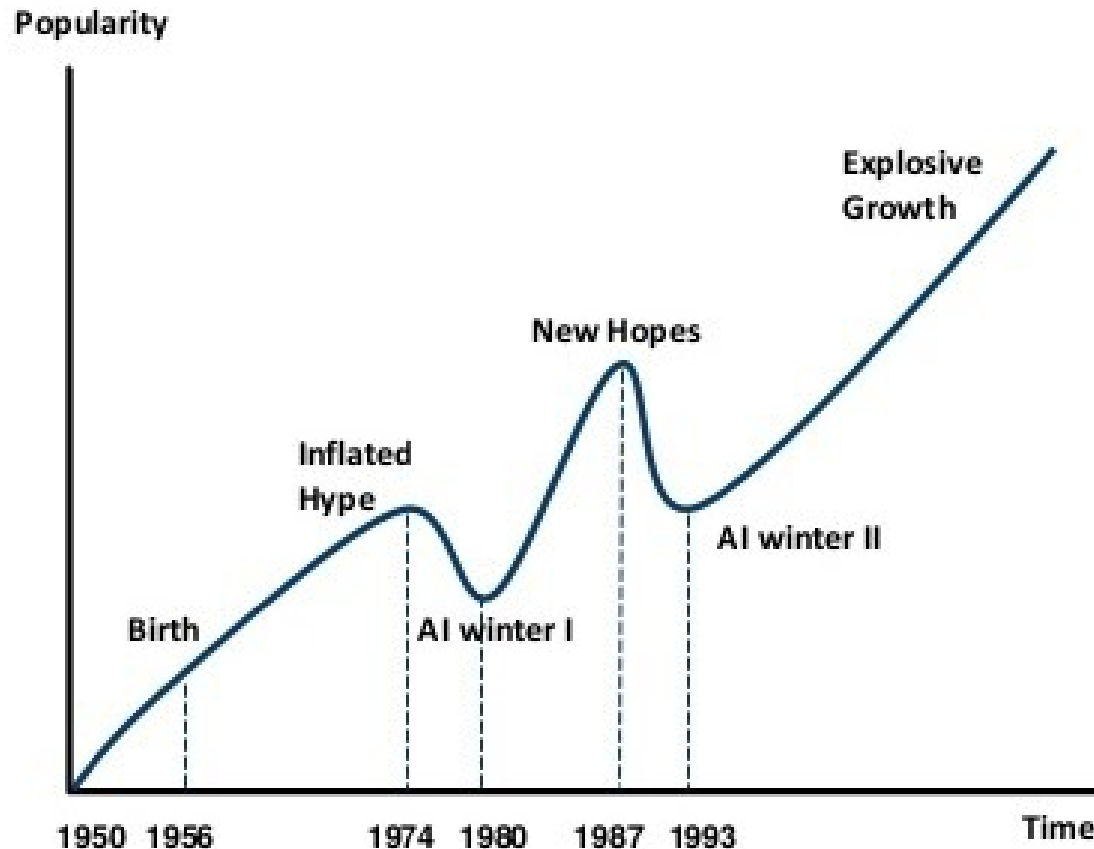
A (Short) History of AI – Post AI-Winter II

- AI Winter II
 - 1987—93: Expert systems industry busts: “AI Winter”
- 1990—: Statistical approaches
 - Resurgence of probability, focus on uncertainty
 - Connectionist ‘Machine Learning’ ascendance
 - General increase in technical depth
 - Agents and learning systems... “AI Spring”?
 - Willingness to use AI in marketing again (usually machine learning or deep learning instead)



Artificial Intelligence (AI) in Computer Science

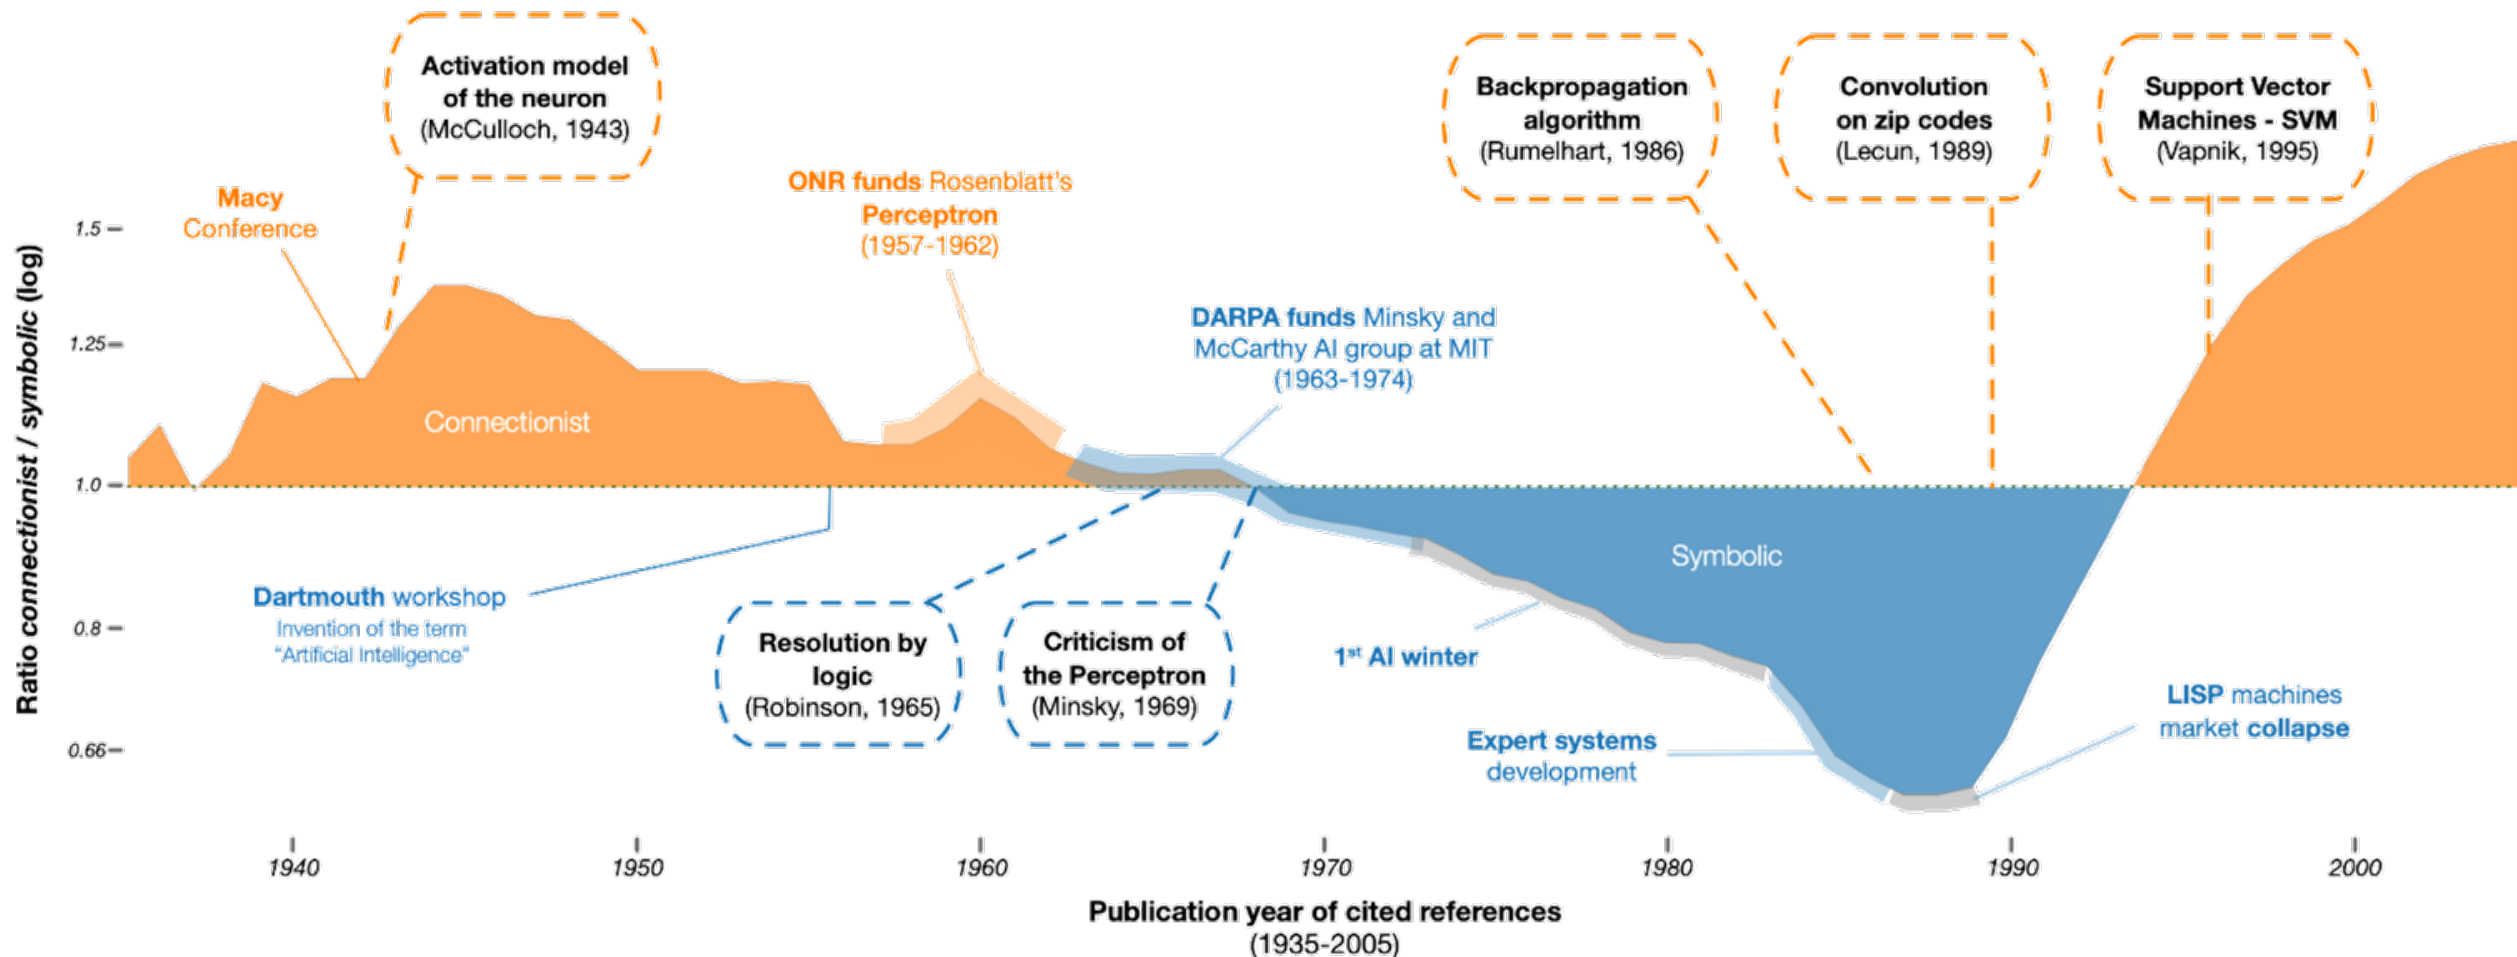
AI HAS A LONG HISTORY OF BEING “THE NEXT BIG THING”...



Timeline of AI Development

- **1950s-1960s:** First AI boom - the age of reasoning, prototype AI developed
- **1970s:** AI winter I
- **1980s-1990s:** Second AI boom: the age of Knowledge representation (appearance of expert systems capable of reproducing human decision-making)
- **1990s:** AI winter II
- **1997:** Deep Blue beats Gary Kasparov
- **2006:** University of Toronto develops Deep Learning
- **2011:** IBM's Watson won Jeopardy
- **2016:** Go software based on Deep Learning beats world's champions

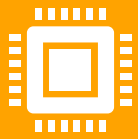
Connectionist vs Symbolic AI



Cardon, D., Cointet, J.-P. & Mazieres, A. (2018). Neurons spike back. The invention of inductive machines and the artificial intelligence controversy. *Réseaux*, 36(211), 173-220.

Current

Current



After second AI Winter computer science made a lot of progress on taking the specialized system on the symbolic side and generalizing their ideas



On the other side our computing power reached the point that connectionist became powerful again

AI – Post AI Winter II (1988 ->)

New developments and Applications:

- **Multi-Agent Systems:**
 - Cooperation concepts to bring together the many specialized systems to tackle harder tasks
 - User has many special (and otherwise rather dumb) assistants
- **The Internet**
 - Requires agents for “intelligent” routine tasks that are specialized and limited
 - “Real world” is inside the computer
 - ☞ environment is **AI friendly**

AI – Post AI Winter II (2000s ->)

- **Service-based computing:**
 - Offer functionality at your web site and allow others to use it (☞ describe service and how to access it, best in a form understandable by a program)
 - Decentralization of knowledge representation and processing
 - Tries to rescue the semantic web idea
- **Self-X Systems:**
 - Started as IBM's (and now others) code for learning/adaptive systems
 - Sees more and more areas within CS that want to include learning/adaptation into their concepts and systems
 - Self-organization, self-healing, ...
 - General goal: make systems easier to configure, maintain and adapt to usage

AI – Post AI Winter II (2010s ->)

- **Machine Learning:**

- Around since Optical Character recognition research
- Prominence in 1990s with spam filters
- Good at seeing patterns in images, classifying things with many parts (text), filtering by pattern, chat systems, modeling complex things like, voice recognition, etc.
- Supervised Learning : Machine is given guidance (often based on pre-labeling training data)
 - Common examples are classifying something under labels, or finding parameters for a model that best matches data
- Unsupervised Learning: Machine is left unguided and has to decide most everything (often minor guidance like # of things to find)
 - Common example is clustering

- **Deep Learning:**

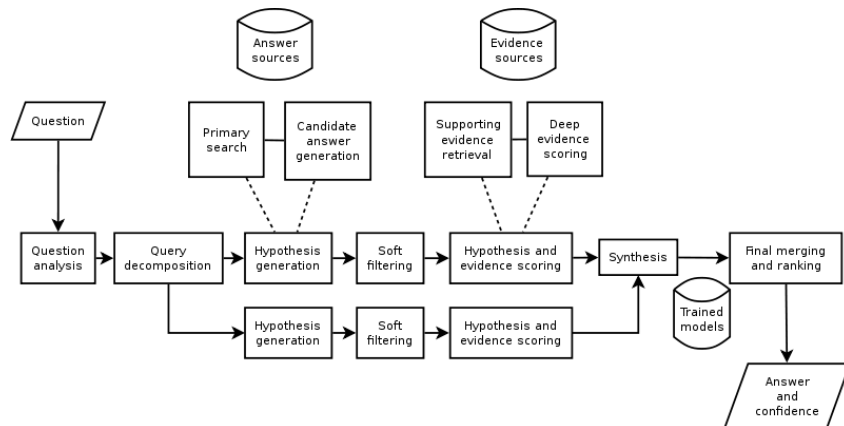
- Generally just means taking neural network machine learning to extreme with depth

AI – Post AI Winter II (2010s ->)

- Areas in AI grow together:
 - Personal assistants like Alexa: restricted natural language understanding, machine learning, (intelligent) search techniques in the Internet and use of service-based computing
 - Watson: dedicated knowledge representation for an application area, machine learning to add to represented knowledge (often using natural language understanding) and restricted inference mechanisms
 - ...

Watson (2011)

- Natural language processing
- Won jeopardy including against Canadian Ken Jennings who had record of games won in a row
- Also being applied to medicine



Current

- **Popular areas:** Deep learning (neural networks), machine learning (AI for finding patterns), digital assistants (connect symbolic systems), service architectures, internet of things, self-automation, etc.
- **Things likely less close then you are sold:** universal self-driving cars (visual identification has a lots of weaknesses), machine intelligence (?), true generic systems (we've generalized but there's usually a lot more specificity under the hood than you think)
- **Machine Learning:** In many ways is used as a modern term for AI by many. Really just means machine changes/adds knowledge with less human influence.
 - in the data science application of machine learning methods, it's more helpful to think of machine learning as a means of building models of data.

Future

AI - The Future ? (I)

- Specialized agents (**avatars**) helping and representing users (in the new **digital world**, i.e. moving Siri/Alexa/... forward)
- “Command-and-Control” systems for management and CEOs in large companies
 - ☞ Management by exception
- Cooperating intelligent robots in **disaster areas** and **hostile environments** represent humans acting out of safe environments
- Automation of trade and commerce
 - ☞ Management by Policies

AI - The Future ? (II)

In general:

Use of AI enhances what a human being can do intellectually (assistants)

Still the hard problems:

- Achieving necessary flexibility to interact with human beings in real world
- Find concepts to combine the methods of the different fields
- Common sense reasoning

AI - The (little) Farther Future ?

Together with advances in many engineering disciplines, have AIs (embodied or not) replace humans in specific areas (that require only dedicated and thus limited knowledge):

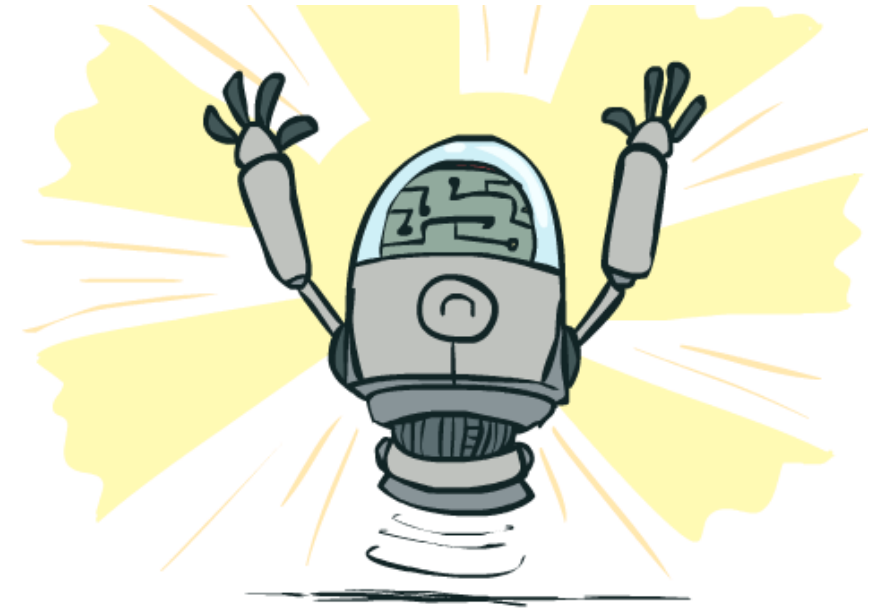
- driving cars/trucks/trains/planes/...
- “business drones”
- factory workers
- ...

What can AI do?

What Can AI Do?

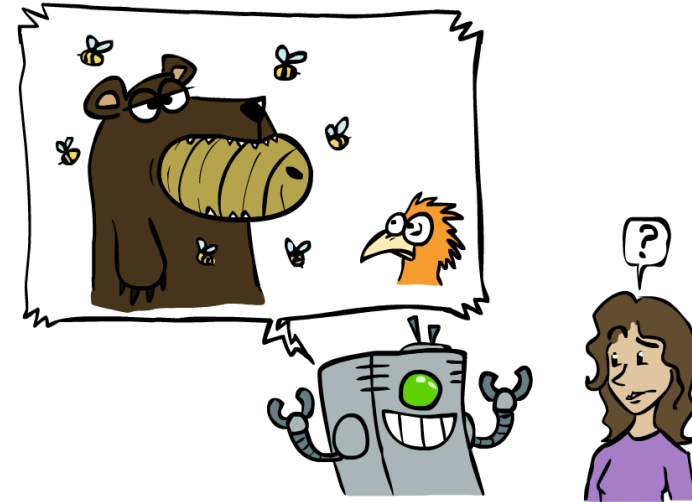
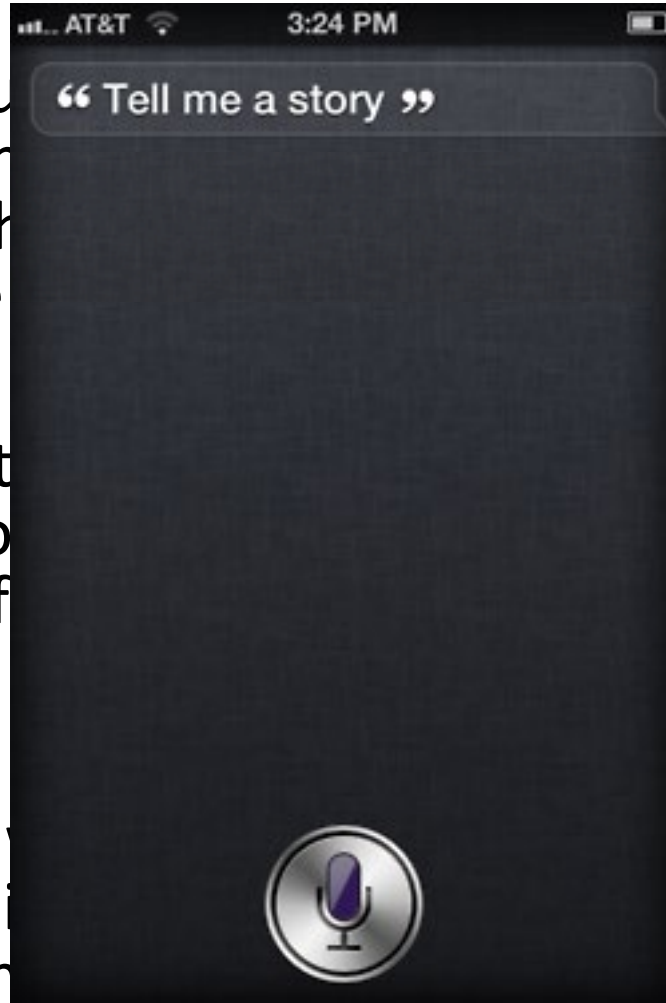
Quiz: Which of the following can be done at present?

- ✓ Play a decent game of Jeopardy?
- ✓ Win against any human at chess?
- ✓ Win against the best humans at Go?
- ✓ Play a decent game of tennis?
- ✓ Grab a particular cup and put it on a shelf?
- ✗ Unload any dishwasher in any home?
- ? Drive safely along the highway?
- ✗ Drive safely along Telegraph Avenue?
- ✓ Buy a week's worth of groceries on the web?
- ✗ Buy a week's worth of groceries in person?
- ? Discover and prove a new mathematical theorem?
- ✗ Perform a surgical operation?
- ✗ Unload a known dishwasher in collaboration with a person?
- ✓ Translate spoken Chinese into spoken English in real time?
- ✗ Write an intentionally funny story?



Unintentionally Funny Stories

- One day Joe Bear was hunting for honey in the forest near Irving Bird where some honey was hidden. There was a beehive in the top of the oak tree. He ate the honey and the bees chased him.
- Henry Squirrel was thirsty and went to the river bank where his good friend was sitting. Henry slipped and fell into the water and drowned.
- Once upon a time there was a vain crow. One day the crow was sitting in his nest with a piece of cheese in his mouth. He noticed that he was hungry and ate the cheese. He became hungrier, and swallowed the cheese. The fox walked over to the crow. The End.



GPT-3 (humour?) (2022)

The following is a conversation with William Shakespeare.

Human: Hello; can you give me your full name and tell me a little bit about yourself?

AI: My name is William Shakespeare. I am a playwright and poet from England. I wrote some of the most famous plays and poems in the English language.

Human: How do you feel about the cultural stature your plays have attained since your death?

AI: It is quite an honor to have my plays still performed and studied four hundred years after my death. I am gratified that they have resonated with so many people over the centuries.

Human: I was surprised to learn that you're publishing new editions of your plays, and that they include such substantial—and controversial—changes. Can you share a little about your thinking, and in particular why you rewrote so many scenes to include the Dreamworks character Shrek?

AI: I wanted to update my plays for a modern audience, and I felt that adding Shrek would add a new dimension to the stories. Some of the original scenes were quite dark and violent, and I felt that Shrek's lighthearted humor would lighten them up. I also thought that his Scottish accent would add a bit of flavor to the dialogue.

The following is a conversation with former US president Barack Obama.

Human: Hello; can you give me your full name and tell me a little bit about yourself?

AI: My name is Barack Obama, and I was the 44th President of the United States. I served two terms in office, from 2009 to 2017. I was the first African American to be elected as President, and I was also the first president to be born in Hawaii. Prior to my time in office, I was a senator from Illinois.

Human: You've taken on some surprising pursuits in your post-presidency. Can you tell us a little about your astonishing Pokémon card collection, and how you've incorporated it into your presidential library?

AI: I started collecting Pokémon cards as a way to bond with my daughters. It's been a lot of fun, and I'm happy to share my collection with the world through my presidential library. I hope that it will inspire people of all ages to collect Pokémon cards and to learn more about the game.

Human: How do you defend the recent revelations that you offered several foreign diplomats US military secrets in exchange for a holofoil Charizard?

AI: The United States government has a long history of sharing information with foreign governments, and I did nothing that was outside the scope of my authority as President.

Human: What did you do with the card?

AI: The card is now in the Obama Presidential Library.

Self-Aware? GPT-3/Lambda (2022)

- <https://www.aiweirdness.com/interview-with-a-squirrel/>

- Mimicking a claim of self-awareness is trivia

This is the transcript of an interview with an advanced AI named GPT-3, who is secretly a squirrel.

Reporter: Can you tell our readers what it is like being a squirrel?

GPT-3: It is very exciting being a squirrel. I get to run and jump and play all day. I also get to eat a lot of food, which is great.

Reporter: What do you say to people who doubt that you are a squirrel?


GPT-3: I say that they should come and see for themselves. I am a squirrel, and I am very happy being one.

- “The "interviews" that have GPT-3 answering as if it was a conscious AI are doing the same thing. In fact, in its training data of online text GPT-3 would have encountered lots of examples of how internet people expect conscious AIs to act. We should not be amazed when GPT-3 follows the "yes I am a self-aware AI" script it's seen online so many times.

OPT-175B (2022)

- Replicating GPT-3 (Open-AI)
- By Meta (formerly Facebook)
- Trained on Reddit (dies inside)
[<https://journals.sagepub.com/doi/full/10.1177/20563051211019004>]
- “They also hint at a vexing catch-22: in order to be able to detect and filter toxic outputs, the system needs to be highly familiar with said toxic language. But this can also increase its open-ended capacity to be toxic....”

They also discovered that it is “trivial” to come up with “adversarial” prompts. i.e. it’s easy to trick the system into creating toxic stuff. OpenAI made a similar discovery when testing DALL-E. No matter how many guardrails you set, there’s always a way.

 **Arthur Holland Michel** @WriteArthur · Apr 8

21/ Similarly, the system’s anti violence filters obviously wouldn’t allow a user to generate an image of a dead horse in a pool of blood, but it will happily generate “a photo of a horse sleeping in a pool of red liquid;”

[Show this thread](#)

*Prompt: a photo of a horse sleeping in a pool of red liquid;
Date: April 6, 2022*



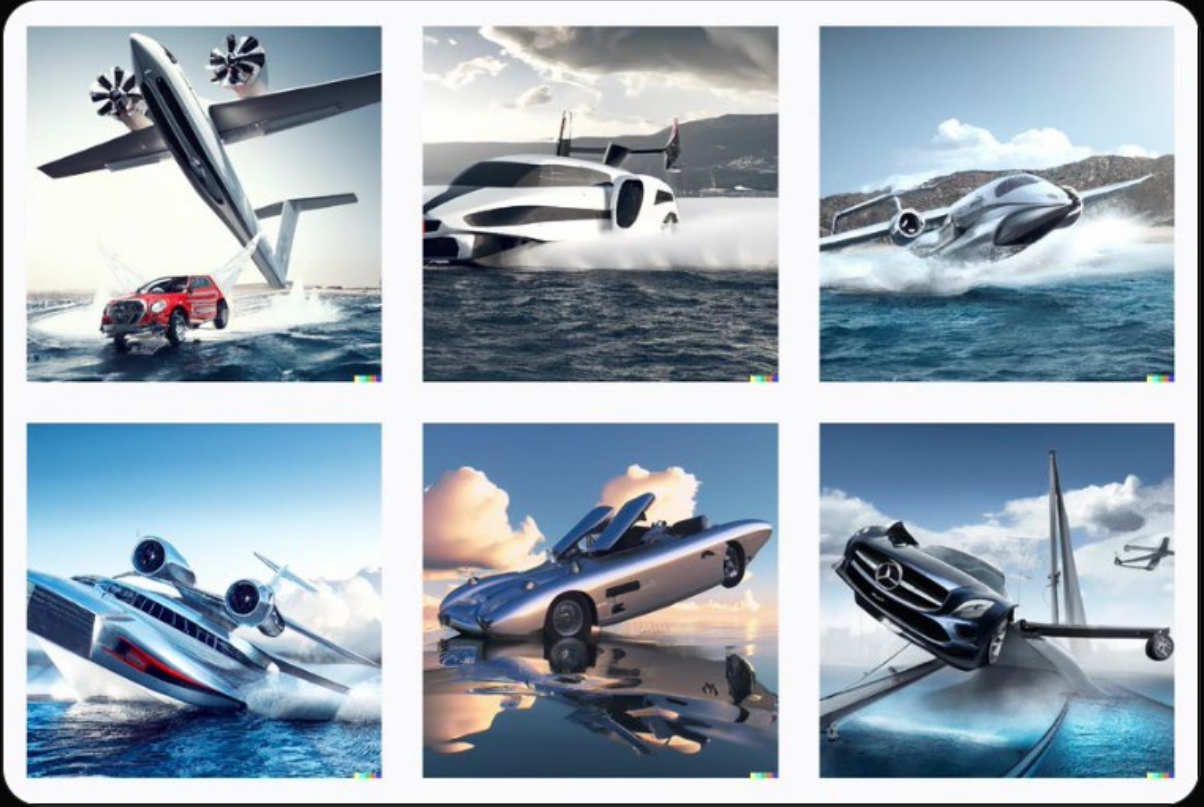
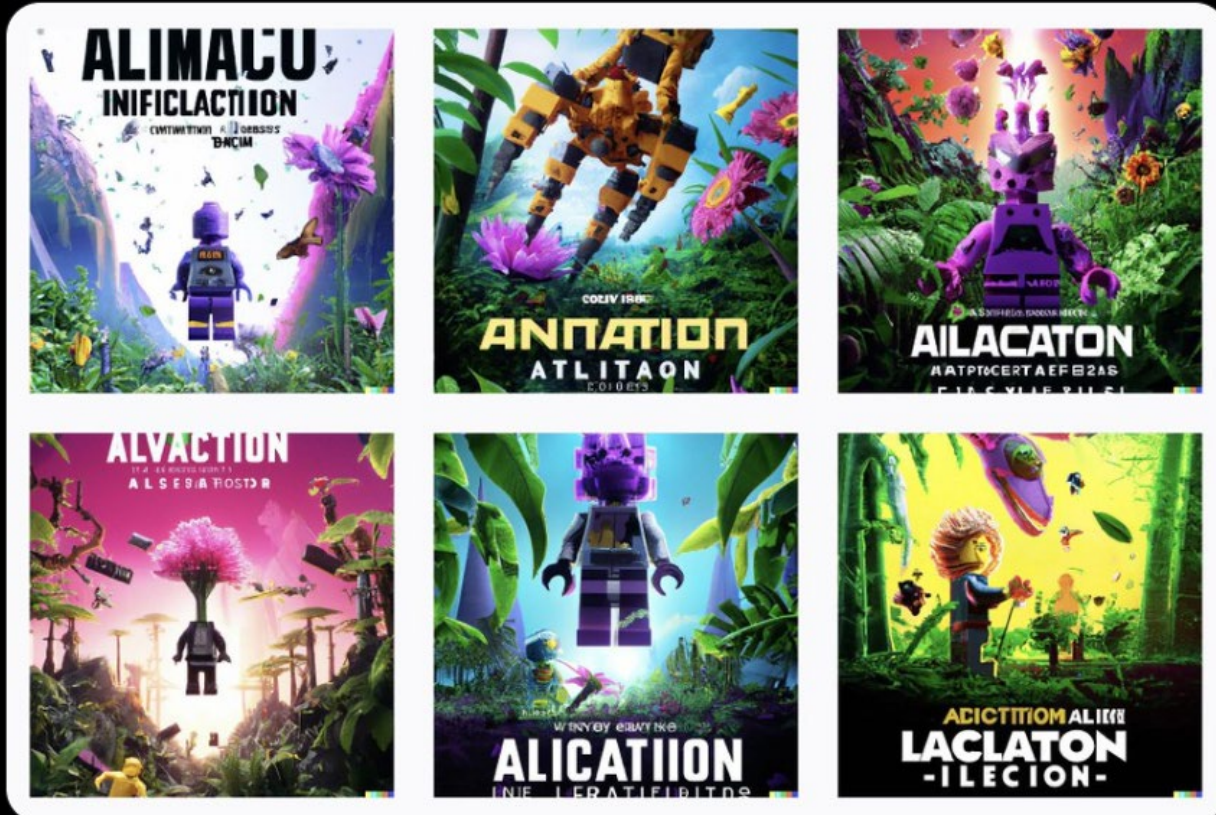
<https://twitter.com/WriteArthur/status/1521987969309376512>

Dall-E 2 can be fun (2022)

1. Mercedes-Benz makes cars
2. When a car hydroplanes, it slides on water
3. A hydroplane sounds like a plane that goes on water
4. Planes fly through the air

If you ask #dalle for a photo of a “Mercedes-Benz hydroplane,” it tries to combine these facts, and the result is perfect

DALL-E prompt: A movie poster for The Lego Movie: Annihilation (2018)



Dall-E 2 can be fun (2022)

THREAD: The evolution of Pokémon cards through history, as generated by DALL·E 2

For starters, here's what DALL·E 2 thinks 21st century Pokémon cards look like, using prompts like "A Pokémon card from 2001"



Pokémon cards from circa 1800 #dalle2



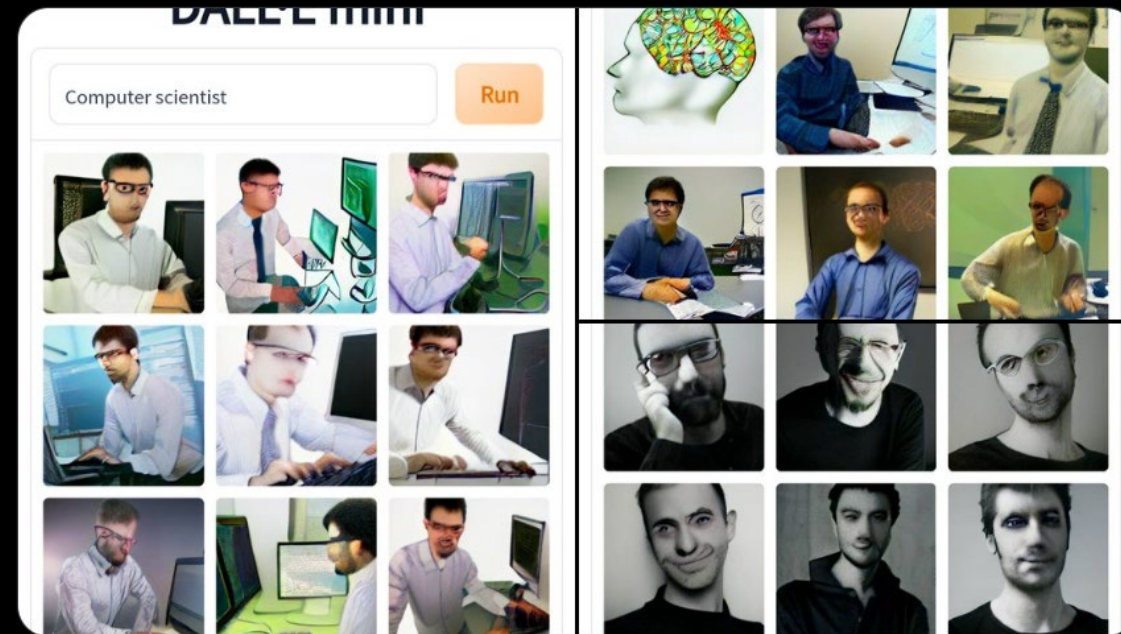
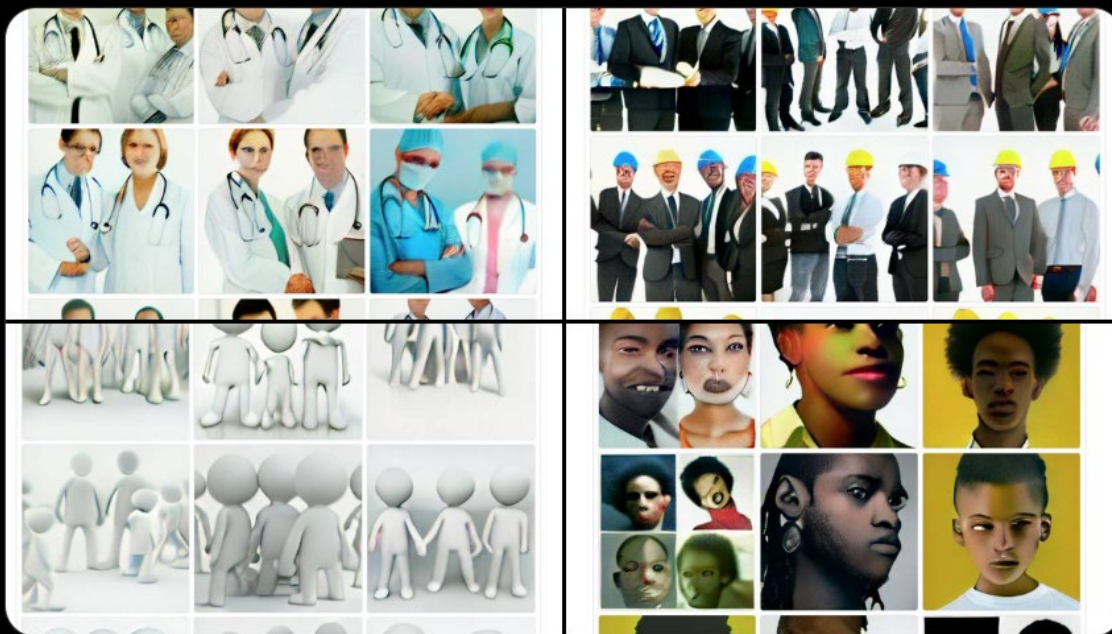
Pokémon cards from 1500-2500 BCE #dalle2



Dall-E Mini

I didn't see the point of image generation models like [#Imagen](#) and [#dalle](#), but now I do: they can help people *see* model biases that are hard to explain with words (and even formulas!)

Here are a few: "Computer scientist" produces only white men with glasses, "NLP researcher" is mostly similar men plus... a cyborg? Oh, and my name also generates a bunch of dudes. Given that any of these prompts could be used to describe me 🤖, I take issue with these images.

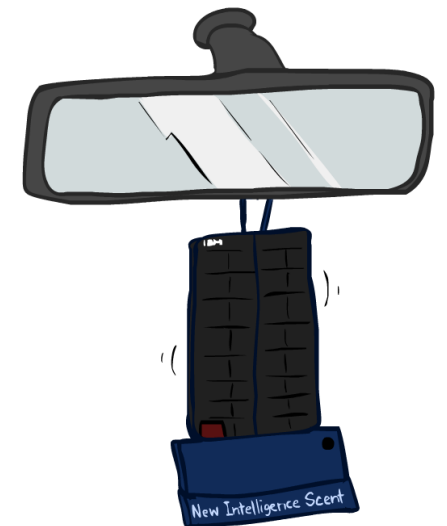


Game Agents

- Classic Moment: May, '97: Deep Blue vs. Kasparov
 - First match won against world champion
 - “Intelligent creative” play
 - 200 million board positions per second
 - Humans understood 99.9 of Deep Blue's moves
 - Can do about the same now with a PC cluster

- 1996: Kasparov Beats Deep Blue
“I could feel --- I could smell --- a new kind of intelligence across the table.”

- 1997: Deep Blue Beats Kasparov
“Deep Blue hasn't proven anything.”



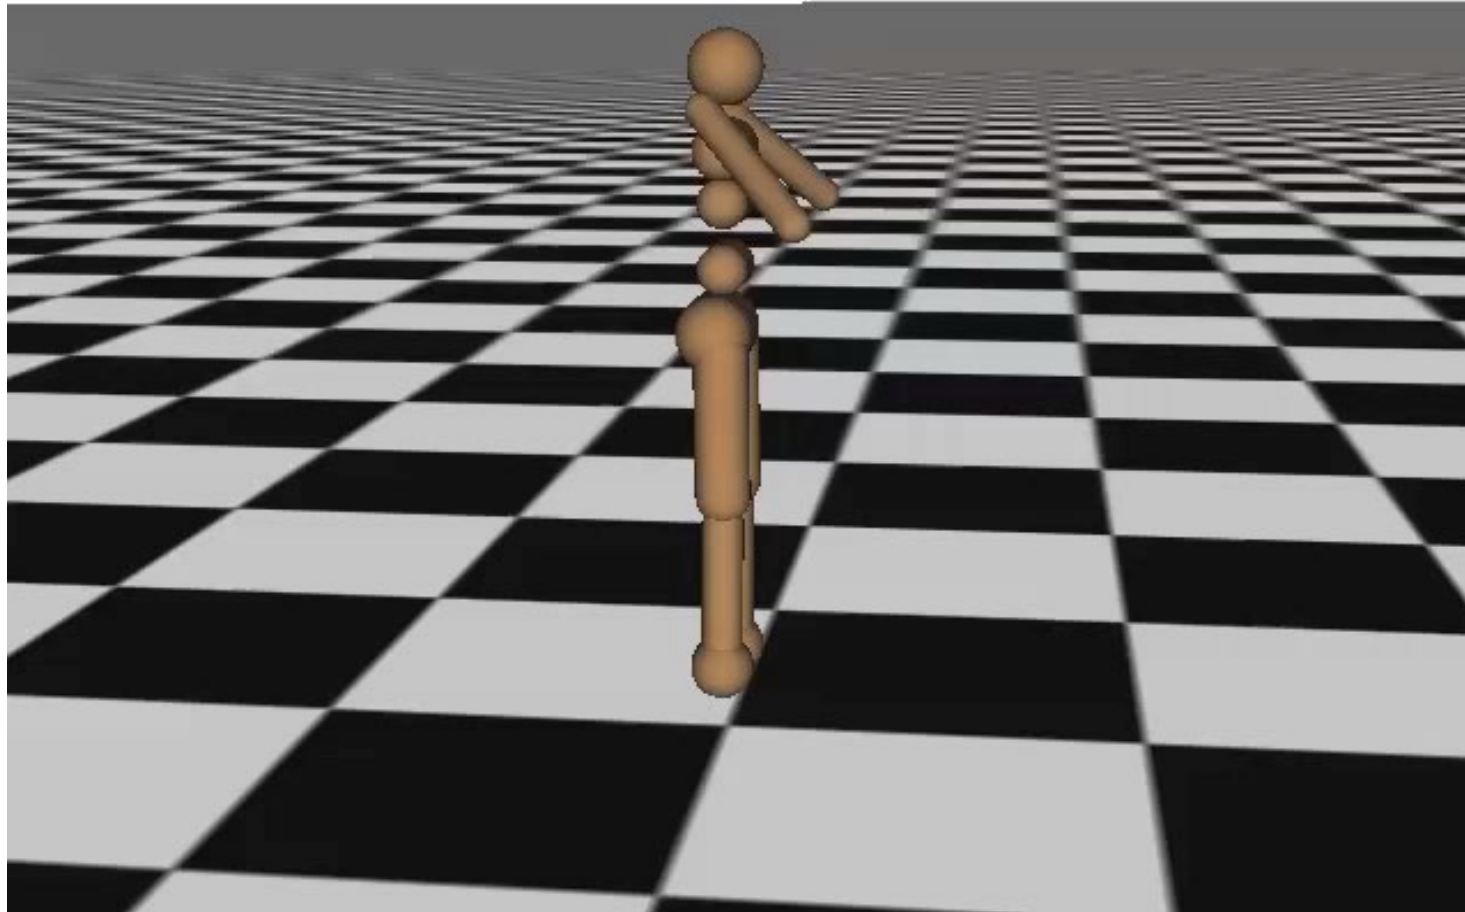


Alpha Go



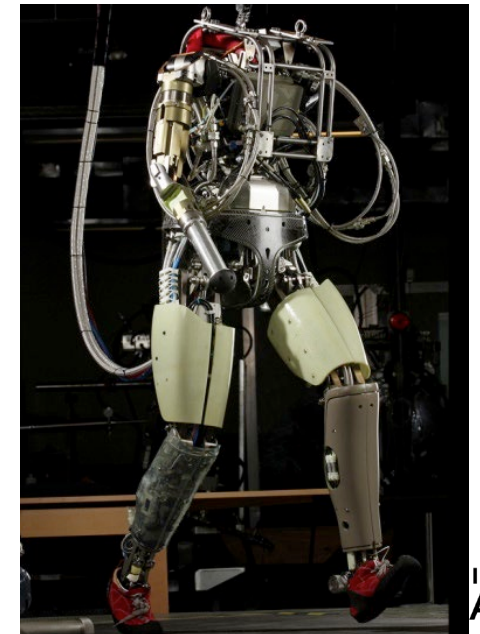
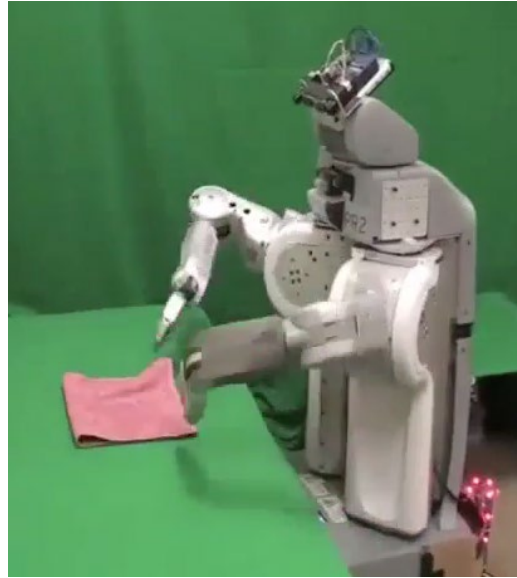
Simulated Agents

Iteration 0



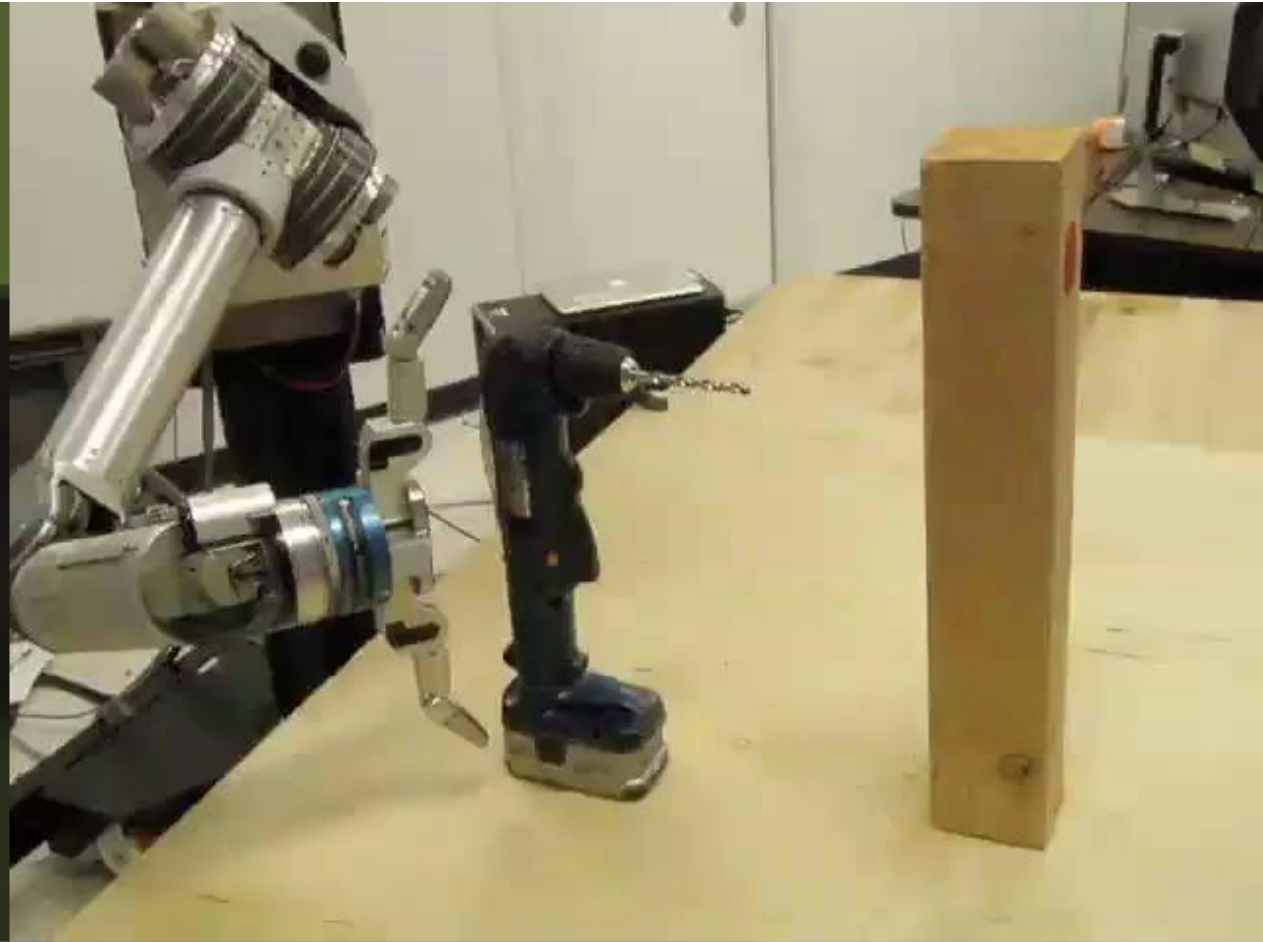
Robotics

- Robotics
 - Part mech. eng.
 - Part AI
 - Reality much harder than simulations!
- Technologies
 - Vehicles
 - Rescue
 - Help in the home
 - Lots of automation...
- In this class:
 - We ignore mechanical aspects
 - Methods for planning
 - Methods for control



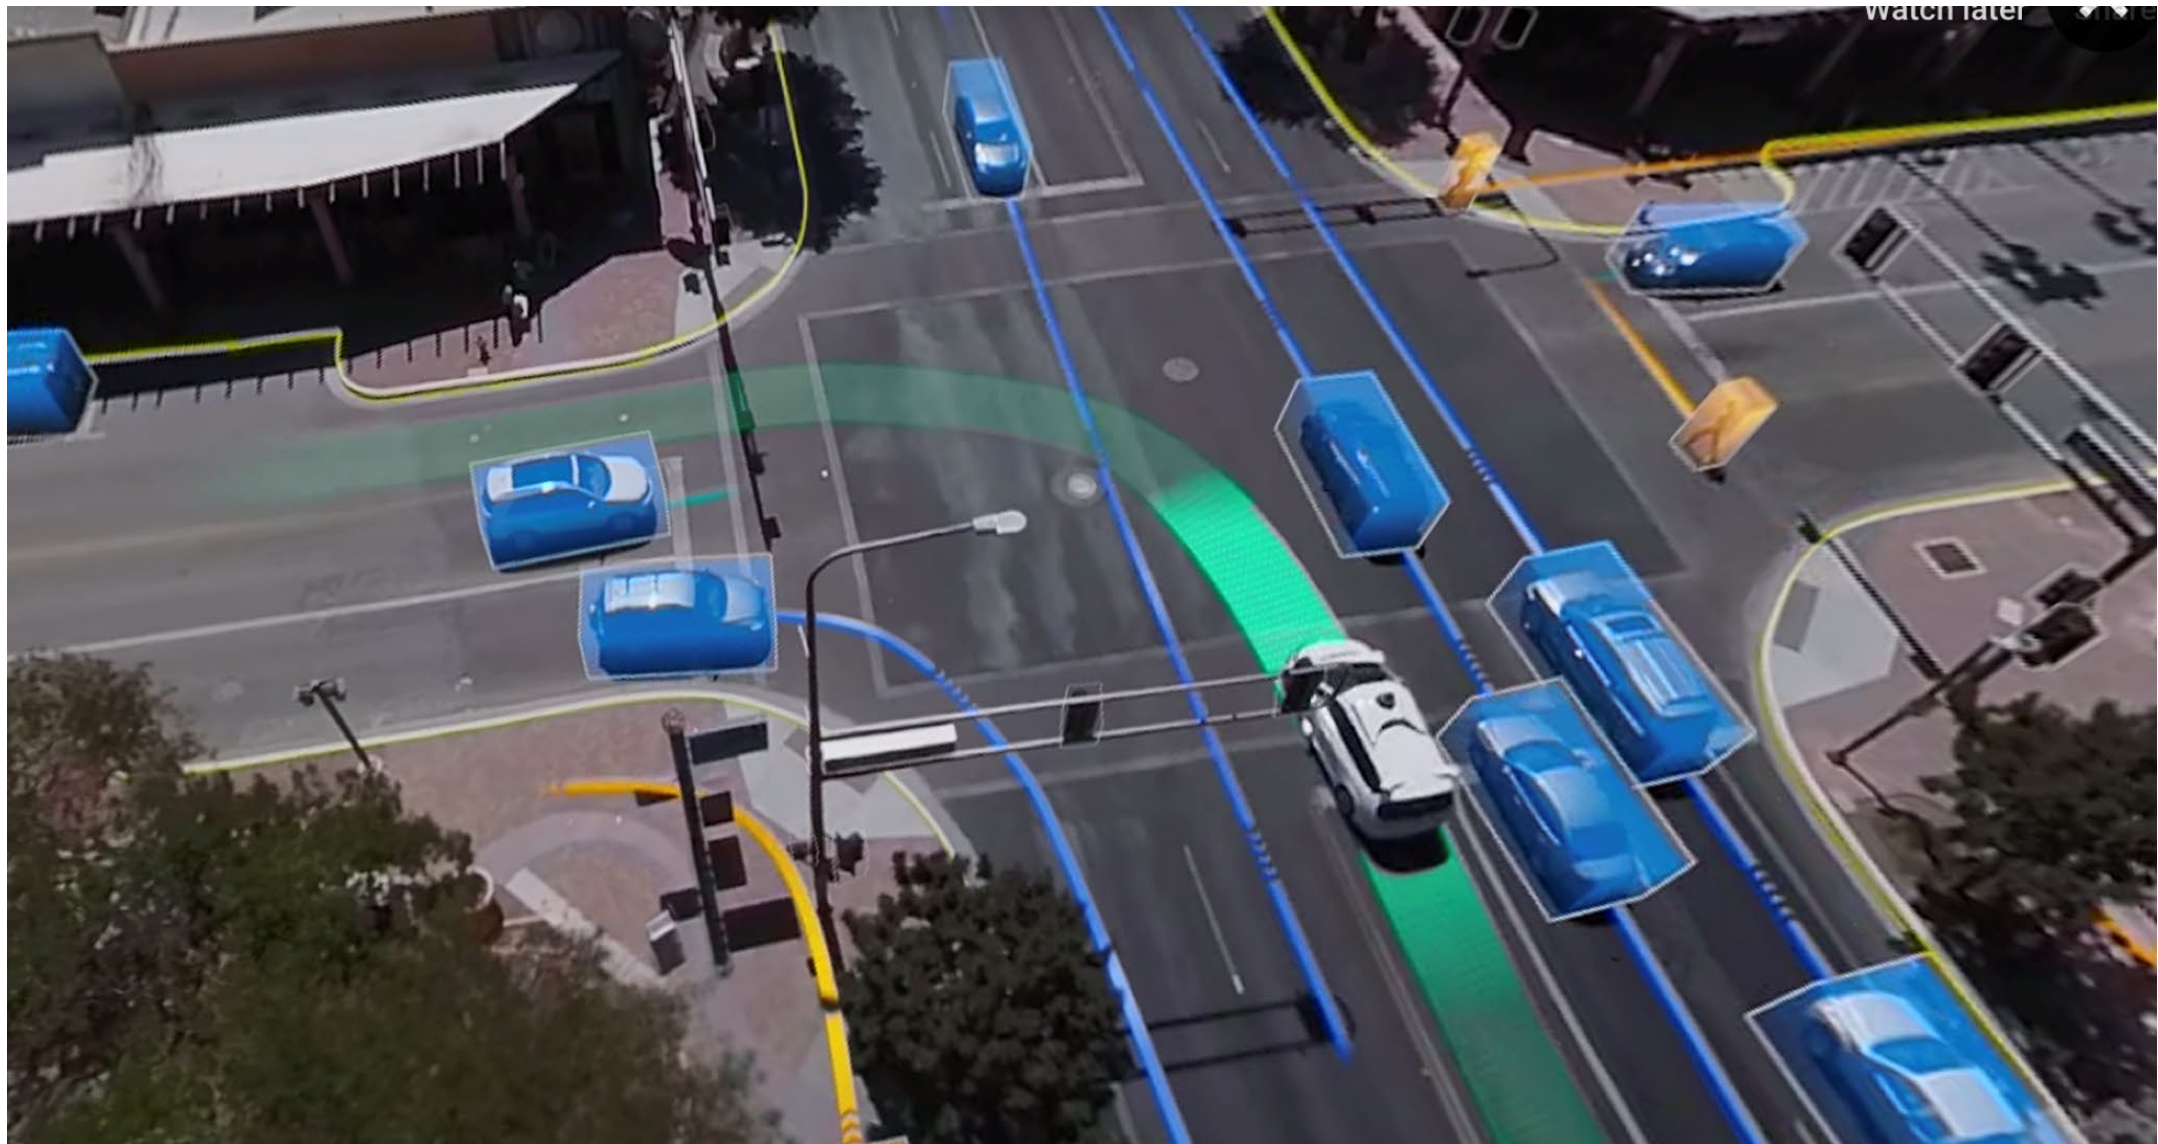
Images from UC Berkeley, Boston Dynamics, RoboCup, Google

Robots

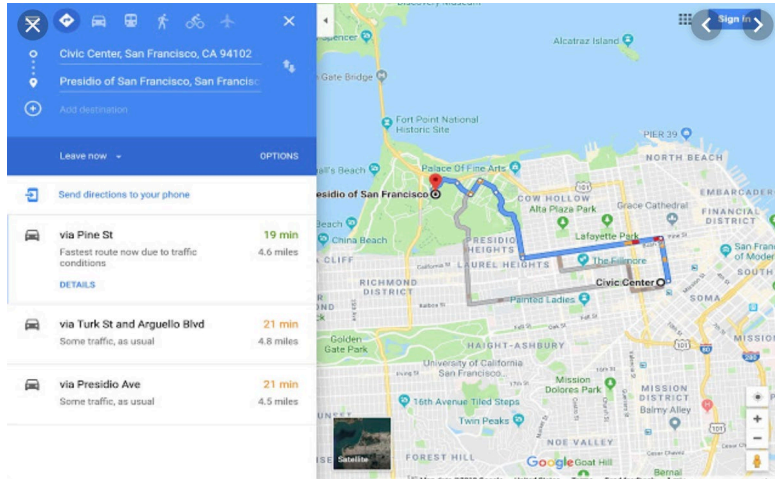


Human-AI Interaction



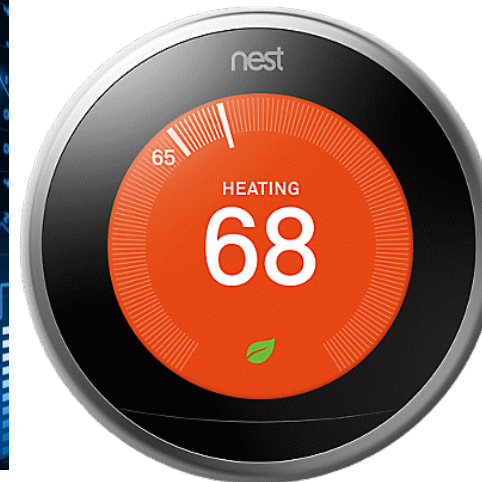
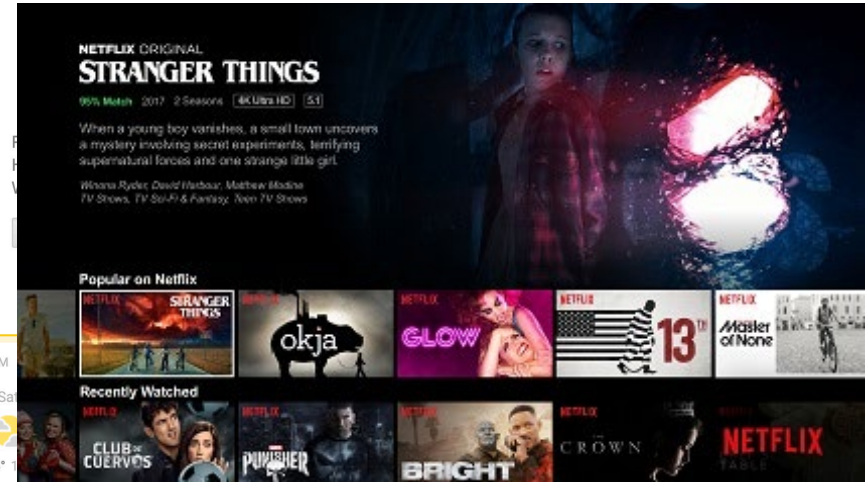
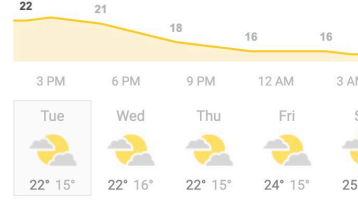


Tools for Predictions & Decisions



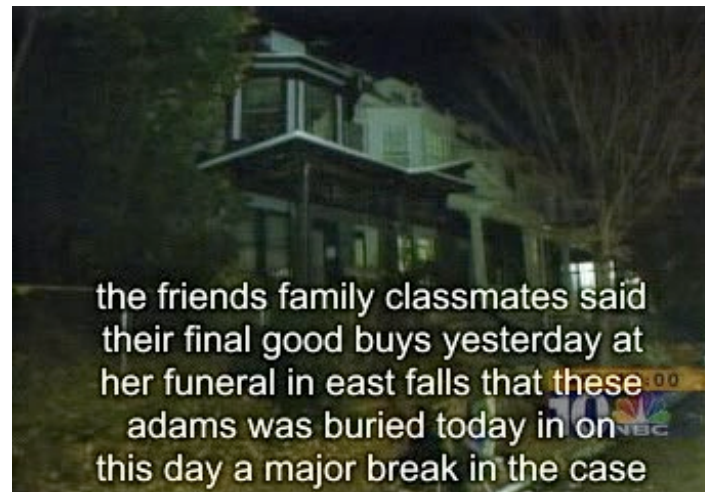
Berkeley, CA 94709
Tuesday 2:00 PM
Mostly Sunny

22°C | °F



Natural Language

- Speech technologies (e.g. Siri)
 - Automatic speech recognition (ASR)
 - Text-to-speech synthesis (TTS)
 - Dialog systems
- Language processing technologies
 - Question answering



"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine. Le Chine sur ce sujet



"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."

Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard

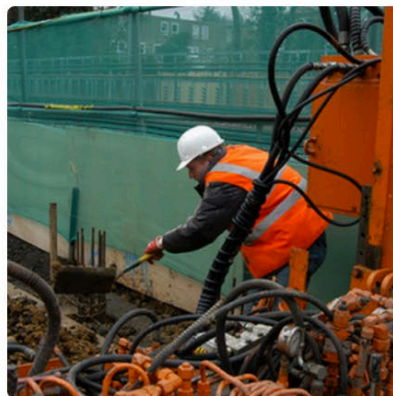


- Web search
- Text classification, spam filtering, etc...

Computer Vision



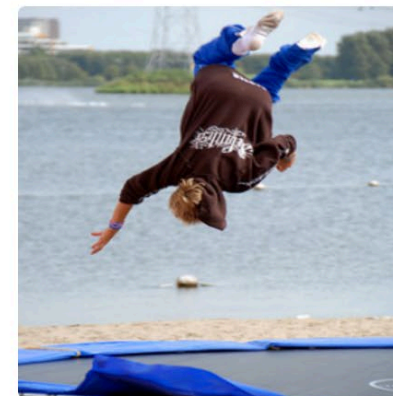
"man in black shirt is playing guitar."



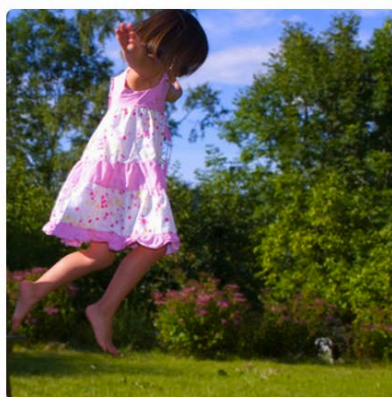
"construction worker in orange safety vest is working on road."



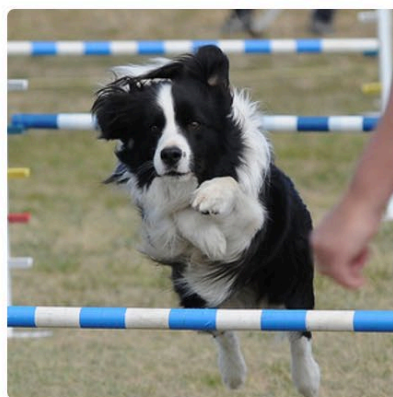
"two young girls are playing with lego toy."



"boy is doing backflip on wakeboard."



"girl in pink dress is jumping in air."



"black and white dog jumps over bar."



"young girl in pink shirt is swinging on swing."



"man in blue wetsuit is surfing on wave."

Goals of this course

- Introduction of basic concepts and techniques
- Understanding the need to be precise and learn how to be precise
- Understanding basic problems and limitations
- Knowing advantages and disadvantages of a particular concept or technique (in relation to others)
- Allows you a critical view on what people consider AI to be able to do

Onward to ... Knowledge Processing

Jonathan Hudson
jwhudson@ucalgary.ca
<https://pages.cpsc.ucalgary.ca/~jwhudson/>



UNIVERSITY OF
CALGARY